

J-series[™] Services Router

Advanced WAN Access Configuration Guide

Release 8.5

Juniper Networks, Inc.

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

www.juniper.net

Part Number: 530-021978-01, Revision 1

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright @ 1996, 1997, Maker Communications, Inc.

Juniper Networks, the Juniper Networks logo, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOS and JUNOSe are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

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

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

J-series[™] Services Router Advanced WAN Access Configuration Guide Release 8.5 Copyright © 2007, Juniper Networks, Inc. All rights reserved. Printed in USA.

Writing: Nidhi Bhargava, Michael Bushong, Maya Devi, Taffy Everts, Walter Goralski, Joshua Kim, Jerry Isaac, Archana Maheshwari, Hareesh Kumar Kozhippurath Narayana Panicker, Laura Phillips, Frank Reade, Hariharan I.S, Selvakumar T. S., Alan Twhigg, and Aiswarya J.Y. Editing: Taffy Everts and Stella Hackell Illustration: Faith Bradford Brown and Nathaniel Woodward Cover Design: Edmonds Design

Revision History 12 October 2007—Revision 1.

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

YEAR 2000 NOTICE

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

SOFTWARE LICENSE

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

End User License Agreement

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Abbreviated Table of Contents

		About This Guide	xvii
Part 1		Configuring Private Communications over Public with MPLS	: Networks
	Chapter 1	Multiprotocol Label Switching Overview	3
	Chapter 2	Configuring Signaling Protocols for Traffic Engineering	19
	Chapter 3	Configuring Virtual Private Networks	31
	Chapter 4	Configuring CLNS VPNs	55
	Chapter 5	Configuring IPSec for Secure Packet Exchange	67
Part 2		Managing Multicast Transmissions	
	Chapter 6	Multicast Overview	101
	Chapter 7	Configuring a Multicast Network	109
Part 3		Configuring DLSw Services	
	Chapter 8	Configuring Data Link Switching	125
Part 4		Configuring a Policy Framework	
	Chapter 9	Policy Framework Overview	149
	Chapter 10	Configuring Routing Policies	169
	Chapter 11	Configuring NAT	185
	Chapter 12	Configuring Stateful Firewall Filters and NAT	205
	Chapter 13	Configuring Stateless Firewall Filters	221
Part 5		Configuring Class of Service	
	Chapter 14	Class-of-Service Overview	263
	Chapter 15	Configuring Class of Service	283
Part 6		Index	
		Index	349

J-series[™] Services Router Advanced WAN Access Configuration Guide

Table of Contents

About This Guide

xvii

Audience How to Use This Guide Document Conventions Related Juniper Networks Documentation Documentation Feedback Requesting Support	xvii xviii xviii xix xx xx xxiii xxiii
Configuring Private Communications over Pub with MPLS	lic Networks
Multiprotocol Label Switching Overview	3
MPLS and VPN Terms	3
MPLS Overview	5
Label Switching	6
Label-Switched Paths	6
Label-Switching Routers	7
Labels	8
Label Operations	8
Penultimate Hop Popping	9
LSP Establishment	9
Static LSPs	9
Dynamic LSPs	9
Traffic Engineering with MPLS	9
Signaling Protocols Overview	
Label Distribution Protocol	
LDP Operation	
LDP Messages	
Resource Reservation Protocol	
RSVP Fundamentals	
Bandwidth Reservation Requirement	
Explicit Route Objects	
Constrained Shortest Path First	
Link Coloring	
VPN Overview	
VPN Components	
VPN Routing Requirements	
	Audience How to Use This Guide Document Conventions Related Juniper Networks Documentation Documentation Feedback Requesting Support Configuring Private Communications over Pubwith MPLS Multiprotocol Label Switching Overview MPLS and VPN Terms MPLS Overview Label Switching Label-Switched Paths Label-Switched Paths Labels Label Operations Penultimate Hop Popping LSP Establishment Static LSPs Dynamic LSPs Signaling Protocol Overview Label Distribution Protocol LDP Operation LDP Messages Resource Reservation Requirement Explicit Route Objects Constrained Shortest Path First Link Coloring VPN Overview VPN Routing Requirements

Objectivesxvii

VPN Routing Information	16
VRF Instances	16
Route Distinguishers	16
Route Targets to Control the VRF Table	16
Types of VPNs	
Layer 2 VPNs	
Layer 2 Circuits	
Layer 3 VPNs	

Configuring Signaling Protocols for Traffic Engineering

19

Signaling Protocol Overview	19
LDP Signaling Protocol	20
RSVP Signaling Protocol	20
Before You Begin	20
Configuring LDP and RSVP with a Configuration Editor	20
Configuring LDP-Signaled LSPs	
Configuring RSVP-Signaled LSPs	23
Verifying an MPLS Configuration	25
Verifying an LDP-Signaled LSP	25
Verifying LDP Neighbors	25
Verifying LDP Sessions	26
Verifying the Presence of LDP-Signaled LSPs	27
Verifying Traffic Forwarding over the LDP-Signaled LSP	27
Verifying an RSVP-Signaled LSP	27
Verifying RSVP Neighbors	28
Verifying RSVP Sessions	28
Verifying the Presence of RSVP-Signaled LSPs	29

Chapter 3

Chapter 2

Configuring Virtual Private Networks

VPN Configuration Overview	
Sample VPN Topology	
Basic Layer 2 VPN Configuration	
Basic Layer 2 Circuit Configuration	
Basic Layer 3 VPN Configuration	
Before You Begin	
Configuring VPNs with a Configuration Editor	
Configuring Interfaces Participating in a VPN	35
Configuring Protocols Used by a VPN	
Configuring MPLS for VPNs	
Configuring a BGP Session	
Configuring Routing Options for VPNs	40
Configuring an IGP and a Signaling Protocol	41
Configuring LDP for Signaling	41
Configuring RSVP for Signaling	43
Configuring a Layer 2 Circuit	44

	Configuring a VPN Routing Instance	45
	Configuring a VPN Routing Policy	47
	Configuring a Routing Policy for Layer 2 VPNs	48
	Configuring a Routing Policy for Layer 3 VPNs	51
	Verifying a VPN Configuration	
	Pinging a Layer 2 VPN	53
	Pinging a Layer 3 VPN	53
	Pinging a Layer 2 Circuit	53
Chapter 4	Configuring CLNS VPNs	55
	CLNC Torme	FF
	CLNS Terms	
	CLINS OVERVIEW	
	Before You Begin	
	Configuring CLNS with a Configuration Editor	
	Configuring a VPN Routing Instance (Required)	
	Configuring ES-IS	
	Configuring IS-IS for CLNS	60
	Configuring CLNS Static Routes	62
	Configuring BGP for CLNS	63
	Verifying CLNS VPN Configuration	63
	Displaying CLNS VPN Configuration	63
Chapter 5	Configuring IPSec for Secure Packet Exchange	67
	IPSec Terms	67
	IPSec Overview	69
	Authentication and Encryption Algorithms in IPSec	69
	Authentication Methods in IPSec	70
	Preshared Keys	70
	Digital Certificates	70
	Certificate Revocation Lists (CRLs)	71
	Traffic Protection in IPSec	71
	Security Associations	72
	Dynamic Security Associations and IKE Protocol	72
	IPSec Modes	73
	Before You Begin	73
	Configuring an IPSec Tunnel with Quick Configuration	73
	Configuring IPSec with a Configuration Editor	75
	Configuring IPSec Manual Security Associations	76
	Configuring IPSec Dynamic Security Associations	77
	Configuring an IKE Proposal	
	Configuring an IKE Policy	
	Configuring an IPSec Proposal	81
	Configuring an IPSec Policy	82
	Configuring IPSec Rules	
	Configuring IPSec Services Interfaces	
	Configuring Service Sets	

Configuring a NAT Pool	90
Configuring Digital Certificates for IPSec Tunnels	91
Configuring a CA Profile with a Configuration Editor	92
Requesting a CA Certificate from a CA	93
Generating a Public and Private Key Pair	94
Generating and Enrolling a Local Digital Certificate	94
Loading a Digital Certificate on a Services Router	95
Applying the Local Digital Certificate to an IPSec Tunnel	96
Deleting a Digital Certificate	97
Verifying the IPSec Tunnel Configuration	98
Verifying IPSec Tunnel Statistics	98

Part 2 Managing Multicast Transmissions

Chapter 6	Multicast Overview	101
	Multicast Terms	101
	Multicast Architecture	103
	Upstream and Downstream Interfaces	103
	Subnetwork Leaves and Branches	104
	Multicast IP Address Ranges	104
	Notation for Multicast Forwarding States	105
	Dense and Sparse Routing Modes	105
	Strategies for Preventing Routing Loops	105
	Reverse-Path Forwarding for Loop Prevention	105
	Shortest-Path Tree for Loop Prevention	106
	Administrative Scoping for Loop Prevention	106
	Multicast Protocol Building Blocks	106
Chapter 7	Configuring a Multicast Network	109
	· · · · · · · · · · · · · · · · · · ·	
	Before You Begin	109
	Configuring a Multicast Network with a Configuration Editor	110
	Configuring SAP and SDP (Optional)	110
	Configuring IGMP (Required)	111
	Configuring the PIM Static RP (Optional)	112
	Filtering PIM Register Messages from Unauthorized Groups and Sou	rces
	(Optional)	114
	Rejecting Incoming PIM Register Messages on an RP Router	115
	Stopping Outgoing PIM Register Messages on a Designated	
	Router	116
	Configuring a PIM RPF Routing Table (Optional)	117
	Verifying a Multicast Configuration	119
	Verifying SAP and SDP Addresses and Ports	119
	Verifying the IGMP Version	119
	Verifying the PIM Mode and Interface Configuration	120
	Verifying the PIM RP Configuration	120
	Verifying the RPF Routing Table Configuration	121

Part 3 Configuring DLSw Services

Chapter 8	Configuring Data Link Switching	125
	DLSw Terms	126
	DI Sw Overview	127
	Switch-to-Switch Protocol for DLSw	127
	DLSw Operational Stages	128
	DLSw Capabilities Exchange	
	DLSw Circuits Establishment	
	Class of Service for DLSw	
	DLSw Ethernet Redundancy	129
	DLSw Peer Preference and Load Balancing	129
	Before You Begin	129
	Configuring DLSw with Quick Configuration	129
	Configuring DLSw with a Configuration Editor	131
	Configuring Basic DLSw (Required)	131
	Configuring LLC Type 2 Properties on an Ethernet Interface	132
	Configuring DLSw on the Local Services Router	132
	Configuring DLSw on the Remote Services Router	134
	Configuring CoS for DLSw (Optional)	134
	Configuring DLSw Ethernet Redundancy (Optional)	136
	Configuring DLSw Peer Preference and Load Balancing (Optional)	139
	Clearing the DLSw Reachability Cache	141
	Verifying DLSw Configuration	142
	Displaying LLC Type 2 Properties on a Fast Ethernet Interface	142
	Displaying DLSw Capabilities	142
	Displaying DLSw Circuit State	143
	Displaying Details of a DLSw Circuit State	143
	Displaying DLSw Peers	144
	Displaying Details of DLSw Peers	144
	Displaying DLSw Reachability Information	145
	Displaying DLSw Ethernet Redundancy Properties	146
	Displaying DLSw Ethernet Redundancy Statistics	146

Part 4 Configuring a Policy Framework

Cha	pter	9
-----	------	---

Policy Framework Overview

Policy Framework Terms 149 Routing Policies 151 Routing Policy Overview 151 Routing Policy Terms 151 Default and Final Actions 151 Applying Routing Policies 151 Routing Policy Match Conditions 152 Routing Policy Actions 153		
Routing Policies 151 Routing Policy Overview 151 Routing Policy Terms 151 Default and Final Actions 151 Applying Routing Policies 151 Routing Policy Match Conditions 152 Routing Policy Actions 153	Policy Framework Terms	149
Routing Policy Overview151Routing Policy Terms151Default and Final Actions151Applying Routing Policies151Routing Policy Match Conditions152Routing Policy Actions153	Routing Policies	151
Routing Policy Terms151Default and Final Actions151Applying Routing Policies151Routing Policy Match Conditions152Routing Policy Actions153	Routing Policy Overview	151
Default and Final Actions	Routing Policy Terms	151
Applying Routing Policies	Default and Final Actions	151
Routing Policy Match Conditions	Applying Routing Policies	151
Routing Policy Actions153	Routing Policy Match Conditions	
	Routing Policy Actions	

Stateful Firewall Filters	
Stateful Firewall Filter Overview	
Stateful Firewall Filter Match Conditions	
Stateful Firewall Filter Actions	
Stateless Firewall Filters	
Stateless Firewall Filter Overview	
Stateless Firewall Filter Terms	
Chained Stateless Firewall Filters	
Planning a Stateless Firewall Filter	
Stateless Firewall Filter Match Conditions	
Stateless Firewall Filter Actions and Action Modifiers	
Network Address Translation	163
NAT Overview	
Source Static NAT	
Source Dynamic NAT with NAPT	
Source Dynamic NAT Without NAPT	
Destination Static NAT	165
Full-Cone NAT (Bidirectional NAT)	165
NAT Components	166
NAT Pools	
NAT Rules	

Chapter 10

Configuring Routing Policies

169

Before You Begin	169
Configuring a Routing Policy with a Configuration Editor	170
Configuring the Policy Name (Required)	170
Configuring a Policy Term (Required)	171
Rejecting Known Invalid Routes (Optional)	172
Injecting OSPF Routes into the BGP Routing Table (Optional)	174
Grouping Source and Destination Prefixes in a Forwarding Class	
(Optional)	176
Configuring a Policy to Prepend the AS Path (Optional)	177
Configuring Damping Parameters (Optional)	179

Chapter 11

Configuring NAT

Before You Begin	185
Configuring NAT with a Configuration Editor	
Configuring Basic Source Static NAT	
Statically Assigning NAT Addresses from a I	Dynamic Pool187
Configuring Full-Cone NAT	
Configuring NAT Rules Without Defining Poo	ols192
Defining an Overload Pool or an Overload P	Prefix193
Defining Rules for Transparent NAT	
Applying NAT to an Interface	
Verifying NAT Configuration	
Displaying NAT Configurations	
Verifying NAT	

Chapter 12	Configuring Stateful Firewall Filters and NAT 205		
	Before You Begin	205	
	Configuring a Stateful Firewall Filter with Ouick Configuration	206	
	Configuring a Stateful Firewall Filter with a Configuration Editor	211	
	Verifying Stateful Firewall Filter Configuration	217	
	Displaying Stateful Firewall Filter Configurations	217	
	Verifying a Stateful Firewall Filter	217	
Chapter 13	Configuring Stateless Firewall Filters	221	
	Before You Begin	221	
	Configuring a Stateless Firewall Filter with Ouick Configuration	2221	
	Configuring IPv4 and IPv6 Stateless Firewall Filters	222	
	Assigning IPV4 and IPV6 Firewall Filters to Interfaces	222	
	Assigning IFV4 and IFV0 Filewall Filter with a Configuration Editor	200	
	Configuring a Stateless Filewall Filter with a Configuration Eultor	208	
	Stateless Firewall Filter Strategies	238	
	Strategy for a Typical Stateless Firewall Filter	238	
	Strategy for Handling Packet Fragments	238	
	Configuring a Routing Engine Firewall Filter for Services and Protoc from Trusted Sources	ols 238	
	Configuring a Routing Engine Firewall Filter to Protect Against TCP a	and	
	ICMP Floods	241	
	Configuring a Routing Engine Firewall Filter to Handle Fragments .	246	
	Applying a Stateless Firewall Filter to an Interface	251	
	Verifying Stateless Firewall Filter Configuration	252	
	Displaying Stateless Firewall Filter Configurations	252	
	Displaying Stateless Firewall Filter Logs	255	
	Displaying Firewall Filter Statistics	256	
	Verifying a Services, Protocols, and Trusted Sources Firewall Filter	257	
	Verifying a TCP and ICMP Flood Firewall Filter	258	
	Verifying a Firewall Filter That Handles Fragments	259	
Part 5	Configuring Class of Service		
Chapter 14	Class-of-Service Overview	263	
		o / -	
	CoS Terms	263	
	Benefits of CoS	264	
	CoS Across the Network	265	
	JUNOS CoS Components	266	
	Code-Point Aliases	266	
	Classifiers	266	
	Behavior Aggregate Classifiers	266	
	Multifield Classifiers	267	
	Forwarding Classes		
	Loss Priorities	267	
	Forwarding Policy Options	267	
		201	

Transmission Queues	268
Schedulers	
Transmit Rate	
Delay Buffer Size	
Scheduling Priority	
Shaping Rate	
RED Drop Profiles	
Virtual Channels	
Policers for Traffic Classes	
Rewrite Rules	
How CoS Components Work	
CoS Process on Incoming Packets	
CoS Process on Outgoing Packets	
Default CoS Settings	
Default CoS Values and Aliases	273
Forwarding Class Queue Assignments	
Scheduler Settings	
Default Behavior Aggregate Classifiers	
CoS Value Rewrites	
Sample Behavior Aggregate Classification	
Transmission Scheduling on J-series Services Routers	

Chapter 15

Configuring Class of Service

Before You Begin	283
Configuring CoS with Quick Configuration	284
Defining CoS Components	284
Defining CoS Value Aliases	286
Defining Forwarding Classes	288
Defining Classifiers	290
Defining Rewrite Rules	292
Defining Schedulers	294
Defining Virtual Channel Groups	300
Assigning CoS Components to Interfaces	302
Configuring CoS Components with a Configuration Editor	305
Configuring a Policer for a Firewall Filter	306
Configuring and Applying a Firewall Filter for a Multifield Classifier	307
Assigning Forwarding Classes to Output Queues	310
Configuring and Applying Rewrite Rules	312
Configuring and Applying Behavior Aggregate Classifiers	315
Configuring RED Drop Profiles for Congestion Control	319
Configuring Schedulers	321
Configuring and Applying Scheduler Maps	324
Configuring and Applying Virtual Channels	327
Configuring and Applying Adaptive Shaping for Frame Relay	331
Configuring Strict High Priority for Queuing with a Configuration Editor	332
Configuring Large Delay Buffers with a Configuration Editor	340
Maximum Delay Buffer Sizes Available to Interfaces	340
Delay Buffer Size Allocation Methods	341

Specifying Delay Buffer Sizes for Queues	342
Configuring a Large Delay Buffer on a Channelized T1 interface	343
Verifying a CoS Configuration	345
Verifying Multicast Session Announcements	345
Verifying a Virtual Channel Configuration	345
Verifying a Virtual Channel Group Configuration	
Verifying an Adaptive Shaper Configuration	

Part 6

Index

|--|

J-series[™] Services Router Advanced WAN Access Configuration Guide

About This Guide

This preface provides the following guidelines for using the *J*-series[™] Services Router Advanced WAN Access Configuration Guide:

- Objectives on page xvii
- Audience on page xvii
- How to Use This Guide on page xviii
- Document Conventions on page xix
- Related Juniper Networks Documentation on page xx
- Documentation Feedback on page xxiii
- Requesting Support on page xxiii

Objectives

This guide contains instructions for configuring Services Routers in virtual private networks (VPNs) and multicast networks, configure data link switching (DLSw) services, and apply routing techniques such as policies, firewall filters, IP Security (IPSec), and class-of-service (CoS) classification for safe, efficient routing.

J-series Services Router operations are controlled by the JUNOS software. You direct the JUNOS software through either a Web browser or a command-line interface (CLI).



NOTE: This guide documents Release 8.5 of the JUNOS software. For additional information about J-series Services Routers—either corrections to or omissions from this guide—see the *J-series Services Router Release Notes* at http://www.juniper.net.

Audience

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

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

 Network administrators who install, configure, and manage products of Juniper Networks

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

How to Use This Guide

J-series documentation explains how to install, configure, and manage J-series routers by providing information about JUNOS implementation specifically on J-series routers. (For comprehensive JUNOS information, see the JUNOS software manuals listed in "Related Juniper Networks Documentation" on page xx.) Table 1 on page xviii shows the location of J-series information, by task type, in Juniper Networks documentation.

Table 1: Location of J-series Information

J-series Tasks	Location of Instruction
Installing hardware and establishing basic connectivity	Getting Started Guide for your router
Configuring interfaces and routing protocols such as RIP, OSPF, BGP, and IS-IS	J-series Services Router Basic LAN and WAN Access Configuration Guide
Configuring advanced features such as virtual private networks (VPNs), IP Security (IPSec), multicast, routing policies, firewall filters, and class of service (CoS)	J-series Services Router Advanced WAN Access Configuration Guide
Managing users and operations, monitoring performance, upgrading software, and diagnosing common problems	J-series Services Router Administration Guide
Using the J-Web interface	J-Web Interface User Guide
Using the CLI	JUNOS CLI User Guide

Typically, J-series documentation provides both general and specific information—for example, a configuration overview, configuration examples, and verification methods. Because you can configure and manage J-series routers in several ways, you can choose from multiple sets of instructions to perform a task. To make best use of this information:

- If you are new to the topic—Read through the initial overview information, keep the related JUNOS guide handy for details about the JUNOS hierarchy, and follow the step-by-step instructions for your preferred interface.
- If you are already familiar with the feature—Go directly to the instructions for the interface of your choice, and follow the instructions. You can choose a J-Web method, the JUNOS CLI, or a combination of methods based on the level of complexity or your familiarity with the interface.

For many J-series features, you can use J-Web Quick Configuration pages to configure the router quickly and easily without configuring each statement individually. For

more extensive configuration, use the J-Web configuration editor or CLI configuration mode commands.

To monitor, diagnose, and manage a router, use the J-Web interface or CLI operational mode commands.

Document Conventions

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

Table 2: Notice Icons

lcon	Meaning	Description
(F	Informational note	Indicates important features or instructions.
<u>!</u>	Caution	Indicates a situation that might result in loss of data or hardware damage.
4	Warning	Alerts you to the risk of personal injury or death.
*	Laser warning	Alerts you to the risk of personal injury from a laser.

Table 3 on page xix defines the text and syntax conventions used in this guide.

Table 3: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command:
		user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
Italic text like this	 Introduces important new terms. Identifies book names. Identifies RFC and Internet draft titles. 	 A policy <i>term</i> is a named structure that defines match conditions and actions. <i>JUNOS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i>

Table 3: Text and Syntax Conventions (continued)

Convention	Description	Examples
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name domain-name
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	 To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric metric="">;</default-metric>
(pipe symbol)	Indicates a choice between the mutually	broadcast multicast
	side of the symbol. The set of choices is often enclosed in parentheses for clarity.	(string1 string2 string3)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [community-ids]
Indention and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options {
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	<pre>static { route default { nexthop address; retain; } } }</pre>
J-Web GUI Conventions		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	 In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select Protocols > Ospf .

Related Juniper Networks Documentation

J-series Services Routers are documented in multiple guides. Although the J-series guides provide instructions for configuring and managing a Services Router with the JUNOS CLI, they are not a comprehensive JUNOS software resource. For complete

documentation of the statements and commands described in J-series guides, see the JUNOS software manuals listed in Table 4 on page xxi.

Chapter in a J-series Guide	Corresponding JUNOS Software Manual
Getting Started Guide for Your Router	
"Services Router User Interface Overview"	■ JUNOS CLI User Guide
"Establishing Basic Connectivity"	■ JUNOS System Basics Configuration Guide
J-series Services Router Basic LAN and WAN Access Co	nfiguration Guide
"Using Services Router Configuration Tools"	 JUNOS CLI User Guide JUNOS System Basics Configuration Guide
"Interfaces Overview"	■ JUNOS Network Interfaces Configuration Guide
"Configuring DS1, DS3, Ethernet, and Serial Interfaces"	■ JUNOS Interfaces Command Reference
"Configuring Channelized T1/E1/ISDN PRI Interfaces"	
"Configuring Digital Subscriber Line Interfaces	
"Configuring Point-to-Point Protocol over Ethernet"	
"Configuring ISDN"	
"Configuring Link Services Interfaces"	 JUNOS Services Interfaces Configuration Guide JUNOS System Basics and Services Command Reference
"Configuring VoIP"	 JUNOS Network Interfaces Configuration Guide JUNOS Interfaces Command Reference
"Configuring uPIMs as Ethernet Switches"	 JUNOS Network Interfaces Configuration Guide JUNOS System Basics Configuration Guide JUNOS System Basics and Services Command Reference
"Routing Overview"	■ JUNOS Routing Protocols Configuration Guide
"Configuring Static Routes"	■ JUNOS Routing Protocols and Policies Command Reference
"Configuring a RIP Network"	
"Configuring an OSPF Network"	
"Configuring the IS-IS Protocol"	
"Configuring BGP Sessions"	

Table 4: J-series Guides and Related JUNOS Software Publications

	Table ⁴	4: J-series	Guides	and Rela	ated JUNOS	Software	Publications	(continued)
--	--------------------	-------------	--------	----------	------------	-----------------	--------------	-------------

Chapter in a J-series Guide	Corresponding JUNOS Software Manual		
"Multiprotocol Label Switching Overview"	 JUNOS MPLS Applications Configuration Guide UNOS Pouting Protocols and Policies Command Reference 		
"Configuring Signaling Protocols for Traffic Engineering"	 JUNOS Kouting Protocols and Policies Command Reference JUNOS VPNs Configuration Guide 		
"Configuring Virtual Private Networks"			
"Configuring CLNS VPNs"			
"Configuring IPSec for Secure Packet Exchange"	 JUNOS System Basics Configuration Guide JUNOS Services Interfaces Configuration Guide JUNOS System Basics and Services Command Reference 		
"Multicast Overview"	JUNOS Multicast Protocols Configuration Guide		
"Configuring a Multicast Network"	 JUNOS Routing Protocols and Policies Command Reference 		
"Configuring Data Link Switching"	 JUNOS Services Interfaces Configuration Guide JUNOS System Basics and Services Command Reference 		
"Policy Framework Overview"	JUNOS Policy Framework Configuration Guide		
"Configuring Routing Policies"	 JUNOS Routing Protocols and Policies Command Reference 		
"Configuring NAT"	■ JUNOS Network Interfaces Configuration Guide		
"Configuring Stateful Firewall Filters and NAT"	 JUNOS Policy Framework Configuration Guide IUNOS Services Interfaces Configuration Guide 		
"Configuring Stateless Firewall Filters"	 Secure Configuration Guide for Common Criteria and JUNOS-FIPS 		
	 JUNOS System Basics and Services Command Reference JUNOS Routing Protocols and Policies Command Reference 		
"Class-of-Service Overview"	■ JUNOS Class of Service Configuration Guide		
"Configuring Class of Service"	■ JUNOS System Basics and Services Command Reference		
J-series Services Router Administration Guide			
"Managing User Authentication and Access"	 JUNOS System Basics Configuration Guide Secure Configuration Guide for Common Criteria and JUNOS-FIPS 		
"Configuring SNMP for Network Management"	JUNOS Network Management Configuration Guide		
"Configuring the Router as a DHCP Server"	JUNOS System Basics Configuration Guide		
"Configuring Autoinstallation"			
"Automating Network Operations and Troubleshooting"	JUNOS Configuration and Diagnostic Automation Guide		

Chapter in a J-series Guide	Corresponding JUNOS Software Manual
"Monitoring the Router and Routing Operations"	 JUNOS System Basics and Services Command Reference JUNOS Interfaces Command Reference JUNOS Routing Protocols and Policies Command Reference
"Monitoring Events and Managing System Log Files"	 JUNOS System Log Messages Reference Secure Configuration Guide for Common Criteria and JUNOS-FIPS
"Configuring and Monitoring Alarms"	JUNOS System Basics Configuration Guide
"Performing Software Upgrades and Reboots"	JUNOS Software Installation and Upgrade Guide
"Using Services Router Diagnostic Tools"	 JUNOS System Basics and Services Command Reference JUNOS Interfaces Command Reference JUNOS Routing Protocols and Policies Command Reference
"Configuring Packet Capture"	JUNOS Services Interfaces Configuration Guide
"Configuring RPM Probes"	JUNOS System Basics and Services Command Reference

Table 4: J-series Guides and Related JUNOS Software Publications (continued)

Documentation Feedback

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

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

Requesting Support

For technical support, open a support case with 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).

J-series[™] Services Router Advanced WAN Access Configuration Guide

Part 1 Configuring Private Communications over Public Networks with MPLS

- Multiprotocol Label Switching Overview on page 3
- Configuring Signaling Protocols for Traffic Engineering on page 19
- Configuring Virtual Private Networks on page 31
- Configuring CLNS VPNs on page 55
- Configuring IPSec for Secure Packet Exchange on page 67

J-series[™] Services Router Advanced WAN Access Configuration Guide

Chapter 1 Multiprotocol Label Switching Overview

Multiprotocol Label Switching (MPLS) provides a framework for controlling traffic patterns across a network. The MPLS framework allows Services Routers to pass traffic through transit networks on paths that are independent of the individual routing protocols enabled throughout the network.

The MPLS framework supports traffic engineering and the creation of virtual private networks (VPNs). Traffic is engineered (controlled) primarily by the use of signaling protocols to establish label-switched paths (LSPs). VPN support includes Layer 2 and Layer 3 VPNs and Layer 2 circuits.

This chapter contains the following topics. For more information, see the *JUNOS Routing Protocols Configuration Guide*, *JUNOS MPLS Applications Configuration Guide*, and *JUNOS VPNs Configuration Guide*.

- MPLS and VPN Terms on page 3
- MPLS Overview on page 5
- Signaling Protocols Overview on page 10
- VPN Overview on page 14

MPLS and VPN Terms

To understand MPLS and VPNs, become familiar with the terms defined in Table 5 on page 3.

Table 5: MPLS and VPN Terms

Term	Definition
color	See link coloring.
Constrained Shortest Path First (CSPF)	MPLS algorithm that has been modified to include specific restrictions for calculating the shortest path across the network.
customer edge (CE) device	Services Router or switch in the customer's network that is connected to a service provider's provider edge (PE) router and participates in a Layer 3 VPN.
Explicit Route Object (ERO)	Extension to the Resource Reservation Protocol (RSVP) that allows an RSVP PATH message to traverse an explicit sequence of routers independently of conventional shortest-path IP routing.

Table 5: MPLS and VPN Terms (continued)

Term	Definition
inbound router	Entry point for a label-switched path (LSP). Each LSP must have exactly one inbound router that is different from the outbound router. Inbound routers are also known as ingress routers. See also <i>outbound router</i> .
label	In Multiprotocol Label Switching (MPLS), a 20-bit unsigned integer in the range 0 through 1,048,575, used to identify a packet traveling along a label-switched path (LSP).
Label Distribution Protocol (LDP)	Protocol for distributing labels in non-traffic-engineered applications. LDP allows Services Routers to establish label-switched paths (LSPs) through a network by mapping Network layer routing information directly to Data Link layer switched paths.
label-switched path (LSP)	Sequence of Services Routers that cooperatively perform Multiprotocol Label Switching (MPLS) operations for a packet stream. The first router in an LSP is called the inbound router, and the last router in the path is called the outbound router. An LSP is a point-to-point, half-duplex connection from the inbound router to the outbound router. (The inbound and outbound routers cannot be the same router.)
label-switching router (LSR)	Any Services Router that is part of an LSP.
Layer 2 circuit	Point-to-point Layer 2 connection transported by means of Multiprotocol Label Switching (MPLS) or another tunneling technology on a service provider's network. Multiple Layer 2 circuits can be transported over a single label-switched path (LSP) tunnel between two provider edge (PE) routers.
Layer 2 VPN	Private network service among a set of customer sites that use a service provider's existing Multiprotocol Label Switching (MPLS) and IP network. One customer's data is separated from another's by software rather than hardware. In a Layer 2 VPN, the Layer 3 routing of customer traffic occurs within the <i>customer</i> network.
Layer 3 VPN	Private network service among a set of customer sites that use a service provider's existing Multiprotocol Label Switching (MPLS) and IP network. One customer's routes and data are separated from another customer's routes and data by software rather than hardware. In a Layer 3 VPN, the Layer 3 routing of customer traffic occurs within the <i>service provider</i> network.
link coloring	In Constrained Shortest Path First (CSPF) routing, a way to group Multiprotocol Label Switching (MPLS) interfaces for CSPF path selection by assigning a color identifier and number to each administrative group.
Multiprotocol Label Switching (MPLS)	Method for engineering network traffic patterns by assigning short labels to network packets that describe how to forward the packets through the network.
multiple push	Addition by a Services Router of up to three labels to a packet as it enters a Multiprotocol Label Switching (MPLS) domain.
outbound router	Exit point for a label-switched path (LSP). Each LSP must have exactly one outbound router that is different from the inbound router. Outbound routers are also called egress routers. See also <i>inbound router</i> .
penultimate hop popping (PHP)	Using the penultimate router rather than the outbound router in a label-switched path (LSP) to remove the Multiprotocol Label Switching (MPLS) label from a packet.
penultimate router	Second-to-last Services Router in an LSP. The penultimate router is responsible for label popping when penultimate hop popping (PHP) is configured.

Term	Definition
рор	Removal by a Services Router of the top label from a packet as it exits the Multiprotocol Label Switching (MPLS) domain.
provider edge (PE) router	Services Router in the service provider network that is connected to a customer edge (CE) device and participates in a virtual private network (VPN).
provider router	Services Router in the service provider's network that does not attach to a customer edge (CE) device.
push	Addition of a label or stack of labels by a Services Router to a packet as it enters a Multiprotocol Label Switching (MPLS) domain.
Resource Reservation Protocol (RSVP)	Resource reservation setup protocol that interacts with integrated services on the Internet.
route distinguisher	A 6-byte virtual private network (VPN) identifier that is prefixed to an IPv4 address to make it unique. The new address is part of the VPN-IPv4 address family, which is a Border Gateway Protocol (BGP) extension. A route distinguisher allows you to configure private addresses within the VPN by preventing any overlap with the private addresses in other VPNs.
routing instance	Collection of routing tables, their interfaces, and the routing protocol parameters that control the information they contain.
swap	Replacement by a Services Router of a label or stack of labels on a packet as it travels through a Multiprotocol Label Switching (MPLS) domain.
swap and push	Replacement and subsequent push by a Services Router of a label or stack of labels on a packet as it travels through a Multiprotocol Label Switching (MPLS) domain.
Traffic engineering (TE)	The techniques and processes used to cause routed traffic to travel through the network on a path other than the one that would have been chosen if standard routing methods had been used.
traffic engineering database (TED)	Database populated by label-switched path (LSP) information such as the network topology, current reservable bandwidth of links, and link colors. The traffic engineering database is used to determine Constrained Shortest Path First (CSPF) path selection.
transit router	Any label-switching router (LSR) between the inbound and outbound Services Router of a label-switched path (LSP).
virtual private network (VPN)	Private data network that uses a public TCP/IP network, typically the Internet, while maintaining privacy with a tunneling protocol, encryption, and security procedures.
VPN routing and forwarding (VRF) instance	Routing instance for a Layer 3 VPN implementation that consists of one or more routing tables, a derived forwarding table, the interfaces that use the forwarding table, and the policies and routing protocols that determine what goes into the forwarding table.

Table 5: MPLS and VPN Terms (continued)

MPLS Overview

Multiprotocol Label Switching (MPLS) is a method for engineering traffic patterns by assigning short labels to network packets that describe how to forward them through

the network. MPLS is independent of routing tables or any routing protocol and can be used for unicast packets.

This overview contains the following topics:

- Label Switching on page 6
- Label-Switched Paths on page 6
- Label-Switching Routers on page 7
- Labels on page 8
- Label Operations on page 8
- Penultimate Hop Popping on page 9
- LSP Establishment on page 9
- Traffic Engineering with MPLS on page 9

Label Switching

In a traditional IP network, packets are transmitted with an IP header that includes a source and destination address. When a router receives such a packet, it examines its forwarding tables for the next-hop address associated with the packet's destination address and forwards the packet to the next-hop location.

In an MPLS network, each packet is encapsulated with an MPLS header. When a router receives the packet, it copies the header as an index into a separate MPLS forwarding table. The MPLS forwarding table consists of pairs of inbound interfaces and path information. Each pair includes forwarding information that the router uses to forward the traffic and modify, when necessary, the MPLS header.

Because the MPLS forwarding table has far fewer entries than the more general forwarding table, the lookup consumes less processing time and processing power. The resultant savings in time and processing are a significant benefit for traffic that uses the network to transit between outside destinations only.

Label-Switched Paths

Label-switched paths (LSPs) are unidirectional routes through a network or autonomous system (AS). In normal IP routing, the packet has no predetermined path. Instead, each router forwards a packet to the next-hop address stored in its forwarding table, based only on the packet's destination address. Each subsequent router then forwards the packet using its own forwarding table.

In contrast, MPLS routers within an AS determine paths through a network through the exchange of MPLS traffic engineering information. Using these paths, the routers direct traffic through the network along an established route. Rather than selecting the next hop along the path as in IP routing, each router is responsible for forwarding the packet to a predetermined next-hop address.

Figure 1 on page 7 shows a typical LSP topology.

Figure 1: Typical LSP Topology



In the topology shown in Figure 1 on page 7, traffic is forwarded from Host C1 to the transit network with standard IP forwarding. When the traffic enters the transit network, it is switched across a preestablished LSP through the network. In this example, an LSP might switch the traffic from Router R4 to Router R2 to Router R1. When the traffic exits the network, it is forwarded to its destination by IP routing protocols.

Label-Switching Routers

Routers that are part of the LSP are label-switching routers (LSRs). Each LSR must be configured with MPLS so that it can interpret MPLS headers and perform the MPLS operations required to pass traffic through the network. An LSP can include four types of LSRs:

- Inbound router—The only entry point for traffic into MPLS. Native IPv4 packets are encapsulated into the MPLS protocol by the inbound router. Each LSP can have only one inbound router.
- Transit router—Any router in the middle of an LSP. An individual LSP can contain between 0 and 253 transit routers. Transit routers forward MPLS traffic along the LSP, using only the MPLS header to determine how the packet is routed.
- Penultimate router—The second-to-last router in the LSP. The penultimate router in an LSP is responsible for stripping the MPLS header from the packet before forwarding it to the outbound router.
- Outbound router—The endpoint for the LSP. The outbound router receives MPLS packets from the penultimate router and performs an IP route lookup. The router then forwards the packet to the next hop of the route. Each LSP can have only one outbound router.

Labels

To forward traffic through an MPLS network, MPLS routers encapsulate packets and assign and manage headers known as labels. The routers use the labels to index the MPLS forwarding tables that determine how packets are routed through the network.

When a network's inbound router receives traffic, it inserts an MPLS label between the IP packet and the appropriate Layer 2 header for the physical link. The label contains an index value that identifies a next-hop address for the particular LSP. When the next-hop transit router receives the packet, it uses the index in the MPLS label to determine the next-hop address for the packet and forwards the packet to the next router in the LSP.

As each packet travels through the transit network, every router along the way performs a lookup on the MPLS label and forwards the packet accordingly. When the outbound router receives a packet, it examines the header to determine that it is the final router in the LSP. The outbound router then removes the MPLS header, performs a regular IP route lookup, and forwards the packet with its IP header to the next-hop address.

Label Operations

Each LSR along an LSP is responsible for examining the MPLS label, determining the LSP next hop, and performing the required label operations. LSRs can perform five label operations:

Push—Adds a new label to the top of the packet. For IPv4 packets arriving at the inbound router, the new label is the first label in the label stack. For MPLS packets with an existing label, this operation adds a label to the stack and sets the stacking bit to 0, indicating that more MPLS labels follow the first.

When it receives the packet, the inbound router performs an IP route lookup on the packet. Because the route lookup yields an LSP next hop, the inbound router performs a label push on the packet, and then forwards the packet to the LSP next hop.

• Swap—Replaces the label at the top of the label stack with a new label.

When a transit router receives the packet, it performs an MPLS forwarding table lookup. The lookup yields the LSP next hop and the path index of the link between the transit router and the next router in the LSP.

Pop—Removes the label from the top of the label stack. For IPv4 packets arriving at the penultimate router, the entire MPLS label is removed from the label stack. For MPLS packets with an existing label, this operation removes the top label from the label stack and modifies the stacking bit as necessary—sets it to 1, for example, if only a single label remains in the stack.

If multiple LSPs terminate at the same outbound router, the router performs MPLS label operations for all outbound traffic on the LSPs. To share the operations among multiple routers, most LSPs use penultimate hop popping (PHP).

 Multiple push—Adds multiple labels to the top of the label stack. This action is equivalent to performing multiple push operations. The multiple push operation is used with label stacking, which is beyond the scope of this guide.

 Swap and push—Replaces the top label with a new label and then pushes a new label to the top of the stack.

The swap and push operation is used with label stacking, which is beyond the scope of this guide.

Penultimate Hop Popping

Multiple LSPs terminating at a single outbound router load the router with MPLS label operations for all their outbound traffic. Penultimate hop popping (PHP) transfers the operation from the outbound router to penultimate routers.

With PHP, the penultimate router is responsible for popping the MPLS label and forwarding the traffic to the outbound router. The outbound router then performs an IP route lookup and forwards the traffic. For example, if four LSPs terminate at the same outbound router and each has a different penultimate router, label operations are shared across four routers.

LSP Establishment

An MPLS LSP is established by one of two methods: static LSPs and dynamic LSPs.

Static LSPs

Like a static route, a static LSP requires each router along the path to be configured explicitly. You must manually configure the path and its associated label values. Static LSPs require less processing by the LSRs because no signaling protocol is used. However, because paths are statically configured, they cannot adapt to network conditions. Topology changes and network outages can create black holes in the LSP that exist until you manually reconfigure the LSP.

Dynamic LSPs

Dynamic LSPs use signaling protocols to establish themselves and propagate LSP information to other LSRs in the network. You configure the inbound router with LSP information that is transmitted throughout the network when you enable the signaling protocols across the LSRs. Because the LSRs must exchange and process signaling packets and instructions, dynamic LSPs consume more resources than static LSPs. However, dynamic LSPs can avoid the network black holes of static LSPs by detecting topology changes and outages and propagating them throughout the network.

Traffic Engineering with MPLS

Traffic engineering facilitates efficient and reliable network operations while simultaneously optimizing network resources and traffic performance. Traffic engineering provides the ability to move traffic flow away from the shortest path selected by the interior gateway protocol (IGP) to a potentially less congested physical path across a network. To support traffic engineering, besides source routing, the network must do the following:

- Compute a path at the source by taking into account all the constraints, such as bandwidth and administrative requirements.
- Distribute the information about network topology and link attributes throughout the network once the path is computed.
- Reserve network resources and modify link attributes.

MPLS traffic engineering uses the following components:

- MPLS LSPs for packet forwarding
- IGP extensions for distributing information about the network topology and link attributes
- CSPF for path computation and path selection
- RSVP extensions to establish the forwarding state along the path and reserve resources along the path

The Services Router also supports traffic engineering across different OSPF regions. For more details, see the *JUNOS MPLS Applications Configuration Guide*.

Signaling Protocols Overview

Two MPLS signaling protocols are used to dynamically establish and maintain LSPs within a network:

- Label Distribution Protocol on page 10
- Resource Reservation Protocol on page 11

Label Distribution Protocol

LDP is a simple, fast-acting signaling protocol that automatically establishes LSP adjacencies within an MPLS network. Routers then share LSP updates such as hello packets and LSP advertisements across the adjacencies.

LDP Operation

Because LDP runs on top of an interior gateway protocol (IGP) such as IS-IS or OSPF, you must configure LDP and the IGP on the same set of interfaces. After both are configured, LDP begins transmitting and receiving LDP messages through all LDP-enabled interfaces.

Because of LDP's simplicity, it cannot perform true traffic engineering like RSVP. LDP does not support bandwidth reservation or traffic constraints.

LDP Messages

When you configure LDP on an LSR, the router begins sending LDP discovery messages out all LDP-enabled interfaces. When an adjacent LSR receives LDP discovery messages, it establishes an underlying TCP session. An LDP session is then created on top of the TCP session. The TCP three-way handshake ensures that the LDP session has bidirectional connectivity. After they establish the LDP session, the LDP neighbors maintain, and terminate, the session by exchanging messages.

LDP advertisement messages allow LSRs to exchange label information to determine the next hops within a particular LSP.

Any topology changes, such as a router failure, generate LDP notifications that can terminate the LDP session or generate additional LDP advertisements to propagate an LSP change.

Resource Reservation Protocol

Resource Reservation Protocol (RSVP) is a signaling protocol that handles bandwidth allocation and true traffic engineering across an MPLS network. Like LDP, RSVP uses discovery messages and advertisements to exchange LSP path information between all hosts. However, RSVP also includes a set of features that control the flow of traffic through an MPLS network.

This section contains the following topics:

- RSVP Fundamentals on page 11
- Bandwidth Reservation Requirement on page 12
- Explicit Route Objects on page 12
- Constrained Shortest Path First on page 13
- Link Coloring on page 14

RSVP Fundamentals

RSVP uses unidirectional and simplex flows through the network to perform its function. The inbound router initiates an RSVP path message and sends it downstream to the outbound router. The path message contains information about the resources needed for the connection. Each router along the path begins to maintain information about a reservation.

When the path message reaches the outbound router, resource reservation begins. The outbound router sends a reservation message upstream to the inbound router. Each router along the path receives the reservation message and sends it upstream, following the path of the original path message. When the inbound router receives the reservation message, the unidirectional network path is established.

The established path remains open as long as the RSVP session is active. The session is maintained by the transmission of additional path and reservation messages that report the session state every 30 seconds. If a router does not receive the maintenance

messages for 3 minutes, it terminates the RSVP session and reroutes the LSP through another active router.

Bandwidth Reservation Requirement

When a bandwidth reservation is configured, reservation messages propagate the bandwidth value throughout the LSP. Routers must reserve the bandwidth specified across the link for the particular LSP. If the total bandwidth reservation exceeds the available bandwidth for a particular LSP segment, the LSP is rerouted through another LSR. If no segments can support the bandwidth reservation, LSP setup fails and the RSVP session is not established.

Explicit Route Objects

Explicit Route Objects (EROs) limit LSP routing to a specified list of LSRs. By default, RSVP messages follow a path that is determined by the network IGP's shortest path. However, in the presence of a configured ERO, the RSVP messages follow the path specified.

EROs consist of two types of instructions: loose hops and strict hops. When a loose hop is configured, it identifies one or more transit LSRs through which the LSP must be routed. The network IGP determines the exact route from the inbound router to the first loose hop, or from one loose hop to the next. The loose hop specifies only that a particular LSR be included in the LSP.

When a strict hop is configured, it identifies an exact path through which the LSP must be routed. Strict-hop EROs specify the exact order of routers through which the RSVP messages are sent.

You can configure loose-hop and strict-hop EROs simultaneously. In this case, the IGP determines the route between loose hops, and the strict-hop configuration specifies the exact path for particular LSP path segments.

Figure 2 on page 13 shows a typical RSVP-signaled LSP that uses EROs.


Figure 2: Typical RSVP-Signaled LSP with EROs

In the topology shown in Figure 2 on page 13, traffic is routed from Host C1 to Host C2. The LSP can pass through Router R4 or Router R7. To force the LSP to use R4, you can set up either a loose-hop or strict-hop ERO that specifies R4 as a hop in the LSP. To force a specific path through Routers R4, R3, and R6, configure a strict-hop ERO through the three LSRs.

Constrained Shortest Path First

Whereas IGPs use the Shortest Path First (SPF) algorithm to determine how traffic is routed within a network, RSVP uses the Constrained Shortest Path First (CSPF) algorithm to calculate traffic paths that are subject to the following constraints:

- LSP attributes—Administrative groups such as link coloring, bandwidth requirements, and EROs
- Link attributes—Colors on a particular link and available bandwidth

These constraints are maintained in the traffic engineering database (TED). The database provides CSPF with up-to-date topology information, the current reservable bandwidth of links, and the link colors.

In determining which path to select, CSPF follows these rules:

- 1. Computes LSPs one at a time, beginning with the highest-priority LSP—the one with the lowest setup priority value. Among LSPs of equal priority, CSPF starts with those that have the highest bandwidth requirement.
- 2. Prunes the traffic engineering database of links that are not full duplex and do not have sufficient reservable bandwidth.
- 3. If the LSP configuration includes the **include** statement, prunes all links that do not share any included colors.
- 4. If the LSP configuration includes the **exclude** statement, prunes all links that contain excluded colors. If a link does not have a color, it is accepted.

- 5. Finds the shortest path toward the LSP's outbound router, taking into account any EROs. For example, if the path must pass through Router A, two separate SPF algorithms are computed: one from the inbound router to Router A and one from Router A to the outbound router.
- 6. If several paths have equal cost, chooses the one with a last-hop address the same as the LSP's destination.
- 7. If several equal-cost paths remain, selects the path with the fewest number of hops.
- 8. If several equal-cost paths remain, applies CSPF load-balancing rules configured on the LSP.

Link Coloring

RSVP allows you to configure administrative groups for CSPF path selection. An administrative group is typically named with a color, assigned a numeric value, and applied to the RSVP interface for the appropriate link. Lower numbers indicate higher priority.

After configuring the administrative group, you can either exclude, include, or ignore links of that color in the traffic engineering database:

- If you exclude a particular color, all segments with an administrative group of that color are excluded from CSPF path selection.
- If you include a particular color, only those segments with the appropriate color are selected.
- If you neither exclude nor include the color, the metrics associated with the administrative groups and applied on the particular segments determine the path cost for that segment.

The LSP with the lowest total path cost is selected into the traffic engineering database.

VPN Overview

Virtual private networks (VPNs) are private networks that use a public network to connect two or more remote sites. In place of dedicated connections between networks, VPNs use virtual connections routed (tunneled) through public networks that are typically service provider networks. The type of the VPN is determined by the connections it uses and whether the customer network or the provider network performs the virtual tunneling.

This overview contains the following topics:

- VPN Components on page 15
- VPN Routing Requirements on page 15
- VPN Routing Information on page 16
- Types of VPNs on page 17

VPN Components

All types of VPNs share certain components. Figure 3 on page 15 shows a typical VPN topology.

Figure 3: Typical VPN Topology



The provider edge (PE) routers in the provider's network connect to the customer edge (CE) devices located at customer sites. PE routers support VPN and MPLS label functionality. Within a single VPN, pairs of PE routers are connected through a virtual tunnel, typically an LSP.

Provider routers within the core of the provider's network are not connected to any routers at a customer site but are part of the tunnel between pairs of PE routers. Provider routers support LSP functionality as part of the tunnel support, but do not support VPN functionality.

Customer edge (CE) devices are the routers or switches located at the customer site that connect to the provider's network. CE devices are typically IP routers, but they can also be Asynchronous Transfer Mode (ATM), Frame Relay, or Ethernet switches.

All VPN functions are performed by the PE routers. Neither CE devices nor provider routers are required to perform any VPN functions.

VPN Routing Requirements

VPNs tunnel traffic as follows from one customer site to another customer site, using a public network as a transit network, when certain requirements are met:

1. Traffic is forwarded by standard IP forwarding from the CE devices to the PE routers.

The CE devices require only a BGP connection to the PE routers.

2. The PE routers establish an LSP through the provider network.

The provider network must be running either OSPF or IS-IS as an IGP, as well as IBGP sessions through either a full mesh or route reflector. IBGP is required so that the PE routers can exchange route information for routes that originate or terminate in the VPN.

3. When the inbound PE router receives traffic, it performs a route lookup. The lookup yields an LSP next hop, and the traffic is forwarded along the LSP.

Either LDP or RSVP must be configured to dynamically set up LSPs through the provider network.

4. When the traffic reaches the outbound PE router, the PE router pops the MPLS label and forwards the traffic with standard IP routing.

Because the tunnel information is maintained at both PE routers, neither the provider core routers nor the CE devices need to maintain any VPN information in their configuration databases.

VPN Routing Information

Routing information, including routes, route distinguishers, and routing policies, is stored in a VPN routing and forwarding (VRF) table. Routers must maintain separate VRF tables for each VPN.

VRF Instances

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The interfaces belong to the routing tables, and the routing protocol parameters control the information in the routing tables. In the case of VPNs, each VPN has a VPN routing and forwarding (VRF) instance.

A VRF instance consists of one or more routing tables, a derived forwarding table, the interfaces that use the forwarding table, and the policies and routing protocols that determine what goes into the forwarding table. Because each instance is configured for a particular VPN, each VPN has separate tables, rules, and policies that control its operation.

A separate VRF table is created for each VPN that has a connection to a CE router. The VRF table is populated with routes received from directly connected CE sites associated with the VRF instance, and with routes received from other PE routers in the same VPN.

Route Distinguishers

Because a typical transit network is configured to handle more than one VPN, the provider routers are likely to have multiple VRF instances configured. As a result, depending on the origin of the traffic and any filtering rules applied to the traffic, the BGP routing tables can contain multiple routes for a particular destination address. Because BGP requires that exactly one BGP route per destination be imported into the forwarding table, BGP must have a way to distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs.

A route distinguisher is a locally unique number that identifies all route information for a particular VPN. Unique numeric identifiers allow BGP to distinguish between routes that are otherwise identical.

Route Targets to Control the VRF Table

On each PE router, you must define routing policies that specify how routes are imported into and exported from the router's VRF table. Each advertisement must have an associated route target that uniquely identifies the VPN for which the advertisement is valid. The route target allows you to keep routing and signaling information for each VPN separate.

Types of VPNs

There are three primary types of VPNs: Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs.

Layer 2 VPNs

In a Layer 2 VPN, traffic is forwarded to the PE router in Layer 2 format, carried by MPLS through an LSP over the service provider network, and then converted back to Layer 2 format at the receiving CE device.

On a Layer 2 VPN, routing occurs on the customer routers, typically on the CE router. The CE router connected to a service provider on a Layer 2 VPN must select the appropriate circuit on which to send traffic. The PE router receiving the traffic sends it across the network to the PE router on the outbound side. The PE routers need no information about the customer's routes or routing topology, and need only to determine the virtual tunnel through which to send the traffic.

Layer 2 Circuits

A Layer 2 circuit is a point-to-point Layer 2 connection that transports traffic by MPLS or another tunneling technology on a service provider network. The Layer 2 circuit creates a virtual connection to direct traffic between two CE routers. The primary difference between a Layer 2 circuit and an Layer 2 VPN is the method of setting up the virtual connection. Like a leased line, a Layer 2 circuit forwards all packets received from the local interface to the remote interface.

Layer 3 VPNs

In a Layer 3 VPN, routing occurs on the service provider's routers. As a result, Layer 3 VPNs require information about customer routes and a more extensive VRF policy configuration to share and filter routes that originate or terminate in the VPN.

Because Layer 3 VPNs require the provider routers to route and forward VPN traffic at the entry and exit points of the transit network, the routes must be advertised and filtered throughout the provider network.

Route advertisements originate at the CE devices and are shared with the inbound PE routers through standard IP routing protocols, typically BGP. Based on the source address, the PE router filters route advertisements and imports them into the appropriate VRF table.

The PE router then exports the route in IBGP sessions to the other provider routers. Route export is governed by any routing policy that has been applied to the particular VRF table. To propagate the routes through the provider network, the PE router must also convert the route to VPN format, which includes the route distinguisher. When the outbound PE router receives the route, it strips off the route distinguisher and advertises the route to the connected CE device, typically through standard BGP IPv4 route advertisements.

Chapter 2 Configuring Signaling Protocols for Traffic Engineering

Signaling protocols are used within a Multiprotocol Label Switching (MPLS) environment to establish label-switched paths (LSPs) for traffic across a transit network. J-series Services Routers support the Label Distribution Protocol (LDP) and the Resource Reservation Protocol (RSVP) as part of their suite of traffic engineering features.

You can use either the J-Web configuration editor or CLI configuration editor to configure signaling protocols.

This chapter contains the following topics. For more information about MPLS traffic engineering, see the *JUNOS MPLS Applications Configuration Guide*.

- Signaling Protocol Overview on page 19
- Before You Begin on page 20
- Configuring LDP and RSVP with a Configuration Editor on page 20
- Verifying an MPLS Configuration on page 25

Signaling Protocol Overview

When transit traffic is routed through an IP network, MPLS is often used to engineer its passage. Although the exact path through the transit network is of little importance to either the sender or the receiver of the traffic, network administrators often want to route traffic more efficiently between certain source and destination address pairs. By adding a short label with specific routing instructions to each packet, MPLS switches packets from router to router through the network rather than forwarding packets based on next-hop lookups. The resulting routes are called label-switched paths (LSPs). LSPs control the passage of traffic through the network and speed traffic forwarding.

You can create LSPs manually, or through the use of signaling protocols. Services Routers support two signaling protocols—the Label Distribution Protocol (LDP) and the Resource Reservation Protocol (RSVP).

LDP Signaling Protocol

The Label Distribution Protocol (LDP) is a signaling protocol that runs on a Services Router configured for MPLS support. The LDP configuration is added to the existing interior gateway protocol (IGP) configuration and included in the MPLS configuration. To configure a network to use LDP for LSP establishment, you first enable MPLS on all transit interfaces in the MPLS network and then enable LDP sessions on the interfaces.

The successful configuration of both MPLS and LDP initiates the exchange of TCP packets across the LDP interfaces. The packets establish TCP-based LDP sessions for the exchange of MPLS information within the network. Enabling both MPLS and LDP on the appropriate interfaces is sufficient to establish LSPs.

RSVP Signaling Protocol

The Resource Reservation Protocol (RSVP) is a more flexible and powerful way to engineer traffic through a transit network. Like LDP, RSVP establishes LSPs within an MPLS network when you enable both MPLS and RSVP on the appropriate interfaces. However, whereas LDP is restricted to using the configured IGP's shortest path as the transit path through the network, RSVP uses a combination of the Constrained Shortest Path First (CSPF) algorithm and Explicit Route Objects (EROs) to determine how traffic is routed through the network.

Basic RSVP sessions are established in exactly the same way that LDP sessions are established. By configuring both MPLS and RSVP on the appropriate transit interfaces, you enable the exchange of RSVP packets and the establishment of LSPs. However, RSVP also lets you configure link authentication, explicit LSP paths, and link coloring. For more information about these topics, see the *JUNOS MPLS Applications Configuration Guide*.

Before You Begin

Before you begin configuring signaling protocols for traffic engineering, complete the following tasks:

- Establish basic connectivity. See the Getting Started Guide for your router.
- Configure network interfaces. See the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.
- Configure an interior gateway protocol (IGP) across your network. See the *J-series* Services Router Basic LAN and WAN Access Configuration Guide. For information about the IS-IS IGP, see the JUNOS Routing Protocols Configuration Guide.

Configuring LDP and RSVP with a Configuration Editor

To configure either LDP or RSVP as a signaling protocol on the Services Router to establish LSPs through an IP network, perform one of the following tasks:

- Configuring LDP-Signaled LSPs on page 21
- Configuring RSVP-Signaled LSPs on page 23

For information about using the J-Web and CLI configuration editors, see the *J-series Services Router Basic LAN and WAN Access Configuration Guide*.

Configuring LDP-Signaled LSPs

Using LDP as a signaling protocol, you create LSPs between Services Routers in an IP network. A sample network is shown in Figure 4 on page 21.

Figure 4: Typical LDP-Signaled LSP



To establish an LSP between Services Routers R6 and R7, you must configure LDP on Services Routers R5, R6, and R7. This configuration ensures that Hosts C1 and C2 use the LDP-signaled LSP when the entry (ingress) router is R6 or R7.

To configure LDP to establish the LSP shown in Figure 4 on page 21, perform these steps:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 6 on page 21.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to "Verifying an LDP-Signaled LSP" on page 25.

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level of the configuration hierarchy	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit interfaces
	2.	Next to Interfaces, click Configure or Edit .	

Table 6: Configuring an LDP-Signaled LSP

Table 6: Configuring an LDP-Signaled LSP (continued)

Task	J-Web Con	figuration Editor	CLI	Configuration Editor
Enable the MPLS family on all transit interfaces on	1. Click want	he transit interface on which you to configure MPLS.	1.	Add the MPLS family to all transit interfaces. For example:
each router in the MPLS network.	In the which	Unit table, click the unit number for you want to enable MPLS.		set ge-0/0/0 unit 0 family mpls
	3. In the box.	Family area, select the Mpls check	2.	Repeat Step 1 for each transit interface on the routers in the MPLS network.
	4. Click	OK.		
	5. Repea interfa netwo	It Steps 1 through 4 for each transit ace on the routers in the MPLS rk.		
Enable the MPLS process on all MPLS interfaces for	1. On the Protoc	e main Configuration page next to cols, click Configure or Edit .	1.	From the [edit] hierarchy level, enter
each router in the MPLS	2. Next t	to Mpls, click Configure or Edit .		edit protocols mpls
	3. Next t	to Interface, click Add new entry .	2.	Enter
(See the interface naming conventions in the <i>J-series</i>	4. In the	Interface name box, type all.		set interface all
Services Router Basic LAN and WAN Access	5. Click	OK.	3.	Repeat Steps 1 and 2 for each transit
and WAN Access Configuration Guide.)	6. Repea interfa netwo	at Steps 1 through 5 for each transit ace on the routers in the MPLS rk.		interface on the routers in the MPLS network.
Create the LDP instance on each Services Router in the	1. On the Protoc	e main Configuration page next to cols, click Configure or Edit .	1.	From the [edit] hierarchy level, enter
MPLS network.	2. Next t	to Ldp, click Configure or Edit .		edit protocols ldp
	3. Next t	to Interface, click Add new entry .	2.	Enable LDP on a transit interface. For example:
	4. In the of a tr	Interface name box, type the name ansit interface—for example,)/0		set interface ge-0/0/0
	5. Click	OK.	3.	Repeat Steps 1 and 2 for each transit interface on the routers in the MPLS
	6. Repea interfa netwo	at Steps 1 through 5 for each transit ace on the routers in the MPLS ork.		network.
Set the keepalive interval	1. In the	Keepalive interval box, type 10.	On	each router in the MPLS network, enter
to 10 seconds.	2. Click	OK.	set	keepalive-interval 10
The keepalive interval specifies the number of seconds between the transmission of keepalive messages along the LDP link.	3. Repea the M	tt Steps 1 and 2 for each router in PLS network.	200	

Configuring RSVP-Signaled LSPs

Using RSVP as a signaling protocol, you create LSPs between Services Routers in an IP network. A sample network is shown in Figure 5 on page 23.

Figure 5: Typical RSVP-Signaled LSP



To establish an LSP between Services Routers R1 and R7, you must configure RSVP on all MPLS transit interfaces in the network. This configuration ensures that Hosts C1 and C2 use the RSVP-signaled LSP corresponding to the network IGP's shortest path. Additionally, this configuration reserves 10 Mbps of bandwidth.

To configure RSVP to establish the LSP shown in Figure 5 on page 23, perform these steps:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 7 on page 23.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to "Verifying an RSVP-Signaled LSP" on page 27.

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level of the configuration hierarchy	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit interfaces
	2.	Next to Interfaces, click Configure or Edit .	

Table 7: Configuring an RSVP-Signaled LSP

Table 7: Configuring an RSVP-Signaled LSP (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Enable the MPLS family on all transit interfaces on each router in the MPLS network.	1. Click the transit interface on which you want to configure MPLS.	1. Add the MPLS family to all transit interfaces. For example:
	2. In the Unit table, click the unit number for which you want to enable MPLS.	r set ge-0/0/0 unit 0 family mpls
	3. In the Family area, select the Mpls chec box.	2. Repeat Step 1 for each transit interface on the routers in the MPLS network.
	4. Click OK .	
	5. Repeat Steps 1 through 4 for each trans interface on the routers in the MPLS network.	t
Enable the MPLS process on all MPLS interfaces for	 On the main Configuration page next to Protocols, click Configure or Edit. 	1. From the [edit] hierarchy level, enter
each router in the MPLS	2. Next to Mpls, click Configure or Edit .	edit protocols mpls
	3. Next to Interface, click Add new entry .	2. Enter
	4. In the Interface name box, type all.	set interface all
	5. Click OK .	3. Repeat Steps 1 and 2 for each transit
	6. Repeat Steps 1 through 5 for each trans interface on the routers in the MPLS network.	interface on the routers in the MPLS t network.
Create the RSVP instance on each Services Router in	 On the main Configuration page next to Protocols, click Configure or Edit. 	1. From the [edit] hierarchy level, enter
the MPLS network.	2. Next to Rsvp, click Configure or Edit .	edit protocols rsvp
(See the interface naming	3. Next to Interface, click Add new entry .	 Enable RSVP on a transit interface. For example:
Services Router Basic LAN and WAN Access	 In the Interface name box, type the nam of a transit interface—for example, ge-0/0/0 	set interface ge-0/0/0
conjugaration Guide.)	5. Click OK .	3. Repeat Steps 1 and 2 for each transit interface on the routers in the MPLS
	6. Repeat Steps 1 through 5 for each trans interface on the routers in the MPLS network.	network. t
On the entry (ingress) router, R1, define the LSP	 On the main Configuration page next to Protocols, click Configure or Edit. 	1. From the [edit] hierarchy level, enter
r1–r7, using Router R7's loopback address (10.0.9.7).	2. Next to Mpls, click Configure or Edit .	edit protocols mpls
	3. Next to Label switched path, click Add new entry .	2. Enter
	4. In the Path name box, type r1-r7.	set aber switched patient - in to 10.0.3.1
	5. In the To box, type 10.0.9.7 .	

Task	J-Web Configuration Editor	CLI Configuration Editor
Reserve 10 Mbps of	1. In the Bandwidth box, click Configure	e. Enter
bandwidth on the LSP.	2. In the Ct0 box, type 10m .	set label-switched-path r1–r7 bandwidth 10m
	3. Click OK .	
Disable the use of the Constrained Shortest Path First (CSPF) algorithm.	 Select the No cspf check box. Click OK. 	Enter set label-switched-path r1-r7 no-cspf
By disabling the CSPF algorithm, you specify that traffic through the LSP is to be routed along the network IGP's shortest path.		

Table 7: Configuring an RSVP-Signaled LSP (continued)

Verifying an MPLS Configuration

The tasks required to verify your MPLS configuration depend on the signaling protocol used. To validate the configuration, perform the appropriate set of tasks:

- Verifying an LDP-Signaled LSP on page 25
- Verifying an RSVP-Signaled LSP on page 27

Verifying an LDP-Signaled LSP

Suppose that LDP is configured to establish an LSP as shown in Figure 4 on page 21.

To verify the LDP configuration, perform these verification tasks:

- Verifying LDP Neighbors on page 25
- Verifying LDP Sessions on page 26
- Verifying the Presence of LDP-Signaled LSPs on page 27
- Verifying Traffic Forwarding over the LDP-Signaled LSP on page 27

Verifying LDP Neighbors

- **Purpose** Verify that each Services Router shows the appropriate LDP neighbors—for example, that Router R5 has both Router R6 and Router R7 as LDP neighbors.
 - Action From the CLI, enter the show ldp neighbor command.

user@r5> s	how ldp neighbor		
Address	Interface	Label space ID	Hold time
10.0.8.5	ge-0/0/0.0	10.0.9.6:0	14
10.0.8.10	ge-0/0/1.0	10.0.9.7:0	11

- **What It Means** The output shows the IP addresses of the neighboring interfaces along with the interface through which the neighbor adjacency is established. Verify the following information:
 - Each interface on which LDP is enabled is listed.
 - Each neighboring LDP interface address is listed with the appropriate corresponding LDP interface.
 - Under Label space ID, the appropriate loopback address for each neighbor appears.
- **Related Topics** For a complete description of **show ldp neighbor** output, see the *JUNOS Routing Protocols and Policies Command Reference.*

Verifying LDP Sessions

- **Purpose** Verify that a TCP-based LDP session has been established between all LDP neighbors. Also, verify that the modified keepalive value is active.
- Action From the CLI, enter the show ldp session detail command.

```
user@r5> show ldp session detail
Address: 10.0.9.7, State: Operational, Connection: Open, Hold time: 28
Session ID: 10.0.3.5:0--10.0.9.7:0
Next keepalive in 3 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Keepalive interval: 10, Connect retry interval: 1
Local - Restart: disabled, Helper mode: enabled
Remote - Restart: disabled, Helper mode: disabled
Local maximum recovery time: 240000 msec
Next-hop addresses received:
10.0.8.10
10.0.2.17
```

- **What It Means** The output shows the detailed information, including session IDs, keepalive interval, and next-hop addresses, for each established LDP session. Verify the following information:
 - Each LDP neighbor address has an entry, listed by loopback address.
 - The state for each session is Operational, and the connection for each session is Open. A state of Nonexistent or a connection of Closed indicates a problem with one of the following:
 - LDP configuration
 - Passage of traffic between the two Services Routers
 - Physical link between the two routers
 - For Keepalive interval, the appropriate value, **10**, appears.

Related Topics For a complete description of **show ldp session detail** output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying the Presence of LDP-Signaled LSPs

- **Purpose** Verify that each Services Router's inet.3 routing table has an LSP for the loopback address on each of the other routers.
- **Action** From the CLI, enter the **show route table inet.3** command.

user@r5> show route table inet.3
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.9.6/32	*[LDP/9/0] 00:05:29, metric 1
	> to 10.0.8.5 via ge-0/0/0.0
10.0.9.7/32	*[LDP/9/0] 00:05:37, metric 1
	> to 10.0.8.10 via ge-0/0/1.0

- **What It Means** The output shows the LDP routes that exist in the inet.3 routing table. Verify that an LDP-signaled LSP is associated with the loopback addresses of the other routers in the MPLS network.
- **Related Topics** For a complete description of **show route table** output, see the *JUNOS Routing Protocols and Policies Command Reference.*

Verifying Traffic Forwarding over the LDP-Signaled LSP

- **Purpose** Verify that traffic between Hosts C1 and C2 is forwarded over the LDP-signaled LSP between Services Router R6 and Services Router R7. Because traffic uses any configured gateway address by default, you must explicitly specify that the gateway address is to be bypassed.
 - Action If Host C1 is a Juniper Networks router, from the CLI enter the traceroute 220.220.0.0 source 200.200.0.1 bypass-routing gateway 172.16.0.1 command.

user@c1> traceroute 220.220.0.0 source 200.200.0.1 bypass-routing gateway 172.16.0.1 traceroute to 220.220.0.1 (172.16.0.1) from 200.200.0.1, 30 hops max, 40 byte packets

172.16.0.1 (172.16.0.1) 0.661 ms 0.538 ms 0.449 ms
10.0.8.9 (10.0.8.9) 0.511 ms 0.479 ms 0.468 ms MPLS Label=100004 CoS=0 TTL=1 S=1
10.0.8.5 (10.0.8.5) 0.476 ms 0.512 ms 0.441 ms
220.220.0.1 (220.220.0.1) 0.436 ms 0.420 ms 0.416 ms

What It Means
The output shows the route that traffic travels between Hosts C1 and C2, without using the default gateway. Verify that traffic sent from C1 to C2 travels through Router R7. The 10.0.8.9 address is the interface address for Router R5.

Related Topics For information about the traceroute command and its output. see the *JUNOS System Basics and Services Command Reference*.

Verifying an RSVP-Signaled LSP

Suppose that RSVP is configured to establish an LSP as shown in Figure 5 on page 23.

To verify the RSVP configuration, perform these verification tasks:

- Verifying RSVP Neighbors on page 28
- Verifying RSVP Sessions on page 28
- Verifying the Presence of RSVP-Signaled LSPs on page 29

Verifying RSVP Neighbors

- **Purpose** Verify that each Services Router shows the appropriate RSVP neighbors—for example, that Router R1 lists both Router R3 and Router R2 as RSVP neighbors.
- **Action** From the CLI, enter the **show rsvp neighbor** command.

user@r1> show rsvp neighbor						
RSVP neighbor: 2	learne	b				
Address	Idle	Up/Dn	LastChange	HelloInt	HelloTx/Rx	
10.0.6.2	0	3/2	13:01	3	366/349	
10.0.3.3	0	1/0	22:49	3	448/448	

- **What It Means** The output shows the IP addresses of the neighboring routers. Verify that each neighboring RSVP router loopback address is listed.
- **Related Topics** For a complete description of **show rsvp neighbor** output, see the *JUNOS Routing Protocols and Policies Command Reference.*

Verifying RSVP Sessions

- **Purpose** Verify that an RSVP session has been established between all RSVP neighbors. Also, verify that the bandwidth reservation value is active.
- Action From the CLI, enter the show rsvp session detail command.

```
user@r1> show rsvp session detail
Ingress RSVP: 1 sessions
```

10.0.9.7

```
From: 10.0.6.1, LSPstate: Up, ActiveRoute: 0
LSPname: r1-r7, LSPpath: Primary
Bidirectional, Upstream label in: -, Upstream label out: -
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 100000
Resv style: 1 FF, Label in: -, Label out: 100000
Time left: -, Since: Thu Jan 26 17:57:45 2002
Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
Port number: sender 3 receiver 17 protocol 0
PATH rcvfrom: localclient
PATH sentto: 10.0.4.13 (ge-0/0/1.0) 1467 pkts
RESV rcvfrom: 10.0.4.13 (ge-0/0/1.0) 1467 pkts
Record route: <self> 10.0.4.13 10.0.2.1 10.0.8.10
```

- **What It Means** The output shows the detailed information, including session IDs, bandwidth reservation, and next-hop addresses, for each established RSVP session. Verify the following information:
 - Each RSVP neighbor address has an entry for each neighbor, listed by loopback address.
 - The state for each LSP session is **Up**.
 - Under **Tspec**, the appropriate bandwidth value, **10Mbps**, appears.
- **Related Topics** For a complete description of **show rsvp session detail** output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying the Presence of RSVP-Signaled LSPs

- **Purpose** Verify that the inet.3 routing table of the entry (ingress) Services Router, R1, has a configured LSP to the loopback address of Router R7.
- **Action** From the CLI, enter the show route table inet.3 command.

user@r1> show route table inet.3
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.9.7/32 *[RSVP/7] 00:05:29, metric 10 > to 10.0.4.17 via ge-0/0/0.0, label-switched-path r1-r7

- **What It Means** The output shows the RSVP routes that exist in the inet.3 routing table. Verify that an RSVP-signaled LSP is associated with the loopback address of the exit (egress) router, R7, in the MPLS network.
- **Related Topics** For a complete description of **show route table** output, see the *JUNOS Routing Protocols and Policies Command Reference.*

J-series[™] Services Router Advanced WAN Access Configuration Guide

Chapter 3 Configuring Virtual Private Networks

You can configure a Services Router to participate in several types of virtual private networks (VPNs). A VPN allows remote sites and users to use a public communication infrastructure to create secure access to an organization's network. VPNs are a cost-effective alternative to expensive dedicated lines.

There are many ways to set up a VPN and direct traffic through it. This chapter describes the most common tasks involved in setting up a basic Layer 2 VPN, Layer 2 circuit, or Layer 3 VPN configuration. For more information about VPNs, including other configurations and advanced or less common tasks, see the *JUNOS VPNs Configuration Guide*.

You can use either the J-Web configuration editor or the CLI configuration editor to configure VPNs.

This chapter contains the following topics:

- VPN Configuration Overview on page 31
- Before You Begin on page 34
- Configuring VPNs with a Configuration Editor on page 34
- Verifying a VPN Configuration on page 52

VPN Configuration Overview

To configure VPN functionality on a Services Router, you must enable support on the provider edge (PE) Services Router as well as configure the Services Router to distribute routing information to other Services Routers in the VPN. The sample configurations in this chapter describe setting up a basic Multiprotocol Label Switching (MPLS) Layer 2 VPN, Layer 3 VPN, and Layer 2 circuit.

This section contains the following topics:

- Sample VPN Topology on page 32
- Basic Layer 2 VPN Configuration on page 32
- Basic Layer 2 Circuit Configuration on page 32
- Basic Layer 3 VPN Configuration on page 33

Sample VPN Topology

Figure 6 on page 32 shows the overview of a basic VPN topology for the sample configurations in this chapter.

Figure 6: Basic VPN Topology



Basic Layer 2 VPN Configuration

Implementing a Layer 2 VPN on the Services Router is similar to implementing a VPN using a Layer 2 technology such as Asynchronous Transfer Mode (ATM) or Frame Relay. However, for a Layer 2 VPN on the Services Router, traffic is forwarded to the router in a Layer 2 format. Traffic is then carried by Multiprotocol Label Switching (MPLS) over the service provider's network, and then converted back to Layer 2 format at the receiving end.

On a Layer 2 VPN, routing occurs on the customer's Services Routers, typically on the customer edge (CE) router. The CE Services Router connected to a service provider on a Layer 2 VPN must select the appropriate circuit on which to send traffic. The provider edge (PE) Services Router receiving the traffic sends it across the service provider's network to the PE Services Router connected to the receiving site. PE Services Routers are not required to learn the customer's routes or routing topology, but they must identify the tunnel through which to send the data.

In this sample Layer 2 VPN configuration, the PE routers use the same autonomous system (AS). Within the AS, routing information is communicated through an interior gateway protocol (IGP). Outside the AS, routing information is shared with other ASs through Border Gateway Protocol (BGP). Each AS has a single routing policy and uses a group of one or more IP prefixes. The PE routers must use the same signaling protocols to communicate.

Each routing instance that you configure on a PE router must have a unique route distinguisher associated with it. VPN routing instances need a route distinguisher to help BGP identify overlapping network layer reachability information (NLRIs) messages from different VPNs.

Basic Layer 2 Circuit Configuration

A Layer 2 circuit is a point-to-point Layer 2 connection that transports traffic by means of Multiprotocol Label Switching (MPLS) or another tunneling technology on the service provider network. The Layer 2 circuit creates a virtual connection to direct

traffic between two CE Services Routers across a service provider network. The main difference between a Layer 2 VPN and a Layer 2 circuit is the method of setting up the virtual connection. As with a leased line, a Layer 2 circuit forwards all packets received from the local interface to the remote interface.

On the interface communicating with the other PE router, you must specify MPLS and IPv4, and include the IP address. For the loopback interface, you must specify inet, and include the IP address. For IPv4, you must designate the loopback interface as primary so it can receive control packets. Because it is always operational, the loopback interface is best able to perform the control function.

On the PE router interface facing the CE router, you must specify a circuit cross-connect (CCC) encapsulation type. The type of encapsulation depends on the interface type. For example, an Ethernet interface uses **ethernet-ccc**. The encapsulation type determines how the packet is constructed for that interface.

On the CE router interface that faces the PE router, you must specify inet (for IPv4), and include the IP address. You also specify a routing protocol such as Open Shortest Path First (OSPF) which specifies the area and IP address of the Services Router interface.

With this information, the Services Routers can send and receive packets across the circuit.

Basic Layer 3 VPN Configuration

A Layer 3 VPN operates at the Layer 3 level of the OSI model, the Network layer. In this configuration, the service provider network must learn the IP addresses of devices sending traffic across the VPN. The Layer 3 VPN requires more processing power on the PE Services Routers, because it has larger routing tables for managing network traffic on the customer sites.

A Layer 3 VPN is a set of sites that share common routing information, and connectivity of the sites is controlled by a collection of policies. The sites making up a Layer 3 VPN are connected over a service provider's existing public Internet backbone.

An interface on each CE Services Router communicates with an interface on a PE Services Router through the external Border Gateway Protocol (EBGP).

On the provider Services Router, you configure two interfaces: one to communicate with each PE Services Router. The interfaces communicate with the PE Services Routers by using IPv4 and MPLS. The provider router is in the same AS as the PE routers, which is typically the case for Layer 3 VPNs.

The provider router uses OSPF and Label Distribution Protocol (LDP) to communicate with the PE Services Routers. For OSPF, the provider Services Router interfaces that communicate with the PE routers are specified, as well as the loopback interface. For the PE routers, the loopback interface is in passive mode, meaning it does not send OSPF packets to perform the control function. In this example, the provider router and PE routers are in the same backbone area. For the LDP configuration, the provider router interfaces that communicate with the PE routers are specified.

Before You Begin

Before you begin configuring VPNs, perform the following tasks:

- Determine which Services Routers are participating in the VPN configuration. This chapter describes configuring an interface for basic VPN connectivity. To configure an interface, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.
- Determine the protocols to use in the VPN configuration. These protocols include
 - MPLS—See "Multiprotocol Label Switching Overview" on page 3 and the JUNOS Routing Protocols Configuration Guide.
 - BGP, EBGP, and internal BGP (IBGP)—See the *J*-series Services Router Basic LAN and WAN Access Configuration Guide and the JUNOS Routing Protocols Configuration Guide.
 - LDP and Resource Reservation Protocol (RSVP)—See "Configuring Signaling Protocols for Traffic Engineering" on page 19 and the *JUNOS MPLS Applications Configuration Guide.*
 - OSPF—See the J-series Services Router Basic LAN and WAN Access Configuration Guide and the JUNOS Routing Protocols Configuration Guide.

Configuring VPNs with a Configuration Editor

To configure a basic Layer 3 VPN, Layer 2 VPN, or Layer 2 circuit, perform the following tasks. Use Table 8 on page 34 to help you select the tasks for your VPN type. For information about using the J-Web and CLI configuration editors, see the *J-series Services Router Basic LAN and WAN Access Configuration Guide*.

- Configuring Interfaces Participating in a VPN on page 35
- Configuring Protocols Used by a VPN on page 37
- Configuring a VPN Routing Instance on page 45
- Configuring a VPN Routing Policy on page 47

Table 8: VPN Configuration Task Summary

Section	Layer 3 VPN	Layer 2 VPN	Layer 2 Circuit
"Configuring Interfaces Participating in a VPN" on page 35	All Services Routers	All Services Routers	All Services Routers
"Configuring Protocols Used by a VPN" on page 37	All Services Routers	All Services Routers	All Services Routers
"Configuring a VPN Routing Instance" on page 45	PE Services Routers	PE Services Routers	N/A

Section	Layer 3 VPN	Layer 2 VPN	Layer 2 Circuit
"Configuring a VPN Routing Policy" on page 47	CE Services Routers (PE Services Routers if you are not using a route target)	PE Services Routers if you are not using a route target	N/A

Table 8: VPN Configuration Task Summary (continued)

Configuring Interfaces Participating in a VPN

Configuring the Services Router interfaces that participate in the VPN is similar to configuring them for other uses, with a few requirements for VPN.

Before following the procedures in this section, make sure you have initially configured the interface as described in the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

To configure an interface for a VPN:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 9 on page 36 for each interface involved in the VPN, except Layer 3 loopback interfaces, which do not require other configuration.
- 3. Go on to "Configuring Protocols Used by a VPN" on page 37.

Table 9: Configuring an Interface for a VPN

Task	J-W	J-Web Configuration Editor		CLI Configuration Editor		
Configure IPv4. (interfaces on all Services	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.		For all interfaces except loopback, and a Layer 2 VPN interface facing a CE router: From the [edit] hierarchy level enter		
Routers)	2.	Next to Interfaces, click Configure or Edit.		edit interfaces interface-name unit		
(See the interface naming conventions in the <i>J-series</i>	3.	In the Interface name column, select the interface.		logical_interface family inet address ipv4_address		
Services Router Basic LAN and WAN Access Configuration Guide.)	 4. 5. 6. 7. 8. 9. 	For Layer 2 VPNs on the interface facing a CE router, select an encapsulation type, such as ethernet-ccc from the Encapsulation list. For Fast Ethernet interfaces, you also must select Vlan tagging from the Vlan tag mode list. In the Interface unit number column, select the logical interface. In the Family group, select Inet and click Edit . Next to Address, click Add new entry In the Source box, type the IPv4 address—for example, 10.49.102.1/30 . For a loopback address on a Layer 2 configuration, select Primary . Click OK to return to the Unit page.	•	For a loopback address on a Layer 2 configuration: From the [edit] hierarchy level, enter edit interfaces IoO unit <i>logical_interface</i> family inet address <i>ipv4_address</i> primary For a Layer 2 VPN interface facing a CE router: From the [edit] hierarchy level, enter set interfaces <i>interface-name</i> vlan-tagging encapsulation vlan-ccc unit <i>logical_interface</i> encapsulation vlan-ccc vlan-id <i>id-number</i>		
Configure the MPLS address family.	On gro	the Unit page, select Mpls in the Family up.	At t	the [edit interfaces interface] level, enter		
(for interfaces on a PE or provider Services Router that communicate with a PE or provider Services Router only, and not for loopback addresses)			set	unit <i>logical_interface</i> family mpls		
For Layer 2 VPNs and circuits, configure encapsulation.	1.	On the Unit page, select an encapsulation type from the Encapsulation list.	1.	At the [edit interfaces interface] level, enter set encapsulation encapsulation_type		
If multiple logical units are configured, the encapsulation type is needed at the interface level only. It is always required at the unit level.	2. 3. 4.	On the Interface page, select an encapsulation type from the Encapsulation list. Click OK until you see the Configuration Interfaces page displaying all interfaces on the router.	2.	Enter set unit logical_interfaceencapsulation encapsulation_type		
(for interfaces on a PE Services Router that communicate with a CE Services Router)						

Configuring Protocols Used by a VPN

The Services Routers in a VPN use a variety of protocols to communicate between PE and provider Services Routers. Use Table 10 on page 37 to help you select the tasks for your VPN type. For more information about configuring routing protocols, see the *JUNOS Routing Protocols Configuration Guide* and the *JUNOS MPLS Applications Configuration Guide*.

This section contains the following topics:

- Configuring MPLS for VPNs on page 37
- Configuring a BGP Session on page 39
- Configuring Routing Options for VPNs on page 40
- Configuring an IGP and a Signaling Protocol on page 41
- Configuring LDP for Signaling on page 41
- Configuring RSVP for Signaling on page 43
- Configuring a Layer 2 Circuit on page 44

Table 10: VPN Protocol Configuration Task Summary

Section	Layer 3 VPN	Layer 2 VPN	Layer 2 Circuit
"Configuring MPLS for VPNs" on page 37	N/A unless you are using RSVP	PE and provider Services Routers	PE Services Routers
"Configuring a BGP Session" on page 39	PE Services Routers	PE Services Routers	PE Services Routers
"Configuring Routing Options for VPNs" on page 40	All Services Routers	All Services Routers	All Services Routers
"Configuring an IGP and a Signaling Protocol" on page 41—one of the following tasks:	PE and provider Services Routers	PE Services Routers	PE Services Routers
 Configuring LDP for Signaling on page 41 Configuring RSVP for Signaling on page 43 			
"Configuring a Layer 2 Circuit" on page 44	N/A	N/A	PE Services Routers

Configuring MPLS for VPNs

For Layer 2 VPN and Layer 2 circuit interfaces that communicate with other PE Services Routers and provider Services Routers, you must advertise the interface using MPLS. Unless you are using RSVP, this section does not apply to Layer 3 VPNs because MPLS is configured on the interface.

For more information about configuring MPLS, see "Multiprotocol Label Switching Overview" on page 3*JUNOS MPLS Applications Configuration Guide*.

To configure MPLS for VPNs:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 11 on page 38 on each PE Services Router and provider Services Router interface that communicates with another PE Services Router.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the configuration, see "Verifying a VPN Configuration" on page 52
- 5. Go on to "Configuring a BGP Session" on page 39.

Table 11: Configuring MPLS for VPNs

Task	J-W	eb Configuration Editor	CLI	I Configuration Editor
Navigate to the top of the configuration hierarchy and specify the interfaces	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	Fro foll wa	om the [edit] hierarchy level, enter the lowing command for each interface you nt to enable:
used for communication between PE routers and	2.	Next to Mpls, click Configure or Edit.	edi	t protocols mpls interface interface name
between PE routers and	3.	Next to Interface, click Configure or Edit .	cui	
(PE and provider Services	4.	In the Interface name box, type interface-name.		
Routers)	5.	Click OK .		
(See the interface naming conventions in the <i>J-series</i> <i>Services Router Basic LAN</i> and WAN Access Configuration Guide.)				
For RSVP only, configure an MPLS label-switched path (LSP) to the destination point on the PE router for LSP. During configuration, you specify	1.	In the MPLS page, click Add New Entry in the Label switched path group.	1.	From the [edit] hierarchy level, enter
	2.	Type a path name in the Path name box and an IP address in the To box.		edit protocols mpls label-switched-path path-name
	3.	Click OK .	2.	Enter
the IP address of the LSP destination point, which is	4.	Next to Interface, click Add New Entry.		set to ip-address
an address on the remote PE router.	5.	Type <i>interface-name</i> in the Interface name box.	3.	Enter up.
The path name is defined on the source Services Router only and is unique between two routers.	6.	Click OK .	4.	Enter
	7.	Repeat Steps 4 through 6 for each interface.		interface interface-name
(PE Services Router interface communicating with another PE Services Router)				

Configuring a BGP Session

You must configure an internal BGP (IBGP) session between PE Services Routers so the Services Routers can exchange information about routes originating and terminating in the VPN. The PE routers use this information to determine which labels to use for traffic destined for remote sites. The IBGP session for the VPN runs through the loopback address. This section is valid for Layer 2 VPNs and Layer 3 VPNs, but not Layer 2 circuits.

For the Layer 3 example, you also configure an EBGP session.

For more information about configuring IBGP sessions, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide and the JUNOS Routing Protocols Configuration Guide.

To configure an IBGP session:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 12 on page 40 on each PE router.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the configuration, "Verifying a VPN Configuration" on page 52.
- 5. Go on to "Configuring Routing Options for VPNs" on page 40.

Table 12: Configuring an IBGP Session

Task	J-W	eb Configuration Editor	CLI	Configuration Editor
Navigate to the top of the configuration hierarchy and configure the IGBP session.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	1.	From the [edit] hierarchy level, enter
(PE Services Router)	2.	Next to Bgp, click Configure or Edit .	2.	edit protocols bgp group group-name
	3.	Next to Group, click Add New Entry.	~.	
	4.	Type a name in the Group name box.		set type internal
	5.	From the Type list, select Internal.	3.	Enter
	6.	In the Local address box, type the local loopback IP address.		set local-address loopback-interface-ip-address
	7.	In the Family group, select L2vpn for a Layer 2 VPN or Inet vpn for a Layer 3 VPN	4.	Enter
	Q	Select Unicast		set farmy farmy-type unicast
	9.	Click OK .		Replace <i>family-type</i> with l2vpn for a Layer 2 VPN or inet–vpn for a Layer 3 VPN.
	10.	In the Neighbor group, click Add new entry	5.	Enter up.
	11.	In the Address box, type the loopback IP address of the neighboring PE router.	6.	Enter the loopback address of the neighboring PE router:
	12.	Click OK until you return to the BGP page.		set neighbor <i>ip-address</i>

Configuring Routing Options for VPNs

The only required routing option for VPNs is the autonomous system (AS) number. You must specify it on each router involved in the VPN.

To configure routing options for a VPN:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration task described in Table 13 on page 41.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the configuration, see "Verifying a VPN Configuration" on page 52
- 5. Go on to "Configuring an IGP and a Signaling Protocol" on page 41.

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Configure the AS number.	1. 2. 3. 4.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration. Next to Routing options, click Configure or Edit. In the AS number box, type the AS number. Click OK.	From the [edit] hierarchy level, enter set routing-options autonomous-system as-number

Table 13: Configuring Routing Options for a VPN

Configuring an IGP and a Signaling Protocol

The PE Services Routers and provider Services Routers must be able to exchange routing information. To enable this exchange, you must configure either an IGP such as OSPF or static routes on these routers. You must configure the IGP at the [edit protocols] level, not within the routing instance at the [edit routing-instances] level.

You can use LDP or RSVP between PE routers and between PE routers and provider routers, but not for interfaces between PE routers and CE routers. LDP routes traffic using IGP metrics. RSVP has traffic engineering that lets you override IGP metrics as needed. For more information about these protocols, see "Signaling Protocols Overview" on page 10.

Each PE Services Router's loopback address must appear as a separate route. Do not configure any summarization of the PE Services Router's loopback addresses at the area boundary.

For more information about configuring IGPs and static routes, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide and the JUNOS Routing Protocols Configuration Guide.

Configure the appropriate signaling protocol for your VPN:

- Configuring LDP for Signaling on page 41
- Configuring RSVP for Signaling on page 43

Configuring LDP for Signaling

You must configure LDP and OSPF on PE and provider routers. For more information about configuring OSPF see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

To configure LDP and OSPF:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 14 on page 42 on PE and provider router interfaces that communicate with a PE router or provider router.

For the protocols to work properly, you also must configure the MPLS address family for each interface that uses LDP or RSVP, as described previously in "Configuring Interfaces Participating in a VPN" on page 35.

- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the configuration, see "Verifying a VPN Configuration" on page 52.
- 5. Go on to "Configuring a VPN Routing Instance" on page 45.

Table 14: Configuring LDP and OSPF for Signaling

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Navigate to the top of the configuration hierarchy and specify the LDP	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter the following command for each interface you want to enable:
interfaces that	2.	Next to Ldp, click Configure or Edit .	edit protocols ldp interface interface-name
communicate with a PE	3.	Next to Interface, click Configure or Edit .	
and the loopback interface of the PE router.	4.	In the Interface name column, type <i>interface-name</i> .	
(PE and provider Services Routers)	5.	Click OK .	
	6.	Repeat Steps 4 and 5 for each interface you want to enable.	
(See the interface naming conventions in the J-series Services Router Basic LAN and WAN Access Configuration Guide.)			

Task	J-Web Configuration Editor			CLI Configuration Editor		
Configure OSPF for each	For	OSPF:	For	OSPF:		
interface that uses LDP.	1.	On the main Configuration page next to Protocols, click Configure or Edit .	1.	From the [edit] hierarchy level, enter the following command for each interface you		
configure at least one area	2.	Next to Ospf, click Configure or Edit .		want to enable:		
router's interfaces. An AS can be divided into	3.	For Layer 2 VPN or circuit, select Traffic engineering.		edit protocols ospf area 0.0.0.0 interface interface-name		
multiple areas. This example uses the backbone area 0.0.0.0 . (PE and provider Services Routers)	4.	Next to Area group, click Add new entry and add the area.	2.	For Layer 2 VPN or circuit, move up to the [edit protocols ospf] level and enter		
	5.	Next to Area group, select the area (0.0.0.0).		set traffic-engineering		
	6.	Next to Interface group, select Add new entry.				
	7.	In the Interface name box, type <i>interface-name</i> .				
	8.	Click OK .				
	9.	Repeat Steps 5 through 7 to enable additional interfaces.				
	10.	Click OK twice to return to the Protocols page.				

Table 14: Configuring LDP and OSPF for Signaling (continued)

Configuring RSVP for Signaling

You must enable RSVP for all connections that participate in the label-switched path (LSP) on PE and provider Services Routers. In addition, you must configure OSPF on various interfaces.

For more information about configuring OSPF see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

To configure RSVP and OSPF:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 15 on page 44 on each PE router and provider router, as specified.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the configuration, see "Verifying a VPN Configuration" on page 52.
- 5. Go on to "Configuring a VPN Routing Instance" on page 45.

Table 15: Configuring RSVP and OSPF for Signaling

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Navigate to the top of the		OSPF, follow these steps:	From the [edit] hierarchy level, enter the
configuration hierarchy and configure OSPF with traffic engineering support.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	following command for each interface you want to enable:
(PE Services Router)	2.	Next to Protocols, click Configure or Edit .	
	3.	Next to Ospf, click Configure or Edit .	
	4.	Select Traffic engineering , and then click Configure .	
	5.	Select Shortcuts.	
	6.	Click OK until you return to the Protocols page.	
Enable RSVP on interfaces that participate in the LSP.	1.	On the main Configuration page next to Protocols, click Configure or Edit .	From the [edit] hierarchy level, enter the following command for each interface you
(PE Services Router) Enable	2.	Next to Rsvp, click Configure or Edit.	want to enable:
interfaces on the source and destination points.	3.	In the Interface group, click Add New Entry .	edit protocols rsvp interface interface-name
(provider Services Router)	4.	Type an interface name.	
Enable interfaces that	5.	Click OK .	
the PE Services Routers.	6.	Repeat Steps 2 through 4 for each interface you want to enable.	
(See the interface naming conventions in the <i>J-series</i> <i>Services Router Basic LAN</i> <i>and WAN Access</i> <i>Configuration Guide.</i>)	7.	Click OK .	

Configuring a Layer 2 Circuit

Each Layer 2 circuit is represented by the logical interface connecting the local PE Services Router to the local CE Services Router. All Layer 2 circuits using a particular remote PE Services Router neighbor is identified by its IP address and is usually the endpoint destination for the LSP tunnel transporting the Layer 2 circuit.

You configure a virtual circuit ID on each interface. Each virtual circuit ID uniquely identifies the Layer 2 circuit among all the Layer 2 circuits to a specific neighbor. The key to identifying a particular Layer 2 circuit on a PE router is the neighbor address and the virtual circuit ID. Based on the virtual circuit ID and the neighbor relationship, an LDP label is bound to an LDP circuit. LDP uses the binding for sending traffic on that Layer 2 circuit to the remote CE router.

To configure a Layer 2 circuit:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 16 on page 45 on each PE router and provider router, as specified.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the configuration, see "Verifying a VPN Configuration" on page 52.

Task	J-Web Configuration Editor		CLI	CLI Configuration Editor	
Navigate to the top of the configuration hierarchy	1.	In the J-Web interface, select Configuration > View and Edit > Edit	1.	From the [edit] hierarchy level, enter	
and enable a Layer 2		Configuration.		edit protocols l2circuit neighbor	
interface.	2.	Next to Protocols, click Configure or Edit .		interface-name interface interface-name	
(DE Comvises Deuter)	3.	Next to L2circuit, click Configure or Edit .		For neighbor, specify the local loopback	
(PE Services Router)	4.	Next to Neighbor, click Add new entry.		address, and for interface, specify the interface name of the remote PE router.	
(See the interface naming conventions in the J-series Services Router Basic LAN and WAN Access Configuration Guide.)	5.	In the Neighbor box, enter the loopback address of the local router.	2.	Enter	
	6.	Next to Interface, click Add new entry.		set virtual-circuit-id id-number	
	7.	In the Interface box, type the interface name of the remote PE router.			
	8.	In the Virtual circuit id box, type an ID number.			
	9.	Click OK until you return to the Protocols page.			

Table 16: Configuring a Layer 2 Circuit

Configuring a VPN Routing Instance

You must configure a routing instance for each VPN on each PE Services Router participating in the VPN. The routing instance has the same name on each PE router. VPN routing instances need a route distinguisher to help BGP distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs. This section does not apply to Layer 2 circuit configurations.

Each routing instance that you configure on a PE router must have a unique route distinguisher. There are two possible formats:

- as-number:number, where as-number is an autonomous system (AS) number (a 2-byte value) in the range 1 through 65,535, and number is any 4-byte value. We recommend that you use an Internet Assigned Numbers Authority (IANA)-assigned, nonprivate AS number, preferably the ISP or the customer AS number.
- *ip-address:number*, where *ip-address* is an IP address (a 4-byte value) and *number* is any 2-byte value. The IP address can be any globally unique unicast address.

We recommend that you use the address that you configure in the router-id statement, which is a public IP address in your assigned prefix range.

The route target defines which route is part of a VPN. A unique route target helps distinguish between different VPN services on the same router. Each VPN also has a policy that defines how routes are imported into the VPN routing and forwarding (VRF) table on the router. A Layer 2 VPN is configured with import and export policies. A Layer 3 VPN uses a unique route target to distinguish between VPN routes.

To configure a VPN routing instance:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 17 on page 46 on each PE router.
- 3. If you are finished configuring the router, commit the configuration.
- To verify the configuration, see "Verifying a VPN Configuration" on page 52. 4.
- 5. Go on to "Configuring a VPN Routing Policy" on page 47.

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Navigate to the top of the configuration hierarchy and create the routing	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit routing-instances routing-instance-name
(PE Services Bouter)	2.	Next to Routing instances, click Configure or Edit .	
	3.	Next to Mpls, click Configure or Edit.	
	4.	In the Instance group, click Add New Entry .	
	5.	Type a name in the Instance name box.	
Specify a text description for the routing instance	In the Description box, type a description.		Enter
This text appears in the output of the show route instance detail command.			set description "text"
(PE Services Router)			
Specify the instance type, either 12vpn for Laver 2	Fro typ	m the Instance type list, select an instance e.	Enter
VPNs or vrf for Layer 3 VPNs.	51		set instance-typeinstance-type
(PE Services Router)			

Tal

Task	J-Web Configuration Editor	CLI Configuration Editor
Specify the interface of the remote PE Services Router.	 Next to Interface group, click Add New Entry. 	Enter
(PE Services Router)	2. In the Interface name box, enter <i>interface-name</i> .	set interface interface-name
(See the interface naming conventions in the J-series Services Router Basic LAN and WAN Access Configuration Guide.)	3. Click OK .	
Specify the route distinguisher.	In the Rd type box, enter a route distinguisher in the format as-number:numberor	Enter one of the following commands:
(PE Services Router)	ip-address:number.	 set route-distinguisheras-number:number set route-distinguisher <i>ip-address:number</i>
Specify the policy for the Layer 2 VRF table.	For the sample Layer 2 VPN configuration, which uses import and export policies:	For the sample Layer 2 VPN configuration, which uses import and export policies, enter
For the Layer 2 VPN example, the routing	1. Next to Vrf export group, select Add new entry.	set vrf-import import-policy-name vrf-export export-policy-name
policies are defined in "Configuring a Routing	2. In the Value box, type the export routing policy name.	
VPNs" on page 48.	3. Click OK .	
(PE Services Router)	4. Next to Vrf import group, click Add new entry.	
	5. In the Value box, type the import routing policy name.	
	6. Click OK .	
Specify the policy for the Layer 3 VRF table.	For the sample Layer 3 VPN configuration, which uses a route target:	For the sample Layer 3 VPN configuration, which uses a route target, enter
For the Layer 3 VPN	1. In the Vrf target box, click Configure .	set vrf-target target:community-id
example, the routing policy is defined in "Configuring a Routing Policy for Layer 3 VPNs" on page 51.	 In the Community box, type the community (target:community-id, where community-id is as-number:number or ip-address:number). 	Replace <i>community-id</i> with either of the following:
(PE Services Router)	3. Click OK .	■ ip-address:number

Table 17: Configuring a VPN Routing Instance (continued)

Configuring a VPN Routing Policy

Layer 2 and Layer 3 VPNs require a routing policy that describes which packets are sent and received across the VPN. Layer 2 circuits do not use a policy, and therefore, Layer 2 circuits send and receive all packets. For Layer 2 VPNs, the routing policy resides on the PE Services Routers. For the Layer 3 VPN example, the routing policy resides on the CE Services Routers.

This section contains the following topics. For more information about configuring routing policies, see "Configuring Routing Policies" on page 169 and the *JUNOS Routing Protocols Configuration Guide*.

- Configuring a Routing Policy for Layer 2 VPNs on page 48
- Configuring a Routing Policy for Layer 3 VPNs on page 51

Configuring a Routing Policy for Layer 2 VPNs

If the routing instance uses a policy for accepting and rejecting packets instead of a route target, you must specify the import and export routing policies and the community on each PE Services Router.

To configure a Layer 2 VPN routing policy on a PE Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 18 on page 48 and Table 19 on page 50 on each PE router.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the configuration, see "Verifying a VPN Configuration" on page 52.

Table 18: Configuring an Import Routing Policy for Layer 2 VPNs

Task	J-W	eb Configuration Editor	CLI Configuration Editor	
Navigate to the top of the configuration hierarchy and configure the import	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit policy-options policy-statement	
routing policy. (PE Services Router)	ng policy. 2. Next to Policy options, click Configure Edit .	import-policy-name		
,	3.	Next to Policy statement, click Add new entry.		
	4.	In the Policy name box, type the policy name—for example, import_vpn.		
Task	J-Web Configuration Editor	CLI Configuration Editor		
--	--	--	--	--
Define the term for	1. Next to Term group, click Add new ent	ry . 1. Enter		
(PE Services Router)	2. In the Term name box, type a term name—for example, 10 .	set termterm-name-accept from protocol bgp community community-name		
``````````````````````````````````````	3. Next to From, click <b>Configure</b> .	2. Enter		
	4. Click Add new entry.			
	5. Click <b>Protocol</b> and select <b>bgp</b> from the Value menu.	set termterm-name-accept then accept		
	6. Click <b>OK</b> .			
	7. Next to Community, click Add new ent	ry.		
	8. Type the <i>community-name</i> value in the Community Name box.			
	9. Click <b>OK</b> .			
	10. Next to Then, click <b>Configure</b> .			
	11. From the Accept reject list, select acce	pt.		
	12. Click <b>OK</b> until you are at the Policy statement page.			
Define the term for rejecting packets.	<ol> <li>Next to the Term group, click Add new entry.</li> </ol>	<b>v</b> Enter		
(PE Services Router)	2. In the Term name box, type a term name—for example, <b>20</b> .	set term term-name-reject then reject		
	3. Next to Then, click <b>Configure</b> .			
	4. From the Accept list, select <b>reject</b> .			
	5. Click <b>OK</b> until you return to the Policy options page.			

#### Table 18: Configuring an Import Routing Policy for Layer 2 VPNs (continued)

After configuring an import routing policy for a Layer 2 VPN, configure an export routing policy for the Layer 2 VPN. The export routing policy defines how routes are exported from the PE Services Router routing table. An export policy is applied to routes sent to other PE Services Routers in the VPN. The export policy must also evaluate all routes received over the routing protocol session with the CE Services Router. The export policy must also contain a second term for rejecting all other routes.

Task	J-Web Configuration Editor	CLI Configuration Editor		
Configure the export routing policy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter edit policy-options policy-statement		
(PE Services Router)	<ol> <li>Next to Policy options, click Configure or Edit.</li> </ol>	export-policy-name		
	3. Next to Policy statement, click <b>Add new</b> entry.			
	<ol> <li>In the Policy name box, type the policy name—for example, export_vpn.</li> </ol>			
Define the term for accepting packets.	1. Next to the Term group, click <b>Add new</b> entry.	1. Enter		
(PE Services Router)	<ol> <li>In the Term name box, type a term name—for example, 10.</li> </ol>	set termterm-name-accept from community add community-name		
	3. Next to From, click <b>Configure</b> .	2. Enter		
	4. Next to Community, click <b>Add new entry</b> .	set termterm-name-accept then accept		
	5. Type the <i>community-name</i> value in the Community Name box.			
	6. Click <b>OK</b> .			
	7. Next to Then, click <b>Configure</b> .			
	8. From the Accept reject list, select <b>accept</b> .			
	9. Click <b>OK</b> twice until you are at the Policy statement page.			
Define the term for rejecting packets.	1. Next to the Term group, click <b>Add new</b> entry.	1. Enter		
(PE Services Router)	<ol> <li>In the Term name box, type a term name—for example, 20.</li> </ol>	set termterm-name-reject from community add community-name		
	3. Next to Then, click <b>Configure</b> .	2. Enter		
	4. From the Accept reject list, select reject.	set termterm-name-reject then reject		
	5. Click <b>OK</b> until you return to the Policy options page.			

# Table 19: Configuring an Export Routing Policy for Layer 2 VPNs

Task	J-Web Con	figuration Editor	CLI Configuration Editor	
Define the community.	1. In the entry	Community group, click <b>Add new</b>	Type the following commands:	
(PE Services Router)	2. In the comm	Community name box, type a nunity name—for example, <b>VPN</b> .	communitycommunity-nametarget:as-number or ip-address:number	
	3. In the entry	Members group, click <b>Add new</b>		
	4. In the where or <i>ip-a</i>	Value box, type target:community-id, e community-id is as-number:number address:number.		
	5. Click option	<b>OK</b> until you return to the Policy ns page.		

## Table 19: Configuring an Export Routing Policy for Layer 2 VPNs (continued)

# **Configuring a Routing Policy for Layer 3 VPNs**

To configure a Layer 3 VPN routing policy on a CE Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 20 on page 51 on each CE Services Router.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the configuration, see "Verifying a VPN Configuration" on page 52.

Task	J-W	/eb Configuration Editor	CLI Configuration Editor	
Navigate to the top of the configuration hierarchy1.In the J-Web interface, selectConfiguration > View and Edit > Edit		In the J-Web interface, select Configuration > View and Edit > Edit	From the [edit] hierarchy level, enter	
and configure the routing policy for the loopback interface.		Configuration.	edit policy-options policy-statement policy-name	
	2.	Next to Policy options, click <b>Configure</b> or <b>Edit</b> .		
(CE Services Router)	3.	Next to Policy statement, click <b>Configure</b> or <b>Edit</b> .		
<ol> <li>In the Policy name box, type the policy name—for example, loopback.</li> </ol>		In the Policy name box, type the policy name—for example, loopback.		

#### Table 20: Configuring a Routing Policy for Layer 3 VPNs

Task	J-Web Configuration Editor	CLI Configuration Editor		
Define the term for	1. In the Term group, click Add new entry.	1. Enter		
accepting packets.	<ol> <li>In the Term name box, type a term name—for example, 1.</li> </ol>	set termterm-name-accept from protocol direct route-filter		
()	3. Next to From, click <b>Configure</b> .	local-loopback-address/netmask exact		
	4. Click protocol, then Add new entry.	2. Enter		
	5. Select <b>direct</b> from the Value menu, and click <b>OK</b> .	set termterm-name-accept then accept		
	7. Next to Route Filter, click <b>Add new entry</b>			
	8. Type <i>local-loopback-address/netmask</i> in the Address box.			
	9. Select <b>exact</b> from the Modifier list.			
	10. Click <b>OK</b> twice.			
	11. Next to Then, click <b>Configure</b> .			
	12. From the Accept reject list, select accept			
	13. Click <b>OK</b> until you are at the Policy statement page.			
Define the term for rejecting packets.	<ol> <li>Next to the Term group, click Add new entry.</li> </ol>	Enter		
(CE Services Router)	<ol> <li>In the Term name box, type a term name—for example, 2.</li> </ol>	set termterm-name-reject then reject		
	3. Next to Then, click <b>Configure</b> .			
	4. From the Accept reject list, select <b>reject</b> .			
	5. Click <b>OK</b> until you return to the Policy options page.			

#### Table 20: Configuring a Routing Policy for Layer 3 VPNs (continued)

# **Verifying a VPN Configuration**

To verify the connectivity of Layer 2 VPNs, Layer 3 VPNs, and Layer 2 circuits, use the **ping mpls** command. This command helps to verify that a VPN or circuit has been enabled. This command tests the integrity of the VPN or Layer 2 circuit connection between the PE Services Routers. It does not test the connection between a PE and a CE Services Router.

This section contains the following topics:

- Pinging a Layer 2 VPN on page 53
- Pinging a Layer 3 VPN on page 53
- Pinging a Layer 2 Circuit on page 53

## Pinging a Layer 2 VPN

To ping a Layer 2 VPN, use one of the following commands:

■ ping mpls l2vpn interfaceinterface-name

Ping an interface configured for the Layer 2 VPN on the PE router.

ping mpls l2vpn instance l2vpn-instance-name local-site-idlocal-site-id-number remote-site-idremote-site-id-number

Ping a combination of the Layer 2 VPN routing instance name, the local site identifier, and the remote site identifier to test the integrity of the Layer 2 VPN connection (specified by identifiers) between the two PE Services Routers.

## **Pinging a Layer 3 VPN**

To ping a Layer 3 VPN, use the following command:

ping mpls I3vpn I3vpn-nameprefixprefix <count count>

Ping a combination of a IPv4 destination prefix and a Layer 3 VPN name on the destination PE Services Router to test the integrity of the VPN connection between the source and destination Services Routers. The destination prefix corresponds to a prefix in the Layer 3 VPN. However, ping tests only whether the prefix is present in a PE VRF table.

## **Pinging a Layer 2 Circuit**

To ping a Layer 2 circuit, use one of the following commands:

■ ping mpls l2circuit interfaceinterface-name

Ping an interface configured for the Layer 2 circuit on the PE Services Router.

ping mpls l2circuit virtual-circuit<prefix> <virtual-circuit-id>

Ping a combination of the IPv4 prefix and the virtual circuit ID on the destination PE Services Router to test the integrity of the Layer 2 circuit between the source and destination Services Routers.

J-series[™] Services Router Advanced WAN Access Configuration Guide

# Chapter 4 Configuring CLNS VPNs

Connectionless Network Service (CLNS) is a Layer 3 protocol similar to IPv4 for linking hosts (end systems) with routers (intermediate systems) in an Open Systems Interconnection (OSI) network. CLNS and its related OSI protocols, Intermediate System-to-Intermediate System (IS-IS) and End System-to-Intermediate System (ES-IS), are International Organization for Standardization (ISO) standards.

You can configure Services Routers as provider edge (PE) routers within a CLNS network. CLNS networks can be connected over an IP MPLS network core using BGP and MPLS Layer 3 virtual private networks (VPNs). For more information, see RFC 2547, *BGP/MPLS VPNs*.

You can use either the J-Web configuration editor or CLI configuration editor to configure CLNS.

This chapter contains the following topics. For more information about CLNS, IS-IS, and ES-IS, see the *JUNOS Routing Protocols Configuration Guide*.

- CLNS Terms on page 55
- CLNS Overview on page 56
- Before You Begin on page 57
- Configuring CLNS with a Configuration Editor on page 57
- Verifying CLNS VPN Configuration on page 63

# **CLNS Terms**

Before configuring CLNS, become familiar with the terms defined in Table 21 on page 55.

### Table 21: CLNS Terms

Term	Definition
CLNS island	Typically one IS-IS level 1 area that is part of a single IGP routing domain. An island can contain more than one area. CLNS islands can be connected by virtual private networks (VPNs).
Connectionless Network Service (CLNS)	Layer 3 protocol similar to IPv4 for linking hosts (end systems) with routers (intermediate systems) in an Open Systems Interconnection (OSI) network, by using network service access points (NSAPs) instead of prefix addresses to specify hosts and routers.

### Table 21: CLNS Terms (continued)

Term	Definition
customer edge (CE) router	Router or switch in the customer's network that is connected to a service provider's provider edge (PE) router and participates in a Layer 3 VPN.
end system	A host in an Open Systems Interconnection (OSI) network.
End System-to-Intermediate System (ES-IS)	Protocol that enables end systems (hosts) and intermediate systems (routers) to discover each other, by a method similar to Address Resolution Protocol (ARP) discovery in an IPv4 network.
intermediate system	A router in an Open Systems Interconnection (OSI) network.
International Organization for Standardization (ISO)	Worldwide federation of standards bodies that promotes international standardization and published international agreements as International Standards.
network layer reachability information (NLRI)	Information about routes exchanged in update messages by Border Gateway Protocol (BGP) systems, to enable routers to determine the relationships among all known BGP autonomous systems.
network services access point (NSAP)	International Standards Organization (ISO) addressing method for identifying hosts (end systems) and routers (intermediate systems) at the data-link layer (Layer 3) in an Open Systems Interconnection (OSI) network. An NSAP is from 8 to 20 bytes long and consists of an area address, a system ID, and an NSAP selector (NSEL) byte.
Open Systems Interconnection (OSI)	Standard reference model for representing the way messages are transmitted between two points on a network.
provider edge (PE) router	Services Router in the service provider network that is connected to a customer edge (CE) device and participates in a virtual private network (VPN).
virtual private network (VPN)	Private data network that uses a public TCP/IP network, typically the Internet, while maintaining privacy with a tunneling protocol, encryption, and security procedures.

# **CLNS Overview**

CLNS uses network service access points (NSAPs), similar to IP addresses found in IPv4, to identify end systems (hosts) and intermediate systems (routers). ES-IS enables the hosts and routers to discover each other. IS-IS is the interior gateway protocol (IGP) that carries ISO CLNS routes through a network.

Depending on your network topology, one or more of the following components are needed to route within a CLNS environment:

ES-IS—Provides the basic interaction between CLNS hosts (end systems) and routers (intermediate systems). Using ES-IS, hosts advertise their ISO NSAP addresses and subnetwork point-of-attachment (SNPA) addresses to other routers and hosts attached to the subnetwork. The resolution of Layer 3 ISO NSAPs to Layer 2 SNPAs by ES-IS is equivalent to ARP within an IPv4 network.

If a CLNS island does not contain any end systems, you do not need to configure ES-IS on a Services Router.

- IS-IS extensions—Provide the basic IGP support for collecting intradomain routing information for CLNS destinations within a CLNS network. Routers learning host addresses through ES-IS can advertise them to other routers (intermediate systems) using IS-IS.
- Static routes—You can configure static routes to exchange CLNS routes within a CLNS island. You can use static routing with or without IS-IS.
- Border Gateway Protocol (BGP) extensions—BGP extensions allow BGP to carry CLNS VPN network layer reachability information (NLRI) between PE routers. Each CLNS route is encapsulated into a CLNS VPN NLRI and propagated between remote sites in a VPN.

For more information about CLNS, see the ISO 8473 standards. For more information about IS-IS, see the ISO 10589 standard. For more information about ES-IS, see the ISO 9542 standard.

# **Before You Begin**

Before you begin configuring CLNS, complete the following tasks:

- Configure IS-IS. See the JUNOS Routing Protocols Configuration Guide.
- Configure the network interfaces. See the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.
- If applicable, configure BGP and VPNs. See the *J-series Services Router Basic LAN* and WAN Access Configuration Guide and "Configuring Virtual Private Networks" on page 31.

# **Configuring CLNS with a Configuration Editor**

To configure CLNS on a Services Router, you must perform the first task and then one or more of the following tasks (depending on your network):

- Configuring a VPN Routing Instance (Required) on page 58
- Configuring ES-IS on page 59
- Configuring IS-IS for CLNS on page 60
- Configuring CLNS Static Routes on page 62
- Configuring BGP for CLNS on page 63



**NOTE:** Many of the configuration statements used in this section can be included at different hierarchy levels in the configuration. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

## **Configuring a VPN Routing Instance (Required)**

You typically configure ES-IS, IS-IS, and CLNS static routes using a VPN routing instance. For more information about routing instances, see "Configuring a VPN Routing Instance" on page 45.

To configure a VPN routing instance:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 22 on page 58.
- 3. Go on to one of the following tasks:
  - Configuring IS-IS for CLNS on page 60
  - Configuring CLNS Static Routes on page 62
  - Configuring BGP for CLNS on page 63
  - Verifying CLNS VPN Configuration on page 63

#### **Table 22: Configuring a VPN Routing Instance for CLNS**

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the top of the configuration hierarchy and create the routing instance aaaa.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [ <b>edit</b> ] hierarchy level, enter
	2.	Next to Routing instances, click <b>Configure</b> or <b>Edit</b> .	edit routing-instances aaaa
	3.	Next to Instance, click Add new entry.	
	4.	In the Instance name box, type aaaa.	
	5.	Click <b>OK</b> .	
Specify the instance type vrf for Layer 3 VPNs.	In	the Instance type list, select <b>vrf</b> .	Enter
			set instance-type vrf

Task	J-Web Configuration Editor	CLI Configuration Editor
Specify the interfaces that belong to the routing instance aaaa—for example, lo0.1, e1–2/0/0.0, and t1–3/0/0.0. (See the interface naming conventions in the <i>J</i> -series Services Router Basic LAN and WAN Access Configuration Guide.)	<ol> <li>Next to Interface, click Add New Entry.</li> <li>In the Interface name box, type lo0.1.</li> <li>Click OK.</li> <li>Next to Interface, click Add New Entry.</li> <li>In the Interface name box, type e1–2/0/0.0.</li> </ol>	Enter 1. set interface lo0.1 2. set interface e1-2/0/0.0 3. set interface t1-3/0/0.0
	<ol> <li>Click OK.</li> <li>Next to Interface, click Add New Entry.</li> <li>In the Interface name box, type t1–3/0/0.0.</li> <li>Click OK.</li> </ol>	
Specify the route distinguisher—for example, <b>10.255.245.1:1</b> .	In the Rd type box, type <b>10.255.245.1:1</b> .	Enter set route-distinguisher 10.255.245.1:1
Specify the policy for the Layer 3 VRF table—for example, target:11111:1.	<ol> <li>Next to Vrf target, click Configure.</li> <li>In the Community box, type target:1111:1.</li> <li>Click OK.</li> </ol>	Enter set vrf-target target:11111:1

#### Table 22: Configuring a VPN Routing Instance for CLNS (continued)

# **Configuring ES-IS**

If a Services Router is a PE router within a CLNS island that contains any end systems, you must configure ES-IS on the Services Router.

To configure ES-IS for the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
- 2. Perform the configuration tasks described in Table 23 on page 60.
- 3. If you are finished configuring the router, commit the configuration.
- 4. If applicable, go on to one of the following tasks:
  - Configuring IS-IS for CLNS on page 60
  - Configuring CLNS Static Routes on page 62
  - Configuring BGP for CLNS on page 63
  - Verifying CLNS VPN Configuration on page 63

#### Table 23: Configuring ES-IS

Task	J-V	/eb Configuration Editor	CLI Configuration Editor	
Navigate to the <b>Routing</b> instances level in the	1.	In the J-Web interface, select <b>Configuration &gt; View and</b> <b>Edit &gt; Edit Configuration</b> .	From the [ <b>edit</b> ] hierarchy level, enter	
configuration hierarchy.	2.	Next to Routing instances, click Configure or Edit.	edit routing-instances aaaa	
	3.	Under Instance name, click <b>aaaa</b> .		
Enable ES-IS on all interfaces.	1.	Next to Protocols, click Configure.	Enter	
		Next to Esis, click Configure.	set protocols esis interface all	
	3. N	Next to Interface, click Add new entry.		
	4.	In the Interface name box, type all.		
	5.	Click <b>OK</b> until you return to the Protocols statement page.		

# **Configuring IS-IS for CLNS**

You can configure IS-IS to exchange CLNS routes within a CLNS island. To export BGP routes into IS-IS, you must configure and apply an export policy. For more information about policies, see "Configuring Routing Policies" on page 169.

If you have a pure CLNS island—an island that does not contain any IP devices—you must disable IPv4 and IPv6 routing.

To configure IS-IS for CLNS:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
- 2. Perform the configuration tasks described in Table 24 on page 60.
- 3. If you are finished configuring the router, commit the configuration.
- 4. If applicable, go on to one of the following tasks:
  - Configuring CLNS Static Routes on page 62
  - Configuring BGP for CLNS on page 63
  - Verifying CLNS VPN Configuration on page 63

#### Table 24: Configuring IS-IS to Exchange CLNS Routes

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Routing</b> instances level in the	1.	In the J-Web interface, select <b>Configuration &gt; View</b> <b>and Edit &gt; Edit Configuration</b> .	From the [edit] hierarchy level, enter
configuration hierarchy.	2.	Next to Routing instances, click <b>Configure</b> or <b>Edit</b> .	edit routing-instances aaaa
	3.	Under Instance name, click <b>aaaa</b> .	

Task	J-Web Configuration Editor		CLI Configuration Editor
Enable CLNS routing.	1.	Next to Protocols, click <b>Configure</b> .	Enter
	2.	Next to Isis, click Configure.	set protocols isis clns-routing
	3.	Next to CLNS routing, select the <b>Yes</b> box.	eer processie is e cuite rooming
Enable IS-IS on all interfaces.	1.	Next to Interface, click Add new entry.	Enter
(See the interface naming	2.	In the Interface name box, type all.	set protocols isis interface all
conventions in the J-series Services Router Basic LAN and WAN Access Configuration Guide.)	3.	Click <b>OK</b> .	
(Optional) To configure a pure	1.	Next to No ipv4 routing, select the <b>Yes</b> box.	Enter
CLNS network, disable IPv4 and IPv6 routing.	2.	Next to No ipv6 routing, select the $\mathbf{Yes}$ box.	set protocols isis no-ipv4-routing
0	3.	Click <b>OK</b> .	no-ipv6-routing
Define the BGP export policy name—for example,	1.	On the main Configuration page next to Policy options, click <b>Configure</b> or <b>Edit</b> .	From the <b>[edit]</b> hierarchy level, enter
dist-bgp—and the family and protocol.	2.	Next to Policy statement, click Add new entry.	set policy-options
	3.	In the Policy name box, type dist-bgp.	policy-statement dist-bgp
	4.	Next to From, click Configure.	from family iso protocol bgp
	5.	In the Family list, select <b>iso</b> .	
	6.	Next to Protocol, click Add new entry.	
	7.	In the Value list, select <b>bgp</b> .	
	8.	Click <b>OK</b> until you return to the Policy statement page.	
Define the action for the export	1.	Next to Then, click <b>Configure</b> .	From the [edit] hierarchy level,
policy.	2.	In the Accept reject list, select <b>accept</b> .	enter
	3.	Click <b>OK</b> until you return to the main Configuration page.	set policy-options policy-statement dist-bgp then accept
Apply the export policy to IS-IS.	1.	On the main Configuration page next to Routing instances, click <b>Configure</b> or <b>Edit</b> .	From the <b>[edit]</b> hierarchy level, enter
	2.	Next to aaaa, click <b>Protocols</b> .	set routing-instances agaa
	3.	Next to Isis, click <b>Edit</b> .	protocols isis export dist-bgp
	4.	Next to Export, click Add new entry.	
	5.	In the Value box, type dist-bgp.	
	6.	Click <b>OK</b> until you return to the Instance page.	

# Table 24: Configuring IS-IS to Exchange CLNS Routes (continued)

# **Configuring CLNS Static Routes**

If some devices in your network do not support IS-IS, you must configure CLNS static routes. You might also consider using static routes if your network is simple.

This procedure, as well as the configuration provided in "Verifying CLNS VPN Configuration" on page 63, uses the following ISO NET address and NSAP prefix:

- 47.0005.80ff.f800.0000.aaaa.1000.1921.6800.4196.00
- 47.0005.80ff.f800.0000.bbbb.1022/104

To configure CLNS static routes:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
- 2. Perform the configuration tasks described in Table 25 on page 62.
- 3. If you are finished configuring the router, commit the configuration.
- 4. If applicable, go on to one of the following tasks:
  - Configuring BGP for CLNS on page 63
  - Verifying CLNS VPN Configuration on page 63

#### **Table 25: Configuring Static CLNS Routes**

Task	J-Web Configuration Editor		CLI Configuration Editor		
Navigate to the <b>Routing</b> <b>instances</b> level in the configuration hierarchy.	1.	In the J-Web interface, select <b>Configuration &gt; View</b> and Edit > Edit Configuration.	From the [edit] hierarchy level, enter		
	2.	Next to Routing instances, click <b>Configure</b> or <b>Edit</b> .	edit routing-instances aaaa		
	3.	Under Instance name, click <b>aaaa</b> .			
Configure the	1.	Next to Routing options, click <b>Configure</b> .	Enter		
next-hop ISO NET address for an NSAP	<ol> <li>Next to</li> <li>In the I</li> </ol>	Next to Rib, click Add new entry.	set routing-options iso-route		
		In the Rib name box, type aaaa.iso.0.	47.0005.80ff.f800.0000.bbbb.1022/104 next-hop		
prenz.	4.	Next to Static, click <b>Configure</b> .	47.0005.80ff.f800.0000.aaaa.1000.1921.6800.4196.00		
	5.	Next to Iso route, click Add new entry.			
	6.	In the Destination box, type 47.0005.80ff.f800.0000.bbbb.1022/104.			
	7.	From the Next hop list, select Next hop.			
	8.	Next to Next hop, click Add new entry.			
	9.	In the Value box, type 47.0005.80ff.f800.0000.aaaa.1000.1921.6800.4196.00.			
	10.	Click <b>OK</b> .			

# **Configuring BGP for CLNS**

To configure BGP to carry CLNS VPN NLRI:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
- 2. Perform the configuration tasks described in Table 26 on page 63.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the configuration, see "Verifying CLNS VPN Configuration" on page 63.

#### Table 26: Configuring BGP to Carry CLNS VPN NLRI Messages

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Bgp</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; Vie and Edit &gt; Edit Configuration.</li> </ol>	• From the [edit] hierarchy level, enter
	2. Next to Protocols, click <b>Configure</b> or <b>Edit</b> .	set protocols bgp
	3. Next to Bgp, click <b>Configure</b> or <b>Edit</b> .	group pedge-pedge neighbor 10.255.245.215
Define a BGP group name—for	1. Next to Group, click <b>Add new entry</b> .	family iso-vpn unicast
example, <b>pedge-pedge</b> .	2. In the Group name box, type <b>pedge-pedge</b> .	
Define a BGP peer neighbor address	1. Next to Neighbor, click Add new entry.	
for the group—for example, 10.255.245.215.	2. In the Address box, type <b>10.255.245.215</b> .	
Define the family.	1. Under Family, next to Iso vpn, click <b>Configure</b> .	
	2. Next to Unicast, select the <b>Yes</b> box.	
	3. Click <b>OK</b> .	

# **Verifying CLNS VPN Configuration**

Verify that the Services Router is configured correctly for CLNS VPNs.

# **Displaying CLNS VPN Configuration**

- **Purpose** Verify the configuration of CLNS VPNs.
  - Action From the J-Web interface, select Configuration > View and Edit > View Configuration Text. Alternatively, from configuration mode in the CLI, enter the show command.

[edit] user@host# **show** interfaces { e1-2/0/0.0 { unit 0 {

```
family inet {
 address 192.168.37.51/31;
 }
 family iso;
 family mpls;
 }
 }
 t1-3/0/0.0 {
 unit 0 {
 family inet {
 address 192.168.37.24/32;
 family iso;
 family mpls;
 }
 }
 100 {
 unit 0 {
 family inet {
 address 127.0.0.1/32;
 address 10.255.245.215/32;
 }
 family iso {
 address 47.0005.80ff.f800.0000.0108.0001.1921.6800.4215.00;
 }
 }
 unit 1 {
 family iso {
 address 47.0005.80ff.f800.0000.0108.aaa2.1921.6800.4215.00;
 }
 }
 }
}
routing-options {
 autonomous-system 230;
}
protocols {
 bgp {
 group pedge-pedge {
 type internal;
 local-address 10.255.245.215;
 neighbor 10.255.245.212 {
 family iso-vpn {
 unicast;
 }
 }
 }
 }
}
policy-options {
 policy-statement dist-bgp {
 from {
 protocol bgp;
 family iso;
 }
 then accept;
```

```
}
 }
 routing-instances {
 aaaa {
 instance-type vrf;
 interface lo0.1;
 interface e1-2/0/0.0;
 interface t1-3/0/0.0;
 route-distinguisher 10.255.245.1:1;
 vrf-target target:11111:1;
 routing-options {
 rib aaaa.iso.0 {
 static {
 iso-route 47.0005.80ff.f800.0000.bbbb.1022/104
 next-hop 47.0005.80ff.f800.0000.aaaa.1000.1921.6800.4196.00;
 }
 }
 }
 protocols {
 esis {
 interface all;
 }
 isis {
 export dist-bgp;
 no-ipv4-routing;
 no-ip64-routing;
 clns-routing;
 interface all;
 }
 }
 }
 }
What It Means
 Verify that the output shows the intended configuration of CLNS VPNs.
```

**Related Topics** For more information about the format of a configuration file, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

J-series[™] Services Router Advanced WAN Access Configuration Guide

# Chapter 5 Configuring IPSec for Secure Packet Exchange

IP security (IPSec) is a framework of open standards for securing Layer 3 IP communications by encrypting and authenticating all IP packets. You can use IPSec to protect one or more paths between a pair of hosts, between a pair of security gateways (such as J-series Services Routers), or between a Services Router security gateway and a host.

You can use either J-Web Quick Configuration or a configuration editor to configure IPSec.

This chapter contains the following topics. For more information about IPSec, see the *JUNOS System Basics Configuration Guide* and the *JUNOS Services Interfaces Configuration Guide*.

- IPSec Terms on page 67
- IPSec Overview on page 69
- Before You Begin on page 73
- Configuring an IPSec Tunnel with Quick Configuration on page 73
- Configuring IPSec with a Configuration Editor on page 75
- Verifying the IPSec Tunnel Configuration on page 98

# **IPSec Terms**

To understand IPSec, you must be familiar with the terms defined in Table 27 on page 67.

### Table 27: IPSec Terms

Term	Definition
Advanced Encryption Standard (AES)	Encryption algorithm that uses a fixed block size of 128 bits, key sizes of 128, 192, or 256 bits, and multiple rounds of processing to encrypt data.
Authentication Header (AH)	Component of the IPSec protocol used to verify that the contents of a data packet have not changed, and to validate the identity of the sender. See also <i>ESP</i> .

# Table 27: IPSec Terms (continued)

Term	Definition
certificate	Secure electronic identifier conforming to the X.509 standard, definitively identifying an individual, system, company, or organization. In addition to identification data, the digital certificate contains a serial number, a copy of the certificate holder's public key, the identity and digital signature of the issuing certificate authority (CA), and an expiration date.
certificate authority (CA)	Third-party organization or company that issues digital certificates used to create digital signatures and public-private key pairs. The CA guarantees the identity of the individual or device that presents the digital certificate.
certificate revocation list (CRL)	Document maintained and published by a CA that lists revoked or suspended certificates.
Data Encryption Standard (DES)	Encryption algorithm that uses a 64-bit key (56 bits for encryption and 8 bits for error checking) to encrypt data. DES is considered a legacy method and insecure for many applications. See <i>3DES</i> and <i>AES</i> .
Diffie-Hellman (DH) protocol	Asymmetric cryptographic key agreement protocol developed by Diffie and Hellman in 1976. The protocol enables two users to exchange a secret key over an insecure medium without any prior secrets. Diffie-Hellman is used by the IKE protocol.
digital signature	A digital code that is attached to an electronically transmitted message to uniquely identify the sender.
Encapsulating Security Payload (ESP)	A protocol for securing packet flows for IPSec using encryption, data integrity checks, and sender authentication, which are added as a header to an IP packet. If an ESP packet is successfully decrypted, and no other party knows the secret key the peers share, the packet was not wiretapped in transit. See also <i>AH</i> .
Hashed Message Authentication Code (HMAC)	Method for message authentication that uses cryptographic hash functions. HMAC can be used with any iterative cryptographic hash function, such as MD5 or SHA-1, in combination with a secret shared key. The cryptographic strength of HMAC depends on the properties of the underlying hash function.
Internet Key Exchange (IKE)	Protocol that provides authentication of the IPSec peers, negotiates security associations (SAs), and establishes IPSec keys.
IP security (IPSec)	Framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. The secure aspects of IPSec are usually implemented in three parts: the Authentication Header (AH), the Encapsulating Security Payload (ESP), and the Internet Key Exchange (IKE).
Message Digest 5 (MD5)	Authentication algorithm that takes a data message of arbitrary length and produces a 128-bit message digest.
Perfect Forward Secrecy (PFS)	Key-establishment protocol used to secure VPN communications. A property which ensures that the compromise of an encryption key does not compromise security of previous or future encrypted sessions, because new keys are negotiated for each exchange and keys are securely deleted after use.
public key infrastructure (PKI)	Framework for public key cryptography on which other applications and network security components are built.
replay attack	Type of network attack in which valid data is maliciously transmitted repeatedly.

#### Table 27: IPSec Terms (continued)

Term	Definition
security association (SA)	In IPSec, an agreement between two network devices about what rules to use for authentication and encryption algorithms, key exchange mechanisms, and secure communications.
security parameter index (SPI)	Unique identifier for a security association (SA) at a network host or routing platform.
Secure Hash Algorithm 1 (SHA-1)	Authentication algorithm that takes a data message of less than 264 bits and produces a 160-bit message digest. SHA-1 is the most commonly used cryptographic function in the SHA family of authentication algorithms.
triple Data Encryption Standard (3DES)	Enhanced DES algorithm that provides 168-bit encryption by processing data three times with three different keys.

# **IPSec Overview**

Designed to address the lack of built-in security for IP traffic in the TCP/IP protocol suite, IPSec provides network-level data integrity, data confidentiality, data origin authentication, and protection from replay. IPSec can protect any protocol running over IP on any medium or a mixture of application protocols running on a complex combination of media.

This overview includes the following topics:

- Authentication and Encryption Algorithms in IPSec on page 69
- Authentication Methods in IPSec on page 70
- Traffic Protection in IPSec on page 71
- Security Associations on page 72
- Dynamic Security Associations and IKE Protocol on page 72
- IPSec Modes on page 73

# Authentication and Encryption Algorithms in IPSec

IPSec uses two types of algorithms: authentication algorithms and encryption algorithms.

IPSec authentication algorithms use a shared key to verify the identity of the sending IPSec device. The IPSec protocol suite defines two authentication algorithms: MD5 and SHA-1. The Services Router uses an HMAC variant of MD5 and SHA-1 algorithms that provide an additional level of hashing.

In an IPSec-enabled network, the Services Router that sends an IP packet computes a MD5 or SHA-1 digital signature, and adds this digital signature to the packet. The Services Router that receives the packet computes the digital signature and compares it with the signature stored in the packet's header. If the digital signatures match, the packet is authenticated. Encryption encodes data into a secure format so that it cannot be deciphered by unauthorized users. Like authentication algorithms, encryption algorithms use a shared key to verify the authenticity of the IPSec devices. The Services Router uses the following encryption algorithms:

- Data Encryption Standard-cipher block chaining (DES-CBC)
- Triple Data Encryption Standard-cipher block chaining (3DES-CBC)
- Advanced Encryption Standard (AES)

## **Authentication Methods in IPSec**

The IPSec implementation in the Services Router allows you to use one of two authentication methods: preshared keys or digital certificates.

When you configure IPSec for secure communications in the network, the peer devices in the network must have at least one common authentication method. Only one authentication method can be used between a pair of devices, regardless of the number of authentication methods configured.

## **Preshared Keys**

Preshared keys are secret passwords shared by the peer devices in an IPSec-enabled network. You must configure these keys on each Services Router in the network before any communication can take place. You can configure the preshared keys on each device manually and use protocols such as IKE to manage the keys dynamically.

## **Digital Certificates**

Certificates are digital identifiers that validate the authenticity of an individual or a device. A digital certificate implementation uses the public key infrastructure (PKI), which requires you to generate a key pair consisting of a public key and a private key. Certificates are issued by certificate authorities (CAs), which are public or private organizations that manage a PKI.

The main function of a digital certificate is to associate a device or user with a public-private key pair. Digital certificates also verify the authenticity of data and indicate privileges and roles within secure communication. A digital certificate consists of data that definitively identifies an individual, system, company, or organization. In addition to identification data, the digital certificate contains a serial number, a copy of the certificate holder's public key, the identity and digital signature of the issuing CA, and an expiration date.

**NOTE:** We recommend that you become familiar with PKI and digital certificates before implementing this feature on a Services Router.

For white papers about digital certificates and additional information about PKI, see the following Web sites:

- http://www.verisign.com
- http://www.thawte.com
- http://www.entrust.com

#### **Certificate Revocation Lists (CRLs)**

During the course of business, circumstances such as the following cause a certificate to become invalid before the validity period expires:

- Change of name
- Change of association between the subject and CA
- Compromise or suspected compromise of the corresponding private key

When events like these occur, the CA revokes or suspends a certificate. Revoked certificates are permanently deactivated, whereas suspended certificates can be reactivated later. Each CA periodically issues a list of revoked certificates, called Certificate Revocation Lists (CRLs). Each revoked certificate is identified in a CRL by the serial number of the certificate. You can automatically access the CA's CRL online at daily, weekly, or monthly intervals or at the default interval set by the CA.

You can configure the Services Router to check the CRLs at specified intervals to verify the validity of certificates. You can download CRLs either automatically using the Lightweight Directory Access Protocol (LDAP) or manually. Only Microsoft and Entrust CAs are supported. For more information about configuring CRLs, see the *JUNOS Services Interfaces Configuration Guide*.

## **Traffic Protection in IPSec**

Ê

IPSec provides a set of cryptographic protections for IP traffic. To provide security for the Layer 3 traffic, IPSec defines two protocols: Authentication Header (AH) and Encapsulating Security Payload (ESP). These protocols provide data and identity protection for each IP packet.

The AH protocol provides data origin authentication, data integrity, and antireplay protection for the entire IP packet, except for the fields in the IP header that are allowed to change in transit. AH protocol does not provide encryption. AH protocol is useful when the requirement is only to verify data integrity, but not to maintain data confidentiality.

The ESP protocol provides data confidentiality with encryption, data origin authentication, data integrity, and antireplay protection. ESP protocol can be implemented without encryption also. Although ESP provides an adequate level of authentication and encryption, it does so only for part of the IP packet, and excludes the IP header.

In addition to AH and ESP, the Services Router allows you to use a hybrid of AH and ESP protocols for protecting traffic. The hybrid of AH and ESP protocols, known as a protocol bundle, allows you to combine the benefits of both protocols and overcome their shortcomings.

## Security Associations

A security association (SA) is a set of IPSec specifications negotiated between devices that are establishing an IPSec relationship. These specifications include preferences for the type of authentication and encryption, and the IPSec protocol that is used to establish the IPSec connection. A security association is uniquely identified by a security parameter index (SPI), an IPv4 or IPv6 destination address, and a security protocol (AH or ESP).

IPSec security associations are established either manually through configuration statements, or dynamically by Internet Key Exchange (IKE) negotiation. In the case of manually configured security associations, the connection is established when both ends of the tunnel are configured, and the connections last until one of the endpoints is taken offline. In the case of dynamic security associations, you can configure when connections are to be established; immediately after both ends of the tunnel are configured, or only when traffic is sent through the tunnel, and dissolve after a preset amount of time or traffic. You can configure unidirectional security associations (one security association for both incoming and outgoing traffic).

### **Dynamic Security Associations and IKE Protocol**

When you deploy and use IPSec on a large scale in the network, manually managing the security associations (SAs) and keys on each device in the network is not practical. You can configure dynamic SAs in such scenarios so that authentication and key negotiation are automated.

To use dynamic SAs in a Services Router, you must configure the Internet Key Exchange (IKE) protocol and IPSec settings under the IPSec-VPN service configuration. IPSec uses the IKE protocol to dynamically negotiate the security association settings and exchange keys.

The IKE negotiation in a Services Router takes place in two phases. Phase 1 establishes a secure channel between the key management processes on the two peers, and phase 2 directly negotiates IPSec security associations. During phase 1, the peers negotiate at minimum an authentication method, an encryption algorithm, a hash algorithm, and a Diffie-Hellman group to create a phase 1 security association. The peers use this information to authenticate each other and compute key material to use for protecting phase 2. Phase 2, also called quick mode, results in an IPSec tuple, one security association for incoming traffic and another for outgoing traffic

Optionally, you can enable perfect forward secrecy (PFS) security for keys so that a shared key is used only once in phase 2 negotiation. PFS requires a Diffie-Hellman exchange to generate the shared key information for each new key.

# **IPSec Modes**

An IPSec mode describes how the original IP packet is transformed into a protected packet. IPSec supports two modes of secure communication: transport mode and tunnel mode.

Transport mode provides a security association (SA) between two hosts. In transport mode, the protocols provide protection primarily for upper-layer protocols.

Tunnel mode helps protect an entire IP packet by treating it as an AH or ESP payload. In tunnel mode, an IP packet is encapsulated with an AH or an ESP header and an additional IP header. The IP addresses of the outer IP header are the local tunnel endpoint and the remote tunnel endpoint. Packets with a destination address matching the private network prefix are encrypted and encapsulated in a tunnel packet that is routable through the outside network. The source address of the tunnel packet is the local gateway, and the destination address is the remote gateway. The IP addresses of the encapsulated IP header are the original source and final destination addresses. Once the encapsulation packet reaches the other side, the remote end determines how to route the packet.

When one side of a security association is a Services Router operating as a security gateway, the security association must use tunnel mode. However, when traffic (for example, SNMP commands or BGP sessions) is destined for the Services Router, the system acts as a host. Transport mode is allowed in this case because the system does not act as a security gateway and does not send or receive transit traffic.

## **Before You Begin**

Before you begin configuring IPSec, you must have completed these tasks:

- Establish basic connectivity. See the Getting Started Guide for your router.
- Configure network interfaces. See the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.
- Configure one or more routing protocols. See the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.
- Ensure that you have connectivity between the two routers in the network segment, and also that the traffic is routed through the routers on which the IPSec tunnel is configured. For example, if you want to send traffic from Router R1 to Router R4 through an IPSec tunnel set up between Routers R2 and R3, you must ensure that connectivity exists between R1 and R4, with traffic passing through R2 and R3.

### **Configuring an IPSec Tunnel with Quick Configuration**

J-Web Quick Configuration allows you to create IPSec tunnels. Figure 7 on page 74 shows the Quick Configuration page for IPSec tunnels.

	ROUTER - J4300						
Monitor Configuration	n Diagnose	Manage	Events	Logged in as: regres	s Help	About	Logout
Quick Configuration 💦 🚴				Configuration >	Quick Configur	ation > IPS	ec Tunnels
View and Edit 🕨 🎽	Quick Cor	nfigurati	on				
History	IPSec Tu	Innels			Add an I	PSec T	unnel
Rescue							
	Tunnel In	formatio	on				
	+ Local 1	Tunnel En	dpoint		•		
	* Remote	Tunnel En	Idpoint		2		
	+	IKE Secr	et Key		\$		
	* Verify	IKE Secr	et Key				
	Pri	ivate Pref	ix List		?		
				Add Delete			
				Add Delete			
	OK	Cancel					
		Cancer					
						un la cett	LANK NAL
Copyright © 2004-2005, J Privacy.	uniper Networks	s, Inc. <u>All Ri</u>	phis Reser	ved. Trademark Notice.	,	uniper	our wet.

#### **Figure 7: Quick Configuration Page for IPSec Tunnels**

To configure an IPSec tunnel with Quick Configuration:

- 1. In the J-Web interface, select **Configuration > Quick Configuration > IPSec Tunnels**.
- 2. In the IPSec Tunnels Quick Configuration main page, click Add.
- 3. Enter information into the Quick Configuration page for IPSec Tunnels, as described in Table 28 on page 75.
- 4. From the IPSec Tunnels Quick Configuration main page, click one of the following buttons:
  - To apply the configuration and stay on the IPSec Tunnels Quick Configuration page, click **Apply**.
  - To apply the configuration and return to the Quick Configuration main page, click **OK**.
  - To cancel your entries and return to the Quick Configuration main page, click **Cancel**.
- 5. To use digital certificates for authentication, see "Configuring Digital Certificates for IPSec Tunnels" on page 91.
- 6. To check the configuration, see "Verifying the IPSec Tunnel Configuration" on page 98.

Field	Function		Your Action		
<b>Tunnel Information</b>					
Local Tunnel Endpoint (required)	Externally routable IP address that is the local endpoint of the IPSec tunnel		Type the IPSec tunnel's local endpoint 32-bit IP address, in dotted decimal notation.		
Remote Tunnel Endpoint (required)	Externally routable IP address that is the peer endpoint of the IPSec tunnel		Type the IPSec tunnel's peer endpoint 32-bit IP address, in dotted decimal notation.		
IKE Secret Key (required)	Internet Key Exchange key (password) that is preshared to ensure authentication across the IPSec tunnel		Type the IKE key to be used for authentication across the IPSec tunnel. Characters are disguised as you type.		
Verify IKE Secret Key (required)	Internet Key Exchange key that is preshared to ensure authentication across the IPSec tunnel	Verify the IKE key by retyping the key to b used for authentication across the IPSec tur Characters are disguised as you type.			
Private Prefix List	List of addresses or address prefixes for which the IPSec tunnel is used. Packets whose destination address matches any of the addresses or prefixes in this list are transported through the		In the text box at the bottom of the list, type an IP address or address prefix. For example:		
	IPSec tunnel to the remote tunnel endpoint.		10.10.10/24		
		2.	Click Add.		
		3.	Click <b>OK</b> .		

#### **Table 28: IPSec Tunnels Quick Configuration Summary**

# **Configuring IPSec with a Configuration Editor**

To configure a Services Router to transport traffic across a secure IPSec connection, you can define the IPSec tunnel with security associations (SAs), services interfaces, IPSec tunnel endpoints, and IPSec rules to direct traffic to the tunnel.

In a network consisting of Services Routers, you can define manual SAs or dynamic SAs. Manual SAs require you to configure all security parameters of the security association, such as authentication and encryptions algorithms, encryptions keys, and the protocols, in the Services Routers at the tunnel endpoints. Dynamic SAs require you to configure the IKE protocol to manage the negotiation and exchange of encryption keys.

For a security association, you can optionally define NAT pools to hide IP addresses from the Internet.

This section contains the following topics:

- Configuring IPSec Manual Security Associations on page 76
- Configuring IPSec Dynamic Security Associations on page 77
- Configuring a NAT Pool on page 90
- Configuring Digital Certificates for IPSec Tunnels on page 91

## **Configuring IPSec Manual Security Associations**

To configure a manual security association (SA) in a Services Router, you must configure an IPSec-VPN rule and specify all the parameters such as authentication and encryptions algorithms, protocols, security parameter index (SPI), and the authentication and encryption keys required for the security association on the Services Routers at both tunnel endpoints. The sample configuration in Table 29 on page 76 configures a manual SA that applies to all inbound traffic on a Services Router.

Repeat the same procedure to define another rule for oubound traffic with the same parameters. Configure a manual SA with the same parameters, authentication and encryption keys, and security parameter index (SPI) on the Services Router at the other endpoint of the tunnel.

To configure a manual SA:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 29 on page 76.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify that IPSec is configured correctly, see "Verifying the IPSec Tunnel Configuration" on page 98.

#### **Table 29: Configuring IPSec Manual SAs**

Task	J-Web Configuration Editor	CLI Configuration Editor		
Navigate to the <b>Services &gt; Ipsec vpn</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter edit services ipsec-vpn		
	2. Next to Services, click <b>Configure</b> or <b>Edit</b> .			
	3. Next to Ipsec vpn, click <b>Configure</b> .			
Configure a rule—for example,	1. Next to Rule, click <b>Add new entry</b> .	Enter		
manualSARule—that applies to all incoming traffic.	2. In the Rule name box, type manualSARule.	set rule manualSARule match-direction input		
	3. In the Match direction box, select <b>input</b> .			
Configure a term ⁻ —for example,	1. Next to Term, click Add new entry.	1. Enter		
manualSATerm—for the rule, and the remote gateway for the IPSec	2. In the Term name box, type manualSATerm.	edit rule manualSARule		
tunnel—for example, 10.90.90.1.	3. Next to Then, select the check box, and click <b>Configure</b> .	2. Enter		
	4. In the Remote gateway box, type <b>10.90.90.1</b> .	set term manualSATerm then remote-gateway 10.90.90.1		

Task	J-Web Configuration Editor	<b>CLI Configuration Editor</b>		
Configure the manual SA, and specify	1. In the Sa choice box, select Manual.	1. Enter		
the direction of traffic to which the SA is applicable—for example, <b>bidirectional</b> .	2. Next to Manual, click <b>Configure</b> .	edit term manualSATerm		
	3. Next to Direction, click <b>Add new entry</b> .	then		
	4. In the Direction box, select <b>bidirectional</b> .	2. Enter		
		set manual direction bidirectional		
Configure the security parameter index	1. In the Spi box, type <b>1024</b> .	1. Enter		
(SPI)—for example, <b>1024</b> —and the IPSec protocol—for example, <b>esp</b> .	2. In the Protocol box, select <b>esp</b> .	edit manual direction bidirectional		
		2. Enter		
		set spi 1024 protocol esp		
Configure the authentication	1. Next to Authentication, click <b>Configure</b> .	Enter		
algorithm—for example, hmac-md5-96—and an authentication	2. In the Algorithm box, select <b>hmac-md5-96</b> .	set authentication algorithm		
key—for example, juniper—to be used while establishing the manual SA	3. Next to Key, click <b>Configure</b> .	hmac-md5-96 key ascii-text		
while establishing the manual SA.	4. In the Key choice box, select <b>Ascii text</b> .	Juniper		
	5. In the Ascii text box, type juniper.			
	<ol> <li>Click <b>OK</b> until you return to the Direction page.</li> </ol>			
Configure an encryption algorithm—for	1. Next to Encryption, click Configure.	Enter		
example, <b>3des-cbc</b> —and an encryption key—for example, <b>juniper123</b> .	2. In the Algorithm box, select <b>3des-cbc</b> .	set encryption algorithm		
	3. Next to Key, click <b>Configure</b> .	3des-cbc key ascii-text		
	4. In the Key choice box, select <b>Ascii text</b> .	Juniper123		
	5. In the Ascii text box, type juniper123.			
	<ol> <li>Click <b>OK</b> until you return to the Ipsec vpn page.</li> </ol>			

#### Table 29: Configuring IPSec Manual SAs (continued)

# **Configuring IPSec Dynamic Security Associations**

Dynamic SAs require you to configure the IKE protocol, which manages the negotiation and exchange of encryption keys. Configuring a dynamic SA involves setting up an IKE IPSec tunnel, which can be activated either on completion of the configuration or when the traffic flow starts. To establish an IKE IPSec tunnel, two phases of negotiation are required:

 In Phase 1, the participants establish a secure channel to negotiate the IPSec SAs. ■ In Phase 2, the participants negotiate the IPSec SAs for encrypting and authenticating the exchanges of user data.

To configure an IPSec dynamic SA, you must complete the following tasks in the Services Routers at both tunnel endpoints:

- Configuring an IKE Proposal on page 78
- Configuring an IKE Policy on page 80
- Configuring an IPSec Proposal on page 81
- Configuring an IPSec Policy on page 82
- Configuring IPSec Rules on page 83
- Configuring IPSec Services Interfaces on page 84
- Configuring Service Sets on page 86

## **Configuring an IKE Proposal**

An IKE proposal determines the authentication method, authentication and encryption algorithms, lifetime for the authentication and encryption keys, and the Diffie-Hellman group that determines the cryptographic strength of the key negotiation. You can configure one or more IKE proposals.

To configure an IKE proposal:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 30 on page 78.
- 3. Go on to "Configuring an IKE Policy" on page 80.

Task		eb Configuration Editor	CLI Configuration Editor	
Navigate to the <b>Services &gt; Ipsec vpn &gt; Ike</b> level in the configuration hierarchy.		In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [ <b>edit]</b> hierarchy level, enter	
	2.	Next to Services, click <b>Configure</b> or <b>Edit</b> .	edit services ipsec-vpn ike	
	3.	Next to Ipsec vpn, click <b>Configure</b> or <b>Edit</b> .		
	4.	Next to Ike, click Configure.		
Configure an IKE proposal—for example, ike-dynamic-proposal—that defines the	1.	Next to Proposal, click <b>Add new</b> entry.	Enter	
authentication method, authentication and encryption algorithms, and the lifetime of the keys.	2.	In the Name box, type ike-dynamic-proposal.	set proposal ike-dynamic-proposal	

#### **Table 30: Configuring IKE Proposal**

# Table 30: Configuring IKE Proposal (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the authentication algorithm—for example, <b>sha1</b> .	In the Authentication algorithm box, select <b>sha1</b> .	Enter
		set proposal ike-dynamic-proposal authentication-algorithm sha1
Configure the authentication method—for example, <b>pre-shared-keys</b> .	In the Authentication method box, select pre-shared-keys.	Enter
<b>NOTE:</b> Alternatively, you can use digital certificates as an authentication method. For details, see "Configuring Digital Certificates for IPSec Tunnels" on page 91.		set proposal ike-dynamic-proposal authentication-method pre-shared-keys
Configure the Diffie-Helman group to be used In the Dh group box, select <b>group1</b> .		Enter
		set proposal ike-dynamic-proposal dh-group group1
Configure an encryption algorithm—for example, <b>3des-cbc</b> .	In the Encryption algorithm box, select <b>3des-cbc</b> .	Enter
		set proposal ike-dynamic-proposal encryption-algorithm 3des-cbc
Configure the lifetime (in seconds) of the encryption and authentication keys—for	1. In the Lifetime seconds box, type 3600.	Enter
example, 3600.	2. Click <b>OK</b> until you return to the Configuration page.	set proposal ike-dynamic-proposal lifetime-seconds 3600

### **Configuring an IKE Policy**

An IKE policy defines a combination of security parameters (IKE proposals) to be used during IKE negotiation. The policy defines a peer address, the preshared key for the given peer, and the proposals needed for that connection. During the IKE negotiation, IKE searches 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 peer policies have a proposal that contains the same configured attributes. If the lifetimes are not identical, the shorter lifetime between the two policies is used. The configured preshared key must also match its peer.



**NOTE:** You can create an IKE access profile that uses the IKE policy 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. You can also use the digital certificate method for IKE authentication with dynamic peers. For more information about IKE access profiles, see the *JUNOS System Basics Configuration Guide*. For detailed information, see the *JUNOS Services Interfaces Configuration Guide*.

To configure an IKE policy:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 31 on page 80.
- 3. Go on to "Configuring an IPSec Proposal" on page 81.

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Services &gt; Ipsec</b> <b>vpn &gt; Ike</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter edit services ipsec-vpn ike
	2. Next to Services, click <b>Configure</b> or <b>Edit</b> .	
	3. Next to Ipsec vpn, click <b>Configure</b> .	
	4. Next to Ike, click <b>Configure</b> .	
Configure an IKE policy—for example, ike-dynamic-policy.	1. Next to Policy, click Add new entry.	Enter
	2. In the Name box, type ike-dynamic-policy.	set policy ike-dynamic-policy

#### **Table 31: Configuring IKE Policy**

Table	31:	Configuring	IKE	Policy	(continued)
-------	-----	-------------	-----	--------	-------------

Task	J-Web Configuration Editor	CLI Configuration Editor		
Configure a local ID for the policy—for	1. Next to Local id, click <b>Configure</b> .	Enter		
example, 10.90.90.2.	2. In the Id type box, select <b>Ipv4 addr</b> .	set policy ike-dynamic-policy local-id		
	<ol> <li>In the Ipv4 addr box, type 10.90.90.2.</li> </ol>	ipv4_addr 10.90.90.2		
Configure a remote ID for the policy—for	1. Next to Remote id click <b>Configure</b> .	Enter		
example, 10.90.90.1.	<ol> <li>Next to Ipv4 addr, click Add new entry.</li> </ol>	set policy ike-dynamic-policy remote-id ipv4_addr 10.90.90.1		
	3. In the Value box, type <b>10.90.90.1</b> .			
Configure a preshared key—for example, \$1991poPPi—for IKE in ASCII format.	<ol> <li>Next to Pre-shared key, click Configure.</li> </ol>	Enter		
<b>NOTE:</b> The IKE preshared key must be configured exactly the same way at both	<ol> <li>In the Key choice box, select Ascii text from the list.</li> </ol>	set policy ike-dynamic-policy pre-shared-key ascii-text \$1991poPPi		
the local and remote endpoints of the IPSec tunnel.	<ol> <li>In the Ascii text box, type the plain text IKE key \$1991poPPi</li> </ol>			
Configure the IKE proposal to be used for the IKE policy—for example,	1. Next to Proposals, click Add new entry.	Enter		
ike-aynamic-proposal.	<ol> <li>In the Value keyword, type ike-dynamic-proposal.</li> </ol>	set policy ike-dynamic-policy proposals ike-dynamic-proposal		
	<ol> <li>Click <b>OK</b> until you return to the main Configuration page.</li> </ol>			

# **Configuring an IPSec Proposal**

An IPSec proposal determines the authentication and encryption algorithms, lifetime for the authentication and encryption keys, and the protocols to be negotiated with the remote IPSec peer.

To configure an IPSec proposal:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 32 on page 82.
- 3. Go on to "Configuring an IPSec Policy" on page 82.

#### **Table 32: Configuring IPSec Proposal**

Task	J-Web Configuration Editor	CLI Configuration Editor	
Navigate to the <b>Services &gt; Ipsec vpn &gt; IPsec</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [ <b>edit</b> ] hierarchy level, enter	
	<ol> <li>Next to Services, click Configure or Edit.</li> </ol>	edit services ipsec-vpn ipsec	
	3. Next to Ipsec vpn, click <b>Configure</b> .		
	4. Next to Ipsec, click <b>Configure</b> .		
Configure an IPSec proposal—for example,	1. Next to Proposal, click Add new entry.	Enter	
authentication and encryption algorithms, the lifetime of the keys, and the protocol.	2. In the Name box, type ipsec-dynamic-proposal.	set proposal ipsec-dynamic-proposal	
Configure the authentication algorithm—for example, hmac-md5-96.	In the Authentication algorithm box, select <b>hmac-md5-96</b> .	Enter	
		set proposal ipsec-dynamic-proposal authentication-algorithm hmac-md5-96	
Configure an encryption algorithm—for example. 3des-cbc.	In the Encryption algorithm box, select <b>3des-cbc</b> .	Enter	
		set proposal ipsec-dynamic-proposal encryption-algorithm 3des-cbc	
Configure the lifetime (in seconds) of the encryption and authentication keys—for	In the Lifetime seconds box, type <b>3600</b> .	Enter	
example, 3600.		set proposal ipsec-dynamic-proposal lifetime-seconds 3600	
Configure the protocol to be used for key	1. In the Protocol box, select <b>esp</b> .	Enter	
negotiations—for example, <b>esp</b> .	2. Click <b>OK</b> until you return to the main Configuration page.	set proposal ipsec-dynamic-proposal protocol esp	

## **Configuring an IPSec Policy**

An IPSec policy defines a combination of security parameters (IPSec proposals) used during IPSec negotiation. 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.

To configure an IPSec policy:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.

- 2. Perform the configuration tasks described in Table 33 on page 83.
- 3. Go on to "Configuring IPSec Rules" on page 83.

#### **Table 33: Configuring IPSec Policy**

Task	J-Web Configuration Editor	CLI Configuration Editor	
Navigate to the <b>Services &gt; Ipsec</b> <b>vpn &gt; Ipsec</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter	
	2. Next to Services, click <b>Configure</b> or <b>Edit</b> .	edit services ipsec-vpn ipsec	
	3. Next to Ipsec vpn, click <b>Configure</b> .		
	4. Next to Ipsec, click <b>Configure</b> .		
Configure an IPSec policy—for example, ipsec-dynamic-policy.	<ol> <li>Next to Policy, click Add new entry.</li> <li>In the Name box, type ipsec-dynamic-policy.</li> </ol>	Enter set policy ipsec-dynamic-policy	
Configure the IPSec proposal to be	1. Next to Proposals, click Add new entry.	Enter	
used for the IPSec policy—for example, ipsec-dynamic-proposal.	<ol> <li>In the Value keyword, type ipsec-dynamic-proposal.</li> </ol>	set policy ipsec-dynamic-policy proposals ipsec-dynamic-proposa	
	<ol> <li>Click <b>OK</b> until you return to the main Configuration page.</li> </ol>		

## **Configuring IPSec Rules**

A rule defines a set of conditions that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. An IPSec rule specifies the traffic that you want to send through the IPSec tunnel using source and destination address combinations, and also specifies the IKE and IPSec policies to be applied on that traffic.

To configure an IPSec rule:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 34 on page 84.
- 3. Go on to "Configuring IPSec Services Interfaces" on page 84.

#### **Table 34: Configuring IPSec Rules**

Task		J-Web Configuration Editor		CLI Configuration Editor		
Navigate to the <b>Services &gt; Ipsec vpn</b> level in the configuration hierarchy.		<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> <li>Next to Services, click Configure or Edit.</li> </ol>		From the <b>[edit]</b> hierarchy level, enter		
				edit services ipsec-vpn		
		Next to Ipsec vpn, click <b>Configure</b> .				
Configure an IPSec rule named	1.	1. Next to Rule, click <b>Add new entry</b> .		Enter		
traffic.	2.	In the Rule name box, type ipsec-dynamic-rule.	set rule ipsec-dynamic-rule match-direction input			
	3.	In the Match direction box, select <b>Input</b> from the list.				
Configure a term—for example, term1,	1.	Next to Term, click Add new entry.	1.	Enter		
and a remote gateway—for example, <b>10.90.90.1</b> .	2.	In the Term name box, type term1.		edit rule ipsec-dynamic-rule		
<b>NOTE:</b> Because the rule applies to all traffic you configure only the action (or	3. Nex clic	Next to Then, select the <b>Yes</b> check box and click <b>Configure</b> .	2.	Enter		
then statement) for the term.	4.	In the Remote gateway box, type <b>10.90.90.1</b> .		set term term1 then remote-gateway 10.90.90.1		
Configure the IPSec rule	the IPSec rule 1.	In the Sa choice box, select <b>Dynamic</b> .	1.	Enter		
ipsec-dynamic-rule to reference the IKE policy ike-dynamic-policy and the IPSec	2.	Next to Dynamic, click <b>Configure</b> .		edit term term1.		
policy <b>ipsec-dynamic-policy</b> for the IPSec	3.	In the Ike policy box, type ike-dynamic-policy.	2.	Enter		
	4.	Click <b>OK</b> until you return to the main Configuration page.		set then dynamic ike-policy ike-dynamic-policy ipsec-policy ipsec-dynamic-policy		

## **Configuring IPSec Services Interfaces**

To enable IPSec on a Services Router, you must configure the services interfaces. In the Services Router, the service interface is always **sp-0/0/0**.*unit*. For the services to be applied, you must first define the logical interfaces to be used. The logical interface must have a unit number other than 0. By default, the J-Web interface uses the unit number **1001** for inside-service logical interfaces, and **2001** for outside-service logical interfaces.

To configure an IPSec tunnel, you must configure the following services interfaces:

 Inside services interface—Logical interface used to apply the service sets that define the behavior of the IPSec tunnel for outbound traffic (traffic whose next hop is inside the IPSec tunnel).
Outside services interface—Logical interface used to apply the service sets that define the behavior of the IPSec tunnel for inbound traffic (traffic whose next hop is outside the IPSec tunnel).

To configure IPSec inside services interfaces and outside services interfaces:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor..
- 2. Perform the configuration tasks described in Table 35 on page 85.
- 3. Go on to "Configuring Service Sets" on page 86.

### **Table 35: Configuring IPSec Service Interfaces**

Navigate to the Interfaces level in the configuration hierarchy.1.In the J-Web interface, select Configuration > View and Edit > Edit Configuration > View and Edit > Edit Configuration 2.From the [edit] hierarchy level, enter edit interfacesConfigure the inside services interface for the IPSec tunnel.1.Next to Interface, click Add new entry.1.Configure the services interface as an inside-service interface: set sp-0/0/0, and click OK.1.Configure the services interface as an inside-service interface: set sp-0/0/0 unit 1001 services Router Basic LANand WAN Access Configuration Guide.)1.Next to Unit, click Add new entry.1.Configure the services interface as an inside5.In the Interface unit number box, type 1001.1.Configure the services interface as an inet interface2.Configure the services interface as an inet interface6.In the Service domain box, select the check box next to inter and click Configure.1.Next to Interface, click sp-0/0/0.2.7.In the Interface, click sp-0/0/0.1.Next to Interface, click sp-0/0/0.3.3.8.Select the Primary check box, and click OK until you return to the Interface spage.1.Configure the services interface as an outside-service interface7.In the Interface, click sp-0/0/0.1.Configure the services interface as an outside-service interface8.Select the Primary check box, and click OK until you return to the Interface outside form the list.1.Configure the services interface as an outside-service interface9.In t	Task	J-W	J-Web Configuration Editor		CLI Configuration Editor		
2.Next to Interfaces, click Configure or Edit.Configure the inside services interface for the IPSec tunnel, (See the interface naming conventions in the <i>J-services</i> <i>Services Router Basic LAN and WAN Access Configuration</i> <i>Guide.</i> )1.Next to Interface, click Add new entry. 2.1.Configure the services interface as an inside-service interface as an inside-service interface as an inside3.In the Interface box, click sp-0/0/0.3.In the Interface box, click sp-0/0/0.1.Configure the services interface as an inside4.Next to Unit, click Add new entry. 5.In the Interface unit number box, type 1001.2.Configure the services interface as an inet interface.5.In the Interface domain box, select inside from the list.7.In the Family box, select the check box next to Inter and click Configure.8.Select the Primary check box, and click OK until you return to the Interfaces page.1.Configure the services interface as an outside-service interfaceConfigure the outside services interface for the IPSec tunnel, interface for the IPSec tunnel, interface for the IPSec tunnel, interface unit number box, type 2001.1.Configure the services interface as an outside-service interface as an outside-service interface2.In the Interface unit number box, type 2001.1.Configure the services interface as an outside3.In the Interface unit number box, type 2001.1.Configure the services interface as an outside4.In the Interface unit number box, type 2001.1.Configure the services interface as an inet interface. </td <td>Navigate to the <b>Interfaces</b> level in the configuration hierarchy.</td> <td>1.</td> <td>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</td> <td>Frc edi</td> <td>om the [edit] hierarchy level, enter</td>	Navigate to the <b>Interfaces</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	Frc edi	om the [edit] hierarchy level, enter		
Configure the inside services interface for the IPSec tunnel.1.Next to Interface, click Add new entry. 2.1.Configure the services interface as an inside-service interface.(See the interface naming conventions in the J-series Services Router Basic LAN and WAN Access Configuration Guide.)1.Next to Unit, click Add new entry. 4.1.Configure the services interface as an inside5.In the Interface unit number box, type 1001. 6.In the Family box, select the check box next to Inter and click Configure.1.Configure the services interface as an inet interface: set sp-0/0/0 unit 1001 family inet6.Next to Unit, click Add new entry. 5.1.Next to Interface, click sp-0/0/0. 		2.	Next to Interfaces, click <b>Configure</b> or <b>Edit</b> .				
<ul> <li>In the Interface naming conventions in the <i>J-series Services Router Basic LAN and Click OK.</i></li> <li>In the Interface box, click sp-0/0/0.</li> <li>In the Interface box, click sp-0/0/0.</li> <li>In the Interface of the IPSec turnel.</li> <li>In the Family box, select the check box next to Interface of the IPSec turnel.</li> <li>In the Interface, click sp-0/0/0.</li> <li>Select the Primary check box, and click OK until you return to the Interface spage.</li> <li>In the Interface unit number box, type 2001.</li> <li>In the Interface unit number box, select outside from the list.</li> <li>In the Interface unit number box, type 2001.</li> <li>In the Interface unit number box, select outside from the list.</li> <li>In the Interface unit number box, type 2001.</li> <li>In the Interface unit number box, select outside from the list.</li> <li>In the Interface unit number box, type 2001.</li> <li>In the Interface unit number box, select outside from the list.</li> <li>In the Interface unit number box, select outside from the list.</li> <li>In the Interface unit number box, select outside from the list.</li> <li>In the Interface unit number box, select outside from the list.</li> <li>In the Interface unit number box, select outside from the list.</li> <li>In the Earnily box, select the check box next to Interface:</li> <li>In the Family box, select the check box next to Interface:</li> <li>In the Family box, select the check box next to Interface:</li> <li>Select the Primary check box, and click OK.</li> </ul>	Configure the inside services	1.	Next to Interface, click Add new entry.	1.	Configure the services interface		
conventions in the <i>J</i> -series Services Router Basic LAN and WAN Access Configuration Guide.)3.In the Interface box, click <b>sp-0/0/0</b> .service-domain inside4.Next to Unit, click <b>Add new entry</b> . 5.In the Interface unit number box, type 1001.2.Configure the services interface as an inet interface:5.In the Service domain box, select <b>inside</b> from the list.5.In the Family box, select the check box next to Inet and click <b>Configure</b> .5.Select the <b>Primary</b> check box, and click <b>OK</b> until you return to the Interface spage.5.1.Configure the outside services interface for the IPSec tunnel.1.Next to Unit, click <b>Add new entry</b> .1.Configure the services interface as an outside-service interface: as an outside-service interfaceConfigure the outside services interface for the IPSec tunnel.1.Next to Unit, click <b>Add new entry</b> .1.Configure the services interface as an outside-service interface: as an outside-service interface: as an inet interface: set sp-0/0/0 unit 2001 service-domain outsideConfigure the services interface1.Next to Linterface unit number box, type 2001.1.Configure the services interface as an inet interface: set sp-0/0/0 unit 2001 service-domain outsideSelect the Primary check box, and click OK.5.In the Family box, select the check box next to Inet and click <b>Configure</b> .2.Configure the services interface as an inet interface: set sp-0/0/0 unit 2001 family inet	(See the interface naming	2.	In the Interface name box, type <b>sp-0/0/0</b> , and click <b>OK</b> .		set sp-0/0/0 unit 1001		
WAN Access Configuration Guide.)4.Next to Unit, click Add new entry.2.Configure the services interface as an inet interface:Guide.)5.In the Interface unit number box, type 1001.6.In the Service domain box, select inside from the list.7.In the Service domain box, select the check box next to Inet and click Configure.8.Select the Primary check box, and click OK until you return to the Interface page.8.Select the Primary check box, and click OK until you return to the Interface page.1.Configure the services interface as an outside-service interfaceConfigure the outside services interface for the IPSec tunnel.1.Next to Unit, click Add new entry.1.Configure the services interface as an outside-service interface2.Next to Unit, click Add new entry.3.In the Interface unit number box, type 2001.1.Configure the services interface as an outside-service interface as an inet interface3.In the Service domain box, select outside from the list.1.Configure the services interface as an inet interface4.In the Service domain box, select outside from the list.2.Configure the services interface as an inet interface5.In the Family box, select the check box next to Inet and click Configure.2.Configure the services interface as an inet interface: set sp-0/0/0 unit 2001 family inet	conventions in the <i>J-series</i>	3.	In the Interface box, click <b>sp-0/0/0</b> .		service-domain inside		
Guide.)       5. In the Interface unit number box, type 1001.       as an inter interface:         6. In the Service domain box, select inside from the list.       set sp-0/0/0 unit 1001 family inet         7. In the Family box, select the check box next to Inet and click Configure.       set sp-0/0/0 unit 1001 family inet         8. Select the Primary check box, and click OK until you return to the Interfaces page.       1. Next to Interface, click sp-0/0/0.         2. Next to Unit, click Add new entry.       3. In the Interface unit number box, type 2001.         4. In the Service domain box, select outside from the list.       set sp-0/0/0 unit 2001 service-domain outside         5. In the Family box, select the check box next to Inet and click Configure.       set sp-0/0/0 unit 2001 service interface:         6. Select the Primary check box, and click OK.       set sp-0/0/0 unit 2001 service interface:	WAN Access Configuration	4.	Next to Unit, click Add new entry.	2.	Configure the services interface		
6.In the Service domain box, select inside from the list.set sp-0/0/0 unit 1001 family inet7.In the Family box, select the check box next to Inet and click Configure.set sp-0/0/0 unit 1001 family inet8.Select the Primary check box, and click OK until you return to the Interfaces page.1.Configure the outside services interface for the IPSec tunnel.1.Next to Interface, click sp-0/0/0.1.2.Next to Unit, click Add new entry.1.Configure the services interface as an outside-service interface:3.In the Interface unit number box, type 2001.set sp-/0/0/0 unit 2001 service-domain outside4.In the Service domain box, select outside from the list.Select the check box next to Inet and click Configure.5.In the Family box, select the check box next to Inet and click Configure.2.6.Select the Primary check box, and click OK.set sp-0/0/0 unit 2001 family inet	Guide.)	5. In the Interface unit number box, type <b>1001</b> .		as an inet interface:			
7.In the Family box, select the check box next to Inet and click Configure.8.Select the Primary check box, and click OK until you return to the Interfaces page.Configure the outside services interface for the IPSec tunnel.1.2.Next to Interface, click sp-0/0/0.2.Next to Unit, click Add new entry.3.In the Interface unit number box, type 2001.4.In the Service domain box, select outside from the list.5.In the Family box, select the check box next to Inet and click Configure.6.Select the Primary check box, and click OK.		6.	In the Service domain box, select <b>inside</b> from the list.		set sp-0/0/0 unit 1001 family inet		
8.Select the Primary check box, and click OK until you return to the Interfaces page.1.Configure the services interface as an outside-service interface as an outside-service interface.Configure the outside services interface for the IPSec tunnel.1.Next to Interface, click <b>sp-0/0/0</b> .1.Configure the services interface as an outside-service interface.2.Next to Unit, click Add new entry.1.In the Interface unit number box, type 2001.1.Set sp-/0/00 unit 2001 service-domain outside4.In the Service domain box, select outside from the list.1.Configure the services interface as an inet interface.5.In the Family box, select the check box next to Inet and click Configure.2.Configure the services interface as an inet interface.6.Select the Primary check box, and click OK.Set sp-0/0/0 unit 2001 family inet		7.	In the Family box, select the check box next to Inet and click <b>Configure</b> .				
Configure the outside services interface for the IPSec tunnel.1.Next to Interface, click <b>sp-0/0/0</b> . 2.1.Configure the services interface as an outside-service interface:2.Next to Unit, click Add new entry.3.In the Interface unit number box, type 2001. 4.1.Set sp-/0/0/0 unit 2001 service-domain outside4.In the Service domain box, select outside from the list.2.Configure the services interface as an outside5.In the Family box, select the check box next to Inet and click Configure.2.Configure the services interface as an inet interface:6.Select the Primary check box, and click OK.set sp-0/0/0 unit 2001 family inet		8.	Select the <b>Primary</b> check box, and click <b>OK</b> until you return to the Interfaces page.				
interface for the IPSec tunnel.2.Next to Unit, click Add new entry.as an outside-service interface:3.In the Interface unit number box, type 2001.set sp-/0/0/0 unit 20014.In the Service domain box, select outside from the list.configure the services interface5.In the Family box, select the check box next to Inet and click Configure.configure.6.Select the Primary check box, and click OK.set sp-0/0/0 unit 2001 family inet	Configure the outside services	1.	Next to Interface, click <b>sp-0/0/0</b> .	1.	Configure the services interface		
<ol> <li>In the Interface unit number box, type 2001.</li> <li>In the Service domain box, select outside from the list.</li> <li>In the Family box, select the check box next to Inet and click Configure.</li> <li>Select the Primary check box, and click OK.</li> <li>set sp-/0/0/0 unit 2001 service-domain outside</li> <li>Configure the services interface as an inet interface:</li> <li>set sp-/0/0/0 unit 2001 service-domain outside</li> </ol>	interface for the IPSec tunnel.	2.	Next to Unit, click Add new entry.		as an outside-service interface:		
<ol> <li>In the Service domain box, select outside from the list.</li> <li>In the Family box, select the check box next to Inet and click Configure.</li> <li>Select the Primary check box, and click OK.</li> <li>service-domain outside</li> <li>Configure the services interface as an inet interface:</li> <li>set sp-0/0/0 unit 2001 family inet</li> </ol>		3.	In the Interface unit number box, type 2001.		set sp-/0/0/0 unit 2001		
<ul> <li>the list.</li> <li>5. In the Family box, select the check box next to Inet and click Configure.</li> <li>6. Select the Primary check box, and click OK.</li> <li>2. Configure the services interface as an inet interface:</li> <li>set sp-0/0/0 unit 2001 family inet</li> </ul>		<ol> <li>In the Service domain box, select outside from the list.</li> <li>In the Family box, select the check box next to Inet and click Configure.</li> </ol>	2.	service-domain outside			
<ul> <li>5. In the Family box, select the check box next to Inet and click Configure.</li> <li>6. Select the Primary check box, and click OK.</li> </ul>				Configure the services interface			
Inet and click Configure.set sp-0/0/0 unit 2001 family6.Select the Primary check box, and click OK.inet				as an met interface.			
6. Select the <b>Primary</b> check box, and click <b>OK</b> . inet				set sp-0/0/0 unit 2001 family			
		6.	Select the <b>Primary</b> check box, and click <b>OK</b> .		inet		

Ê

# **Configuring Service Sets**

To use dynamic SAs on the Services Router, you must create service sets to define the following information for IPSec service:

• The local gateway. If the IKE gateway IP address is in a VPN routing and forwarding (VRF) instance, you must configure the routing instance.

**NOTE:** You can configure Internet Key Exchange (IKE) gateway IP addresses that are present in a VPN routing and forwarding (VRF) instance as long as the peer is reachable through the VRF instance. For next-hop service sets, the key management process (kmd) places the IKE packets in the routing instance that contains the **outside-service-interface** value you specify. For interface service sets, the services interface (the interface on which the service set is applied) determines the VRF.

- A next-hop service set that defines which services interface to use for all inside-service next hops and all outside-service next hops (traffic inside the network and outside the network). Alternatively, you can create an interface service set that defines the services interface to be used for all IPSec traffic.
- An IPSec rule to act on input traffic, set the remote gateway on all traffic, and reference an IKE policy.

This configuration allows you to set the remote gateway address and perform IKE validation on all incoming traffic through the IPSec tunnel.

To configure a service set, you must complete the following tasks:

- Configure a gateway. See "Configuring a Local Gateway" on page 86
- Define a services interface. See either of the following tasks:
  - Configuring Next-Hop Services Interfaces on page 87
  - Configuring Interface Service Sets on page 88
- Apply a rule. See "Applying IPSec Rules to Service Sets" on page 89

## **Configuring a Local Gateway**

The sample service set configuration in Table 36 on page 87 configures the IPSec service set **ipsec-dynamic** and sets the local gateway to **10.90.90.2**.

To configure a local gateway for the service set:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 36 on page 87.
- 3. Go on to one of the following:
  - Configuring Next-Hop Services Interfaces on page 87
  - Configuring Interface Service Sets on page 88

## **Table 36: Configuring a Local Gateway**

Task		eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Services</b> level in the configuration hierarchy.		In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit services
	2.	Next to Services, click <b>Configure</b> or <b>Edit</b> .	
Configure the service set ipsec-dynamic.		Next to Service set, click <b>Add new</b> entry.	Enter
	2.	In the Service set name box, type ipsec-dynamic.	set service-set ipsec-dynamic
	3.	Click OK.	
Configure the IP address of the local gateway for the IPSec service set to the	1.	In the Service set list, click ipsec-dynamic.	Enter
local tunnel endpoint—for example, <b>10.1.15.1</b> .	2.	Next to Ipsec vpn options, click <b>Configure</b> .	set service-set ipsec-dynamic ipsec-vpn-options local-gateway 10.1.15.1
		In the Local gateway box, type <b>10.1.15.1</b> .	
	4.	Click <b>OK</b> until you return to the Services page.	

## **Configuring Next-Hop Services Interfaces**

The sample next-hop configuration in Table 37 on page 87 adds the next-hop services interfaces to the IPSec service set **ipsec-dynamic** created in Table 36 on page 87. This sample next-hop configuration sets the inside services interface to **sp-0/0/0.1001**, and sets the outside services interface (facing the remote IPSec site) to **sp-0/0/0.2001**.

To configure next-hop services interfaces:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 37 on page 87.
- 3. Go on to "Applying IPSec Rules to Service Sets" on page 89.

Table 37: Configuring	(Next-Hop Se	rvices Interfaces
-----------------------	--------------	-------------------

Task	J-Web Configuration Editor	CLI Configuration Editor	
Navigate to the <b>Services</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter edit services	
	<ol> <li>Next to Services, click Configure or Edit.</li> </ol>		

## Table 37: Configuring Next-Hop Services Interfaces (continued)

Task	J-W	eb Configuration Editor	CLI	Configuration Editor
Configure the next-hop service set for the IPSec tunnel.	1.	In the Service set list, click <b>ipsec-dynamic</b> .	1.	Enter
You must include an interface name and unit number for the inside-service interface and the outside-service interface. By default, the J-Web interface uses the following values:		In the Service type choice box, select <b>Next hop service</b> from the list.		set service-set ipsec-dynamic next-hop-service inside-service-interface sp-0/0/0.1001
		Next to Next hop service, click <b>Configure</b> .	2.	Enter
■ For the inside-service interface—sp-0/0/0.1001		In the Inside service interface box, type <b>sp-0/0/0.1001</b> .		set service-set ipsec-dynamic next-hop-service
■ For the outside-service interface—sp-0/0/0.2001	5.	In the Outside service interface box, type <b>sp-0/0/0.2001</b> .		outside-service-interface sp-0/0/0.2001
	6.	Click <b>OK</b> until you return to the Services page.		

## **Configuring Interface Service Sets**

The sample interface service set configuration in Table 38 on page 88 adds the interface service-set configuration to the IPSec service set **ipsec-dynamic** created in Table 36 on page 87. This sample interface service-set configuration sets the services interface **sp-0/0/0**.

To configure interface service sets:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 38 on page 88.
- 3. Go on to "Applying IPSec Rules to Service Sets" on page 89.

## **Table 38: Configuring Interface Service Sets**

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Services</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter edit services
	2. Next to Services, click <b>Configure</b> or <b>Edit</b> .	

Task	J-Web Configuration Editor	CLI Configuration Editor		
Configure the interface service set and specify <b>sp-0/0/0</b> as the services interface	<ol> <li>In the Service set list, click ipsec-dynamic.</li> </ol>	Enter		
to be used for IPSec traffic.	2. In the Service type choice box, select <b>Interface service</b> from the list.	set service-set ipsec-dynamic interface-service service-interface sp-0/0/0		
	<ol> <li>Next to Interface service, click Configure.</li> </ol>			
	4. In the Service interface box, type sp-0/0/0.			
	5. Click <b>OK</b> until you return to the Services page.			

## Table 38: Configuring Interface Service Sets (continued)

## **Applying IPSec Rules to Service Sets**

The sample configuration in Table 39 on page 89 configures the service set **ipsec-dynamic** configured in Table 36 on page 87 to use the IPSec rule **ipsec-dynamic-rule** defined in Table 34 on page 84.

To apply an IPSec rule to a service set:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 39 on page 89.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to the optional task "Configuring a NAT Pool" on page 90.
- 5. To check the configuration, see "Verifying the IPSec Tunnel Configuration" on page 98.

## **Table 39: Applying IPSec Rules to Service Sets**

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Services</b> level in the configuration hierarchy.	1.	In the J-Web interface, select <b>Configuration &gt; View</b> <b>and Edit &gt; Edit Configuration</b> .	From the [ <b>edit]</b> hierarchy level, enter
	2.	Next to Services, click <b>Configure</b> or <b>Edit</b> .	edit services
Apply the IPsec rule	1.	In the Service set list, click ipsec-dynamic.	Enter
<b>ipsec-dynamic-rule</b> to all traffic through the service set.	2.	In the Ipsec vpn rules choice box, select <b>Ipsec vpn rules</b> .	set service-set ipsec-dynamic ipsec-vpn-rules ipsec-dynamic-rule
		Next to Ipsec vpn rules, click Add new entry.	
		In the Rule name box, type ipsec-dynamic-rule.	
	5.	Click OK.	

# **Configuring a NAT Pool**

To hide internal IP addresses from the rest of the Internet, you configure the local tunnel endpoint as the only address in a Network Address Translation (NAT) pool, to ensure that it is the address used for address translation.

For more information about NAT, see "Network Address Translation" on page 163.

To configure a NAT pool for IPSec:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 40 on page 90.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to one of the following procedures:
  - To use digital certificates for authentication, see "Configuring Digital Certificates for IPSec Tunnels" on page 91.
  - To check the configuration, see "Verifying the IPSec Tunnel Configuration" on page 98.

### **Table 40: Configuring a NAT Pool for IPSec**

Task	J-W	eb Configuration Editor	CLI	Configuration Editor
Configure the NAT pool from which the addresses for Network Address Translation are taken.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	1.	From the <b>[edit]</b> hierarchy level, enter
Name the NAT need with any	2.	Next to Services, click Configure or Edit.		edit services nat
unique string of fewer than	3.	Next to Nat, click <b>Configure</b> or <b>Edit</b> .	2.	Add the local tunnel endpoint to the NAT address pool:
of characters.	4.	Next to Pool, click Add new entry.		-
Provide the IP address of the local tunnel endpoint—for example,	5.	In the Pool name box, type the name of the NAT pool.		set pool <i>pool-name</i> address 1.1.1.1
1.1.1.1.1.	6.	From the the Address choice list, select <b>Address</b> .		
	7.	In the Address box, type <b>1.1.1.1</b> .		

Task	Web Configuration Editor	CLI	CLI Configuration Editor		
Configure the router so that all outgoing traffic is matched against the IP address of the local tunnel	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	1.	From the [edit] hierarchy level, enter		
endpoint.	Next to Services, click Configure or Edi	it.	edit services nat		
Use any unique string for the NAT	. Next to Nat, click Configure or Edit.	2.	Configure a NAT rule and		
rule name and for the name of the term in the rule.	Next to Rule, click Add new entry.		apply it to an output traine.		
The source address must be the IP	In the Rule name box, type the name o rule.	of the	set rule <i>rule-name</i> match-direction output		
endpoint—for example, <b>1.1.1.1</b> .	. From the Match direction list, select <b>Ou</b>	atput. ^{3.}	Configure the rule to match traffic with a source address		
,	. Next to Term, click Add new entry.		that is the same as the local		
8	In the Term name box, type the name term.	of the	tunnel endpoint: set rule <i>rule-name</i> term		
(	Click <b>From</b> .		term-name from source-address		
	0. Next to Source address, click Add new e	entry.	1.1.1.1		
	1. From the address list, select Enter spec value.	cific			
	2. In the Address box, type <b>1.1.1.1</b> .				
	3. Click <b>OK</b> .				
Configure the router so that the source address for traffic through	On the main Configuration page next to Services, click <b>Configure</b> or <b>Edit</b> .	o 1.	From the [edit] hierarchy level, enter		
the local endpoint is translated to ,	. Next to Nat, click Configure or Edit.		edit services nat rule rule name		
, , , , , , , , , , , , , , , , , , ,	. Under Rule name, click the name of the	e rule.	term term-name		
	. Under Term name, click the name of the	term. 2.	Configure the source pool:		
5	. Click <b>Then</b> .		set then translated source-pool		
(	Click Translated.		pool-name		
	In the Source pool box, type the name	of the 3. point	Configure the type of translation:		
	is configured.				
8	<ul><li>is configured.</li><li>From the Source list, select Static.</li></ul>		set then translated translation-type source static		

## Table 40: Configuring a NAT Pool for IPSec (continued)

# **Configuring Digital Certificates for IPSec Tunnels**

Digital certificates are digitally signed statements providing independent confirmation of a network public key. Most digital certificates are issued by trusted third parties such as governments, financial institutions, or certificate authority (CA) companies specializing in certificate services.

A certificate authority (CA) is a location on a network that issues and manages security credentials and public keys for data encryption. As part of a public key infrastructure (PKI), a CA checks with a registration authority (RA) to verify information provided by the requestor of a digital certificate. If the RA verifies the requestor's information, the CA can issue a certificate.

The digital certificate is installed locally on the Services Router and used to encrypt and decrypt data on a network with IPSec peers configured for digital certificates. This section contains the following topics:

- Configuring a CA Profile with a Configuration Editor on page 92
- Requesting a CA Certificate from a CA on page 93
- Generating a Public and Private Key Pair on page 94
- Generating and Enrolling a Local Digital Certificate on page 94
- Loading a Digital Certificate on a Services Router on page 95
- Applying the Local Digital Certificate to an IPSec Tunnel on page 96
- Deleting a Digital Certificate on page 97

## **Configuring a CA Profile with a Configuration Editor**

The CA profile contains the name and the URL of the CA as well as a public key and additional information. The sample configuration in Table 41 on page 92 configures a CA profile **ca-profile-ipsec**.

To configure a CA profile:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor..
- 2. Perform the tasks described in Table 41 on page 92.
- 3. Go on to "Requesting a CA Certificate from a CA" on page 93.

Table	41:	<b>Configuring</b> a	I CA	Profile
-------	-----	----------------------	------	---------

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Security &gt; Pki</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter edit security pki
	2. Next to Security, click <b>Configure</b> of <b>Edit</b> .	r
	<ol> <li>Next to Pki, select the check box, and click Configure.</li> </ol>	

## Table 41: Configuring a CA Profile (continued)

Task	J-Web Configuration Editor CLI Configuration Editor	
Add a new CA profile to the Services Router.	1. Next to Ca profile, click <b>Add new</b> entry.	Enter set ca-profile ca-profile-ipsec ca-identity
Configure the profile name and the CA authority identification—for example, ca-profile-ipsec and versign.	<ol> <li>In the Ca profile name box, type ca-profile-ipsec.</li> <li>In the Ca identity box, type verisign.</li> </ol>	verisign
<ul> <li>Configure the following enrollment options:</li> <li>Enrollment retry—Number of attempts at online enrollment with the CA profile to allow for a router certificate, if enrollment fails—for example, 10. The range is from 0 through 100 attempts.</li> <li>Enrollment retry-interval—Length of time, in seconds, to allow between enrollment attempts—for example, 60. The range is from 0 through 3600 seconds.</li> <li>Enrollment URL—URL where the Simple Certificate Enrollment Protocol (SCEP) request is sent to the certification authority configured in this profile—for example, http://pilotonsiteipsec.verisign.com /cgi-bin/pkiclient.exe.</li> </ul>	<ol> <li>Next to Enrollment, click Configure.</li> <li>In the Retry box, type 10.</li> <li>In the Retry interval box, type 60.</li> <li>In the Url box, type http://pilotonsiteipsec.verisign.com /cgi-bin/pkiclient.exe.</li> <li>Click OK until you return to the main Configuration page.</li> </ol>	Enter set ca-profile ca-profile-ipsec enrollment retry 10 retry-interval 60 url http://pilotonsiteipsec.verisign.com /cgi-bin/pkiclient.exe

## **Requesting a CA Certificate from a CA**

CA certificates can be requested either manually or online. To request a certificate online, you can use the Simple Certificate Enrollment Protocol (SCEP) to contact the CA.

You can request a CA certificate in CLI operational mode only. To request a CA certificate:

- 1. Enter the CLI operational mode.
- 2. Perform the tasks described in Table 42 on page 94.
- 3. Go on to "Generating a Public and Private Key Pair" on page 94.

## Table 42: Requesting a CA Certificate from a CA

Task	CLI Operational Mode
Using the CA profile <b>ca-profile-ipsec</b> configured in Table 41 on page 92, contact the CA to request a CA certificate.	Enter
	request security pki ca-certificate enroll ca-profile ca-profile-ipsec

## **Generating a Public and Private Key Pair**

Every digital certificate has a pair consisting of an associated private key and public key. You must generate a public and private key pair to use digital certificates. A larger key pair is more secure than a smaller key pair. The available sizes, in bits, are as follows:

- 512
- 1024
- **2048**

Generating public and private key pairs can be performed in the CLI operational mode only. The sample configuration in Table 43 on page 94 generates a public and private key pair of 1024 bits for the certificate ID **local-verisign**.

To generate a public and private key pair:

- 1. Enter the CLI operational mode.
- 2. Perform the tasks described in Table 43 on page 94.
- 3. Go on to "Generating and Enrolling a Local Digital Certificate" on page 94.

## Table 43: Generating a Public and Private Key Pair

Task	CLI Operational Mode
Generate a public and private key pair.	Enter
The certificate ID is a unique ID that you create to identify all related files including the key pair, the certificate, and the certificate request files.	request security pki generate-key-pair certificate-id local-verisign size 1024

# **Generating and Enrolling a Local Digital Certificate**

Each Services Router is initially enrolled manually with the CA and then obtains the router certificate for its identity. This certificate is sent to the remote peer router during the Internet Key Exchange (IKE) negotiation.

You can generate and enroll a local digital certificate in the CLI operational mode only. To generate and enroll a local digital certificate:

1. Enter the CLI operational mode.

- 2. Perform the tasks described in Table 44 on page 95.
- 3. Go on to "Loading a Digital Certificate on a Services Router" on page 95.

## **Table 44: Generating and Enrolling a Local Certificate**

Task C	CLI Operational Mode	
Generate a local digital certificate.	Enter	
<ul> <li>The certificate has the following parameters:</li> <li>Certificate ID—Unique ID used to identify all of the related key pairs, certificates, and PKCS-10 certificate request files—for example, local-verisign</li> <li>CA profile—Name of the configured certificate authority profile—for example, ca-profile-ipsec</li> <li>Subject—Common name (CN), department or organizational unit name (OU), company name (O), state (ST), and country (C)for the digital certificate</li> <li>Domain name—Fully qualified domain name that identifies the certificate owner during IKE negotiations</li> <li>Challenge password—Password used by the CA for certificate enrollment and revocation</li> <li>IP address (Optional)—IP address if the Services Router has a static IP address</li> <li>Validity start time (Optional)—Length of time that a certificate is valid</li> </ul>	request security pki local-certificate enroll certificate-id local-verisign Enter request security pki local-certificate enroll ca-profile ca-profile-ipsec subject subject-distinguished-name domain-name domain-name challenge-password challenge-password ip-address <i>ip-address</i> validity-start-time start-time validity-end-time end-time	

# Loading a Digital Certificate on a Services Router

A CA certificate can be manually loaded onto the router from the certificates file.

You can load a local digital certificate in the CLI operational mode only. To load a local certificate:

- 1. Enter the CLI operational mode.
- 2. Perform the tasks described in Table 45 on page 95.
- 3. Go on to "Applying the Local Digital Certificate to an IPSec Tunnel" on page 96.

## Table 45: Loading a Certificate on a Services Router

Task	CLI Operational Mode
Load a certificate from an external file. You must specify the certificate ID—for example, <b>local-verisign</b> —to keep the proper linkage between	Enter
the private and public key pair.	request security pki local-certificate load certificate-id local-verisign filename <i>file-path</i>

## Table 45: Loading a Certificate on a Services Router (continued)

Task	CLI Operational Mode
Load a CA certificate from an external file. You must specify the CA profile—for example, <b>ca-profile-ipsec</b> .	Enter
	request security pki ca-certificate load ca-profile ca-profile-ipsec filename <i>file-path</i>

## Applying the Local Digital Certificate to an IPSec Tunnel

You can add a digital certificate to the IPSec tunnel using the J-Web configuration editor or the CLI configuration editor. To apply a certificate to an IPSec tunnel:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the tasks described in Table 46 on page 96.
- 3. If you are finished configuring the router, commit the configuration.

#### Task J-Web Configuration Editor **CLI Configuration Editor** Navigate to the Services level of 1. In the J-Web interface, select From the [edit] hierarchy level, the configuration hierarchy. Configuration > View and Edit > Edit enter Configuration. Use any unique string for the edit services service-set 2. Next to Services, click Configure or Edit. service set name. service-set-name 3. Next to Service set, click Add new entry. In the Service set name box, type a service set 4. name. Configure the IPSec VPN options 1. Next to Ipsec vpn options, click Configure. Enter for the services set. 2 In the Local gateway box, type an IP address. edit services service-set Use the CA profile you created in Next to Trusted ca, click Configure. 3. service-set-nameipsec-vpn-options Table 41 on page 92. In the Trusted ca profile box, type ca-profile-ipsec. 4. Enter Click OK until you return to the Services page. 5. set local-gateway ip-address Enter set trusted-ca ca-profile-ipsec

#### Table 46: Applying the Local Digital Certificate to an IPSec Tunnel

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Configure the IPSec VPN policy.	1.	Next to Ipsec vpn, click Configure.	Return to the [edit services]
Use the certificate ID you created in Table 44 on page 95.	2.	Next to Ike, click Configure.	hierarchy.
	3.	Next to Policy, click Add new entry.	Enter
	4.	In the Name box, type the policy name.	set ipsec-vpn ike policy policy-name
	5.	In the Local certificate box, type local-verisign.	local-certificate local-verisign
	6.	Click <b>OK</b> .	
Configure the IPSec VPN	1.	Next to Proposal, click Add new entry.	Enter
proposal.	2.	In the Name box, type the proposal name.	set ipsec-vpn ike proposal
	3.	From the Authentication method list, select <b>rsa-signatures</b> .	proposal-name authentication-method
	4.	Click <b>OK</b> .	rsa-signatures

## Table 46: Applying the Local Digital Certificate to an IPSec Tunnel (continued)

# **Deleting a Digital Certificate**

You can delete digital certificates using the CLI operational mode only. To delete certificates:

- 1. Enter the CLI operational mode.
- 2. Perform one of the tasks described in Table 47 on page 97.
- 3. If you are finished configuring the router, commit the configuration.

# Table 47: Deleting Digital Certificates on a Services Router

Task	CLI Operational Mode
Deleting all digital certificates for all service sets from the Services Router.	To delete all digital certificates for all service sets from the cache, enter
	clear services ipsec-vpn certificates service-set all
Deleting all digital certificates for a specific service set—for example <b>ipsec-dynamic</b> —from the Services Router.	To delete all digital certificates for the service set <b>ipsec-dynamic</b> from the cache, enter
	clear services ipsec-vpn certificates service-set ipsec-dynamic
Deleting the digital certificate that matches a specified certificate cache entry number—for example, <b>3</b> —for all service sets from the Services Router.	To delete the digital certificate that matches the certificate cache entry number <b>3</b> , enter
<b>NOTE:</b> To view the certificate cache entry numbers, issue the show services ipsec-vpn certificates command.	clear services ipsec-vpn certificates service-set certificate-cache-entry 3

## Table 47: Deleting Digital Certificates on a Services Router (continued)

Task	CLI Operational Mode
Deleting the digital certificate that matches a specified certificate cache entry number—for example, <b>3</b> —for a specified service set—for example, <b>ipsec-dynamic</b> from the Services	To delete the digital certificate that matches the certificate cache entry number <b>3</b> for the service set <b>ipsec-dynamic</b> , enter
Router.	clear services ipsec-vpn certificates service-set ipsec-dynamic certificate-cache-entry 3

## **Verifying the IPSec Tunnel Configuration**

To verify the IPSec tunnel configuration, perform the following task.

## **Verifying IPSec Tunnel Statistics**

**Purpose** Verify that traffic is being sent through the configured IPSec tunnel.

Action From the CLI, enter the show services ipsec-vpn ipsec statistics command.

# user@host> show services ipsec-vpn ipsec statistics PIC: sp-0/0/0, Service set: service-set-1

local gateway: 1.1.1.1. Remote gateway: 2.2.2.2. Tunnel index: 1

Local galeway. I.I.I.I, Remote	yateway.	2.2.2.2,	runner	muex.
ESP Statistics:				
Encrypted bytes:	0			
Decrypted bytes:	0			
Encrypted packets:	0			
Decrypted packets:	0			
AH Statistics:				
Input bytes:	0			
Output bytes:	0			
Input packets:	0			
Output packets:	0			
Errors:				
AH authentication failures: 0	, Replay (	errors: 0		
ESP authentication failures: (	), Decryp [.]	tion erro	rs: 0	
Bad headers: 0 Bad trailers: (	C			

- **What It Means** The output shows the statistics for the particular service set that defines the IPSec tunnel, including the local and remote gateway addresses, the number of packets that have been encrypted and transported, and the number of errors and failures. Verify the following information:
  - The local and remote tunnel endpoints are configured correctly.
  - The number of Authentication Header (AH) and Encapsulation Security Payload (ESP) errors is zero. If these numbers are nonzero, the Services Router might be having a problem either transmitting or receiving encrypted packets through the IPSec tunnel.

# **Related Topics** For a complete description of **show services ipsec-vpn ipsec statistics** output, see the *JUNOS System Basics and Services Command Reference*.

# Part 2 Managing Multicast Transmissions

- Multicast Overview on page 101
- Configuring a Multicast Network on page 109

J-series[™] Services Router Advanced WAN Access Configuration Guide

# Chapter 6 Multicast Overview

Multicast traffic lies between the extremes of unicast (one source, one destination) and broadcast (one source, all destinations). Multicast is a "one source, many destinations" method of traffic distribution, meaning that the destinations needing to receive the information from a particular source receive the traffic stream.

IP network destinations (clients) do not often communicate directly with sources (servers), so the routers between source and destination must be able to determine the topology of the network from the unicast or multicast perspective to avoid routing traffic haphazardly. The multicast router must find multicast sources on the network, send out copies of packets on several interfaces, prevent routing loops, connect interested destinations with the proper source, and keep the flow of unwanted packets to a minimum. Standard multicast routing protocols provide most of these capabilities.

This chapter contains the following topics. For more information about multicast, see the *JUNOS Multicast Protocols Configuration Guide*. For configuration instructions, see "Configuring a Multicast Network" on page 109.

- Multicast Terms on page 101
- Multicast Architecture on page 103
- Dense and Sparse Routing Modes on page 105
- Strategies for Preventing Routing Loops on page 105
- Multicast Protocol Building Blocks on page 106

# **Multicast Terms**

To understand multicast routing, you must be familiar with the terms defined in Table 48 on page 101. See Figure 8 on page 104 for a general view of some of the elements commonly used in an IP multicast network architecture.

### **Table 48: Multicast Terms**

Term	Definition
administrative scoping	Multicast routing strategy that limits the routers and interfaces used to forward a multicast packet by reserving a range of multicast addresses.
Auto-RP	Cisco multicast routing protocol that allows sparse-mode routing protocols to find rendezvous points (RPs) within a routing domain.

## Table 48: Multicast Terms (continued)

Term	Definition
bootstrap router (BSR)	Multicast mechanism that allows routers running PIM sparse mode to find rendezvous points (RPs) within a routing domain.
branch	Part of a multicast network that is formed when a leaf subnetwork is joined to the multicast distribution tree. Branches with no interested receivers are pruned from the tree so that multicast packets are no longer replicated on the branch.
broadcast routing protocol	Protocol that distributes traffic from a particular source to all destinations.
dense mode	Multicast routing mode appropriate for LANs with many interested receivers.
Designated Router (DR)	Router on a subnet that is selected to control multicast routes for the sources and receivers on the subnet. When more than one multicast-enabled router is located on a subnet, the selected DR is the router with the highest priority. If the DR priorities match, the router with the highest IP address is selected as the DR.
	The source's DR sends PIM register messages from the source network to the rendezvous point (RP). The receiver's DR sends PIM join and PIM prune messages from the receiver network toward the RP.
Distance Vector Multicast Routing Protocol (DVMRP)	Distributed multicast routing protocol that dynamically generates IP multicast distribution trees using reverse-path multicasting (RPM) to forward multicast traffic to downstream interfaces.
distribution tree	Path linking multicast receivers (listeners) to sources. The root of the tree is at the source, and the branches connect subnetworks of interested receivers (leaves). Multicast packets are replicated only where a distribution tree branches. To shorten paths to a source at the edge of a network, sparse mode multicast protocols can use a <i>shared</i> distribution tree located more centrally in the network backbone.
downstream interface	Interface on a multicast router that is leading toward the receivers. You can configure all the logical interfaces except one as downstream interfaces.
group address	Multicast destination address. A multicast network uses the Class D IP address of a logical group of multicast receivers to identify a destination. IP multicast packets have a multicast group address as the destination address and a unicast source address.
Internet Group Management Protocol (IGMP)	Multicast routing protocol that runs between receiver hosts and routers to determine whether group members are present. Services Routers support IGMPv1, IGMPv2, and IGMPv3.
leaf	IP subnetwork that is connected to a multicast router and that includes at least one host interested in receiving IP multicast packets. The router must send a copy of its multicast packets out on each interface with a leaf, and its action is unaffected by the number of leaves on the interface.
listener	Another name for a receiver in a multicast network.
multicast routing protocol	Protocol that distributes traffic from a particular source to only the destinations needing to receive it. Typical multicast routing protocols are the Distance Vector Multicast Routing Protocol (DVMRP) and Protocol Independent Multicast (PIM).
Multicast Source Discovery Protocol (MSDP)	Multicast routing protocol that connects multicast routing domains and allows them to find rendezvous points (RPs).

## Table 48: Multicast Terms (continued)

Term	Definition
Pragmatic General Multicast (PGM)	Special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic.
Protocol Independent Multicast (PIM) protocol	Protocol-independent multicast routing protocol that can be used in either sparse or dense mode. In sparse mode, PIM routes to multicast groups that might span WANs and interdomain Internets. In dense mode, PIM is a flood-and-prune protocol.
pruning	Removing from a multicast distribution tree branches that no longer include subnetworks with interested hosts. Pruning ensures that packets are replicated only as needed.
reverse-path forwarding (RPF)	Multicast routing strategy that allows a router to receive packets through an interface if it is the same interface a unicast packet uses as the shortest path back to the source.
rendezvous point (RP)	Core router operating as the root of a shared distribution tree in a multicast network.
Session Announcement Protocol (SAP)	Multicast routing protocol used with other multicast protocols—typically Session Description Protocol (SDP)—to handle session conference announcements.
Session Description Protocol (SDP)	Session directory protocol that advertise multimedia conference sessions and communicates setup information to participants who want to join the session.
shortest-path tree (SPT)	Multicast routing strategy for sparse mode multicast protocols. SPT uses a shared distribution tree rooted in the network backbone to shorten paths to sources at the edge of a network.
source-specific multicast (SSM)	Service that allows a client to receive multicast traffic directly from the source, without the help of a rendezvous point (RP).
sparse mode	Multicast routing mode appropriate for WANs with few interested receivers.
unicast routing protocol	Protocol that distributes traffic from one source to one destination.
upstream interface	Interface on a multicast router that is leading toward the source. To minimize bandwidth use, configure only one upstream interface on a router receiving multicast packets.

# **Multicast Architecture**

Multicast-capable routers replicate packets on the multicast network, which has exactly the same topology as the unicast network it is based on. Multicast routers use a multicast routing protocol to build a distribution tree that connects receivers (also called listeners) to sources.

# **Upstream and Downstream Interfaces**

A single upstream interface on the router leads toward the source to receive multicast packets. The downstream interfaces on the router lead toward the receivers to transmit packets. A router can have as many downstream interfaces as it has logical interfaces, minus 1. To prevent looping, the router's upstream interface must never receive copies of its own downstream multicast packets.

## **Subnetwork Leaves and Branches**

On a multicast router, each subnetwork of hosts that includes at least one interested receiver is a leaf on the multicast distribution tree (see Figure 8 on page 104). The router must send out a copy of the IP multicast packet on each interface with a leaf. When a new leaf subnetwork joins the tree, a new branch is built so that the router can send out replicated packets on the interface. The number of leaves on an interface does not affect the router. The action is the same for one leaf or a hundred.

A branch that no longer has leaves is pruned from the distribution tree. No multicast packets are sent out on a router interface leading to an IP subnetwork with no interested hosts. Because packets are replicated only where the distribution tree branches, no link ever carries a duplicate flow of packets.

In IP multicast networks, traffic is delivered to multicast groups based on an IP multicast group address instead of a unicast destination address. The groups determine the location of the leaves, and the leaves determine the branches on the multicast network.



#### **Figure 8: Multicast Elements in an IP Network**

## **Multicast IP Address Ranges**

Multicast uses the Class D IP address range (224.0.0.0 through 239.255.255.255). Multicast addresses usually have a prefix length of /32, although other prefix lengths are allowed. Multicast addresses represent logical groupings of receivers and not physical collections of devices, and can appear only as the destination in an IP packet, never as the source address.

# **Notation for Multicast Forwarding States**

The multicast forwarding state in a router is usually represented by one of the following notations:

- (S,G) notation—S refers to the unicast IP address of the source for the multicast traffic and G refers to the particular multicast group IP address for which S is the source. All multicast packets sent from this source have S as the source address and G as the destination address.
- (*, G) notation—The asterisk (*) is a wildcard for the address of any multicast application source sending to group G. For example, if two sources are originating exactly the same content for multicast group 224.1.1.2, a router can use (*, 224.1.1.2) to represent the state of a router forwarding traffic from both sources to the group.

# **Dense and Sparse Routing Modes**

To keep packet replication to a minimum, multicast routing protocols use the two primary modes shown in Table 49 on page 105.

# <u>!</u>

**CAUTION:** A common multicast guideline is *not to run dense mode on a WAN under any circumstances.* 

## **Table 49: Primary Multicast Routing Modes**

Multicast Mode	Description	Appropriate Network for Use			
Dense mode	Network is flooded with traffic on all possible branches, then pruned back as branches explicitly (by message) or implicitly (time-out silence) eliminate themselves.	LANs—Networks in which all possible subnets are likely to have at least one receiver.			
Sparse mode	Network establishes and sends packets only on branches that have at least one leaf indicating (by message) a need for the traffic.	WANs—Network in which very few of the possible receivers require packets from this source.			

# **Strategies for Preventing Routing Loops**

Routing loops are disastrous in multicast networks because of the risk of repeatedly replicated packets, which can overwhelm a network. One of the complexities of modern multicast routing protocols is the need to avoid routing loops, packet by packet, much more rigorously than in unicast routing protocols. Three multicast strategies—reverse-path forwarding (RPF), shortest-path tree (SPT), and administrative scoping—help prevent routing loops by defining routing paths in different ways.

# **Reverse-Path Forwarding for Loop Prevention**

The router's multicast forwarding state runs more logically based on the reverse path, from the receiver back to the root of the distribution tree. In reverse-path

forwarding (RPF), every multicast packet received must pass an RPF check before it can be replicated or forwarded on any interface. When it receives a multicast packet on an interface, the router verifies that the *source* address in the multicast IP packet is the *destination* address for a unicast IP packet back to the source.

If the outgoing interface found in the unicast routing table is the same interface that the multicast packet was received on, the packet passes the RPF check. Multicast packets that fail the RPF check are dropped, because the incoming interface is not on the shortest path back to the source. Routers can build and maintain separate tables for RPF purposes.

# **Shortest-Path Tree for Loop Prevention**

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT), but this path can be long if the source is at the periphery of the network. Providing a *shared tree* on the backbone as the distribution tree locates the multicast source more centrally in the network. Shared distribution trees with roots in the core network are created and maintained by a multicast router operating as a rendezvous point (RP), a feature of sparse mode multicast protocols.

## Administrative Scoping for Loop Prevention

Scoping limits the routers and interfaces that can forward a multicast packet. Multicast scoping is *administrative* in the sense that a range of multicast addresses is reserved for scoping purposes, as described in RFC 2365, *Administratively Scoped IP Multicast*. Routers at the boundary must filter multicast packets and ensure that packets do not stray beyond the established limit.

# **Multicast Protocol Building Blocks**

Multicast is not a single protocol, but a collection of protocols working together to form trees, prune branches, locate sources and groups, and prevent routing loops:

- Distance Vector Multicast Routing Protocol (DVMRP) and Protocol Independent Multicast (PIM) operate between routers. PIM can operate in dense mode and sparse mode.
- Three versions of the Internet Group Management Protocol (IGMP) run between receiver hosts and routers.
- Several other routing mechanisms and protocols enhance multicast networks by providing useful functions not included in other protocols. These include the bootstrap router (BSR) mechanism, Auto-RP protocol, Multicast Source Discovery Protocol (MSDP), Session Announcement Protocol (SAP) and Session Discovery Protocol (SDP), and Pragmatic General Multicast (PGM) protocol.

Table 50 on page 107 lists and summarizes these protocols.

Multicast Protocol	Description	Uses
DVMRP	Dense-mode-only protocol that uses the flood-and-prune or implicit join method to deliver traffic everywhere and then determine where the uninterested receivers are. DVRMP uses source-based distribution trees in the form (S,G) and builds its own multicast routing tables for RPF checks.	Not appropriate for large-scale Internet use.
PIM dense mode	Sends an <i>implicit</i> join message, so routers use the flood-and-prune method to deliver traffic everywhere and then determine where the uninterested receivers are.	Most promising multicast protocol in use for LANs.
	PIM dense mode uses source-based distribution trees in the form (S,G), and also supports sparse-dense mode, with mixed sparse and dense groups. Both PIM modes use unicast routing information for RPF checks.	
PIM sparse mode	Sends an <i>explicit</i> join message, so routers determine where the interested receivers are and send join messages upstream to their neighbors, building trees from receivers to a rendezvous point (RP) router, which is the initial source of multicast group traffic.	Most promising multicast protocol in use for WANs.
	PIM sparse mode builds distribution trees in the form (*,G), but migrates to an (S,G) source-based tree if that path is shorter than the path through the RP router for a particular multicast group's traffic. Both PIM modes use unicast routing information for RPF checks.	
PIM source-specific multicast (SSM)	Enhancement to PIM sparse mode that allows a client to receive multicast traffic directly from the source, without the help of a rendezvous point (RP).	Used with IGMPv3 to create a shortest-path tree between receiver and source.
IGMPv1	The original protocol defined in RFC 1112, <i>Host Extensions for IP Multicasting</i> . IGMPv1 sends an explicit join message to the router, but uses a time-out to determine when hosts leave a group.	
IGMPv2	Defined in RFC 2236, <i>Internet Group</i> <i>Management Protocol, Version 2.</i> Among other features, IGMPv2 adds an explicit leave message to the join message.	Used by default.

# Table 50: Multicast Protocol Building Blocks

# Table 50: Multicast Protocol Building Blocks (continued)

Multicast Protocol	Description	Uses
IGMPv3	Defined in RFC 3376, <i>Internet Group</i> <i>Management Protocol, Version 3.</i> Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or <i>source-specific</i> <i>multicast (SSM)</i> .	Used with PIM SSM to create a shortest-path tree between receiver and source.
BSR	Allow sparse-mode routing protocols to find rendezvous points (RPs) within the	
Auto-RP	routing domain (autonomous system, or AS). RP addresses can also be statically configured.	
MSDP	Allows groups located in one multicast routing domain to find rendezvous points (RPs) in other routing domains. MSDP is not used on an RP if all receivers and	Typically runs on the same router as PIM sparse mode rendezvous point (RP).
	sources are located in the same routing domain.	Not appropriate if all receivers and sources are located in the same routing domain.
SAP and SDP	Display multicast session names and correlate the names with multicast traffic. SDP is a session directory protocol that advertises multimedia conference sessions and communicates setup information to participants who want to join the session. A client commonly uses SDP to announce a conference session by periodically multicasting an announcement packet to a well-known multicast address and port using SAP.	
PGM	Special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic. PGM allows a receiver to detect missing information in all cases and request replacement information if the receiver application requires it.	

# Chapter 7 Configuring a Multicast Network

You configure a router network to support multicast applications with a related family of protocols. To use multicast, you must understand the basic components of a multicast network and their relationships, and then configure the J-series Services Router to act as a node in the network.



**NOTE:** The J-series Services Router supports both Protocol Independent Multicast (PIM) version 1 and PIM version 2. In this chapter, the term *PIM* refers to both versions of the protocol.

You use either the J-Web configuration editor or CLI configuration editor to configure multicast protocols. The J-Web interface does not include Quick Configuration pages for multicast protocols.

This chapter contains the following topics. For more information about multicast, see the *JUNOS Multicast Protocols Configuration Guide*.

- Before You Begin on page 109
- Configuring a Multicast Network with a Configuration Editor on page 110
- Verifying a Multicast Configuration on page 119

# **Before You Begin**

Before you begin configuring a multicast network, complete the following tasks:

- If you do not already have a basic understanding of multicast, read "Multicast Overview" on page 101.
- Determine whether the Services Router is directly attached to any multicast sources. Receivers must be able to locate these sources.
- Determine whether the Services Router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
- Determine whether to use the sparse, dense, or sparse-dense mode of multicast operation. Each mode has different configuration considerations.
- Determine the address of the rendezvous point (RP) if sparse or sparse-dense mode is used.
- Determine whether to locate the RP with the static configuration, bootstrap router (BSR), or Auto-RP method.
- Determine whether to configure multicast to use its own reverse-path forwarding (RPF) routing table when configuring PIM in sparse, dense, or sparse-dense modes.

# **Configuring a Multicast Network with a Configuration Editor**

To configure the Services Router as a node in a multicast network, you must perform the following tasks marked *(Required)*. For information about using the J-Web and CLI configuration editors, see the *J-series Services Router Basic LAN and WAN Access Configuration Guide*.

- Configuring SAP and SDP (Optional) on page 110
- Configuring IGMP (Required) on page 111
- Configuring the PIM Static RP (Optional) on page 112
- Filtering PIM Register Messages from Unauthorized Groups and Sources (Optional) on page 114
- Configuring a PIM RPF Routing Table (Optional) on page 117

## **Configuring SAP and SDP (Optional)**

Multicast session announcements are handled by two protocols, the Session Announcement Protocol (SAP) and Session Description Protocol (SDP). These two protocols display multicast session names and correlate the names with multicast traffic. Enabling SDP and SAP allows the router to receive announcements about multimedia and other multicast sessions from sources. Enabling SAP automatically enables SDP.

For more information on SAP and SDP, see the *JUNOS Multicast Protocols Configuration Guide*.

The Services Router listens for session announcements on one or more addresses and ports. By default, the router listens to address and port **224.2.127.254:9875**.

To configure SAP and SDP for the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 51 on page 111.
- 3. Go on to "Configuring IGMP (Required)" on page 111.

### **Table 51: Configuring SAP and SDP**

Task	J-W	J-Web Configuration Editor		Configuration Editor
Navigate to the <b>Listen</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	Frc edi	om the [edit] hierarchy level, enter
	<ol> <li>Next to Protocols, click Configure or Edit.</li> </ol>			
		Next to Sap, click <b>Configure</b> or <b>Edit</b> .		
		Click Add new entry next to Listen.		
(Optional) Enter one or more addresses and ports for the Services Router to listen to session announcements on. By default, the Services Router listens to address and port 224.2.127.254:9875.	1.	In the Address box, type the multicast address the Services Router can listen to session announcements on, in dotted decimal notation. In the Port box, type the port number in decimal notation.	1.	Set the <b>address</b> value to the IP address that the Services Router can listen to session announcements on, in dotted decimal notation. For example: set listen 224.2.127.254
	3.	Click <b>OK</b> .	2.	Set the <b>port</b> value to the number of the port that the Services Router can listen to session announcements on, in decimal notation. For example:
				set listen 224.2.127.254 port 9875.

# **Configuring IGMP (Required)**

The Internet Group Management Protocol (IGMP) manages the membership of hosts and routers in multicast groups. IGMP is an integral part of IP and must be enabled on all routers and hosts that need to receive IP multicasts. IGMP is automatically enabled on all broadcast interfaces when you configure PIM or DVMRP.

For more information on IGMP, see JUNOS Multicast Protocols Configuration Guide.

By default, the Services Router runs IGMPv2. However, you might still want to set the IGMP version explicitly on an interface, or all interfaces. Routers running different versions of IGMP negotiate the lowest common version of IGMP supported by hosts on their subnet. One host running IGMPv1 forces the Services Router to use that version and lose features important to other hosts.

To explicitly configure the IGMP version, perform these steps on each Services Router in the network:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 52 on page 112.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to one of the following procedures:
  - To configure PIM sparse mode, see "Configuring the PIM Static RP (Optional)" on page 112.
  - To check the configuration, see "Verifying a Multicast Configuration" on page 119.

### Table 52: Explicitly Configuring the IGMP version

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Interface</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter edit protocols igmp
	2. Next to Protocols, click <b>Configure</b> or <b>Edit</b> .	
	3. Next to Igmp, click <b>Configure</b> or <b>Edit</b> .	
	4. Next to Interface, click <b>Add new</b> entry.	
Set the IGMP version. By default, the Services Router uses IGMPv2, but this	1. In the Interface name box, type the name of the interface, or all.	1. Set the interface value to the interface name, or all. For example:
negotiation with hosts unless explicitly configured.	2. In the Version box, type the version number: 1, 2, or 3.	set igmp interface all
(See the interface naming conventions in the I-series Services Router Basic LAN	3. Click <b>OK</b> .	<ol> <li>Set the version value to 1, 2, or 3. For example:</li> </ol>
and WAN Access Configuration Guide.)		set igmp interface all version 2

# **Configuring the PIM Static RP (Optional)**

Protocol Independent Multicast (PIM) sparse mode is the most common multicast protocol used on the Internet. PIM sparse mode is the default mode whenever PIM is configured on any interface of the Services Router. However, because PIM must not be configured on the network management interface of the Services Router, you must disable it on that interface.

Each any-source multicast (ASM) group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The rendezvous point (RP) router is the root of this shared tree and receives the multicast traffic from the source. To receive multicast traffic from the groups served by the RP, the Services Router must determine the IP address of the RP for the source.

One common way for the Services Router to locate RPs is by static configuration of the IP address of the RP. For information about alternate methods of locating RPs, see the *JUNOS Multicast Protocols Configuration Guide*.

To configure PIM sparse mode, disable PIM on **ge-0/0/0**, and configure the IP address of the RP perform these steps on each Services Router in the network:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 53 on page 113.
- 3. Go on to "Configuring a PIM RPF Routing Table (Optional)" on page 117.

#### Table 53: Configuring PIM Sparse Mode and the RP

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Interface</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter edit protocols pim
	2. Next to Protocols, click <b>Configure</b> or <b>Edit</b> .	
	<ol> <li>Next to Pim, click Configure or Edit.</li> </ol>	
	<ol> <li>Next to Interface, click Add new entry.</li> </ol>	
Enable PIM on all network interfaces.	In the Interface name box, type all.	Set the <b>interface</b> value to <b>all</b> . For example:
(See the interface naming conventions in the J-series Services Router Basic LAN and WAN Access Configuration Guide.)		set pim interface all
Apply your configuration changes.	Click <b>OK</b> to apply your entries to the configuration.	Changes in the CLI are applied automatically when you execute the <b>set</b> command.
Remain at the <b>Interface</b> level in the configuration hierarchy.	Click <b>Add new entry</b> next to Interface.	Remain at the [edit protocols pim interface] hierarchy level.
Disable PIM on the network management interface.	1. In the Interface name box, type ge-0/0/0.	Disable the <b>ge-0/0/0</b> interface:
	2. Select the check box next to Disable.	set pim interface ge-0/0/0 unit 0 disable
Apply your configuration changes.	Click <b>OK</b> to apply your entries to the configuration.	Changes in the CLI are applied automatically when you execute the <b>set</b> command.

## Table 53: Configuring PIM Sparse Mode and the RP (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Rp</b> level in the configuration hierarchy.	<ol> <li>On the main Configuration page next to Protocols, click Configure</li> </ol>	From the [edit] hierarchy level, enter
	or <b>Edit</b> .	edit protocols pim rp
	2. Next to Pim, click <b>Configure</b> or <b>Edit</b> .	
	3. Next to Rp, click <b>Configure</b> or <b>Edit</b> .	
Configure the IP address of the RP—for	1. Click <b>Configure</b> next to Static.	Set the address value to the IP address
example, <b>192.168.14.27</b> .	2. Click <b>Add new entry</b> next to	of the RP:
	Address.	set static address 192 168 14 27
	<ol> <li>In the Addr box, type 192.168.14.27.</li> </ol>	
	4. Click <b>OK</b> .	

# Filtering PIM Register Messages from Unauthorized Groups and Sources (Optional)

When a source in a multicast network becomes active, the source's designated router (DR) encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the rendezvous point (RP) router.

To prevent unauthorized groups and sources from registering with an RP router, you can define a routing policy to reject PIM register messages from specific groups and sources and configure the policy on the designated router or the RP router. For information about routing policies, see the *JUNOS Policy Framework Configuration Guide* 

- If you configure the reject policy on an RP router, it rejects incoming PIM register messages from the specified groups and sources. The RP router also sends a register stop message by means of unicast to the designated router. On receiving the register stop message, the designated router sends periodic null register messages for the specified groups and sources to the RP router.
- If you configure the reject policy on a designated router, it stops sending PIM register messages for the specified groups and sources to the RP router.

(¥

**NOTE:** If you have configured the reject policy on an RP router, we recommend that you configure the same policy on all the RP routers in your multicast network.

# (¥

**NOTE:** If you delete a group and source address from the reject policy configured on an RP router and commit the configuration, the RP router will register the group and source only when the designated router sends a null register message.

This section contains the following topics:

- Rejecting Incoming PIM Register Messages on an RP Router on page 115
- Stopping Outgoing PIM Register Messages on a Designated Router on page 116

## **Rejecting Incoming PIM Register Messages on an RP Router**

To reject incoming PIM register messages on an RP router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 54 on page 115.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To check the configuration, see "Verifying a Multicast Configuration" on page 119.

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Policy</b> options level in the	1. In the J-Web interface, select <b>Configuration &gt;</b> and Edit > Edit Configuration.	View From the [edit] hierarchy level, enter
configuration hierarchy.	2. Next to Policy options, click <b>Configure</b> or <b>Ed</b>	it. edit policy-options
Define a policy to reject	1. Next to Policy statement, click Add new ent	ry. 1. Set the match condition for the
PIM register messages from a group and source address.	2. In the Policy name box, type the name of the p statement—for example, reject-pim-register-m	sg-rp. set policy statement
	3. Next to From, click <b>Configure</b> .	reject-pim-register-msg-rp from
	4. Next to Route filter, click <b>Add new entry</b> .	route-filter 224.1.1.1/32 exact
	5. In the Address box, type the address of the group—for example, <b>224.1.1.1/32</b> .	2. Set the match condition for the address of a source in the group:
	6. From the Modifier list, select <b>Exact</b> .	set policy statement
	7. Click OK.	source-address-filter 10.10.10.1/32
	8. Next to Source address filter, click <b>Add new e</b>	entry. exact
	9. In the Address box, type the address of the source—for example, <b>10.10.1/32</b> .	<ol> <li>Set the match action to reject PIM register messages from the group and source address:</li> </ol>
	10. From the Modifier list, select Exact.	and source address.
	11. Click <b>OK</b> until you return to the Policy staten page.	nent set policy statement reject-pim-register-msg-rp then reject
	12. Next to Then, click <b>Configure</b> .	
	13. From the Accept reject list, select <b>Reject</b> .	

	Table 5	4: Re	jecting	Incoming	PIM	Register	Messages	on an	RP	Route
--	---------	-------	---------	----------	-----	----------	----------	-------	----	-------

Task	J-W	/eb Configuration Editor	CL	I Configuration Editor
Configure the       1.       On the main Configuration page next to Protocols, click Configure or Edit.		1.	From the [edit] hierarchy level, enter	
policy on the RP router.	2. 3.	Next to Pim, click Configure.		edit protocols pim rp
		Next to Rp, click <b>Configure</b> .	2.	Assign the policy on the RP:
4. Next to Rp register policy, click <b>Add new entry</b> .				
	5.	In the Value box, type the name of the policy—reject-pim-register-msg-rp.		set rp-register-policy reject-pim-register-msg-rp
	6.	Click <b>OK</b> .		

## Table 54: Rejecting Incoming PIM Register Messages on an RP Router (continued)

# **Stopping Outgoing PIM Register Messages on a Designated Router**

To stop outgoing PIM register messages on a designated router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 55 on page 116.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To check the configuration, see "Verifying a Multicast Configuration" on page 119.

## Table 55: Stopping Outgoing PIM Register Messages on a Designated Router

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Policy</b> options level in the	1.	In the J-Web interface, select <b>Configuration &gt; View</b> <b>and Edit &gt; Edit Configuration</b> .	From the [edit] hierarchy level, enter
configuration hierarchy.	2.	Next to Policy options, click <b>Configure</b> or <b>Edit</b> .	edit policy-options

Task	J-W	/eb Configuration Editor	CLI Configuration Editor		
Define a policy to not send PIM register messages for a group and source address.	1.	Next to Policy statement, click Add new entry.		Set the match condition for the	
	2.	In the Policy name box, type the name of the policy statement—for example, <b>stop-pim-register-msg-dr</b> .		group address:	
	3.	Next to From, click Configure.		stop-pim-register-msg-dr from	
	4.	Next to Route filter, click Add new entry.	-	route-filter 224.2.2.2/32 exact	
	5.	In the Address box, type the address of the group—for example, <b>224.2.2.2/32</b> .	2.	Set the match condition for the address of a source in the group:	
	6.	From the Modifier list, select Exact.		set policy statement stop-pim-register-msg-dr from source-address-filter 20.20.20.1/32 exact	
	7.	Click OK.			
	8.	Next to Source address filter, click Add new entry.			
	9.	In the Address box, type the address of the source—for example, 20.20.20.1/32.	3.	Set the match action to not send PIM register messages for the group and source address:	
	10.	From the Modifier list, select Exact.		and source address.	
		Click <b>OK</b> until you return to the Policy statement page.		set policy statement stop-pim-register-msg-dr then reject	
	12	Next to Then, click <b>Configure</b> .			
	13	From the Accept reject list, select <b>Reject</b> .			
	14	Click <b>OK</b> .			
Configure the stop-pim-register-msg-dr policy on the designated router.	1.	On the main Configuration page, next to Protocols, click <b>Configure</b> or <b>Edit</b> .	1.	From the [edit] hierarchy level, enter	
	2.	Next to Pim, click Configure.		edit protocols pim rp	
	3.	Next to Rp, click Configure.	2.	Assign the policy on the designated	
	4.	Next to Dr register policy, click Add new entry.		104(0).	
	5.	In the Value box, type the name of the policy—for example, <b>stop-pim-register-msg-dr</b> .		set dr-register-policy stop-pim-register-msg-dr	
	6.	Click <b>OK</b> .			

#### Table 55: Stopping Outgoing PIM Register Messages on a Designated Router (continued)

# **Configuring a PIM RPF Routing Table (Optional)**

By default, PIM uses inet.0 as its reverse-path forwarding (RPF) routing table group. PIM uses an RPF routing table group to resolve its RPF neighbor for a particular multicast source address and for the RP address. PIM can optionally use inet.2 as its RPF routing table group. The inet.2 routing table is organized more efficiently for RPF checks.

Once configured, the RPF routing table must be applied to PIM as a routing table group.

To configure and apply a PIM RPF routing table, perform these steps on each Services Router in the network:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 56 on page 118.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To check the configuration, see "Verifying a Multicast Configuration" on page 119.

Task		eb Configuration Editor	CLI Configuration Editor	
Navigate to the <b>Routing options</b> level in the configuration hierarchy.		In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter	
	2.	Next to Routing options, click <b>Configure</b> or <b>Edit</b> .	edit routing-options	
Configure a new group for the RPF routing table.		xt to Rib groups, click Add new entry.	Enter	
			edit rib-groups	
Configure a name for the new RPF		In the Ribgroup name box, type	Enter	
multicast-rfp-rib—and use inet.2 for its export routing table.	2.	In the Export rib box, type inet.2.	set multicast-rpf-rib export-rib inet.2	
Configure the new RPF routing table group to use inet.2 for its import routing table.		Click Add new entry next to Import rib.	Enter	
		In the Value box, type inet.2.	set multicast-rpf-rib import-rib	
		Click <b>OK</b> three times.	inet.2	
Navigate to the <b>Rib group</b> level in the configuration hierarchy.	1.	On the main Configuration page next to Protocols, click <b>Configure</b> or <b>Edit</b> .	From the [edit] hierarchy level, enter	
		Next to Pim, click Configure or Edit.	edit protocols pim	
	3.	Next to Rib group, click Configure or Edit.		
Apply the new RPF routing table to PIM.	1.	In the Inet box, type the name of the RPF routing table group—multicast-rpf-rib.	Enter	
		Click <b>OK</b> three times.	set rib-group multicast-rpf-rib	
Create a routing table group for the interface routes.	1.	On the main Configuration page next to Routing options, click <b>Configure</b> or <b>Edit</b> .	From the [edit] hierarchy level, enter	
	2.	Next to Rib groups, click Add new entry.	edit routing-options rib-groups.	
Configure a name for the RPF routing table group—for example, if-rib—and use inet.2 and inet.0 for its import routing tables.		In the Ribgroup name box, type if-rib.	Enter	
		Click Add new entry next to Import rib.	set if-rib import-rib inet.2	
		In the Value box, type inet.2 inet.0.		

set if-rib import-rib inet.0

## **Table 56: Configuring a PIM RPF Routing Table**

4. Click OK twice.

## Table 56: Configuring a PIM RPF Routing Table (continued)

Task	J-W	eb Configuration Editor	CLI Configuration Editor	
Add the new interface routing table group to the interface routes.		On the Routing options page next to Interface routes, click <b>Configure</b> or <b>Edit</b> .	From the <b>[edit]</b> hierarchy level, enter	
	2.	Next to Rib group, click <b>Configure</b> or <b>Edit</b> .	edit routing-options	
	3.	In the Inet box, type if-rib.	interface-routes	
	4.	Click <b>OK</b> .	set rib-group inet if-rib	

# **Verifying a Multicast Configuration**

To verify a multicast configuration, perform these tasks:

- Verifying SAP and SDP Addresses and Ports on page 119
- Verifying the IGMP Version on page 119
- Verifying the PIM Mode and Interface Configuration on page 120
- Verifying the PIM RP Configuration on page 120
- Verifying the RPF Routing Table Configuration on page 121

## **Verifying SAP and SDP Addresses and Ports**

- **Purpose** Verify that SAP and SDP are configured to listen on the correct group addresses and ports.
- **Action** From the CLI, enter the **show sap listen** command.

user@host> show sap listen Group Address Port 224.2.127.254 9875

- **What It Means** The output shows a list of the group addresses and ports that SAP and SDP listen on. Verify the following information:
  - Each group address configured, especially the default **224.2.127.254**, is listed.
  - Each port configured, especially the default **9875**, is listed.
- **Related Topics** For a complete description of **show sap listen** output, see the *JUNOS Routing Protocols and Policies Command Reference*.

## **Verifying the IGMP Version**

- **Purpose** Verify that IGMP version 2 is configured on all applicable interfaces.
- Action From the CLI, enter the show igmp interface command.

	user@host> <b>show igmp interface</b> Interface: ge-0/0/0.0 Querier: 192.168.4.36 State: Up Timeout: 197 Version: 2 Groups: 0
	Configured Parameters: IGMP Query Interval: 125.0 IGMP Query Response Interval: 10.0 IGMP Last Member Query Interval: 1.0 IGMP Robustness Count: 2
	Derived Parameters: IGMP Membership Timeout: 260.0 IGMP Other Querier Present Timeout: 255.0
What It Means	<ul><li>The output shows a list of the Services Router interfaces that are configured for IGMP.</li><li>Verify the following information:</li><li>Each interface on which IGMP is enabled is listed.</li></ul>
	<ul> <li>Next to Version, the number 2 appears.</li> </ul>
Related Topics	For a complete description of <b>show igmp interface</b> output, see the <i>JUNOS Routing Protocols and Policies Command Reference</i> .

# **Verifying the PIM Mode and Interface Configuration**

**Purpose** Verify that PIM sparse mode is configured on all applicable interfaces.

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

user@host> <b>show pim interfaces</b>						
Instance: PIM.master						
Name	Stat Mode	IP V State C	Count DR address			
100.0	Up Sparse	4 2 DR	0 127.0.0.1			
pime.32769	Up Sparse	4 2 P2P	0			

- **What It Means** The output shows a list of the Services Router interfaces that are configured for PIM. Verify the following information:
  - Each interface on which PIM is enabled is listed.
  - The network management interface, either ge–0/0/0 or fe–0/0/0, is not listed.
  - Under Mode, the word Sparse appears.
- **Related Topics** For a complete description of **show pim interfaces** output, see the *JUNOS Routing Protocols and Policies Command Reference*.

# **Verifying the PIM RP Configuration**

**Purpose** Verify that the PIM RP is statically configured with the correct IP address.

**Action** From the CLI, enter the **show pim rps**command.
	user@host> <b>shov</b> Instance: PIM.m Address family	<b>v pim rps</b> naster INET				
	RP address 192.168.14.27	Type static	Holdtime 0	Timeout Activ None	ve groups Group prefixes 2 224.0.0.0/4	
What It Means	The output shows a list of the RP addresses that are configured for PIM. At least one RP must be configured. Verify the following information:					
	■ The configu	ured RP is	listed with	the proper IF	address.	
	■ Under Type	, the word	d static appe	ears.		

**Related Topics** For a complete description of **show pim rps** output, see the *JUNOS Routing Protocols and Policies Command Reference*.

## Verifying the RPF Routing Table Configuration

**Purpose** Verify that the PIM RPF routing table is configured correctly.

Action From the CLI, enter the show multicast rpf command.

user@host> show multicast rpf
Multicast RPF table: inet.0 , 2 entries...

- **What It Means** The output shows the multicast RPF table that is configured for PIM. If no multicast RPF routing table is configured, RPF checks use inet.0. Verify the following information:
  - The configured multicast RPF routing table is **inet.0**.
  - The inet.0 table contains entries.
- **Related Topics** For a complete description of **show multicast rpf** output, see the *JUNOS Routing Protocols and Policies Command Reference.*

J-series[™] Services Router Advanced WAN Access Configuration Guide

## Part 3 Configuring DLSw Services

• Configuring Data Link Switching on page 125

J-series[™] Services Router Advanced WAN Access Configuration Guide

## Chapter 8 Configuring Data Link Switching

Data link switching (DLSw) was developed in the early 1990s as a method to transport IBM System Network Architecture (SNA) over a WAN. To route traffic over a WAN link or the Internet, DLSw encapsulates the SNA network traffic in IP. The Services Router supports DLSw as part of an SNA implementation.



**NOTE:** You must have a license to configure DLSw. For license details, see the *J*-series Services Router Administration Guide.

You can use either J-Web Quick Configuration or a configuration editor to configure DLSw. For more information about DLSw, see the *JUNOS Services Interfaces Configuration Guide*.

To monitor DLSw on a Services Router, you can use J-Web or CLI monitoring tools or SNMP.

- For information about J-Web or CLI monitoring, see the *J*-series Services Router Administration Guide.
- For SNMP monitoring with the DLSw MIB (defined in RFC 2024), you must configure SNMP on the router. For SNMP configuration instructions, see the *J-series Services Router Administration Guide*. For information about the DLSw MIB, see the *JUNOS Network Management Configuration Guide*.

This chapter contains the following topics.

- DLSw Terms on page 126
- DLSw Overview on page 127
- Before You Begin on page 129
- Configuring DLSw with Quick Configuration on page 129
- Configuring DLSw with a Configuration Editor on page 131
- Clearing the DLSw Reachability Cache on page 141
- Verifying DLSw Configuration on page 142

## **DLSw Terms**

Before configuring DLSw on a Services Router, become familiar with the terms defined in Table 57 on page 126.

#### Table 57: DLSw Terms

Term	Definition
circuit cost	Value you assign to a remote peer to indicate the relative preference for establishing a circuit through the specified peer. The lower the cost, the higher the preference.
circuit weight	Value you assign to a remote peer to indicate the extent to which the specified peer can participate in establishing circuits. The higher the circuit weight, the greater the percentage of total circuits established with this remote peer.
destination service access point (DSAP)	Service access point (SAP) that identifies the destination for which a logical link control protocol data unit (LPDU) is intended.
DLSw circuit	Path formed by establishing a data link control (DLC) connection between each locally configured SNA end system and a local router configured for DLSw. A DLSw circuit is identified by the circuit ID, which includes the SNA end system MAC address, local service access point (LSAP), destination MAC address, and destination service access point (DSAP). Multiple DLSw circuits can operate over the same DLSw connection.
DLSw connection	Set of TCP connections between two DLSw peers that is established after the initial handshake and successful capabilities exchange.
explorer timeout	Number of seconds a DLSw router waits for a response from its peers to its explorer requests.
I-frame	Information frame used to transfer sequentially numbered logical link control protocol data units (LPDUs) between link stations.
Logical Link Control (LLC)	Data-link layer protocol used on a LAN. LLC1 provides connectionless data transfer, and LLC type 2 provides connection-oriented data transfer.
LLC protocol data unit (LPDU)	Logical link control (LLC) frame on a DLSw network.
local reachability cache	Cache of pairs of local media access control (MAC) addresses and local Logical Link Control (LLC) IP addresses, maintained on a DLSw router for a specified number of seconds. The router uses the local cache to determine whether a local SNA host is reachable through any of the router's LLC interface.
preemption	Process by which a master router takes over from a backup router after recovering from a failure incident.
priority-cost	Value that is deducted from the priority value of a router to determine when it takes over for a master router.
redundancy group	Group of DLSw peer routers on the same Ethernet segment of a network.
remote reachability cache	Cache of pairs of remote media access control (MAC) addresses and remote peer IP addresses, maintained on a DLSw router for a specified number of seconds. The router uses the remote cache to determine whether a remote SNA host is reachable through any of the router's remote peers.

#### Table 57: DLSw Terms (continued)

Term	Definition
service access point (SAP)	OSI term for the component of a network address that identifies the individual application sending or receiving a packet on a host.
source service access point (SSAP)	Service access point (SAP) that identifies the origin of an LPDU on a DLSw network.
Switch-to-Switch Protocol (SSP)	Protocol implemented between two DLSw routers that establishes connections, locates resources, forwards data, and handles error recovery and flow control.

#### **DLSw Overview**

Data link switching (DLSw) was developed in the 1990s as a method to transport IBM Systems Network Architecture (SNA) traffic over an IP WAN network. Switch-to-Switch Protocol (SSP) is used to forward network traffic between routers configured for DLSw (DLSw peers). Then, to route traffic over a WAN link or the Internet, DLSw encapsulates the SNA network traffic into IP packets.

DLSw was developed as a forwarding mechanism for IBM Systems Network Architecture (SNA) protocol. Although DLSw does not provide full routing capabilities, it provides switching at the data link layer and encapsulation in TCP/IP for transport over the Internet.

Because DLSw provides support for SNA, a connection-oriented protocol, the Services Router supports Logical Link Control (LLC) type 2 as part of the DLSw implementation. Figure 9 on page 127 shows a possible DLSw network.

#### **Figure 9: Sample DLSw Network** IBM Internet (WAN) PC mainframe PC 206.142.22.3 64.128.10.1 IBM J-series J-series IP traffic mainframe PC PC DLSw DLSw SNA (LLC2) SNA (LLC2)

### Switch-to-Switch Protocol for DLSw

Switch-to-Switch Protocol (SSP) is used between DLSw peers to establish connections, locate resources, forward data, and handle error recovery as well as flow control. Generally, SSP does not provide full routing between peers, because routing is typically handled by common routing protocols such as OSPF or BGP. Instead, packets are switched at the SNA data link layer and encapsulated in TCP/IP for transport over IP-based networks. TCP is used as reliable transport method between DLSw peers.

g017134

#### **DLSw Operational Stages**

There are several operational stages that take place in DLSw connections. First, two DLSw peers establish a TCP connection with each other. After the connection is established, each peer router exchanges supported capabilities with the other router. The TCP connection ensures reliable and guaranteed delivery of IP traffic, and also ensures the integrity and delivery of traffic encapsulated in the IP protocol. After capability information is exchanged, the DLSw peers establish circuits between SNA end systems and begin transmitting information frames (I-frames) over the network.

#### **DLSw Capabilities Exchange**

DLSw capabilities exchange is based on a switch-to-switch protocol message describing the capabilities of the sending data-link switch. Sent just after the DLSw peers establish a connection, a capabilities exchange control message communicates the following operational parameters between the two peers:

- DLSw version number
- Initial pacing window size (receive window size)
- List of supported link SAPs (LSAPs)
- Number of supported TCP sessions
- Lists of media access control (MAC) addresses

## **DLSw Circuits Establishment**

Establishing DLSw circuits is a process in which local and remote DLSw peers locate each other and set up data link control (DLC) connections between the remote router and a local router. The specific details of establishing circuits are determined by the traffic type, but the process is the same for all types of traffic.

The first step in the process enables the SNA devices on a LAN to find other SNA devices by sending out an explorer frame with the MAC address of the target SNA device. When a DLSw peer receives the explorer frame, it sends a canureach message frame to each of its DLSw peer connections. The canureach message frame queries the DLSw peers to determine if one of the peers can locate the target SNA device. If one of the DLSw peers can reach the target SNA device, it returns an icanreach message frame to the originating DLSw peer to indicate that it can provide a path to the SNA device in question.

After canureach and icanreach message frames are exchanged, the two DLSw peers establish a circuit consisting of a DLC connection between each router and the local SNA end system and a TCP connection between the two DLSw peers. The resulting circuit is uniquely identified by source and destination circuit IDs. Each SNA DLSw circuit ID includes the following information:

- MAC address of the SNA end system
- Link service access point (LSAP)
- DLC port ID

Circuit priority is negotiated when the circuit is set up on the network.

## **Class of Service for DLSw**

You can use the class-of-service (CoS) features on a Services Router to classify DLSw packets and assign them to queues by a type-of-service (TOS) precedence value.

For more information, see "Configuring CoS for DLSw (Optional)" on page 134.

#### **DLSw Ethernet Redundancy**

When more than one DLSw router is configured on the same LAN segment, the DLSw design limits redundancy and load sharing. To ensure a recovery point in case of router failure, DLSw Ethernet redundancy supports parallel paths between two points in an Ethernet environment. You can assign priorities to enable one DLSw router to operate as the master router.

For more information, see "Configuring DLSw Ethernet Redundancy (Optional)" on page 136.

#### **DLSw Peer Preference and Load Balancing**

When more than one remote DLSw peer provides a path to a WAN destination, you can assign a relative cost to each peer to establish preferred DLSw circuits. In addition, you can assign a relative weight to each circuit to balance the number of circuits going to each peer.

For more information, see "Configuring DLSw Peer Preference and Load Balancing (Optional)" on page 139.

## **Before You Begin**

Before you begin configuring DLSw, complete the following tasks:

- Establish basic connectivity. See the Getting Started Guide for your router.
- Configure network interfaces. See the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.
- If you do not already have an understanding of DLSw, read "DLSw Overview" on page 127.

## **Configuring DLSw with Quick Configuration**

You can use the DLSw Quick Configuration page to configure DLSw on a Services Router. The Quick Configuration page allows you to designate the peer routers that make up the DLSw network.

Figure 10 on page 130 shows the DLSw Quick Configuration page.

#### Figure 10: DLSw Quick Configuration Page

			ROUT	'ER - J63	00				
Monitor Configuration	n Diagnose	Manage	Events	Alarms	Logged	in as: regre	ss Help	About L	ogout
Guish Coofiguration → View and Edit ► History Rescue	Quick Cor Routing	nfigurati and Pro	on otocols		Configu	ration > Quick	Configuration >	Routing and P	
	DLSW Cor Connecti Enable Pro	nfigurati on Idle Ti miscuous Loca Remote	on meout Mode Peer Peer		Add D	? ? ?			
	OK	LLC 1 Cancel	ype 2 Apply	Interface Confi	with LLC2 igured	-> (=	Interface w Confi fe-0/0/0	vithout LLC2 igured	•
Converse t @ 2004-2005	luniner Networks		the Reserv	ed Trademark	Notice Privacy		Juniper 14	our Net.	

To configure DLSw with Quick Configuration:

- 1. In the J-Web interface, select **Configuration > Quick Configuration > Routing** and **Protocols > DLSw Protocol**.
- 2. Enter information into the DLSw Quick Configuration page, as described in Table 58 on page 131.
- 3. Click one of the following buttons on the DLSw Quick Configuration page:
  - To apply the configuration and stay in the DLSw Quick Configuration page, click **Apply**.
  - To apply the configuration and return to the Routing and Protocols Quick Configuration page, click **OK**.
  - To cancel your entries and return to the Routing and Protocols Quick Configuration page, click **Cancel**.
- 4. To verify the configuration, see "Verifying DLSw Configuration" on page 142.

Field	Function	Your Action
Connection Idle Timeout	Specifies the length of time, in seconds, a remote DLSw Services Router can be idle before the network connection times out.	Type a value between 0 and 60000.
Enable Promiscuous Mode	Enables or disables promiscuous mode. If enabled, the Services Router accepts all incoming DLSw connections.	To enable promiscuous mode, select <b>Enable</b> <b>Promiscuous Mode</b> .
		To disable promiscuous mode, clear the <b>Enable</b> <b>Promiscuous Mode</b> check box.
Local Peer	Adds the IP address of the local DLSw Services Router.	Type the IPv4 address of the local router in the <b>Local Peer</b> box.
Remote Peer	Configures the IP addresses of the remote DLSw Services Routers.	Type the IPv4 address of a remote router in the IP address box. Click <b>Add</b> to add each remote router.
Interface with LLC2 Configured	Sets or deletes LLC type 2 properties for an Ethernet interface on a DLSw Services Router.	To set LLC type 2 properties on an Ethernet interface, select it, and click the left arrow.
Interface without LLC2 Configured		To delete LLC type 2 properties on an Ethernet interface, select it, and click the right arrow.

### Table 58: DLSw Quick Configuration Page Summary

## **Configuring DLSw with a Configuration Editor**

To configure basic DLSw on a Services Router, perform the following task marked *(Required)*:

- Configuring Basic DLSw (Required) on page 131
- Configuring CoS for DLSw (Optional) on page 134
- Configuring DLSw Ethernet Redundancy (Optional) on page 136
- Configuring DLSw Peer Preference and Load Balancing (Optional) on page 139

**NOTE:** To configure other properties for DLSw, see the *JUNOS Services Interfaces Configuration Guide*.

## **Configuring Basic DLSw (Required)**

To configure basic DLSw on a Services Router, perform the following tasks:

- Configuring LLC Type 2 Properties on an Ethernet Interface on page 132
- Configuring DLSw on the Local Services Router on page 132
- Configuring DLSw on the Remote Services Router on page 134

### **Configuring LLC Type 2 Properties on an Ethernet Interface**

Before configuring DLSw on the Services Router, you must configure the LLC type 2 properties on the Ethernet interfaces of the router. The Logical Link Control (LLC) layer is one of two sublayers into which the OSI data link layer is subdivided for data link protocols used on the LAN. LLC type 2 is implemented anytime SNA is running on a LAN or virtual LAN.

**NOTE:** LLC type 2 properties must be configured on the local Services Router and the remote Services Router.

To configure LLC type 2 properties:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 59 on page 132.
- 3. Go on to one of the following required configurations:
  - To configure DLSw on the local Services Router, go on to "Configuring DLSw on the Local Services Router" on page 132.
  - To configure DLSw on the remote Services Router, go on to "Configuring DLSw on the Remote Services Router" on page 134.
- 4. To verify the basic DLSw properties, see "Verifying DLSw Configuration" on page 142.

#### Table 59: Configuring LLC Type 2 Properties on a Fast Ethernet Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Interfaces</b> level in the configuration hierarchy and select a Fast Ethernet interface—for example <b>fe</b> -3/0/1	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit Configuration.</li> <li>Next to Interfaces, click Configuration</li> </ol>	From the [edit] hierarchy level, enter > Edit edit interfaces fe-3/0/1 ure or Edit.
3. Click <b>fe-3/0/1</b> .		
Configure LLC type 2 properties on the fe-3/0/1 interface.	<ol> <li>Under Unit and Interface unit no</li> <li>Under Family, select Llc2.</li> <li>Click OK until you return to the Configuration page.</li> </ol>	umber, click <b>0</b> . 1. Enter edit unit <b>0</b> main 2. Enter set family llc2

## **Configuring DLSw on the Local Services Router**

To configure DLSw on the local Services Router, you do the following:

Define a local peer.

- Define a remote peer.
- Finally, define connection behavior.

The example in this section shows how to configure DLSw on the local and remote Services Routers with IP addresses listed in Table 60 on page 133. The remote Services Router initiates the peer connection.

#### Table 60: Sample DLSw Peer Router Values

Option	Value		
remote-peer	217.110.111.134		
local-peer	110.0.10.1		

In this example, the local router is configured with **remote-peer** settings because the local router is initiating the connection for SNA traffic over the WAN interface. The remote router is accepting DLSw connections from any DLSw peers.

To configure basic DLSw on the local router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 61 on page 133.
- 3. Go on to "Configuring DLSw on the Remote Services Router" on page 134.

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Dlsw</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter
	<ol> <li>Next to Protocols, click <b>Configure</b> or <b>Edit</b>.</li> </ol>	
	3 Next to Disw make sure the check hox is	
	selected, and click <b>Configure</b> or <b>Edit</b> .	
Configure the local router	In the Local peer box, type <b>110.0.10.1</b> .	Enter
properties.		set local-peer 110.0.10.1
Configure the remote peer	1. Next to Remote peer, click <b>Configure</b> .	Enter
settings.	2. Click Add new entry.	set remote-peer 217,110,111,134
Because the remote router	3. In the Peer ip box, type <b>217.110.111.134</b> .	
is initiating the peer connection, configure the <b>remote-peer</b> setting.	4. Click <b>OK</b> until you return to the Protocols pag	ge.

#### **Table 61: Configuring DLSw on the Local Router**

#### **Configuring DLSw on the Remote Services Router**

To configure DLSw on the remote Services Router, you do the following:

- Define a local peer.
- Define a remote peer.
- Finally, define the connection behavior.

To configure DLSw on a remote router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 62 on page 134.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the DLSw configuration, see "Verifying DLSw Configuration" on page 142.

#### **Table 62: Configuring DLSw on the Remote Router**

Task	J-We	b Configuration Editor	CLI	Configuration Editor
Navigate to the <b>Dlsw</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	Fro edi	m the [edit] hierarchy level, enter
	2.	Next to Protocols, click <b>Configure</b> or <b>Edit</b> .		
	3.	Next to Dlsw, make sure the check box is selected, and click <b>Configure</b> or <b>Edit</b> .		
Configure the local router	1.	In the Local peer box, type 217.110.111.134.	1.	Enter
properties.	2.	Next to Promiscuous, select Yes.		set local-peer 217.110.111.134
promiscuous—Allows all incoming peer	3.	Click <b>OK</b> .	2.	Enter
connections.				set promiscuous



**NOTE:** If the values connection-idle-timeout, dlsw-cos, local-peer, multicast-address, promiscuous, and receive-initial-pacing are modified, any existing DLSw peer connection is torn down. If remote-peer *peer-address* is added or removed, only that remote peer and its associated circuits are affected.

## Configuring CoS for DLSw (Optional)

The J-series Services Router CoS features provide differentiated services when best-effort traffic delivery is not enough. You can use CoS to classify DLSw packets. The packets are sent to a logical tunnel interface on the router, where they are classified and queued based on the configured type-of-service (ToS) value. For information about CoS, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide or the JUNOS Class of Service Configuration Guide.

To configure CoS for DLSw on the Services Router, you do the following:

- Configure the logical tunnel It-0/0/0 interface.
- Configure the CoS classifier on the lt-0/0/0 interface.
- Configure the DLSw type-of-service (ToS) precedence on the lt-0/0/0 interface.

To configure CoS classification for DLSw on a router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 63 on page 135.
- 3. If you are finished configuring the router, commit the configuration.

Task	J-W	eb Configuration Editor	CL	Configuration Editor
Navigate to the <b>Interfaces</b> level in the configuration hierarchy.	1.	<ul> <li>In the J-Web interface, select</li> <li>Configuration &gt; View and Edit &gt; Edit</li> <li>Configuration.</li> </ul>		m the [edit] hierarchy level, enter
	2.	Next to Interfaces, click <b>Configure</b> or <b>Edit</b> .		
Configure the first logical	1.	Next to Interface, click Add new entry.	1.	Enter
interface.	2.	In the Interface name box, type $lt-0/0/0$ .		set unit 0
(See the interface naming	3.	Click OK.	2.	Enter
conventions in the <i>J-series</i> Services Router Basic LAN	4.	Next to lt-0/0/0, click <b>Edit</b> .		
	5.	Next to Unit, click Add new entry.	_	set dici 10
Configuration Guide.)	6.	In the Interface unit number box, type $0$ .	3.	Enter
	7.	In the Dlci box, type <b>10</b> .		set encapsulation frame-relay
	8.	From the Encapsulation list, select <b>frame-relay</b> .	4.	Enter
	9.	In the Peer unit box, type 1.		set peer-unit 1
	10. Under Family, select Inet.		5.	Enter
	11.	Click <b>OK</b> .		set family inet

#### Table 63: Configuring CoS for DLSw on the Remote Router

#### Table 63: Configuring CoS for DLSw on the Remote Router (continued)

Task	-Web Configuration Editor	CLI Configuration Editor
Configure the second	I. Next to Unit, click Add new er	try. 1. Enter
logical unit on the lt-0/0/0 interface.	2. In the Interface unit number	box, type 1. set unit 1
	3. In the Dlci box, type <b>10</b> .	2. Enter
	4. From the Encapsulation list, se	elect frame-relay.
	5. In the Peer unit box, type <b>0</b> .	set dici 10
	6. Under Family, select Inet.	3. Enter
	7. Click <b>OK</b> until you return to t	ne main set encapsulation frame-relay
	Configuration page.	4. Enter
		set peer-unit 0
		5. Enter
		set family inet
Configure the default CoS classifier on the lt-0/0/0	I. On the main Configuration p of service, click <b>Edit</b> .	age next to Class From the [edit] hierarchy level, enter
interface.	2. Next to Interfaces, click Add	new entry. edit class-of-service interfaces It-0/0/0 unit
	3. In the Interface name box, ty	pe It-0/0/0.
	4. Next to Unit, click Add new	entry. Enter
	5. In the Unit number box, type	1. set classifiers dscp default
	6. Next to Classifiers, click <b>Con</b>	igure.
	7. Under Dscp, in the Classifier default.	name box, type
	<ol> <li>Click <b>OK</b> until you return to t Configuration page.</li> </ol>	ne main
Configure the type-of-service precedence	I. On the main Configuration p Protocols, click <b>Configure</b> or	age next to 1. From the [edit] hierarchy level, enter Edit.
value for DLSw packets—for example, 192.	2. Next to Dlsw, make sure the selected, and click <b>Configure</b>	check box is edit protocols dlsw dlsw-cos or Edit. 2. Enter
	3. Next to Dlsw cos, click <b>Confi</b>	gure or Edit.
	<ol> <li>In the Destination interface b It-0/0/0.0.</li> </ol>	ox, type type-of-service 192
	5. In the Type of service box, ty	pe <b>192</b> .
	5. Click <b>OK</b> .	

## **Configuring DLSw Ethernet Redundancy (Optional)**

When more than one DLSw router is connected on the same LAN segment, there are DLSw design limitations for providing redundancy and load sharing. When DLSw

Ethernet redundancy is configured on the network, it enables DLSw to support parallel paths between two points in an Ethernet environment, ensuring a recovery point in the case of router failure.

When DLSw Ethernet redundancy is configured on a LAN segment, one router (DLSw peer), is selected to act as the master router, and other routers become backup routers, depending on the configured priority, in a group of DLSw peers. Only the master router establishes circuits and connections on the LAN and maintains a database of known DLSw peers on the network. By maintaining a circuit database, the master router prevents duplicate circuits from being created for the same SNA session. In addition, only the master router accepts incoming LLC connections while the backup routers simply drop the connections.

When the master router fails, all incoming connections cease, and the backup router with a higher priority than other backup routers becomes the master router and begins handling all connections.

Figure 11 on page 137 shows a typical use of Ethernet LAN redundancy in a DLSw network.



#### Figure 11: DLSw Ethernet Redundancy Network Topology

In Figure 11 on page 137, the local hosts share the same destination MAC address of **00:22:22:22:22:22** and send DLSw traffic to the remote host with a MAC address of **00:30:48:84:99:45**. Router 1 and Router 2 are configured as a DLSw redundancy group and map the local destination MAC address to the remote MAC address. Router 1 is the designated master and if Router 1 becomes unavailable, Router 2 takes over as the master router.

The priority cost feature is used to determine the effective priority by subtracting the priority cost from the configured priority when a tracked event occurs, such as the unavailability of a remote DLSw peer.

To configure DLSw Ethernet redundancy on the DLSw peer Services Router, you do the following:

- Define the redundancy groups on each peer.
- Define the redundancy group options on each peer.
- Finally, define the priority cost of each redundancy group option.

To configure DLSw Ethernet redundancy on a DLSw peer:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 64 on page 138.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the DLSw configuration, see "Verifying DLSw Configuration" on page 142.

Task		/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Interfaces</b> level in the configuration hierarchy.		In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit interfaces fe-1/0/0 unit 0 family llc2
	2.	Next to Interfaces, click <b>Configure</b> or <b>Edit</b> .	
Edit the LLC type 2 properties	1.	Next to the interface fe-1/0/0, click Edit.	_
on a Fast Ethernet interface—for example,	2.	Next to Unit, click <b>Edit</b> .	
fe-1/0/0.	3.	Under Family, select <b>Llc2</b> , and then click <b>Configure</b> .	
Create a redundancy group—for example <b>100</b> .	1.	Next to Redundancy group, click <b>Add new</b> entry.	Enter set redundancy-group 100
	2.	In the Group Id box, type 100.	
Map a local peer MAC address	1.	Next to Map, select Yes.	Enter
to a remote peer MAC address. For instance, the local peer	2.	Click <b>Configure</b> .	set redundancy-group 100 map
MAC address is	3.	Next to Local mac, click Add new entry.	local-mac 00:22:22:22:22:22
00:22:22:22:22:22 and the remote peer MAC address is 00:30:48:84:99:45.	4.	In the Local address box, type 00:22:22:22:22:22.	remote-mac 00:30:48:84:99:45
	5.	In the Remote mac box, type 00:30:48:84:99:45.	
	6.	Click <b>OK</b> .	
Configure a priority value		the Priority box, type <b>250</b> .	Enter
group. The default value is 100.			set redundancy-group 100 priority 250
The priority value determines which DLSw peer becomes the master router during master router selection			

#### Table 64: Configuring DLSw Ethernet Redundancy on a DLSw Peer Router

Task	J-W	eb Configuration Editor	CLI Configuration Editor	
Configure tracking options for	1.	Next to Track, click <b>Configure</b> .	Enter	
the remote peer and destination.	2.	Next to DLSw, click <b>Configure</b> .	set redundancy-group 100 track dlsw	
The track perspector is used to	3.	Next to Destination, click Add new entry.	destination 00:22:22:22:22:22 priority-cost	
track events such as the	4.	In the Mac address box, type	30	
unavailability of a remote		00:30:48:84:99:45.	Enter	
DLSw peer.	5.	In the Priority cost box, type <b>50</b> .	sat radundanay group 00.20.48.84.90.45	
Priority cost is subtracted from	6.	Click OK.	track dlsw peer 10.10.10.1 priority-cost 30	
the priority value when remote peer connectivity is lost, and	7.	Next to Peer, click Add new entry.		
has a value between 1 and 254.	8.	In the Ip address box, type the IP address of the remote peer—for example, <b>10.10.10.1</b> .		
	9.	In the Priority cost box, type <b>30</b> .		
	10.	Click <b>OK</b> until you return to the Redundancy group page.		
Configure advertisement of DLSw peers on the network.	1.	From the Advertisement type list, select Advertise interval.	Enter	
Advertise interval has a value between 1 and 255 seconds.	2.	In the Advertise interval box, type 1.	set redundancy-group 100 advertise-interval	
The default value is 1.	3.	From the Preemption type list, select ${f no}$	_	
The preempt parameter		preempt.	Enter	
determines if a higher-priority backup router takes over for a lower-priority master router.	4.	Click <b>OK</b> .	set redundancy-group group 100 no-preempt	

#### Table 64: Configuring DLSw Ethernet Redundancy on a DLSw Peer Router (continued)

## **Configuring DLSw Peer Preference and Load Balancing (Optional)**

For a DLSw J-series router, when more than one remote DLSw peer provides alternate paths to a remote destination on a WAN, you can specify preferences by assigning costs among the available routers (peers) or enable load balancing for lowest equal-cost alternatives. The DLSw router maintains a reachablity cache of paired MAC address and IP address entries to determine whether an SNA host can be reached by means of any of the peers the router has information about.

Consider a WAN in which the DLSw Services Router R1 has a peer relationship with more than one peer routers as shown in Figure 12 on page 140. The peer routers R2 and R3 are manufactured by vendors other than Juniper Networks.



#### Figure 12: DLSw Peer Preference and Load-Balancing Network Topology

As shown in Figure 12 on page 140, the far-end routers R2 and R3 provide alternate paths to Host H2 from Router R1. Router R2 has an IP address of **192.168.17.2**, and Router R3 has an IP address of **192.168.18.2**. A DLSw circuit between the local host H1 and the remote host H2 can be established through either R2 or R3.

By default, a Services Router has no preference for a next-hop router among its DLSw peers. Router R1 checks its reachability cache for entries. If none exist, R1 sends a canureach message to peers R2 and R3 and selects the first responding router as the next hop to the destination host H2.

You can specify preferences among peers R2 and R3 by assigning a different cost to each. For example, if you assign a cost of 50 to R2 and a cost of 60 to R3, Router R2 is the preferred next-hop peer. Then, Router R1 waits for a specified period of time to get a response from R2. If both R2 and R3 respond, the circuit is routed through R2. If R2 does not respond in the specified time, and R3 responds, then the DLSw router R1 accepts R3's response and the circuit is routed through R3.

To ensure load balancing among peers, you must assign the least cost for the peer routers, and additionally assign them different circuit weights. Assigning circuit weights ensures that the number of circuits going through each peer is balanced according to the circuit weight configured on each peer. For example, if R2 and R3 both have a cost of 50, but R3 can handle more DLSw traffic, then you can assign a circuit weight of 1 to R2 and a circuit weight of 2 to R3 to ensure that twice as much DLSw traffic is routed to Router R3.

To configure DLSw load balancing:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 65 on page 141.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify the DLSw configuration, see "Verifying DLSw Configuration" on page 142.

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Dlsw</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter edit protocols dlsw
	2. Next to Protocols, click <b>Configure</b> or <b>Edit</b> .	
	<ol> <li>Next to Dlsw, make sure the check box is selected, and click Configure or Edit.</li> </ol>	
	<b>NOTE:</b> You can also navigate through the navigation hierarchy in the left pane.	
Configure the load-balancing settings for the first remote DLSw peer:	<ol> <li>Next to Remote peer, click Configure.</li> </ol>	1. Enter
■ IP address—for example,	2. Click Add new entry.	set remote-peer 192.168.17.2
<ul><li>192.168.17.2</li><li>Circuit weight of between 1 and</li></ul>	<ol> <li>In the Peer ip box, type 192.168.17.2.</li> </ol>	2. Enter
■ Circuit cost of between 0 and	4. In the Circuit weight box, type 1.	3 Enter
127—for example, <b>50</b>	5. In the Cost box, type <b>50</b> .	J. Enter
<ul> <li>Keepalive interval of between 0 and 4294967295 seconds—for</li> <li>example 20 The default interval</li> </ul>	<ol> <li>In the Keepalive interval box, type 20.</li> </ol>	set cost 50 4. Enter
is 10 seconds. Setting an interval of 10 seconds ensures that the	<ol> <li>Click <b>OK</b> until you return to the DLSw page.</li> </ol>	set keepalive-interval 20
circuit is always available.	8. Repeat Steps 1 through 7 for the second remote peer.	5. Repeat Steps 1 through 4 for the second remote peer.
remote peer, using an IP address of <b>192.168.18.2</b> and a circuit weight of <b>2</b> .		
Configure the interval during which the DLSw router waits for a response to its	<ol> <li>In the Explorer wait time box, type</li> <li>5.</li> </ol>	1. From the edit protocols dlsw hierarchy level, enter
explorer requests from the peer routers. The interval ranges from 5 through 60 seconds, and the default value is 10	2. In the Reachability cache timeout box, type <b>300</b> .	set explorer-wait-time 5
seconds.	3. Click <b>OK</b> to return to the	2. Enter
Configure the interval for retaining entries in the reachability cache. The interval ranges from 100 through 3600 seconds, and the default value is 900 seconds.	Configuration Protocols page.	set reachability-cache-timeout 300

#### Table 65: Configuring DLSw Peer Preference and Load Balancing on DLSw and Peer Routers

## **Clearing the DLSw Reachability Cache**

You can delete all the entries from the reachability cache for the DLSw load-balancing feature by applying the **clear** command. From the CLI, enter the **clear dlsw reachability** command.

#### user@host> clear dlsw reachability

## **Verifying DLSw Configuration**

To verify DLSw configuration, perform these tasks:

- Displaying LLC Type 2 Properties on a Fast Ethernet Interface on page 142
- Displaying DLSw Capabilities on page 142
- Displaying DLSw Circuit State on page 143
- Displaying Details of a DLSw Circuit State on page 143
- Displaying DLSw Peers on page 144
- Displaying Details of DLSw Peers on page 144
- Displaying DLSw Reachability Information on page 145
- Displaying DLSw Ethernet Redundancy Properties on page 146
- Displaying DLSw Ethernet Redundancy Statistics on page 146

## **Displaying LLC Type 2 Properties on a Fast Ethernet Interface**

**Purpose** Verify the configuration of LLC type 2 properties on a Fast Ethernet interface.

Action From the J-Web interface, select Configuration > View and Edit > View Configuration Text. Alternatively, from configuration mode in the CLI, enter the show interfaces fe-3/0/0 command.

user@host# show interfaces fe-3/0/0
fe-3/0/0 {
unit 0 {
family inet{
address 172.5.20.1/24;
}
family IIc2}
}
}

What It Means	Verify that the	output shows	the intended	LLC type 2	configuration.
---------------	-----------------	--------------	--------------	------------	----------------

**Related Topics** For more information about the format of a configuration file, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

## **Displaying DLSw Capabilities**

	user@host> show dlsw capabilities
Action	From the CLI, enter the show dlsw capabilities command.
Purpose	Verify DLSw capabilities of remote DLSw peers.

	Peer: 50.50.50
	Vendor ID :000585
	Version number :0200
	Initial pacing window size :32
	Version String
	Juniper Networks, Inc. j2300 internet router
	JUNOS Software Release 7.4IO [builder]
	Build date: 2005-07-15 07:13:17 UTC
	Copyright (c) 1996-2005 Juniper Networks,Inc.
	Compiled Wed 26-Jan-05 02:49 by pwade
What It Means	Verify that the output correctly displays the capabilities of remote DLSw peers.

**Related Topics** For a complete description of **show dlsw capabilities** output, see the *JUNOS System Basics and Services Command Reference*.

## **Displaying DLSw Circuit State**

Purpose	Display DLSw circuits currently established after configuration in "Configuring Basic DLSw (Required)" on page 131.		
Action	From the CLI, enter the show dlsw circuits command.		
	user@host> <b>show dlsw circuits</b> Local address LSAP Remote address DSAP Peer Uptime 22:22:00:00:06 04 44:44:00:00:00:06 04 18.255.18.2 00:06:42		
What It Means	The output shows a summary of DLSw circuits. Verify that the information is correct for your DLSw network.		
	■ Local address—MAC address of the local DLSw peer		
	■ LSAP—Number of the local service access point		
	■ <b>Remote address</b> —MAC address of the remote DLSw peer		
	■ DSAP—Number of the destination service access point		
	■ Peer (or remote peer address)—IP address of the remote DLSw peer		
	<ul> <li>Uptime—How long the circuit has been established</li> </ul>		
<b>Related Topics</b>	For a complete description of <b>show dlsw circuits</b> output, see the <i>JUNOS System Basics and Services Command Reference</i> .		

### **Displaying Details of a DLSw Circuit State**

- **Purpose** Display the details of DLSw circuits currently established after configuration in "Configuring Basic DLSw (Required)" on page 131.
  - Action From the CLI, enter the show dlsw circuits detail command.

user@host> show dlsw circuits detail

```
Circuit ID: 9ad20498aa04
 Local address: 22:22:00:00:00:06, LSAP: 04
 Remote address: 44:44:00:00:00:06, DSAP: 04
 Remote peer address: 18.255.18.2
 Circuit state: Connected
 Uptime: 00:09:02
 Max BTU size: 1466
 Circuit priority: 3
 Statistics:
 I-frames received
 : 0
 I-frames sent
 : 0
 Bytes in I-frames received : 0
 Bytes in I-frames sent : 0
I frames rejected : 0
 I frames rejected
 : 0
 Bytes in I-frames rejected : 0
 I-frames retransmitted
 : 0
 Bytes in retransmitted I-frames : 0
 Reject frames received : 0
 Reject frames sent
XID frames received
XID frames sent
 : 0
 : 2
 XID frames sent
 : 2
```

- **What It Means** In addition to the local and remote MAC addresses, the priority, the maximum basic transmission unit (BTU) size, and the statistics are displayed.
- **Related Topics** For a complete description of **show dlsw circuits detail** output, see the *JUNOS System Basics and Services Command Reference.*

#### **Displaying DLSw Peers**

- **Purpose** Display information about the DLSw peers on the network.
- Action From the CLI, enter the show dlsw peers brief command.

#### user@host> show dlsw peers brief

Peer	State	Circuits	Uptime
17.255.17.2	Connected	0	00:00:00
18.255.18.2	Connected	1	00:12:03

- What It Means The output displays the number of active or inactive DLSw peers.
- **Related Topics** For a complete description of **show dlsw peers brief** output, see the *JUNOS System Basics and Services Command Reference.*

#### **Displaying Details of DLSw Peers**

Purpose	Display	detailed	informatior	n about DLSw	peers on a	ı network
---------	---------	----------	-------------	--------------	------------	-----------

**Action** From the CLI, enter the show dlsw peers detail command.

user@host> show dlsw peers detail

```
Peer: 18.255.18.2
 State: Connected, Circuits: 1, Local address: 10.255.4.50
 Uptime: 00:15:05
 Receive initial pacing: 20, No circuits timeout: 0
 Type-of-service value: 0
 Peer cost: 100, Load balancing: Circuit Weight
 Circuit weight: 2
 Statistics:
 Data packets received : 0
 Data packets sent : 0
Data bytes received : 0
 Data bytes sent : 0
 Control packets received : 7
 Control packets sent : 8
 CANUREACH_ex received : 0
 CANUREACH_ex sent
 : 1
 ICANREACH_ex received : 1
ICANREACH_ex sent : 0
```

**What It Means** The output displays the DLSw peer state and the following statistics:

- Packets received—Number of packets received from DLSw peers
- Packets sent—Number of packets sent to the DLSw peers
- Bytes received—Number of bytes received from DLSw peers
- Bytes sent—Number of bytes sent to the DLSw peers
- CANUREACH_ex received—Number of exploratory messages received from remote DLSw peers
- CANUREACH_ex sent—Number of exploratory messages sent to remote DLSw peers
- ICANREACH_ex received—Number of confirmation messages received from remote DLSw peers
- ICANREACH_ex sent—Number of confirmation messages sent to remote DLSw peers
- **Related Topics** For a complete description of **show dlsw peers detail** output, see the *JUNOS System Basics and Services Command Reference.*

### **Displaying DLSw Reachability Information**

- **Purpose** Display information about the MAC cache entries and peer IP addresses currently maintained on the DLSw router.
- Action From the CLI, enter the show dlsw reachability command.

user@host> **show dlsw reachability** 

MAC	index	MAC address	Location	Peer/Interface
0	44:4	4:00:00:00:06	remote	192.168.17.2
				192.168.18.2

	1	22:22:00:00:00:06	local	ge-0/0/1.0	
What It Means	The	output displays the DLS	Sw reachat	pility details:	
	•	MAC index—Number as	signed to t	he DLSw peer	
	•	MAC address—MAC add	dress of th	e DLSw peer	
	•	Location—Local or rem	ote peer		
	•	Peer/interface—Interface remote DLSw peer	ce location	of the local DLSw peer or IP address of th	ıe
Related Topics	For	a complete description	of the <b>sho</b> y	w disw reachability command see the <i>IUN</i>	0.9

## **Related Topics** For a complete description of the **show dlsw reachability** command, see the *JUNOS System Basics and Services Command Reference*.

## **Displaying DLSw Ethernet Redundancy Properties**

**Purpose** Display information about the DLSw Ethernet redundancy state.

**Action** From the CLI, enter the **show llc2 redundancy brief** command.

user@host> **show llc2 redundancy brief** Interface Unit Group Int state ER state ge-0/0/0.0 0 0 up backup

- **What It Means** The output displays the state of the group and the interface. It also indicates if the router is the master router or the backup router.
- **Related Topics** For a complete description of **show llc2 redundancy** output, see the *JUNOS System Basics and Services Command Reference.*

#### **Displaying DLSw Ethernet Redundancy Statistics**

- **Purpose** Display statistics about the number of keepalives sent and received as well as errors detected.
  - Action From the CLI, enter the show llc2 redundancy interface statistics command.

user@host> **show llc2 redundancy interface statistics** Interface: ge-0/0/0.0, Index: 68, Group:0 Interface ERED PDU statistics Advertisement sent :0 Advertisement received :33240 Interface ERED PDU error statistics Invalid ERED TTL value received :0

- **What It Means** The output displays the number of advertisements sent and received as well as any invalid Ethernet redundancy time-to-live (TTL) packets.
- **Related Topics** For a complete description of show llc2 redundancy interface statistics output, see the *JUNOS System Basics and Services Command Reference.*

# Part 4 Configuring a Policy Framework

- Policy Framework Overview on page 149
- Configuring Routing Policies on page 169
- Configuring NAT on page 185
- Configuring Stateful Firewall Filters and NAT on page 205
- Configuring Stateless Firewall Filters on page 221

J-series[™] Services Router Advanced WAN Access Configuration Guide

## Chapter 9 Policy Framework Overview

To control the way routing information and data packets are handled, a Services Router uses the JUNOS policy framework. This framework consists of routing and firewall filter policies. Although these policies share fundamental similarities, they are different in their functionality and application. The routing policies control how route information is imported to and exported from the routing tables. Firewall filters examine data packets at the entry (ingress) and exit (egress) points of the Services Router, filtering router traffic.



**NOTE:** For readability, the firewall filter policy is often referred to as firewall filter in this guide.

To manage the flow of information into and out of a Services Router, you must understand the fundamentals of routing and firewall filter policies. This chapter provides a brief overview of the policy fundamentals, under the following topics. For more information about routing policies and stateless firewall filters, see the *JUNOS Policy Framework Configuration Guide*. For more information about stateful firewall filters and Network Address Translation (NAT), see the *JUNOS Services Interfaces Configuration Guide*.

If the router is operating in a Common Criteria environment, see the Secure Configuration Guide for Common Criteria and JUNOS-FIPS.

- Policy Framework Terms on page 149
- Routing Policies on page 151
- Stateful Firewall Filters on page 155
- Stateless Firewall Filters on page 157
- Network Address Translation on page 163

## **Policy Framework Terms**

Before configuring routing policies or firewall filters on a Services Router, you must become familiar with the terms defined in Table 66 on page 150.

#### **Table 66: Policy Framework Terms**

Term	Definition
action	Operation performed if a route or packet matches all criteria defined in a match condition. Actions are configured in terms. You can specify one or more actions in a term. See also <i>match condition</i> ; <i>term</i> .
firewall filter	See stateful firewall filter; stateless firewall filter.
match condition	Criteria that an incoming or an outgoing route or packet on a Services Router must match for an action to occur. Match conditions are specified in terms. If you specify more than one match condition, all the conditions must match in a route or packet for an action to occur. See also <i>action; term</i> .
multifield (MF) classifier	Firewall filter that scans through a variety of packet fields to determine the forwarding class and loss priority for a packet and polices traffic to a specific bandwidth and burst size. Typically, a classifier performs matching operations on the selected fields against a configured value.
Network Address Port Translation (NAPT)	Method of concealing a set of host ports on a private network behind a pool of public addresses. NAPT can be used as a security measure to protect the host ports from direct targeting in network attacks.
Network Address Translation (NAT)	Method of concealing a set of host addresses on a private network behind a pool of public addresses. NAT can be used as a security measure to protect the host addresses from direct targeting in network attacks.
policer	Component of firewall filters that limits the amount of traffic passing into or out of an interface to thwart denial-of-service (DoS) attacks. A policer applies rate limits on bandwidth and burst size for traffic on a particular Services Router interface.
service set	Collection of services. Examples of services include stateful firewall filters and Network Address Translation (NAT).
stateful firewall filter	Type of firewall filter that evaluates the context of connections, permits or denies traffic based on the context, and updates this information dynamically. The context includes IP source and destination addresses, TCP port numbers, TCP sequencing information, and TCP connection flags.
stateless firewall filter	Type of firewall filter that statically evaluates the contents of packets transiting the router and packets originating from, or destined for, the router. Information about connection states is not maintained.
term	Component of a routing policy or firewall filter that defines its criteria (match conditions) and results (actions). A routing policy or firewall filter can have one or multiple terms. See also <i>match condition</i> ; <i>action</i> .
trusted network	Network from which all originating traffic can be trusted—for example, an internal enterprise LAN. Stateful firewall filters allow traffic to flow from trusted to untrusted networks.
untrusted network	Network from which all originating traffic cannot be trusted—for example, a WAN. Unless configured otherwise, stateful firewall filters do not allow traffic to flow from untrusted to trusted networks.

## **Routing Policies**

This section contains the following topics:

- Routing Policy Overview on page 151
- Routing Policy Match Conditions on page 152
- Routing Policy Actions on page 153

#### **Routing Policy Overview**

Routing protocols send information about routes to a router's neighbors. This information is processed and used to create routing tables, which are then distilled into forwarding tables. Routing policies control the flow of information between the routing protocols and the routing tables and between the routing tables and the forwarding tables. Using policies, you can determine which routes are advertised, specify which routes are imported into the routing table, and modify routes to control which routes are added to the forwarding table. For more information, see the *JUNOS Policy Framework Configuration Guide*.

Routing policies are made up of one or more terms, each of which contains a set of match conditions and a set of actions. Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route. These actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

#### **Routing Policy Terms**

Generally, a Services Router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of **accept** or **reject** is taken. If none of the terms in the policy match the route, the Services Router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

#### **Default and Final Actions**

If none of the terms' match conditions evaluate to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

#### **Applying Routing Policies**

Once a policy is created, it must be applied before it is active. You apply routing policies using the import and export statements at the **Protocols** > *protocol-name* level in the configuration hierarchy.

In the **import** statement, you list the name of the routing policy to be evaluated when routes are imported into the routing table from the routing protocol.

In the **export** statement, you list the name of the routing policy to be evaluated when routes are being exported from the routing table into a dynamic routing protocol. Only active routes are exported from the routing table.

To specify more than one policy and create a policy chain, you list the policies using a space as a separator. If multiple policies are specified, the policies are evaluated in the order in which they are specified. As soon as an accept or reject action is executed, the policy chain evaluation ends.

### **Routing Policy Match Conditions**

A match condition defines the criteria that a route must match for an action to take place. Each term can have one or more match conditions. If a route matches all the match conditions for a particular term, the actions defined for that term are processed.

Each term can consist of two statements, to and from, that define match conditions:

- In the from statement, you define the criteria that an *incoming* route must match. You can specify one or more match conditions. If you specify more than one, all conditions must match the route for a match to occur.
- In the to statement, you define the criteria that an *outgoing* route must match. You can specify one or more match conditions. If you specify more than one, all conditions must match the route for a match to occur.

The order of match conditions in a term is not important, because a route must match all match conditions in a term for an action to be taken.

Table 67 on page 152 summarizes key routing policy match conditions.

Match Condition	Description
aggregate-contributor	Matches routes that are contributing to a configured aggregate. This match condition can be used to suppress a contributor in an aggregate route.
area area-id	Matches a route learned from the specified OSPF area during the exporting of OSPF routes into other protocols.
as-path name	Matches the name of an autonomous systems (AS) path regular expression. BGP routes whose AS path matches the regular expression are processed.
color preference	Matches a color value. You can specify preference values that are finer-grained than those specified in the <i>preference</i> match conditions. The <b>color</b> value can be a number from 0 through 4,294,967,295 ( $2^{32} - 1$ ). A lower number indicates a more preferred route.
community	Matches the name of one or more communities. If you list more than one name, only one name needs to match for a match to occur. (The matching is effectively a logical OR operation.)
external [type metric-type]	Matches external OSPF routes, including routes exported from one level to another. In this match condition, <b>type</b> is an optional keyword. The <b>metric-type</b> value can be either <b>1</b> or <b>2</b> . When you do not specify <b>type</b> , this condition matches all external routes.

#### **Table 67: Summary of Key Routing Policy Match Conditions**

Match Condition	Description
interface interface-name	Matches the name or IP address of one or more router interfaces. Use this condition with protocols that are interface-specific. For example, do not use this condition with internal BGP (IBGP).
	Depending on where the policy is applied, this match condition matches routes learned from or advertised through the specified interface.
internal	Matches a routing policy against the internal flag for simplified next-hop self policies.
level level	Matches the IS-IS level. Routes that are from the specified level or are being advertised to the specified level are processed.
local-preference value	Matches a BGP local preference attribute. The preference value can be from 0 through 4,294,967,295 ( $2^{32}$ – 1).
metric <i>metric</i> metric2 <i>metric</i>	Matches a metric value. The <b>metric</b> value corresponds to the multiple exit discriminator (MED), and <b>metric2</b> corresponds to the interior gateway protocol (IGP) metric if the BGP next hop runs back through another route.
neighbor address	Matches the address of one or more neighbors (peers).
	For BGP export policies, the address can be for a directly connected or indirectly connected peer. For all other protocols, the address is for the neighbor from which the advertisement is received.
next-hop address	Matches the next-hop address or addresses specified in the routing information for a particular route. For BGP routes, matches are performed against each protocol next hop.
origin value	Matches the BGP origin attribute, which is the origin of the AS path information. The value can be one of the following:
	■ egp—Path information originated from another AS.
	■ igp—Path information originated from within the local AS.
	■ incomplete—Path information was learned by some other means.
preference preference	Matches the preference value. You can specify a primary preference value ( <b>preference</b> ) and a secondary preference value ( <b>preference2</b> ). The preference value can be a number
preference2 preference	from 0 through 4,294,967,295 ( $2^{32} - 1$ ). A lower number indicates a more preferred route.
protocol protocol	Matches the name of the protocol from which the route was learned or to which the route is being advertised. It can be one of the following: aggregate, bgp, direct, dvmrp, isis, local, ospf, pim-dense, pim-sparse, rip, ripng, or static.
route-type value	Matches the type of route. The value can be either external or internal.

#### Table 67: Summary of Key Routing Policy Match Conditions (continued)

## **Routing Policy Actions**

An action defines what the Services Router does with the route when the route matches all the match conditions in the **from** and **to** statements for a particular term.

If a term does not have **from** and **to** statements, all routes are considered to match and the actions apply to all routes.

Each term can have one or more of the following types of actions. The actions are configured under the **then** statement.

- Flow control actions, which affect whether to accept or reject the route and whether to evaluate the next term or routing policy
- Actions that manipulate route characteristics
- Trace action, which logs route matches

Table 68 on page 154 summarizes the routing policy actions.

If you do not specify an action, one of the following results occurs:

- The next term in the routing policy, if one exists, is evaluated.
- If the routing policy has no more terms, the next routing policy, if one exists, is evaluated.
- If there are no more terms or routing policies, the accept or reject action specified by the default policy is executed.

#### **Table 68: Summary of Key Routing Policy Actions**

Action	Description
Flow Control Actions	These actions control the flow of routing information into and out of the routing table.
accept	Accepts the route and propagates it. After a route is accepted, no other terms in the routing policy and no other routing policies are evaluated.
reject	Rejects the route and does not propagate it. After a route is rejected, no other terms in the routing policy and no other routing policies are evaluated.
next term	Skips to and evaluates the next term in the same routing policy. Any <b>accept</b> or <b>reject</b> action specified in the <b>then</b> statement is ignored. Any actions specified in the <b>then</b> statement that manipulate route characteristics are applied to the route.
next policy	Skips to and evaluates the next routing policy. Any <b>accept</b> or <b>reject</b> action specified in the <b>then</b> statement is ignored. Any actions specified in the <b>then</b> statement that manipulate route characteristics are applied to the route.
Route Manipulation Actions	These actions manipulate the route characteristics.
as-path-prepend as-path	Appends one or more autonomous system (AS) numbers at the beginning of the AS path. If you are specifying more than one AS number, include the numbers in quotation marks.
	The AS numbers are added after the local AS number has been added to the path. This action adds AS numbers to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the appended AS numbers are placed within a confederation sequence. Otherwise, the appended AS numbers are placed with a nonconfederation sequence.

#### Table 68: Summary of Key Routing Policy Actions (continued)

Action	Description
as-path-expand last-as count n	Extracts the last AS number in the existing AS path and appends that AS number to the beginning of the AS path $n$ times. Replace $n$ with a number from <b>1</b> through <b>32</b> .
	The AS numbers are added after the local AS number has been added to the path. This action adds AS numbers to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the appended AS numbers are placed within a confederation sequence. Otherwise, the appended AS numbers are placed with a nonconfederation sequence.
class class-name	Applies the specified class-of-service (CoS) parameters to routes installed into the routing table.
color preference	Sets the preference value to the specified value. The color and color2 preference values can be a number from 0 through 4,294,967,295 ( $2^{32} - 1$ ). A lower number indicates
color2 preference	a more preferred route.
damping name	Applies the specified route-damping parameters to the route. These parameters override BGP's default damping parameters.
	This action is useful only in import policies.
local-preference value	Sets the BGP local preference attribute. The preference can be a number from 0 through $4,294,967,295$ ( $2^{32} - 1$ ).
metric metric	Sets the metric. You can specify up to four metric values, starting with metric (for the first metric value) and continuing with metric2, metric3, and metric4.
metric2 metric	
metric3 metric	For BGP routes, metric corresponds to the MED, and metric2 corresponds to the IGP metric if the BGP next hop loops through another router.
metric4 metric	
next-hop address	Sets the next hop.
	If you specify <i>address</i> as <i>self</i> , the next-hop address is replaced by one of the local router's addresses. The advertising protocol determines which address to use.

## **Stateful Firewall Filters**

This section contains the following topics:

- Stateful Firewall Filter Overview on page 155
- Stateful Firewall Filter Match Conditions on page 156
- Stateful Firewall Filter Actions on page 156

## **Stateful Firewall Filter Overview**

In a *stateful* firewall filter, all packets flowing from a trusted network to an untrusted network are allowed. Packets flowing from an untrusted network to a trusted network

are allowed only if they are responses to a session originated by the trusted network, or if they are explicitly accepted by a term in the stateful firewall filter rule.

When Network Address Translation (NAT) is enabled, the source address of a packet flowing from a trusted network to an untrusted network is replaced with an address chosen from a specified range, or *pool*, of addresses. In addition, you can configure the Services Router to dynamically translate the source port of the packet—a process called Network Address Port Translation (NAPT). For more information about NAT, see "Network Address Translation" on page 163.

All stateful firewall filters contain one or more terms, and each term consists of two components—match conditions and actions. The match conditions define the values or fields that the packet must contain to be considered a match. If a packet is a match, the corresponding action is taken. By default, a packet that does not match a firewall filter is discarded.



**NOTE:** A firewall filter with a large number of terms can adversely affect both the configuration commit time and the performance of the Routing Engine.

For more information about stateful firewall filters, see the *JUNOS Services Interfaces Configuration Guide*.

#### **Stateful Firewall Filter Match Conditions**

Table 69 on page 156 lists the match conditions you can specify in stateful firewall filter and terms.

For more information about configuring applications and application sets for stateful firewall filters, see the *JUNOS Services Interfaces Configuration Guide*.

Match Condition	Description
application-sets [set-names]	Matches a list of application set names. For more information about application sets, see the <i>JUNOS Services Interfaces Configuration Guide</i> .
applications [application-names]	Matches a list of applications. For more information about applications, see the <i>JUNOS Services Interfaces Configuration Guide</i> .
destination-address address	Matches the IP destination address field.
source-address address	Matches the IP source address field.

#### **Table 69: Stateful Firewall Filter Match Conditions**

#### **Stateful Firewall Filter Actions**

Table 70 on page 157 and Table 75 on page 167 list actions you can specify in stateful firewall filter terms.
Actions	Description
accept	Accepts the packet and send it to its destination.
allow-ip-options [ values ]	Accepts the packet if the IP Option header of the packet contains a value that matches one of the specified values. If this action is not included, only packets without IP options are accepted. This action can be specified only with the <b>accept</b> action.
	You can specify the IP option as text or a numeric value: <b>any</b> (0), <b>ip-security</b> (130), <b>ip-stream</b> (8), <b>loose-source-route</b> (3), <b>route-record</b> (7), <b>route-ralert</b> (148), <b>strict-source-route</b> (9), and <b>timestamp</b> (4).
discard	Does not accept the packet, and do not process it further.
reject	Does not accept the packet, and sends a rejection message. UDP sends an ICMP unreachable code and RCP sends RST. Rejected packets can be logged or sampled.
syslog	Records information in the system logging facility. This action can be used with all options except <b>discard</b> .

## **Table 70: Stateful Firewall Filter Actions**

# **Stateless Firewall Filters**

This section contains the following topics:

- Stateless Firewall Filter Overview on page 157
- Planning a Stateless Firewall Filter on page 158
- Stateless Firewall Filter Match Conditions on page 159
- Stateless Firewall Filter Actions and Action Modifiers on page 162

# **Stateless Firewall Filter Overview**

A *stateless* firewall filter can filter packets transiting the Services Router from a source to a destination, or packets originating from, or destined for, the Routing Engine. Stateless firewall filters applied to the Routing Engine interface protect the processes and resources owned by the Routing Engine.

You can apply a stateless firewall filter to an input or output interface, or to both. Every packet, including fragmented packets, is evaluated against stateless firewall filters.

# **Stateless Firewall Filter Terms**

All stateless firewall filters contain one or more terms, and each term consists of two components—match conditions and actions. The match conditions define the values or fields that the packet must contain to be considered a match. If a packet is a match, the corresponding action is taken. By default, a packet that does not match a firewall filter is discarded.



**NOTE:** A firewall filter with a large number of terms can adversely affect both the configuration commit time and the performance of the Routing Engine.

### **Chained Stateless Firewall Filters**

On a Services Router, you can configure a stateless firewall filter within the term of another filter. This method enables you to add common terms to multiple filters without having to modify all filter definitions. You can configure one filter with the desired common terms, and configure this filter as a term in other filters. Consequently, to make a change in these common terms, you need to modify only one filter that contains the common terms, instead of multiple filters. For more information about how to configure a filter within a filter, see the *JUNOS Policy Framework Configuration Guide*.

### **Planning a Stateless Firewall Filter**

Before creating a stateless firewall filter and applying it to an interface, determine what you want the firewall filter to accomplish and how to use its match conditions and actions to achieve your goal. Also, make sure you understand how packets are matched and the default action of the resulting firewall filter.



**CAUTION:** If a packet does not match any terms in a stateless firewall filter rule, the packet is discarded. Take care that you do not configure a firewall filter that prevents you from accessing the Services Router after you commit the configuration. For example, if you configure a firewall filter that does not match HTTP or HTTPS packets, you cannot access the router with the J-Web interface.

To configure a stateless firewall filter, determine the following:

- Purpose of the firewall filter—for example, to limit traffic to certain protocols, IP source or destination addresses, or data rates, or to prevent denial-of-service (DoS) attacks.
- Appropriate match conditions. The packet header fields to match—for example, IP header fields (such as source and destination IP addresses, protocols, and IP options), TCP header fields (such as source and destination ports and flags), and ICMP header fields (such as ICMP packet type and code).
- Action to take if a match occurs—for example, accept, discard, or evaluate the next term.
- (Optional) Action modifiers. Additional actions to take if a packet matches—for example, count, log, rate limit, or police a packet.
- Interface on which the firewall filter is applied. The input or output side, or both sides, of the Routing Engine interface or a non-Routing Engine interface.

For more information about what a stateless firewall filter can include, see "Stateless Firewall Filter Match Conditions" on page 159. For more information about stateless firewall filters, see the *JUNOS Policy Framework Configuration Guide*.

# **Stateless Firewall Filter Match Conditions**

Table 71 on page 159 lists the match conditions you can specify in stateless firewall filter terms. Some of the numeric range and bit-field match conditions allow you to specify a text synonym. For a complete list of the synonyms, do any of the following:

- If you are using the J-Web interface, select the synonym from the appropriate list.
- If you are using the CLI, type a question mark (?) after the from statement.
- See the JUNOS Policy Framework Configuration Guide.

To specify a bit-field match condition with values, such as **tcp-flags**, you must enclose the values in quotation marks (""). You can use bit-field logical operators to create expressions that are evaluated for matches. For example, if the following expression is used in a filter term, a match occurs if the packet is the initial packet of a TCP session:

# tcp-flags "syn & lack"

Table 72 on page 162 lists the bit-field logical operators in order of highest to lowest precedence.

You can use text synonyms to specify some common bit-field matches. In the previous example, you can specify **tcp-initial** to specify the same match condition.



**NOTE:** When the Services Router compares the stateless firewall filter match conditions to a packet, it compares only the header fields specified in the match condition. There is no implied protocol match. For example, if you specify a match of **destination-port ssh**, the Services Router checks for a value of **0x22** in the 2-byte field that is two bytes after the IP packet header. The protocol field of the packet is not checked.

### **Table 71: Stateless Firewall Filter Match Conditions**

Match Condition	Description
Numeric Range Match Conditions	
keyword-except	Negates a match—for example, destination-port-except number.
	The following keywords accept the <b>-except</b> extension: <b>destination-port</b> , <b>dscp</b> , <b>esp-spi</b> , <b>forwarding-class</b> , fragment-offset, icmp-code, icmp-type, interface-group, ip-options, <b>packet-length</b> , <b>port</b> , <b>precedence</b> , <b>protocol</b> and <b>source-port</b> .
destination-port number	Matches a TCP or User Datagram Protocol (UDP) destination port field. You cannot specify both the <b>port</b> and <b>destination-port</b> match conditions in the same term. Normally, you specify this match in conjunction with the <b>protocol tcp</b> or <b>protocol udp</b> match statement to determine which protocol is being used on the port.
	In place of the numeric value, you can specify a text synonym. For example, you can specify <b>telnet</b> or <b>23</b> .

# Table 71: Stateless Firewall Filter Match Conditions (continued)

Match Condition	Description
esp-spi spi-value	Matches an IPSec encapsulating security payload (ESP) security parameter index (SPI) value. Match on this specific SPI value. You can specify the ESP SPI value in either hexadecimal, binary, or decimal form.
forwarding-class class	Matches a forwarding class. Specify assured-forwarding, best-effort, expedited-forwarding, or network-control.
fragment-offset number	Matches the fragment offset field.
icmp-code number	Matches the ICMP code field. Normally, you specify this match condition in conjunction with the <b>protocol icmp</b> match statement to determine which protocol is being used on the port.
	This value or keyword provides more specific information than <b>icmp-type</b> . Because the value's meaning depends on the associated <b>icmp-type</b> , you must specify <b>icmp-type</b> along with <b>icmp-code</b> .
	In place of the numeric value, you can specify a text synonym. For example, you can specify ip-header-bad or $0$ .
icmp-type number	Matches the ICMP packet type field. Normally, you specify this match condition in conjunction with the <b>protocol icmp</b> match statement to determine which protocol is being used on the port.
	In place of the numeric value, you can specify a text synonym. For example, you can specify time-exceeded or 11.
interface-group group-number	Matches the interface group on which the packet was received. An interface group is a set of one or more logical interfaces. For information about configuration interface groups, see the <i>JUNOS Policy Framework Configuration Guide</i> .
packet-length bytes	Matches the length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.
port number	Matches a TCP or UDP source or destination port field. You cannot specify both the <b>port</b> match and either the <b>destination-port</b> or <b>source-port</b> match conditions in the same term. Normally, you specify this match condition in conjunction with the <b>protocol tcp</b> or <b>protocol udp</b> match statement to determine which protocol is being used on the port.
	In place of the numeric value, you can specify a text synonym. For example, you can specify bgp or 179.
precedence ip-precedence-field	Matches the IP precedence field. You can specify precedence in either hexadecimal, binary, or decimal form.
	In place of the numeric value, you can specify a text synonym. For example, you can specify immediate or 0x40.
protocol number	Matches the IP protocol field. In place of the numeric value, you can specify a text synonym. For example, you can specify <b>ospf</b> or <b>89</b> .

Match Condition	Description
source-port number	Matches the TCP or UDP source port field. You cannot specify the <b>port</b> and <b>source-port</b> match conditions in the same term. Normally, you specify this match condition in conjunction with the <b>protocol tcp</b> or <b>protocol udp</b> match statement to determine which protocol is being used on the port.
	In place of the numeric value, you can specify a text synonym. For example, you can specify http or 80.
Address Match Conditions	
address prefix	Matches the IP source or destination address field. You cannot specify both the <b>address</b> and the <b>destination-address</b> or <b>source-address</b> match conditions in the same term.
destination-address prefix	Matches the IP destination address field. You cannot specify the <b>destination-address</b> and <b>address</b> match conditions in the same term.
destination-prefix-list prefix-list	Matches the IP destination prefix list field. You cannot specify the <b>destination-prefix-list</b> and <b>prefix-list</b> match conditions in the same term.
prefix-list prefix-list	Matches the IP source or destination prefix list field. You cannot specify both the <b>prefix-list</b> and the <b>destination-prefix-list</b> or <b>source-prefix-list</b> match conditions in the same term.
source-address prefix	Matches the IP source address field. You cannot specify the <b>source-address</b> and <b>address</b> match conditions in the same rule.
source-prefix-list prefix-list	Matches the IP source prefix list field. You cannot specify the <b>source-prefix-list</b> and <b>prefix-list</b> match conditions in the same term.
Bit-Field Match Conditions with Va	lues
fragment-flags number	Matches an IP fragmentation flag. In place of the numeric value, you can specify a text synonym. For example, you can specify more-fragments or 0x2000.
ip-options number	Matches an IP option. In place of the numeric value, you can specify a text synonym. For example, you can specify <b>record-route</b> or <b>7</b> .
tcp-flags number	Matches a TCP flag. Normally, you specify this match condition in conjunction with the <b>protocol tcp</b> match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify a text synonym. For example, you can specify <b>syn</b> or <b>0x02</b> .
Bit-Field Text Synonym Match Cond	litions
first-fragment	Matches the first fragment of a fragmented packet. This condition does not match unfragmented packets.
is-fragment	Matches the trailing fragment of a fragmented packet. It does not match the first fragment of a fragmented packet. To match both first and trailing fragments, you can use two terms, or you can use <b>fragment-offset 0-8191</b> .
tcp-established	Matches a TCP packet other than the first packet of a connection. This match condition is a synonym for "(ack   rst)".
	This condition does not implicitly check that the protocol is TCP. To do so, specify the <b>protocol tcp</b> match condition.

# Table 71: Stateless Firewall Filter Match Conditions (continued)

Table	71:	Stateless	<b>Firewall</b>	Filter	Match	Conditions	(continued)	)
-------	-----	-----------	-----------------	--------	-------	------------	-------------	---

Match Condition	Description
tcp-initial	Matches the first TCP packet of a connection. This match condition is a synonym for "(syn & !ack)".
	This condition does not implicitly check that the protocol is TCP. To do so, specify the <b>protocol tcp</b> match condition.

# Table 72: Stateless Firewall Filter Bit-Field Logical Operators

Logical Operator	Description
()	Grouping
!	Negation
& or +	Logical AND
or ,	Logical OR

# **Stateless Firewall Filter Actions and Action Modifiers**

Table 73 on page 162 lists the actions and action modifiers you can specify in stateless firewall filter terms.

Table 13. Stateless i newall i filer Activity and Activit mouthers	Table	73:	Stateless	Firewall	Filter	Actions	and	Action	Modifiers
--------------------------------------------------------------------	-------	-----	-----------	----------	--------	---------	-----	--------	-----------

Action or Action Modifier	Description
accept	Accepts a packet. This is the default if the packet matches. However, we strongly recommend that you always explicitly configure an action in the <b>then</b> statement.
discard	Discards a packet silently, without sending an Internet Control Message Protocol (ICMP) message. Packets are available for logging and sampling before being discarded.
next term	Continues to the next term for evaluation.
reject <message-type></message-type>	Discards a packet, sending an ICMP destination unreachable message. Rejected packets are available for logging and sampling. You can specify one of the following message types: administratively-prohibited (default), bad-host-tos, bad-network-tos, host-prohibited, host-unknown, host-unreachable, network-prohibited, network-unknown, network-unreachable, port-unreachable, port-unreachable, port-cutoff, precedence-violation, protocol-unreachable, source-host-isolated, source-route-failed, or tcp-reset. If you specify tcp-reset, a TCP reset is returned (indicating the end of a TCP flow), if the packet is a TCP packet. Otherwise, nothing is returned.
routing-instance routing-instance	Routes the packet using the specified routing instance.
Action Modifiers	

Action or Action Modifier	Description
count counter-name	Counts the number of packets passing this term. The name can contain letters, numbers, and hyphens (-), and can be up to 24 characters long. A counter name is specific to the filter that uses it, so all interfaces that use the same filter increment the same counter.
forwarding-class class-name	Classifies the packet to the specified forwarding class.
log	Logs the packet's header information in the Routing Engine. You can access this information by entering the <b>show firewall log</b> command at the CLI.
loss-priority priority	Sets the scheduling priority of the packet. The priority can be low or high.
policer policer-name	Applies rate limits to the traffic using the named policer.
sample	Samples the traffic on the interface. Use this modifier only when traffic sampling is enabled. For more information, see the <i>JUNOS Policy Framework Configuration Guide</i> .
syslog	Records information in the system logging facility. This action can be used in conjunction with all options except <b>discard</b> .

### Table 73: Stateless Firewall Filter Actions and Action Modifiers (continued)

## **Network Address Translation**

This section contains the following topics:

- NAT Overview on page 163
- NAT Components on page 166

# **NAT Overview**

Network Address Translation (NAT) allows multiple hosts on a private internal network to access the public external network using a small pool of NAT addresses. Only addresses from this pool are visible to the external network. Between the internal and external network, a router is configured to rewrite the source or destination addresses of IP packets passing through it.

Services Routers support four types of NAT processing: source static NAT, source dynamic NAT *with* Network Address Port Translation (NAPT), source dynamic *without* NAPT, and destination static NAT.

## **Source Static NAT**

Source static NAT translates an internal source address to a NAT address from the referenced pool on a one-to-one basis. Source static NAT is easy to implement and is useful in a situation when the available pool of addresses is equal to or greater than the number of source addresses to be translated.

In the sample source static NAT scenario shown in Figure 13 on page 164, the defined prefix **192.168.1.0/24** is mapped one-to-one to the defined source address pool

121.0.1.0/24. Hence the source address 192.168.1.1 always translates to 121.0.1.1, the source address 192.168.1.2 always translates to 121.0.1.2, and so on.

### **Figure 13: Sample Source Static NAT**



### Source Dynamic NAT with NAPT

Typically, source dynamic NAT implements address translation for source traffic with Network Address Port Translation (NAPT). For each outgoing packet, the source address is replaced by a NAT address from a defined address pool and a port is assigned to it either automatically by the NAT router or from a port pool that you define. A NAT address that is assigned to a host is used for all concurrent sessions from that host. The address is released to the pool only after all the sessions for that host expire. Because all the private hosts might not simultaneously create sessions, they can share a few NAT addresses.

In the sample source dynamic NAT scenario shown in Figure 14 on page 164, the source address **192.168.1.1** is translated to address **121.0.1.1** from the defined NAT pool, and is assigned port **20001** from the defined port pool. The NAT address **121.0.1.1** is reused for source address **192.168.1.2** with a different port, **20002**.

A dynamic NAT pool with NAPT supports address ranges with a maximum of 32 addresses.

### Figure 14: Sample Source Dynamic NAT with NAPT



### Source Dynamic NAT Without NAPT

Alternatively, a Services Router supports source dynamic NAT without NAPT. This technique, also known as oversubscribed NAT, allows NAT addresses from the referenced pool to be assigned dynamically. Assigning addresses dynamically also allows a few public IP addresses to be used by several private hosts in contrast with an equal sized pool required by source static NAT.

A dynamic NAT pool with no address port translation supports address ranges with a maximum of 65,535 addresses.

# **Destination Static NAT**

Destination static NAT translates the destination address for external traffic to an address specified in a destination pool. The destination pool contains one address and no port configuration.

In the destination static NAT scenario shown in Figure 15 on page 165, when the NAT router receives a packet with destination address **121.0.1.1**, it replaces this destination address with the associated local host address **192.168.1.1**. Only the address defined in the destination address pool (**121.0.1.1**) is visible to the external router and not the local host address (**192.168.1.1**).

#### Figure 15: Sample Destination Static NAT



### Full-Cone NAT (Bidirectional NAT)

With *full-cone* NAT, all requests from the same internal IP address and port are mapped to the same external IP address and port. In addition, any external host can send a packet to the internal host by sending it to the mapped external address. Full-cone NAT is useful if you want to allow external hosts from the public network to connect to internal hosts using public IP addresses. However, we recommend that you use this feature along with strict firewall rules that allow only the intended traffic from the public network to reach the customer-edge router.

When the internal host terminates its connection to the external host, any new connection initiation from any external host to the internal host on the public IP network is not permitted. All existing connections from external to internal hosts are not affected. Full-Cone NAT allows connections between externel and internal

hosts to take place independent of the source or destination port and is application independent. A full-cone NAT is enabled or disabled by configuration .

The router handles the connection between the external host and the internal host like any other connection. This feature is available for both source static and source dynamic NAT.



**NOTE:** There is no support for IPv6 or PAT/NAPT.

For more information, see "Configuring Full-Cone NAT" on page 190.

# **NAT Components**

NAT can be configured independently or with stateful firewall filters. For information about configuring NAT independently, see "Configuring NAT" on page 185. For information about configuring NAT with stateful firewall filters, see "Configuring Stateful Firewall Filters and NAT" on page 205.

To configure NAT, you must define a NAT pool, define a NAT rule or rule set, and apply this NAT rule or rule set to an interface.

## **NAT Pools**

You define a pool of source or destination addresses that are used as translated addresses for NAT. In a pool you can specify one or more addresses, prefixes, or address ranges.

When defining a NAT pool, make sure that it meets the following requirements:

- No more than 10 address ranges, prefixes, or a combination of address ranges and prefixes are in the pool.
- The ranges of addresses and prefixes defined in the pool do not overlap.
- In an address range, the low value is a lower number than the high value.

If you have configured multiple address ranges and prefixes, the prefixes are depleted first, followed by the address ranges.

(¥

**NOTE:** Multiple addresses, prefixes, and address ranges are not supported for destination static NAT. Only one address is allowed in the destination address pool.

# NAT Rules

You can define a set of rules or a single rule. To define a rule you must define the following components:

- Term—Named structure in which match conditions and actions are defined.
- Match condition—Criteria against which a route or packets are compared. You can configure one or more criteria. If all criteria match, one or more actions are applied. Table 74 on page 167 summarizes a list of key NAT match conditions.
- Action—What happens when all the specified conditions match. You can configure one or more actions. Table 75 on page 167 summarizes a list of key NAT actions.
- Match direction—Direction in which the match is applied—input or output. For more information about match direction, see the JUNOS Services Interfaces Configuration Guide.

### **Table 74: NAT Match Conditions**

Match Condition	Description
application-sets [set-names]	Matches a list of application set names. For more information about application sets, see the <i>JUNOS Services Interfaces Configuration Guide</i> .
applications [application-names]	Matches a list of applications. For more information about applications, see the <i>JUNOS Services Interfaces Configuration Guide</i> .
destination-address ( <i>address</i>   any-unicast) <i>except</i>	Matches the IP destination address field.
destination-address-range low minimum-value high maximum-value except	Matches the IP destination address range field
destination-prefix-list list-name except	Matches the prefix list of the IP destination.
source-address ( <i>address</i>   any-unicast) except	Matches the IP source address field.
source-address-range low minimum-value high maximum-value except	Matches the IP source address range field
source-prefix-list list-name except	Matches the prefix list of the IP source.

### **Table 75: NAT Actions**

Actions	Description
no-translation	Enables you to specify addresses that you want to exclude from NAT.
syslog	Records information in the system logging facility.
translated source-pool nat-pool-name	Translates the source address using the specified pool.
translated source-prefix source-prefix	Translates the source address using the specified source prefix.

## Table 75: NAT Actions (continued)

Actions	Description				
translated translation-type	Translates the destination and source port using the specified type:				
(destination type   source type)	destination static—Translates the destination address without port mapping. This type requires the size of the source address space to be the same as the size of the destination address space. You must specify a destination-pool name. The referenced pool must contain exactly one address and no port configuration.				
	■ <b>source dynamic</b> —Translates the source address with port mapping by means of NAPT. You must specify a <b>source-pool</b> name. The referenced pool must include a <b>port</b> configuration.				
	source static—Translates the source address without port mapping. This type requires the size of the source address space to be the same as the size of the destination address space. You must specify a source-pool name. The referenced pool must contain exactly one address and no port configuration.				

# Chapter 10 Configuring Routing Policies

Use routing policies as filters to control the information from routing protocols that a Services Router imports into its routing table and the information that the router exports (advertises) to its neighbors. To create a routing policy, you configure criteria against which routes are compared, and the action that is performed if the criteria are met.

You use either the J-Web configuration editor or CLI configuration editor to configure a routing policy.

This chapter contains the following topics. For more information about routing policies, see the *JUNOS Policy Framework Configuration Guide*.

- Before You Begin on page 169
- Configuring a Routing Policy with a Configuration Editor on page 170

# **Before You Begin**

Before you begin configuring a routing policy, complete the following tasks:

- If you do not already have a basic understanding of routing policies, read "Routing Policies" on page 151.
- Determine what you want to accomplish with the policy, and thoroughly understand how to achieve your goal using the various match conditions and actions.
- Make certain that you understand the default policies and actions for the policy you are configuring.
- Configure an interface on the router. See the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.
- Configure an Interior Gateway Protocol (IGP) and Border Gateway Protocol (BGP), if necessary. See the *J-series Services Router Basic LAN and WAN Access Configuration Guide*.
- Configure the router interface to reject or accept routes, if necessary. See "Configuring Stateless Firewall Filters" on page 221.
- Configure static routes, if necessary. See the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

# **Configuring a Routing Policy with a Configuration Editor**

A routing policy has a major impact on the flow of routing information or packets within and through the Services Router. The match conditions and actions allow you to configure a customized policy to fit your needs.

To configure a routing policy, you must perform the following tasks marked (*Required*). Perform additional tasks as needed for your router. For information about using the J-Web and CLI configuration editors, see the *J-series Services Router Basic LAN and WAN Access Configuration Guide*.

- Configuring the Policy Name (Required) on page 170
- Configuring a Policy Term (Required) on page 171
- Rejecting Known Invalid Routes (Optional) on page 172
- Injecting OSPF Routes into the BGP Routing Table (Optional) on page 174
- Grouping Source and Destination Prefixes in a Forwarding Class (Optional) on page 176
- Configuring a Policy to Prepend the AS Path (Optional) on page 177
- Configuring Damping Parameters (Optional) on page 179

### **Configuring the Policy Name (Required)**

Each routing policy is identified by a policy name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each routing policy name must be unique within a configuration.

To configure the policy name:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 76 on page 170.
- 3. Go on to "Configuring a Policy Term (Required)" on page 171.

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Policy statement</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter
	2.	Next to Policy options, click <b>Configure</b> or <b>Edit</b> .	edit policy-options
	3.	Next to Policy statement, click Add new entry.	

#### **Table 76: Configuring the Policy Name**

### Table 76: Configuring the Policy Name (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Enter the policy name—for example, <b>policy1</b> .	1. In the Policy name box, type <b>policy1</b> .	Type the <b>policy-name</b> value:
	2. Click <b>OK</b> .	set policy-statement policy1

# **Configuring a Policy Term (Required)**

Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

To configure a policy term:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 77 on page 171.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To configure additional routing policy features, go on to one of the following procedures:
  - To remove useless routes, see "Rejecting Known Invalid Routes (Optional)" on page 172.
  - To advertise additional routes, see "Injecting OSPF Routes into the BGP Routing Table (Optional)" on page 174.
  - To create a forwarding class, see "Grouping Source and Destination Prefixes in a Forwarding Class (Optional)" on page 176.
  - To make a route less preferable to BGP, see "Configuring a Policy to Prepend the AS Path (Optional)" on page 177.
  - To suppress route information, see "Configuring Damping Parameters (Optional)" on page 179.

Table 7	7:	Config	guring	a P	olicy	Term
---------	----	--------	--------	-----	-------	------

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Policy statement</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the <b>[edit]</b> hierarchy level, enter
	2.	Next to Policy options, click <b>Configure</b> or <b>Edit</b> .	edit policy-options policy-statement policy1
	3.	Under Policy name, click <b>policy1</b> .	

### Table 77: Configuring a Policy Term (continued)

Task	J-W	eb Configuration Editor	CLI Configuration Editor	
Create and name a policy	1.	In the Term box, click Add new entry.	Create and name a policy term:	
term—for example, <b>term1</b> .	2.	In the Term name box, type term1.	set term term1	
	3.	Click <b>OK</b> .		

# **Rejecting Known Invalid Routes (Optional)**

You can specify known invalid ("bad") routes to ignore by specifying matches on destination prefixes. When specifying a destination prefix, you can specify an exact match with a specific route, or a less precise match by using match types. You can configure either a common reject action that applies to the entire list, or an action associated with each prefix. Table 78 on page 172 lists route list match types.

### **Table 78: Route List Match Types**

Match Type	Match Conditions
exact	The route shares the same most-significant bits (described by <i>prefix-length</i> ), and <i>prefix-length</i> is equal to the route's prefix length.
longer	The route shares the same most-significant bits (described by <i>prefix-length</i> ), and <i>prefix-length</i> is greater than the route's prefix length.
orlonger	The route shares the same most-significant bits (described by <i>prefix-length</i> ), and <i>prefix-length</i> is equal to or greater than the route's prefix length.
prefix-length-range prefix-length2-prefix-length3	The route shares the same most-significant bits (described by <i>prefix-length</i> ), and the route's prefix length falls between <i>prefix-length2</i> and <i>prefix-length3</i> , inclusive.
through destination-prefix	All the following are true:
	<ul> <li>The route shares the same most-significant bits (described by prefix-length) of the first destination prefix.</li> </ul>
	The route shares the same most-significant bits (described by prefix-length) of the second destination prefix for the number of bits in the prefix length.
	<ul> <li>The number of bits in the route's prefix length is less than or equal to the number of bits in the second prefix.</li> </ul>
	You do not use the <b>through</b> match type in most routing policy configurations. For more information, see the <i>JUNOS Policy Framework Configuration Guide</i> .
upto prefix-length2	The route shares the same most-significant bits (described by <i>prefix-length</i> ) and the route's prefix length falls between <i>prefix-length</i> and <i>prefix-length2</i> .

For example, you can create a policy named **rejectpolicy1** to reject routes with a mask of /8 and greater (/8, /9, /10, and so on) that have the first 8 bits set to 0, and to accept routes less than 8 bits in length.

To create rejectpolicy1:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 79 on page 173.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To configure additional routing policy features, go on to one of the following procedures:
  - To advertise additional routes, see "Injecting OSPF Routes into the BGP Routing Table (Optional)" on page 174.
  - To create a forwarding class, see "Grouping Source and Destination Prefixes in a Forwarding Class (Optional)" on page 176.
  - To make a route less preferable to BGP, see "Configuring a Policy to Prepend the AS Path (Optional)" on page 177.
  - To suppress route information, see "Configuring Damping Parameters (Optional)" on page 179.

### Table 79: Creating a Policy to Reject Known Invalid Routes

Task	J-M	Web Configuration Editor CLI Configuration Editor		
Navigate to the <b>Policy statement</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter	
	2.	Next to Policy options, click Configure or Edit.	edit policy-options policy-statement	
	3.	Next to Policy statement, click Add new entry.		
Create a rejection policy and	1.	In the Policy name box, type rejectpolicy1.	Enter	
and rejectterm1.	2.	Next to Term, click Add new entry.	set rejectpolicy1 term rejectterm1	
		In the Term name box, type rejectterm1.		
Specify the routes to accept—for	1.	Next to From, click <b>Configure</b> .	Accept routes less than 8 bits in	
example, routes with a mask of 0/0 up to /7.		Next to Route filter, click Add new entry.	length:	
		In the Address box, type $0/0$ .	set from route-filter 0/0 up to /7	
		From the Modifier list, select Upto.	accept	
	5.	In the Upto box, type /7.		
	6.	From the Accept reject list, select <b>accept</b> .		
	7.	Click OK.		

|--|

J-W	eb Configuration Editor	CL	Configuration Editor
1.	Next to Route filter, click <b>Add new entry</b> .	1.	Specify routes less than 8 bits in length:
2. 3.	From the Modifier list, select <b>Orlonger</b> .		set from route-filter /8
4. 5	From the Accept reject list, select <b>reject</b> .	2.	Reject these routes:
5.			set then reject
	J-W 1. 2. 3. 4. 5.	J-Web Configuration Editor1.Next to Route filter, click Add new entry.2.In the Address box, type /8.3.From the Modifier list, select Orlonger.4.From the Accept reject list, select reject.5.Click OK.	J-Web Configuration EditorCL1.Next to Route filter, click Add new entry.1.2.In the Address box, type /8.1.3.From the Modifier list, select Orlonger.4.4.From the Accept reject list, select reject.2.5.Click OK.2.

# Injecting OSPF Routes into the BGP Routing Table (Optional)

You can specify a match condition for policies based on protocols by naming a protocol from which the route is learned or to which the route is being advertised. You can specify one of the following protocols: aggregate, BGP, direct, DVMRP, IS-IS, local, OSPF, PIM-dense, PIM-sparse, RIP, or static

For example, you can inject or redistribute OSPF routes into the BGP routing table by creating a routing policy.

To create a routing policy named injectpolicy1 that redistributes OSPF routes from Area 1 only into BGP and does not advertise routes learned by BGP:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 80 on page 175.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To configure additional routing policy features, go on to one of the following procedures:
  - To create a forwarding class, see "Grouping Source and Destination Prefixes in a Forwarding Class (Optional)" on page 176.
  - To make a route less preferable to BGP, see "Configuring a Policy to Prepend the AS Path (Optional)" on page 177.
  - To suppress route information, see "Configuring Damping Parameters (Optional)" on page 179.

Table ov. cleating a runcy to inject vorr noutes into bur	Table	80:	Creating	a Polic	y to	Inject	<b>OSPF</b>	Routes	into	BGP
-----------------------------------------------------------	-------	-----	----------	---------	------	--------	-------------	--------	------	-----

Task	J-N	/eb Configuration Editor	CLI Configuration Editor		
Navigate to the <b>Policy statement</b> level in the configuration hierarchy.		In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit policy-options policy-statement		
	2.	Next to Policy options, click <b>Configure</b> or <b>Edit</b> .			
	3.	Next to Policy statement, click <b>Add new entry</b> .			
Create an injection policy and term—for example, injectpolicy1 and injectterm1.	1.	In the Policy name box, type injectpolicy1.	Enter		
	2.	Next to Term, click Add new entry.	set injectpolicy1 term injectterm1		
	3.	In the Term name box, type injectterm1.			
Specify the OSPF routes.	1.	In the From option, click <b>Configure</b> .	Specify the OSPF match condition:		
	2.	In the Protocol box, click <b>Add new</b> entry.	set from ospf		
	3.	In the Value drop box, select <b>ospf</b> .			
	4.	Click <b>OK</b> .			
Specify the routes from a particular	1.	In the Area box, type 1.	Specify Area 1 as a match condition:		
OSPF area—for example, Area 1.	2.	Click <b>OK</b> .	set from area 1		
Specify that the route is to be accepted	1.	Next to Then, click Configure.	Specify the action to accept:		
Set the default option to reject other	2.	From the Accept reject list, Select <b>accept</b> .	set then accept		
OSPF routes.	3.	From the Default action list, Select <b>reject</b> .			
	4.	Click <b>OK</b> until you return to the main Configuration page.			
Navigate to the <b>Bgp</b> level in the	1.	On the main Configuration page	From the [edit] hierarchy level, enter		
computation metalony.		or Edit.	edit protocols bgp		
	2.	Next to Bgp, click <b>Configure</b> or <b>Edit</b> .			
Apply the routing policy injectpolicy1 to BGP.	1.	Next to Export, click <b>Add new</b> entry.	Specify the OSPF match condition:		
	2.	In the Value option, type injectpolicy1.	set export injectpolicy1		
	3.	Click <b>OK</b> .			

# Grouping Source and Destination Prefixes in a Forwarding Class (Optional)

Create a forwarding class called **forwarding-class1** that includes packets based on both the destination address and the source address in the packet.

To configure and apply the routing policy **policy1**, which you configured in Table 76 on page 170 and Table 77 on page 171, to group source and destination prefixes in a forwarding class:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 81 on page 176.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To configure additional routing policy features, go on to one of the following procedures:
  - To make a route less preferable to BGP, see "Configuring a Policy to Prepend the AS Path (Optional)" on page 177.
  - To suppress route information, see "Configuring Damping Parameters (Optional)" on page 179.

Task	J-M	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>term1</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit policy-options policy-statement
	2.	Next to Policy options, click <b>Configure</b> or <b>Edit</b> .	policy1 term term1
	3.	Under Policy name, click policy1.	
	4.	Under Term name, click <b>term1</b> .	
Specify the routes to include in the	1.	Next to From, click <b>Configure</b> .	Specify the source routes for the
route filter. For example:	2.	Next to Route filter, click Add new entry.	route filter:
■ Source routes greater than or equal to 10.210.0.0/16	3.	In the Address box, type <b>10.210.0.0/16</b> .	set from route-filter 10.210.0.0/16
<ul> <li>Destination routes greater than</li> </ul>	4.	From the Modifier list, select Orlonger.	orlonger
or equal to 10.215.0.0/16	5.	Click <b>OK</b> to return to the From page.	
	1.	Next to Route filter, click Add new entry.	Specify the destination routes for the
	2.	In the Address box, type <b>10.215.0.0/16</b> .	route filter:
	3.	From the Modifier list, select Orlonger.	set from route-filter 10.215.0.0/16
	4.	Click <b>OK</b> until you return to the Term page.	orionger

#### Table 81: Creating a Policy to Group Source and Destination Prefixes in a Forwarding Class

Task	J-M	/eb Configuration Editor	CLI Configuration Editor
Group the source and destination		Next to Then, click Configure.	Specify the forwarding class name:
prefixes into a forwarding class—for example, forwarding-class1.	2.	In the Forwarding class box, type forwarding-class1.	set then forwarding class forwarding-class1
	3.	Click <b>OK</b> .	
Navigate to the <b>Forwarding table</b> level in the configuration hierarchy.	1.	On the main Configuration page next to Routing options, click <b>Configure</b> or <b>Edit</b> .	From the [edit] hierarchy level, enter
	2.	Next to Forwarding table, click <b>Configure</b> or <b>Edit</b> .	edit routing-options forwarding-table
Apply the <b>policy1</b> policy to the	1.	Next to Export, click Add new entry.	Specify the routing policy to apply:
forwarding table.		In the Value box, type <b>policy1</b> .	set export policy1
The routing policy is evaluated when routes are being exported from the routing table into the forwarding table. Only active routes are exported from the routing table.	3.	Click <b>OK</b> .	You can refer to the same routing policy one or more times in the same or a different <b>export</b> statement.

### Table 81: Creating a Policy to Group Source and Destination Prefixes in a Forwarding Class (continued)

# **Configuring a Policy to Prepend the AS Path (Optional)**

You can *prepend* or add one or more autonomous system (AS) numbers at the beginning of an AS path. The AS numbers are added after the local AS number has been added to the path. Prepending an AS path makes a shorter AS path look longer and therefore less preferable to the Border Gateway Protocol (BGP).

For example, from AS 1, there are two equal paths (through AS 2 and AS 3) to reach AS 4. You might want packets from certain sources to use the path through AS 2. Therefore, you must make the path through AS 3 look less preferable so that BGP chooses the path through AS 2. In AS 1, you can prepend multiple AS numbers.

To create a routing policy **prependpolicy1** that prepends multiple AS numbers:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 82 on page 178.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To suppress route information, see "Configuring Damping Parameters (Optional)" on page 179.

# **Table 82: Creating a Policy to Prepend AS Numbers**

Task	J-M	eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Policy statement</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration	From the [edit] hierarchy level, enter
	2.	Next to Policy options, click Configure or Edit.	eur policy-options policy-statement
	3.	Next to Policy statement, click <b>Add new entry</b> .	
Create a prepend policy and term—for example, <b>prependpolicy1</b> and	1.	In the Policy name box, type prependpolicy1.	Enter
prependterm1.	2.	Next to Term, click Add new entry.	set prependpolicy1 term prependterm1
	3.	In the Term name box, type prependterm1.	
Specify the routes to prepend AS	1.	Next to From, click Configure.	Specify the first routes to prepend:
<ul> <li>numbers to. For example:</li> <li>Routes greater than or equal to</li> </ul>	2.	Next to Route filter, click <b>Add new</b> entry.	set from route-filter 172.16.0.0/12 orlonger
<ul> <li>172.16.0.0/12</li> <li>Routes greater than or equal to</li> <li>100.100.0.0/40</li> </ul>	3.	In the Value box, type <b>172.16.0.0/12</b> .	
<ul> <li>Routes greater than or equal to 10.0.0.0/8</li> </ul>	4.	From the Modifier list, select <b>Orlonger</b> .	
	5.	Click <b>OK</b> .	
	1.	Next to From, click Configure.	Specify the next routes to prepend:
	2.	Next to Route filter, click <b>Add new</b> entry.	set from route-filter 192.168.0.0/16 orlonger
	3.	In the Value box, type 192.168.0.0/16.	
	4.	From the Modifier list, select Orlonger.	
	5.	Click <b>OK</b> .	
	1.	Next to From, click Configure.	Specify the last routes to prepend:
	2.	Next to Route filter, click <b>Add new</b> entry.	set from route-filter 10.0.0.0/8 orlonger
	3.	In the Value box, type 10.0.0/8.	
	4.	From the Modifier list, select Orlonger.	
	5.	Click <b>OK</b> until you return to the Term page.	

I-Web Configuration Editor	CLI Configuration Editor
1. Next to Then, click <b>Configure</b> .	Specify the AS numbers to prepend, and
<ol> <li>In the AS path prepend box, typ 1111.</li> </ol>	enclose them inside double quotation marks:
3. Click <b>OK</b> .	set then as-path-prepend "1 1 1 1"
<ol> <li>On the main Configuration page next to Protocols, click Configuration</li> </ol>	From the [edit] hierarchy level, enter re
or <b>Edit</b> .	edit protocols bgp
<ol> <li>Next to Bgp, click Configure or Edit.</li> </ol>	
1. Next to Import, click Add new entry.	Apply the policy:
<ol> <li>In the Value box, type prependpolicy1.</li> </ol>	set import prependpolicy1
3. Click <b>OK</b> .	You can refer to the same routing policy one or more times in the same or a different import statement.
	<ol> <li>J-Web Configuration Editor</li> <li>Next to Then, click Configure.</li> <li>In the AS path prepend box, typ 11111.</li> <li>Click OK.</li> <li>On the main Configuration page next to Protocols, click Configure or Edit.</li> <li>Next to Bgp, click Configure or Edit.</li> <li>Next to Import, click Add new entry.</li> <li>In the Value box, type prependpolicy1.</li> <li>Click OK.</li> </ol>

### Table 82: Creating a Policy to Prepend AS Numbers (continued)

# **Configuring Damping Parameters (Optional)**

Flap damping reduces the number of update messages by marking routes as ineligible for selection as the active or preferable route. Marking routes in this way leads to some delay, or *suppression*, in the propagation of route information, but the result is increased network stability. You typically apply flap damping to external BGP (EBGP) routes (routes in different ASs). You can also apply flap damping within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.)

You can specify one or more of the damping parameters described in Table 83 on page 179. If you do not specify a damping parameter, the default value of the parameter is used.

Damping Parameter	Description	Default Value	Possible Values
half-life minutes	Decay half-life—Number of minutes after which an arbitrary value is halved if a route stays stable.	15 (minutes)	1 through 4
max-suppress minutes	Maximum hold-down time for a route, in minutes.	60 (minutes)	1 through 720
reuse	Reuse threshold—Arbitrary value below which a suppressed route can be used again.	750	1 through 20000

### **Table 83: Damping Parameters**

### Table 83: Damping Parameters (continued)

Damping Parameter	Description	Default Value	Possible Values
suppress	Cutoff (suppression) threshold—Arbitrary value above which a route can no longer be used or included in advertisements.	3000	1 through 20000

To change the default BGP flap damping values, you define actions by creating a named set of damping parameters and including it in a routing policy with the damping action. For the damping routing policy to work, you also must enable BGP route flap damping.

To configure damping with a policy named **dampenpolicy1**, perform these steps:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 84 on page 180.
- 3. If you are finished configuring the router, commit the configuration.

### **Table 84: Creating a Policy to Accept and Apply Damping on Routes**

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Policy statement</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter
	2.	Next to Policy options, click <b>Configure</b> or <b>Edit</b> .	
	3.	Next to Policy statement, click Add new entry.	
Create a damping policy and term—for example, dampenpolicy1 and	1.	In the Policy name box, type dampenpolicy1.	Enter
dampenterm1.	2.	Next to Term, click Add new entry.	set dampenpolicy1 term dampenterm1
	3.	In the Term name box, type dampenterm1.	

Task		/eb Configuration Editor	CLI Configuration Editor
Specify the routes to dampen and	1.	Next to From, click <b>Configure</b> .	Specify the first routes to dampen:
associate each group of routes with a group name. For example:	2.	Next to Route filter, click <b>Add new</b> entry.	set from route-filter 172.16.0.0/12 orlonger
■ group1—Routes greater than or equal to 172.16.0.0/12	3.	In the Address box, type 172.16.0.0/12.	ramping group T
<ul> <li>group2—Routes greater than or equal to 192.168.0.0/16</li> </ul>	4.	In the Damping box, type group1.	
■ group3—Routes greater than or equal to 10.0.0.0/8	5.	From the Modifier list, select <b>Orlonger</b> .	
	6.	Click <b>OK</b> .	
	1.	Next to Route filter, click <b>Add new</b> entry.	Specify the next routes to dampen:
	2.	In the Address box, type <b>192.168.0.0/16</b> .	set from route-filter 192.168.0.0/16 orlonger
	3.	In the Damping box, type group2.	
	4.	From the Modifier list, select <b>Orlonger</b> .	
	5.	Click <b>OK</b> .	
	1.	Next to Route filter, click <b>Add new</b> entry.	Specify the last routes to dampen:
	2.	In the Address box, type 10.0.0/8.	set from route-filter 10.0.0.0/8 orlonger
	3.	In the Damping box, type group3.	
	4.	From the Modifier list, select Orlonger.	
	5.	Click <b>OK</b> until you return to the Policy options page.	

# Table 84: Creating a Policy to Accept and Apply Damping on Routes (continued)

# Table 84: Creating a Policy to Accept and Apply Damping on Routes (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Create three damping parameter	For each damping group:	Create and configure the damping
groups with different damping actions. For example:	<ol> <li>Next to Damping, click Add new entry.</li> </ol>	parameter groups:
<ul> <li>group1—Increases the half-life to 30 minutes. All other parameters are left at their default values.</li> </ul>	<ol> <li>In the Damping object name box, type the name of a damping group—for example group1</li> </ol>	edit damping group1 half-life 30 max-suppress 60 reuse 750 suppress 3000 edit damping group2 half-life 40
<ul> <li>group2—Increases the half-life to 40 minutes, decreases the maximum hold-down time for a</li> </ul>	<ol> <li>In the Half life box, type the half-life duration, in minutes:</li> </ol>	edit damping group2 disable
route to <b>45</b> minutes, increases the reuse value to <b>1000</b> , and	■ For group1—30	our authing Broado alogolo
reduces the cutoff (suppression) threshold to <b>400</b> .	■ For group2—40	
■ <b>group3</b> —Disables route damping.	<ol> <li>In the Max suppress box, type the maximum hold-down time, in minutes:</li> </ol>	
	<ul><li>■ For group1—60 (the default)</li><li>■ For group2—45</li></ul>	
	5. In the Reuse box, type the reuse threshold, for this damping group:	
	<ul><li>For group1—750 (the default)</li><li>For group2—1000</li></ul>	
	6. In the Suppress box, type the cutoff threshold, for this damping group:	
	<ul><li>For group1—3000 (the default)</li><li>For group2—400</li></ul>	
	<ol> <li>To disable damping for the group3 damping group, select the Disable check box.</li> </ol>	
	8. Click <b>OK</b> when you finish configuring each group.	
Navigate to the <b>Bgp</b> level in the configuration hierarchy.	1. On the main Configuration page next to Protocols, click <b>Configure</b> or <b>Edit</b>	From the [edit] hierarchy level, enter
	<ol> <li>Next to Bgp, click <b>Configure</b> or <b>Edit</b>.</li> </ol>	
Enable damping.	1. Select the <b>Damping</b> check box.	Enable damping:
	2. Click <b>OK</b> .	set damping
Navigate to the <b>Neighbor</b> level in the configuration hierarchy, for the BGP	1. On the main Configuration page next to Protocols, click <b>Edit</b> .	From the [edit] hierarchy level, enter
neighbor to which you want to apply the damping policy—for example, the	2. Next to Bgp, click <b>Edit</b> .	edit protocols bgp group groupA neighbor 172.16.15.14
neighbor at IP address 172.16.15.14.	3. Under Group name, click groupA.	
	<ol> <li>Under Neighbor Address, click 172.16.15.14.</li> </ol>	

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Apply the policy as an import policy for the BGP neighbor.	1.	Next to Import, click <b>Add new</b> entry	Apply the policy:
The routing policy is evaluated when routes are imported to the routing table.	2. 3.	In the Value box, type the name of the policy. Click <b>OK</b> .	Set import dampenpolicy1 You can refer to the same routing policy one or more times in the same or a different import statement.

# Table 84: Creating a Policy to Accept and Apply Damping on Routes (continued)

J-series[™] Services Router Advanced WAN Access Configuration Guide

# Chapter 11 Configuring NAT

Network Address Translation (NAT) enables multiple hosts on a local network to access the external (public) network by using a single IP address from their private internal network. The main benefits of NAT include efficient use of IP addresses, ease of administration, and security. On a J-series Services Router, NAT can be configured in different ways. For information about the types of NAT supported on Services Routers, see "Network Address Translation" on page 163.

You can use either the J-Web configuration editor or CLI configuration editor to configure NAT. NAT can be configured independently or with stateful firewall filters. For information about configuring NAT with stateful firewall filters, see "Configuring Stateful Firewall Filters and NAT" on page 205.

This chapter contains the following topics. For more information about NAT see the *JUNOS Services Interfaces Configuration Guide*.

- Before You Begin on page 185
- Configuring NAT with a Configuration Editor on page 185
- Verifying NAT Configuration on page 200

# **Before You Begin**

Before you begin configuring NAT, complete the following tasks:

- If you do not already have an understanding of NAT, read "Network Address Translation" on page 163.
- Before you begin configuring NAT, you must configure the interfaces on which to apply these services. To configure an interface, see the *J-series Services Router Basic LAN and WAN Access Configuration Guide.*

# **Configuring NAT with a Configuration Editor**

This section contains the following topics:

- Configuring Basic Source Static NAT on page 186
- Statically Assigning NAT Addresses from a Dynamic Pool on page 187
- Configuring Full-Cone NAT on page 190

- Configuring NAT Rules Without Defining Pools on page 192
- Defining an Overload Pool or an Overload Prefix on page 193
- Defining Rules for Transparent NAT on page 196
- Applying NAT to an Interface on page 198

# **Configuring Basic Source Static NAT**

To configure NAT you must define a NAT pool that specifies the address to be used for network address translation. Next, you must define a NAT rule and then apply this rule to an interface. Each NAT rule consists of a set of terms that contain match conditions and actions. For a description of NAT match conditions and actions, see "Network Address Translation" on page 163.

The example in this section shows a basic NAT configuration. It shows how to create the pool **nat-pool** and define the rule **nat-rule** for source static NAT.

To configure basic NAT:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 85 on page 186.
- 3. Apply the NAT configuration to an interface. See "Applying NAT to an Interface" on page 198.

### **Table 85: Configuring Basic Source Static NAT**

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Nat</b> level in the configuration hierarchy	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration	From the [edit] hierarchy level, enter
	2.	Next to Services, click <b>Configure</b> or <b>Edit</b> .	
	3.	Next to Nat, click <b>Configure</b> or <b>Edit</b> .	
Define nat-pool and assign	1.	Next to Pool, click Add new entry.	Set the NAT pool name and the address:
it an address to be used for network address	2.	In the Pool Name box, type nat-pool.	set pool nat-pool address 121.0.1.0/24
translation.	3.	Next to Address, click Add new entry.	
	4.	In the Prefix box, type 121.0.1.0/24.	
	5.	Click <b>OK</b> twice.	
Define <b>nat-rule</b> and set its match direction.	1.	On the Nat page, next to Rule, click <b>Add new entry</b> .	Set the rule name and its match direction:
	2.	In the Rule name box, type nat-rule.	set rule nat-rule match-direction output
	3.	From the Match direction list, select <b>output</b> .	

Task	J-W	eb Configuration Editor	CLI Configuration Editor	
Define <b>nat-term</b> for <b>nat-rule</b> and specify its match	1.	On the Rule page, next to Term, select <b>Add new entry</b> .	Set the term name and its match condition:	
condition—source address <b>10.0.1.0/24</b> .	2.	In the Term name box, type nat-term.	set rule nat-rule term nat-term from source-address 10.0.1.0/24	
	3.	Next to From, click Configure.		
	4.	Next to Source Address, click <b>Add new</b> entry.		
	5.	From the Address list, select <b>Enter</b> Specific Value.		
	6.	In the Prefix box, type 10.0.1.0/24.		
	7.	Click <b>OK</b> twice.		
Specify the referenced pool	1.	Next to Then, select <b>Configure</b> .	Set the pool and action for the term:	
for nat-term and set its action—to translate the source addresses to	2. 3.	From the Designation list, select <b>Translated</b> .	set rule nat-rule term nat-term then translated	
addresses from the		Next to Translated, click Configure.		
one-to-one basis.	4.	From the Source pool choice list, select <b>Source pool</b> .		
	5.	In the Source pool box, type nat-pool.		
	6.	Click <b>OK</b> .		

### Table 85: Configuring Basic Source Static NAT (continued)

# Statically Assigning NAT Addresses from a Dynamic Pool

On a Services Router you can statically assign addresses from a pool that is being used for dynamic NAT. This approach enables you to advertise one subnet representing the NAT pool and use addresses within the subnet for static rules. However, you cannot reuse these statically assigned addresses for dynamic assignment.



**NOTE:** The addresses assigned statically from the dynamic pool can be used only for source static NAT and not for destination static NAT.

The example in this section shows how to create two pools—**static-pool** and **dynamic-pool**—and statically assign NAT addresses from a dynamic NAT pool with the terms described in Table 86 on page 188.

### Table 86: Sample Terms for Statically Assigned NAT Addresses

Term	Purpose
static-pool-term	Statically assigns addresses to translate the source address <b>10.10.10.2</b> . The translated address is an address within the static pool <b>121.0.1.10</b> through <b>121.0.1.12</b> . This static pool is a subnet from the dynamic pool.
dynamic-pool-term	Dynamically assigns addresses for translation of source addresses of all addresses not specified in <b>static-pool-term</b> . The translated address is within the dynamic pool <b>121.0.1.0/24</b> . The addresses <b>121.0.1.10</b> , <b>121.0.1.11</b> and <b>121.0.1.12</b> (reserved for the static pool) are excluded from the dynamic pool.

To statically assign NAT addresses from a dynamic pool:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 87 on page 188.
- 3. Apply the NAT configuration to an interface. See "Applying NAT to an Interface" on page 198.

### Table 87: Statically Assigning NAT Addresses from Dynamic NAT Pool

J-Web Configuration Editor		CLI Configuration Editor
1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit services nat
2.	Next to Services, click Configure or Edit.	
3.	Next to Nat, click <b>Configure</b> or <b>Edit</b> .	
1.	Next to Pool, click Add new entry.	Set the NAT pool name and the address:
2.	In the Pool Name box, type dynamic-pool.	set pool dynamic-pool address 121.0.1.0/24
3.	Next to Address, click Add new entry.	
4.	In the Prefix box, type <b>121.0.1.0/24</b> .	
5.	Click <b>OK</b> twice.	
1.	Next to Pool, click Add new entry.	Set the NAT pool name and the address range:
2.	In the Pool Name box, type static-pool.	set pool static-pool address-range low 121.0.1.10
3.	Next to Address range, click <b>Add new</b> entry.	high 121.0.1.12
4.	In the High box, type <b>121.0.1.12</b> .	
5.	In the Low box, type <b>121.0.1.10</b> .	
6.	Click <b>OK</b> .	
	J-W 1. 2. 3. 1. 2. 3. 4. 5. 1. 2. 3. 4. 5. 6.	J-Web Configuration Editor1.In the J-Web interface, select Configuration > View and Edit > Edit Configuration.2.Next to Services, click Configure or Edit.3.Next to Nat, click Configure or Edit.1.Next to Pool, click Add new entry.2.In the Pool Name box, type dynamic-pool.3.Next to Address, click Add new entry.4.In the Prefix box, type 121.0.1.0/24.5.Click OK twice.1.Next to Pool, click Add new entry.2.In the Prefix box, type 121.0.1.0/24.5.Click OK twice.1.Next to Pool, click Add new entry.2.In the Prefix box, type 121.0.1.0/24.5.Click OK twice.4.In the Pool Name box, type static-pool.5.In the Pool Name box, type static-pool.6.Click OK.

Task	J-Web Configuration Editor		CLI Configuration Editor
Define static-in-dynamic-rule and set its match direction.	1.	On the Nat page, next to Rule, click <b>Add new entry</b> .	Set the rule name and its match direction:
	2.	In the Rule name box, type static-in-dynamic-rule.	set rule static-in-dynamic-rule match-direction input
	3.	From the Match direction list, select <b>input</b> .	
Define static-pool-term for static-in-dynamic-rule and specify its match condition—source address 10.10.10.2	1.	On the Rule page, next to Term, select <b>Add new entry</b> .	Set the term name and its match condition:
	2.	In the Term name box, type static-pool-term.	set rule static-in-dynamic-rule term static-pool-term from source-address 10.10.10.2
	3.	Next to From, click Configure.	
	4.	Next to Source Address, click <b>Add new</b> entry.	
	5.	From the Address list, select <b>Enter</b> Specific Value.	
	6.	In the Prefix box, type 10.10.10.2.	
	7.	Click <b>OK</b> twice.	
Specify the referenced pool	1.	Next to Then, select Configure.	Set the pool and action for the term:
for static-pool-term and set its action—translation type as source static.	2.	From the Designation list, select <b>Translated</b> .	set rule static-in-dynamic-rule term static-pool-term then translated source-pool static-pool
	3.	Next to Translated, click <b>Configure</b> .	translation-type source static
	4.	From the Source pool choice list, select <b>Source pool</b> .	
	5.	In the Source pool box, type static-pool.	
	6.	Click <b>OK</b> .	
Define dynamic-pool-term	1.	Next to Term, click Add new entry.	Set the name of the term, its reference pool
for static-in-dynamic-rule. Specify the pool to be used for address translation and the term's action—to dynamically assign addresses for source address translation.	2.	In the Term name box, type dynamic-pool-term.	set rule static-in-dynamic-rule term
	3.	Next to Then, click Configure.	dynamic-pool-term then translated source-pool
	4.	From the Designation list select <b>Translated</b> .	dynamic-pool translation-type source dynamic
The action is taken on packets not matching static-pool-term.	5.	Next to Translated, click Configure.	
	6.	From the Source pool choice list, select <b>Source pool</b> .	
	7.	In the Source pool box, type dynamic-pool.	
	8.	From the Source translation type list, select <b>dynamic</b> .	
	9.	Click OK.	

# Table 87: Statically Assigning NAT Addresses from Dynamic NAT Pool (continued)

# **Configuring Full-Cone NAT**

To configure full-cone NAT, you must define a NAT pool that specifies the address to be used for network address translation. Next, you must define a NAT rule and then apply this rule to an interface. Each NAT rule consists of a set of terms that contain match conditions and actions. For a description of NAT match conditions and actions, see "Network Address Translation" on page 163.

The example in this section shows a full-cone NAT configuration. It shows how to create the pool **nat-pool** and define the rule **nat-rule** for full-cone NAT.

To configure full-cone NAT:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 88 on page 190.
- 3. Apply the NAT configuration to an interface. See "Applying NAT to an Interface" on page 198.

### **Table 88: Configuring Full-Cone NAT**

J-Web Configuration Editor		CLI Configuration Editor
1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter
2.	Next to Services, click <b>Configure</b> or <b>Edit</b> .	
3.	Next to Nat, click <b>Configure</b> or <b>Edit</b> .	
1.	Next to Pool, click Add new entry.	Set the NAT pool name and the address:
2.	In the Pool Name box, type nat-pool.	set pool nat-pool address 121.0.1.0/24
3.	Next to Address, click Add new entry.	
4.	In the Prefix box, type 121.0.1.0/24.	
5.	Click <b>OK</b> twice.	
1.	On the Nat page, next to Rule, click <b>Add new entry</b> .	Set the rule name and its match direction:
2.	In the Rule name box, type nat-rule.	set rule nat-rule match-direction output
3.	From the Match direction list, select <b>output</b> .	
	J-W 1. 2. 3. 1. 2. 3. 4. 5. 1. 2. 3. 3. 3. 3. 1. 3. 5. 1. 3. 5. 1. 3. 5. 5. 1. 5. 5. 5. 5. 5. 5. 5. 5. 5. 5	J-Web Configuration Editor1.In the J-Web interface, select Configuration > View and Edit > Edit Configuration.2.Next to Services, click Configure or Edit.3.Next to Nat, click Configure or Edit.1.Next to Pool, click Add new entry.2.In the Pool Name box, type nat-pool.3.Next to Address, click Add new entry.4.In the Prefix box, type 121.0.1.0/24.5.Click OK twice.1.On the Nat page, next to Rule, click Add new entry.2.In the Rule name box, type nat-rule.3.From the Match direction list, select output.

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Define <b>nat-term</b> for <b>nat-rule</b> and specify its match condition—source address 10.0.1.0/24.	1.	On the Rule page, next to Term, select <b>Add new entry</b> .	Set the term name and its match condition:
	2.	In the Term name box, type nat-term.	set rule nat-rule term nat-term from application-sets nat-application
· · · · · · · · · · · · · · · · · · ·	3.	Next to From, click Configure.	
	4.	Next to Application sets, click <b>Add new</b> entry.	set rule nat-rule term nat-term from source-address 10.100.136.5/24
	5.	In the Application set name box, type nat-application.	set rule nat-rule term nat-term from source-address-range 10.100.136.1/24
	6.	Next to Applications, click Add new entry.	10.100.136.5/24
	7.	In the Application name box, type nat-application.	set rule nat-rule term nat-term from source-prefix-list nat-source
	8.	Next to Destination address, click <b>Add new entry</b> .	set rule nat-rule term nat-term then translated source-pool nat-pool
	9.	From the Address list, select <b>Enter</b> Specific Value or any-unicast.	
	10.	If you have selected <b>Enter Specific Value</b> , then in the Address box, type <b>10.100.136.1/24</b> .	
	11.	Next to Destination address range, click <b>Add new entry</b> .	
	12.	Next to Low box, type <b>10.100.136.1/24</b> .	
	13.	Next to High box, type <b>10.00.136.5/24</b> .	
	14.	Next to Destination prefix list, click <b>Add new entry</b> .	
	15.	In the Name box, type nat-destination.	
	16.	Next to Source Address, click <b>Add new</b> entry.	
	17.	From the Address list, select <b>Enter</b> <b>Specific Value</b> or <b>any-unicast</b> .	
	18.	If you have selected <b>Enter Specific Value</b> , then in the Prefix box, type <b>10.100.136.1/24</b> .	
	19.	Next to Source address range, click <b>Add new entry</b> .	
	20.	Next to Low box, type 10.100.136.1/24.	
	21.	Next to High box, type <b>10.100.136.5/24</b> .	
	22.	Next to Source prefix list, click <b>Add new entry</b> .	
	23.	In the Name box, type nat-source.	
	24.	Click <b>OK</b> twice.	

# Table 88: Configuring Full-Cone NAT (continued)

### Table 88: Configuring Full-Cone NAT (continued)

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Specify the referenced pool for nat-term and set its action—to translate the source addresses to addresses from the referenced pool on a one-to-one basis.	1.	From the Nat Type choice list, select <b>full-cone</b> .	Set the pool and action for the term:
	2.	Next to Then, select Configure.	set rule nat-rule term nat-term then translated source-pool nat-pool translation-type source static
	3.	From the Designation list, select <b>Translated</b> .	set rule nat-rule term nat-term from
	4.	Next to Translation type, click Configure.	destination-address 10.100.136.1/24
	5.	From the Source pool choice list, select <b>Source pool</b> .	set rule nat-rule term nat-term from destination-address-range 10.100.136.1/24
	6.	In the Source pool box, type <b>nat-pool</b> .	10.100.136.5/24
	7.	Click <b>OK</b> .	set rule nat-rule term nat-term from destination-prefix-list nat-destination
			set rule nat-rule term nat-term then nat-group1
			set rule nat-rule term nat-term nat-type full-cone
Specify the groups for		pand the Advanced option.	Set the group and group exceptions for NAT:
which this NAT configuration is applicable and the exception groups.	1.	Next to the Apply groups, click <b>Add new</b> entry.	set rule nat-rule term nat-term then translated nat-group
	2.	In the Value box, type <b>nat-group</b> .	
	3.	Next to the Apply groups except, click <b>Add new entry</b> .	set rule nat-rule term nat-term then translated nat-group1
	4.	In the Value box, type <b>nat-group1</b> .	
	5.	Click <b>OK</b> twice.	

# **Configuring NAT Rules Without Defining Pools**

For host-to-host NAT, you can define a NAT rule without having to specify a pool. Instead, you specify the translated address directly in a NAT rule.

The example in this section shows how to create a term **no-pool-term** to dynamically assign the translated address from the prefix **121.0.1.0/24** for source address translation. You do not have to specify the referenced pool in the term. Similarly, you can configure destination static NAT by defining a destination prefix in the term instead of defining the destination pool.

To configure NAT rules without defining pools:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 89 on page 193.
- 3. Apply the NAT configuration to an interface. See "Applying NAT to an Interface" on page 198.
| Task                                                                              | J-W | eb Configuration Editor                                                                  | CLI Configuration Editor                                                           |
|-----------------------------------------------------------------------------------|-----|------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| Navigate to the <b>Nat</b> level in the configuration hierarchy.                  | 1.  | In the J-Web interface, select<br>Configuration > View and Edit > Edit<br>Configuration. | From the [edit] hierarchy level, enter edit services nat                           |
|                                                                                   | 2.  | Next to Services, click Configure or Edit.                                               |                                                                                    |
|                                                                                   | 3.  | Next to Nat, click <b>Configure</b> or <b>Edit</b> .                                     |                                                                                    |
| Define <b>no-pool-rule</b> and set its match direction.                           | 1.  | On the Nat page, next to Rule, click <b>Add new entry</b> .                              | Set the rule name and match direction:                                             |
|                                                                                   | 2.  | In the Rule name box, type no-pool-rule.                                                 | set rule no-pool-rule match-direction input                                        |
|                                                                                   | 3.  | From the Match direction list, select <b>input</b> .                                     |                                                                                    |
| Define <b>no-pool-term</b> and set                                                | 1.  | Next to Term, click Add new entry.                                                       | Set the term name and translation type:                                            |
| its translation<br>type—dynamic.                                                  | 2.  | In the Term name box, type no-pool-term.                                                 | set rule no-pool-rule term no-pool-term then                                       |
|                                                                                   | 3.  | Next to Then, click <b>Configure</b> .                                                   | translated translation-type source dynamic                                         |
|                                                                                   | 4.  | From the Designation list, select <b>Translated</b> .                                    |                                                                                    |
|                                                                                   | 5.  | Next to Translated, click <b>Configure</b> .                                             |                                                                                    |
| Define an action for <b>no-pool-term</b> —source                                  | 1.  | From the Source pool choice list, on the Translated page, select <b>Source prefix</b> .  | Set the source prefix:                                                             |
| prefix. This prefix is used<br>for network address<br>translation, and you do not | 2.  | In the Source prefix box, type 121.0.1.0/24.                                             | set rule no-pool-rule term no-pool-term then translated source-prefix 121.0.1.0/24 |
| have to specify a referenced pool.                                                | 3.  | Click <b>OK</b> .                                                                        |                                                                                    |

#### **Table 89: Defining NAT Rules Without NAT Pools**

## **Defining an Overload Pool or an Overload Prefix**

On the Services Router, you can configure an oversubscribed NAT pool to fall back on Network Address Port Translation (NAPT), also known as Port Address Translation (PAT). An overload NAPT pool provides additional NAT sessions when all the addresses in the source pool are in use. You can use one public address multiple times by assigning different port numbers to it.

Alternatively, for an oversubscribed NAT pool, you can configure an overload prefix to be used when the address pool is exhausted.

This example shows how to define an overload pool or an overload prefix. The terms used in the example are described in Table 90 on page 194.



**NOTE:** An overload prefix is an alternative to an overload pool. Define either over-pool-term or over-prefix-term, not both.

### Table 90: Sample Terms for Defining an Overload Pool or Prefix

Term	Purpose
over-pool-term	Dynamically translates the source address (10.10.10.0/24) to an address within the pool 121.0.1.2 through 121.0.1.20. After the addresses from the pool are used, the system uses the NAPT pool (pat-pool) 121.0.1.21 through 121.0.1.22 for address translation in combination with dynamically assigned ports by means of NAPT.
over-prefix-term	Dynamically translates the source address (10.10.10.0/24) to an address within the pool 121.0.1.2 through 121.0.1.20. After these addresses are used, the system uses the prefix 123.0.1.0/24.

To define an overload pool or prefix:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 91 on page 194.
- 3. Apply the NAT configuration to an interface. See "Applying NAT to an Interface" on page 198.

### **Table 91: Defining an Overload Pool or Prefix**

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Nat</b> level in the configuration	1.	In the J-Web interface, select Configuration > View and Edit > Edit	From the [edit] hierarchy level, enter
hierarchy.		Configuration	edit services nat
	2.	Next to Services, click <b>Configure</b> or <b>Edit</b> .	
	3.	Next to Nat, click <b>Configure</b> or <b>Edit</b> .	
Define nat-pool and assign	1.	Next to Pool, click Add new entry.	Set the NAT pool name and the address range:
used for network address	2.	In the Pool Name box, type nat-pool.	set pool nat-pool address-range high 121.0.1.20
translation.	3.	Next to Address range, click <b>Add new</b> entry.	low 121.0.1.2
	4.	In the High box, type <b>121.0.1.20</b> .	
	5.	In the Low box, type <b>121.0.1.2</b>	
	6.	Click <b>OK</b> twice.	
Define <b>pat-pool</b> and assign it an address range to be	1.	On the Nat page, next to Pool, click <b>Add new entry</b> .	Set the NAPT pool and address range:
used after addresses from nat-pool are fully used.	2.	In the Pool name box, type pat-pool.	set pool pat-pool address-range high 121.0.1.22 low 121.0.1.21
	3.	Next to Address range, click <b>Add new</b> entry.	
	4.	In the High box, type <b>121.0.1.22</b> .	
	5.	In the Low box, type <b>121.0.1.21</b> .	
	6.	Click <b>OK</b> .	

Table 9	<b>)1: Defining</b>	an Overload	Pool or	Prefix (	(continued)
---------	---------------------	-------------	---------	----------	-------------

Task	J-Web Configuration Editor	CLI Configuration Editor
Specify the NAT port to be automatically assigned by	<ol> <li>On the Pool page, next to Port, click Configure.</li> </ol>	Set the NAT port to be assigned automatically:
the router.	2. From the Port choice list select <b>Automatic</b> .	set pool pat-pool port automatic
	3. Click <b>OK</b> twice.	
Define <b>over-pool-rule</b> and set its match direction.	1. On the Nat page, next to Rule, click <b>Add new entry</b> .	Set the rule and its match direction:
	2. In the Rule name box, type over-pool-rule.	set rule over-pool-rule match-direction input
	3. From the Match direction list, select <b>input</b> .	
<ul> <li>Define one of the following terms for over-pool-rule:</li> <li>For an overload pool—over-pool-term</li> <li>For an overload prefix—over-perfix-term</li> </ul>	<ol> <li>Next to Term, click Add new entry.</li> <li>In the Term name box, type the appropriate name:         <ul> <li>over-pool-term</li> <li>over-prefix-term</li> </ul> </li> </ol>	<ul> <li>Set the appropriate term for the rule:</li> <li>For an overload pool: set rule over-pool-rule term over-pool-term</li> <li>For an overload prefix: set rule over-pool-rule term over-prefix-term</li> </ul>
Define a match condition—the source address 10.10.10.0/24— for the term (over-pool-term or over-prefix-term).	<ol> <li>Next to From, click Configure.</li> <li>Next to Source address, click Add new entry.</li> <li>From the Address list, select Enter Specific Value.</li> <li>In the Prefix box, type 10.10.10.0/24.</li> <li>Click OK twice.</li> </ol>	<ul> <li>Set the match condition for the term, as appropriate:</li> <li>For an overload pool: set rule over-pool-rule term over-pool-term from source-address 10.10.10.0/24</li> <li>For an overload prefix: set rule over-pool-rule term over-prefix-term from source-address 10.10.10.0/24</li> </ul>

Task	J-Web Configuration Editor	CLI Configuration Editor
<ul> <li>Task</li> <li>Define an action for the term:</li> <li>For over-pool-term, define a translation type, the source pool (nat-pool) and the overload pool (pat-pool).</li> <li>For over-prefix-term, define a translation type, the source pool (nat-pool) and the overload prefix (123.0.1.0/24).</li> </ul>	<ol> <li>J-Web Configuration Editor</li> <li>Next to Then, click Configure.</li> <li>From the Designation list select Translated.</li> <li>Next to Translated, click Configure.</li> <li>From the Source translation type list, select dynamic.</li> <li>From the Source pool choice list, select Source pool.</li> <li>In the Source pool box, type nat-pool.</li> <li>From the Overload pool choice list, select the appropriate choice:         <ul> <li>Overload pool</li> <li>Overload prefix</li> </ul> </li> <li>Do one of the following:         <ul> <li>In the Overload pool box, type nat-pool.</li> </ul> </li> </ol>	<ul> <li>CLI Configuration Editor</li> <li>Set the appropriate action for the term:         <ul> <li>For an overload pool: set rule over-pool-rule term over-pool-term then translated translation-type source dynamic</li> <li>set rule over-pool-rule term over-pool-term then translated source-pool nat-pool set rule over-pool-rule term over-pool-term then translated overload-pool pat-pool</li> </ul> </li> <li>For an overload prefix: set rule over-pool-rule term over-prefix-term then translated translation-type source dynamic</li> <li>set rule over-pool-rule term over-prefix-term then translated source-pool nat-pool</li> <li>set rule over-pool-rule term over-prefix-term then translated source-pool nat-pool set rule over-pool-rule term over-prefix-term then translated source-pool nat-pool</li> </ul>
	<ul> <li>In the Overload pool box, type pat-pool.</li> <li>In the Overload prefix box, type 123.0.1.0/24.</li> <li>Click OK.</li> </ul>	then translated overload-prefix 123.0.1.0/24

### **Defining Rules for Transparent NAT**

On the Services Router, you can define a rule to perform NAT selectively. This method is useful when you want to perform NAT on a large prefix that includes a few addresses that you do not want to translate. Instead of defining multiple terms to specify source addresses for translation, you can define two terms—one to specify the source prefix for translation and the other to specify source addresses in this prefix that are to be skipped.

This example shows how to define rules to perform NAT selectively by using the terms described in Table 92 on page 196.

Table 92:	Sample	<b>Terms for</b>	Defining	<b>Rules for</b>	Transparent NAT
-----------	--------	------------------	----------	------------------	-----------------

Term	Purpose
selective-term	Skips source prefix <b>192.168.1.1/24</b> from network address translation.
accept-all-term	Dynamically translates all addresses besides prefix <b>192.168.1.1/24</b> to an address from the defined source pool.

To define a rule for transparent NAT:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 93 on page 197.
- 3. Apply the NAT configuration to an interface. See "Applying NAT to an Interface" on page 198.

Task	J-Web Configuration Editor		CLI Configuration Editor
Navigate to the <b>Nat</b> level in the configuration		In the J-Web interface, select Configuration > View and Edit > Edit	From the [edit] hierarchy level, enter
hierarchy.		Configuration.	edit services nat
	2.	Next to Services, click <b>Configure</b> or <b>Edit</b> .	
	3.	Next to Nat, click <b>Configure</b> or <b>Edit</b> .	
Define nat-pool and assign	1.	Next to Pool, click Add new entry.	Set the address pool name and the address
it an address range to be used for network address	2.	In the Pool Name box, type nat-pool.	range:
translation.	3.	Next to Address range, click <b>Add new</b> entry.	set pool nat-pool address-range high 10.10.10.16 low 10.10.10.1
	4.	In the High box, type 10.10.10.16.	
	5.	In the Low box, type <b>10.10.10.1</b> .	
	6.	Click <b>OK</b> .	
Specify the source port pool to be automatically		On the Pool page, next to Port, click <b>Configure</b> .	Configure the source port translation to be automatic:
assigned by the router.	2.	From the Port choice list, select <b>Automatic</b> .	set pool nat-pool port automatic
	3.	Click <b>OK</b> twice.	
Define <b>selective-rule</b> and set its match direction.	1.	On the Nat page, next to Rule, click <b>Add new entry</b> .	Set the rule and its match direction:
	2.	In the Rule name box, type selective-rule.	set rule selective-rule match-direction input
	3.	From the Match direction list, select <b>input</b> .	
Define selective-term for	1.	Next to Term, click Add new entry.	Set the term:
selective-rule.	2.	In the Term name box, type selective-term.	set rule selective-rule term selective-term
Define the match condition	1.	Next to From, click <b>Configure</b> .	Set the match condition for the term:
for selective-term—the source prefix 192.168.1.1/24.	2.	Next to Source address, click <b>Add new</b> entry.	set rule selective-rule term selective-term from source-address 192.168.1.1/24
	3.	From the Address list, select <b>Enter</b> Specific Value.	
	4.	In the Prefix box, type <b>192.168.1.1/24</b> .	
	5.	Click <b>OK</b> twice.	

### **Table 93: Defining Rules for Transparent NAT**

#### Table 93: Defining Rules for Transparent NAT (continued)

Task	J-Web Configuration Editor		CLI Configuration Editor
Define an action for selective-term—no translation. The packets coming from the prefix	1.	Next to Then, click <b>Configure</b> .	Set the action for selective-term:
	2.	From the Designation list, select <b>No</b> translation.	set rule selective-rule term selective-term then no-translation
192.168.1.1/24 are skipped and not translated.	3.	Click <b>OK</b> twice.	
Define accept-all-term for	1.	Next to Term, click Add new entry.	Specify a term for selective-rule:
selective-rule.	2.	In the Term name box, type accept-all-term.	set rule selective-rule term accept-all-term
Define an action for accept-all-term and set the translation type for it.	1.	Next to Then, click Configure.	Set the action for accept-all-term:
	2.	From the Designation list, select <b>Translated</b> .	set rule selective-rule term accept-all-term then translated translation-type source dynamic
	3.	Next to Translated, click Configure.	
	4.	From the Source Translation Type list, select <b>dynamic</b> .	set rule selective-rule term accept-all-term then translated source-pool nat-pool
	5.	From the Source pool choice list, select <b>Source pool</b> .	
		In the Source pool box, type nat-pool.	
	7.	Click <b>OK</b> .	

### Applying NAT to an Interface

To enable the NAT services on an interface, you assign the defined NAT rules to a service set and apply the service set to an interface. For more information about applying services to an interface, see the *JUNOS Services Interfaces Configuration Guide*.

You enable NAT services on an interface as follows:

- Define a service set.
- Assign the NAT rule that you have already defined to the service set. You can
  include one or more rules or one rule set for one service type. The rules are
  applied in the order that they are configured.
- Define a service set type for the service set and assign a virtual interface sp-0/0/0 as the service interface for this set. You can configure two types of service sets—interface service sets or next-hop service sets.
- Apply this service interface to the physical interface on which NAT is to be enabled. You assign the defined service set to the input and output sides of the physical interface.

To apply NAT to an interface:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 94 on page 199.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To verify NAT, see "Verifying NAT Configuration" on page 200.

Task	J-Web Configuration Editor		CLI Configuration Editor
Navigate to the <b>Services</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter
	2.	Next to Services, click <b>Configure</b> or <b>Edit</b> .	
Define a service set—for	1.	Next to Service set, click Add new entry.	Set the service set and assign the NAT rule to
example, nat-service-set.	2.	In the Service set name box, type nat-service-set.	II: set service-set service-set-name nat-rules
rule to the service set—for example, nat-rule.	3.	From the Nat rules choice list, select <b>Nat rules</b> .	nat-rule-name
	4.	Next to Nat rules, click Add new entry.	
	5.	In the Rule name box, type the name of the defined NAT rule—for example, nat-rule.	
	6.	Click <b>OK</b> .	
Define a service set type and virtual service interface sp-0/0/0 as the service interface for nat-service-set.	1.	From the Service type choice list, select <b>Interface service</b> .	Define the service set type and the service interface:
	2.	Next to Interface service, click <b>Configure</b> .	set service-set nat-rule-set interface-service
	3.	In the Service interface box, type <b>sp-0/0/0</b> .	service-interface sp-0/0/0
	4.	Click <b>OK</b> .	
Navigate to the <b>Interfaces</b> level in the configuration	the <b>Interfaces</b> On the main Configuration page next to		From the [edit] hierarchy level, enter
hierarchy.			edit interface
Configure the sp-0/0/0	1.	Next to Interface, click Add new entry.	Set the service interface:
service interface.	2.	In the Interface name box, type <b>sp-0/0/0</b> .	set interfaces sp-0/0/0 unit 0 family inet
(See the interface naming	3.	Click <b>OK</b> .	
conventions in the <i>J-series</i> Services Router Basic LAN	4.	Click <b>sp-0/0/0</b> .	
and WAN Access Configuration Guide )	5.	Next to Unit, click Add new entry.	
conjugaration datae.)	6.	In the Interface unit number box, type $0$ .	
	7.	Next to Inet, select the check box.	

8. Click OK.

### Table 94: Applying NAT to an Interface

#### Table 94: Applying NAT to an Interface (continued)

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Apply <b>nat-service-set</b> to the input and output sides of the physical interface on which NAT is to be	1.	On the main Configuration page next to Interfaces, click <b>Edit</b> .	From the [edit] hierarchy level, apply the service set to the interface:
	2.	Under Interface name, click <b>t1-0/0/0</b> .	set interfaces $t1-0/0/0$ unit 0 family inet service
enabled—for example	3.	Under Interface unit number, click $0$ .	input service-set nat-service-set
t1-0/0/0.	4.	Under Family, make sure the Inet check box is selected, and click <b>Configure</b> or <b>Edit</b> .	set interfaces t1–0/0/0 unit 0 family inet service output service-set nat-service-set
	5.	Next to Service, click <b>Configure</b> .	
		Next to Input, click Configure.	
	7.	Next to Service set, click Add new entry.	
	8.	In the Service set name box, type nat-service-set.	
	9.	Click <b>OK</b> twice.	
	10.	Next to Output, click <b>Configure</b> .	
	11.	Next to Service set, click Add new entry.	
	12.	In the Service set name box, type nat-service-set.	
	13.	Click <b>OK</b> .	

## **Verifying NAT Configuration**

NAT is configured independently and with stateful firewall filters. Some **show** commands used for verification are common for the stateful firewall filters and NAT. For verifying NAT configured with stateful firewall filters, see "Verifying Stateful Firewall Filter Configuration" on page 217.

To verify a NAT configuration, perform these tasks:

- Displaying NAT Configurations on page 200
- Verifying NAT on page 202

### **Displaying NAT Configurations**

**Purpose** Verify NAT configuration.

Action From the J-Web interface, select Configuration > View and Edit > View Configuration Text.

Alternatively, from configuration mode in the CLI perform the following tasks:

- Enter the **show services** command to display the complete NAT configuration.
- Enter the **show interfaces** command to display the interface configuration.

The sample output in this section displays the NAT configurations provided in "Configuring Basic Source Static NAT" on page 186.

```
[edit]
user@r1# show services
nat {
 pool nat-pool {
 address {
 121.0.1.0/24;
 }
 }
 rule nat-rule {
 match-direction output;
 term nat-term {
 nat-type (symmetric|full-cone)
 from {
 source-address {
 10.0.1.0/24;
 }
 }
 then {
 translated {
 translation-type {
 source-pool nat-pool;
 translation-type source (static|dynamic);
 }
 }
 }
 }
}
service-set nat-service-set {
 nat-rules nat-rule;
 interface-service {
 service-interface sp-0/0/0;
 }
}
[edit]
user@r1# show interfaces
t3-1/0/0 {
 description "t3-1/0/0 on r1";
 unit 0 {
 family inet {
 service {
 input {
 service-set nat-service-set;
 }
 output {
 service-set nat-service-set;
 }
 }
 }
}
```

**What It Means** Verify that the output shows the intended NAT and interface configurations.

**Related Topics** For more information about the format of a configuration file, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

### **Verifying NAT**

**Purpose** Verify the NAT configured in "Configuring Basic Source Static NAT" on page 186.

**Action** Take the following actions:

To verify that the network address is translated as configured, create a traffic flow between two routers—an internal router r1 and an external router r2. On r1, configure NAT as shown in "Configuring Basic Source Static NAT" on page 186 and apply the defined nat-service-set on an interface. Configure loopback address 10.0.1.2 on r1 and loopback address 24.40.80.2 on r2.

```
(¥
```

**NOTE:** You are configuring loopback addresses in this example for verification purposes only. If you have the network set up and the source address **10.0.1.2** is configured on a host, ping an external router from the host. In this case, you do not need to configure the loopback address.

- Use the ping command to verify that a connection is established between the two routers used in this sample.
- From the CLI, enter the show services stateful-firewall conversations command to display the flow conversations.

```
user@r1> ping 24.40.80.2 source 10.0.1.2
 PING 24.40.80.2 (24.40.80.2): 56 data bytes
 64 bytes from 24.40.80.2: icmp_seq=0 ttl=64 time=6.669 ms
 64 bytes from 24.40.80.2: icmp_seq=1 ttl=64 time=40.441 ms
 . . .
 user@r1> show services stateful-firewall conversations extensive
 Interface: sp-0/0/0, Service set: nat-service-set
 Conversation: ALG protocol: icmp
 Number of initiators: 1, Number of responders: 1
 State Dir Frm count
 Flow.
 10.0.1.2:52499 -> 24.40.80.2 Watch 0
 TCMP
 2
 NAT source 10.0.1.2:52499 -> 121.0.1.2:52499
 Byte count: 84
 Flow role: Master, Timeout: 30, Protocol detail: echo request
 ICMP
 24.40.80.2:52499 -> 121.0.1.2
 Watch I
 2
 NAT dest
 121.0.1.2:52499 -> 10.0.1.2:0
 Byte count: 84
 Flow role: Responder, Timeout: 30, Protocol detail: echo reply
What It Means
 Verify the following information:
```

• A ping request from r1 returns a ping response from r2. The sample ping command output shows a series of replies, indicating that the connection is

working and traffic is transmitted between the two routers. If there is no connection, a "host unreachable" message is displayed.

The source address is translated to an address from the configured NAT address pool. The sample output shows the flow from r1 to r2 and its response. In the flow from r1 to r2, the source address 10.0.1.2 is translated to address 121.0.1.2 from the configured NAT address pool (121.0.1.0/24). The response flow correctly shows reverse translation from 121.0.1.2 to 10.0.1.2.

Alternatively, you can use the **show services stateful-firewall flows** command to display the NAT flows. The **show services stateful-firewall conversations** command is easier to use for verification because it displays corresponding NAT flows together instead of a random listing of all flows.

**Related Topics** For detailed descriptions of the show services stateful-firewall conversations and show services stateful firewall flows commands and output, see the *JUNOS System Basics* and Services Command Reference.

For information about using the J-Web interface to ping a host, see the *J-series Services Router Administration Guide*.

J-series[™] Services Router Advanced WAN Access Configuration Guide

# Chapter 12 Configuring Stateful Firewall Filters and NAT

A *stateful* firewall filter inspects traffic flowing between a trusted network and an untrusted network. In contrast to a *stateless* firewall filter that inspects packets in isolation, a stateful firewall filter provides an extra layer of security by using state information derived from past communications and other applications to make dynamic control decisions.

On the Services Router you can configure Network Address Translation (NAT) either independently or with a stateful firewall filter. For information on configuring NAT independently, see "Configuring NAT" on page 185.

You can use either J-Web Quick Configuration or a configuration editor to configure stateful firewall filters and NAT.

This chapter contains the following topics. For more information about stateful firewall filters and NAT, see the *JUNOS Services Interfaces Configuration Guide*. To configure a *stateless* firewall filter, see "Configuring Stateless Firewall Filters" on page 221.

- Before You Begin on page 205
- Configuring a Stateful Firewall Filter with Quick Configuration on page 206
- Configuring a Stateful Firewall Filter with a Configuration Editor on page 211
- Verifying Stateful Firewall Filter Configuration on page 217

# **Before You Begin**

Before you begin configuring stateful firewall filters, complete the following tasks:

- If you do not already have an understanding of stateful firewall filters, read "Stateful Firewall Filters" on page 155.
- Before you begin configuring stateful firewall filters and NAT, you must configure the interfaces on which to apply these services. To configure an interface, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.



**CAUTION:** If a packet does not match any terms in a firewall filter rule, the packet is discarded. Take care you do not configure a stateful firewall filter that prevents

you from accessing the Services Router after you commit the configuration. For example, if you configure a firewall filter that does not match HTTP or HTTPS packets, you cannot access the router with the J-Web interface.

# **Configuring a Stateful Firewall Filter with Quick Configuration**

You can use the Firewall/NAT Quick Configuration pages to configure a stateful firewall filter and NAT. These Quick Configuration pages allow you to designate the interfaces that make up the untrusted network. In addition, you can designate the applications that are allowed to operate from the untrusted network to the trusted network.

Figure 16 on page 207 and Figure 17 on page 208 show the Firewall/NAT Quick Configuration main and application pages.

Fl.d	Firewall	/NIAT	0	0 and duration	Main	Dere
rigure 10:	rirewali	/ NAI	QUICK	configuration	<b>Wall</b>	гаge

	ROUTER - J4300
Monitor Configuratio	n Diagnose Manage Events Logged in as: regress Help About Logout
Quick Configuration 🔷 🕨	Configuration > Quick Configuration > Firewall/NAT
View and Edit	Quick Configuration
History	Firewall/NAT
Rescue	
	Stateful Firewall
	Stateful firewall inspects traffic flowing between a trusted network and an untrusted network. All packets flowing from a trusted network to an untrusted network are allowed. Packets flowing from an untrusted network to a trusted network are allowed only if they are responses to a session originated by the trusted network.
	Enable Stateful Firewall 🗌
	Trusted Interfaces
	Select the interfaces to be part of a trusted network. Stateful firewall is applied to the untrusted interfaces.
	Untrusted Interfaces
	-> te=0/0/0.2
	Network Address Translation (NAT)
	When NAT is enabled, the source address of a packet flowing from a trusted network to an untrusted network is replaced with an address choosen from the specified range. The source port of the packet is also replaced with a dynamically chosen port.
	Enable NAT
	Low Address in Address Range 10 255:4:36
	High Address in Address Range
	Outside Applications Allowed
	The following applications are allowed to operate from the untrusted network to the trusted network.
	No applications are allowed from the untrusted network onto the trusted network.
	Add
	OK Cancel Apply
Copyright @ 2004-2005	luniner Networks, Inc. All Rights Reserved, Trademark Notice, Privacy, Juniper Vaour Net.

	ROUTER - J6300
Monitor Configuratio	n Diagnose Manage Events Alarms Logged in as: regress Help About Logout
Quick Configuration 👘	Configuration > Quick Configuration > Firewall/NAT
View and Edit. 🕨 🕨	Quick Configuration
History	Firewall/NAT         Allow an Application Through the Firewall
Rescue	
	Application
	• Application bgp
	Source Address
	Any Unicast WAN Address 🛛 🖂
	Source Addresses and Prefixes
	Destination Address
	Any Unicast LAN Address 🔽
	Destination Addresses and Prefixes
	Add Defete
	OK Cancel
Copyright @ 2004-2005, .	I Juniper Netvorks, Inc. <u>All Right's Reserved</u> . <u>Trademark Notice</u> . <u>Privacy</u> . <b>Juniper your Net.</b>

Figure 17: Firewall/NAT Quick Configuration Application Page

To configure a stateful firewall filter and NAT with Quick Configuration:

- 1. In the J-Web interface, select **Configuration > Firewall/NAT**.
- 2. Enter information into the Firewall/NAT Quick Configuration pages, as described in Table 95 on page 209.
- 3. Click one of the following buttons on the Firewall/NAT Quick Configuration main page:
  - To apply the configuration and stay in the Firewall/NAT Quick Configuration main page, click **Apply**.
  - To apply the configuration and return to the Quick Configuration page, click **OK**.
  - To cancel your entries and return to the Quick Configuration page, click **Cancel**.
- 4. Go on to one of the following procedures:

- To display the configuration, see Displaying Stateful Firewall Filter Configurations on page 217.
- To verify a stateful firewall filter, see Verifying a Stateful Firewall Filter on page 219.

### Table 95: Firewall/NAT Quick Configuration Pages Summary

Field	Function	Your Action			
Stateful Firewall					
Enable Stateful Firewall	Enables stateful firewall filter configuration.	To enable stateful firewall filter configuration, select the check box.			
Trusted Interfaces					
Trusted Interfaces	Designates the trusted and untrusted router interfaces. The stateful firewall filter is applied to the untrusted interfaces.	The Trusted Interfaces box displays a list of all the interfaces configured on the router. Do either of the following:			
		• To <i>apply</i> a stateful firewall filter to an interface, click the interface in the Trusted Interfaces box to highlight it, and click the left arrow to add the interface to the Untrusted Interfaces list. You can select multiple interfaces by pressing Ctrl while you click the interface.			
		• To <i>remove</i> a stateful firewall filter from an interface, click the interface in the Untrusted Interfaces box to highlight it, and click the right arrow to add the interface to the Trusted Interfaces list. You can select multiple interfaces by pressing Ctrl while you click the interface.			
Network Address Translat	tion (NAT)				
Enable NAT	Enables NAT configuration.	To enable NAT configuration, select the check box.			
Low Address in Address Range (required)	Specifies the lowest address in the NAT pool address range. If a range of addresses is not specified, you can specify a single address or an IP prefix.	Type an IP address or prefix.			
High Address in Address Range	Specifies the highest address in the NAT pool address range.	Type an IP address. The total range of addresses in the pool must be limited to a maximum of <b>32</b> .			
Outside Applications Allowed					
	Add or delete applications that are allowed to operate from the untrusted network to the trusted network.	Click <b>Add</b> to move to the Firewall/NAT Quick Configuration application page. When you have finished entering information into this page, click <b>OK</b> to save it.			
		To cancel your entries, click <b>Cancel</b> .			
Application					

### Table 95: Firewall/NAT Quick Configuration Pages Summary (continued)

Field	Function	Your Action
Application (required)	Designate which applications are allowed to operate from the untrusted network to the trusted network.	From the list, select the application you want to operate from the untrusted network to the trusted network.
Source Address		
Any Unicast WAN Address	Specifies that any unicast source address is allowed from the untrusted network.	To allow any unicast source address, select the check box.
Source Addresses and Prefixes	Designates the source addresses and prefixes that are allowed from the untrusted network.	To add an IP address and prefix, type them in the boxes above the <b>Add</b> button, then click <b>Add</b> .
		To delete an IP address and prefix, select them in the Source Addresses and Prefixes box, then click <b>Delete</b> .
<b>Destination Address</b>		
Any Unicast LAN Address	Specifies that any unicast destination address is allowed from the untrusted network.	To allow any unicast destination address, select the check box.
Destination Addresses and Prefixes	Designates the destination addresses and prefixes that are allowed from the untrusted network.	To add an IP address and prefix, type them in the boxes above the <b>Add</b> button, then click <b>Add</b> .
		To delete an IP address and prefix, select them in the Destination Addresses and Prefixes box, then click <b>Delete</b> .

# **Configuring a Stateful Firewall Filter with a Configuration Editor**

To configure a stateful firewall filter and NAT with a configuration editor, you do the following:

 Define the stateful firewall filter output and input rules. You must define an output rule that allows all traffic (application and nonapplication) to flow from the trusted network to the untrusted network.

To define the match condition in the term that allows application traffic to flow from the trusted network to the untrusted network, we recommend you specify the JUNOS default group junos-algs-outbound as the application set. To view the configuration of this group, enter the show groups junos-defaults applications application-set junos-algs-outbound configuration mode command. For more information about JUNOS default groups, see the *JUNOS CLI User Guide*.

You also must define an input rule to discard all traffic from the untrusted network that is not a response to a session originated by the trusted network.

- Define an address pool and port pool for NAT.
- Define NAT input and output rules.
- Define a service set that includes all stateful firewall filter and NAT rules and the service interface. You must specify the service interface as sp-0/0/0. This service interface is a virtual interface that must be included at the [edit interfaces] hierarchy level to support stateful firewall filter and NAT services.
- Finally, apply the service set to any interfaces on the Services Router that lead to or from the untrusted network.



**NOTE:** Do not apply the service set to the **sp-0/0/0** interface.

The example in this section shows how to create a stateful firewall filter and NAT with the rules described in Table 96 on page 211.

#### Table 96: Sample Stateful Firewall Filter and NAT Rules

Rule	Туре	Term or Terms	
to-wan-rule	Output	<ul> <li>app-term—Accepts packets from any of the applications defined by the JUNOS default group junos-algs-outbound application set.</li> </ul>	
		■ accept-all-term—Accepts packets that do not match app-term.	
from-wan-rule	Input	<ul> <li>wan-src-addr-term—Accepts input packets with a source prefix of 192.168.33.0/24.</li> </ul>	
		■ discard-all-term—Discards all packets.	
nat-to-wan-rule	Output	private-public-term—Translates the source address to an address within the pool <b>10.148.2.1</b> through <b>10.148.2.32</b> and dynamically translates the source port to a router-assigned port by means of NAPT	

The example also assigns the name **public-pool** to the NAT address pool and NAPT router-assigned port.

In addition, the example creates the service set wan-service-set that includes the stateful firewall filter and NAT services and defines sp-0/0/0 as its service interface. Finally, wan-service-set is applied to the WAN interface to the untrusted network, t1-0/0/0.

For stateful firewall match conditions, see "Stateful Firewall Filter Match Conditions" on page 156 and for stateful firewall actions, see "Stateful Firewall Filter Actions" on page 156.

To configure a stateful firewall filter and NAT and apply them to the WAN interface:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 97 on page 212.
- 3. To apply the stateful firewall filter and NAT to the interface, perform the configuration tasks described in Table 98 on page 215.
- 4. If you are finished configuring the router, commit the configuration.
- 5. Go on to one of the following procedures:
  - To display the configuration, see Displaying Stateful Firewall Filter Configurations on page 217.
  - To verify the stateful firewall filter, see Verifying a Stateful Firewall Filter on page 219.

#### **Table 97: Configuring a Stateful Firewall Filter and NAT**

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Stateful</b> <b>firewall</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit services stateful-firewall.
	2.	Next to Services, click Configure or Edit.	
	3.	Next to Stateful firewall, click <b>Configure</b> or <b>Edit</b> .	

Task	J-Web Configuration Editor		CLI Configuration Editor
Define to-wan-rule and set	1.	Next to Rule, click Add new entry.	Set the rule name, match direction, term name,
its match direction.	2.	In the Rule name box, type to-wan-rule.	and match condition:
	3.	From the Match direction list, select <b>output</b> .	set rule to-wan-rule match-direction output term app-term from application-sets junos-algs-outbound
Define app-term for the	1.	Next to Term, click Add new entry.	-
to-wan-rule rule.	2.	In the Term name box, type app-term.	
Define the match condition	1.	Next to From, click <b>Configure</b> .	-
for app-term—the default junos-algs-outbound application set.	2.	Next to Application sets, click <b>Add new</b> entry.	
-F.F.	3.	In the Application set name box, type junos-algs-outbound.	
	4.	Click <b>OK</b> twice.	
Define an action for app-term.	1.	On the Term <b>app-term</b> page, next to Then, click <b>Configure</b> .	Set the action:
	2.	In the Designation list, select Accept.	set rule to-wan-rule term app-term then accept
	3.	Click <b>OK</b> twice.	
Define accept-all-term for to-wan-rule.	1.	On the Rule <b>to-wan-rule</b> page, next to Term, click <b>Add new entry</b> .	Set the term name and the action:
	2.	In the Term name box, type accept-all-term.	set rule to-wan-rule term accept-all-term then accept
Define an action for	1.	Next to Then, click Configure.	-
accept-all-term. The action is taken only if a packet does not match app-term.	2.	From the Designation list, select Accept.	
	3.	Next to Accept, select the check box.	
	4.	Click <b>OK</b> three times.	

# Table 97: Configuring a Stateful Firewall Filter and NAT (continued)

### Table 97: Configuring a Stateful Firewall Filter and NAT (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define <b>from-wan-rule</b> and set its match direction.	1. On the Rule page, next to Rule, click <b>Add new entry</b> .	Set the rule name, match direction, term name, and the match condition:
	2. In the Rule name box, type from-wan-rule.	set rule from-wan-rule match-direction input term
	3. From the Match direction list, select <b>input</b> .	wan-src-addr-term from source-address 192.168.33.0/24
Define wan-src-addr-term for	1. Next to Term, click <b>Add new entry</b> .	
the from-wan-rule rule.	<ol> <li>In the Term name box, type wan-src-addr-term.</li> </ol>	
Define the match condition	1. Next to From, click <b>Configure</b> .	_
for wan-src-addr-term.	2. Next to Source address, click <b>Add new</b> entry.	
	<ol> <li>From the Address list, select Enter Specific Value— &gt; .</li> </ol>	
	4. In the Prefix box, type <b>192.168.33.0/24</b> .	
	5. Click <b>OK</b> twice.	
Define an action for wan-src-addr-term.	<ol> <li>On the Term wan-src-addr-term page, next to Then, click Configure.</li> </ol>	Set the action:
	2. In the Designation list, select Accept.	set rule from-wan-rule term wan-src-addr-term then accept
	3. Click <b>OK</b> twice.	
Define discard-all-term for from-wan-rule.	<ol> <li>On the Rule from-wan-rule page, next to Term, click Add new entry.</li> </ol>	Set the term name and the action:
	2. In the Term name box, type discard-all-term.	set rule from-wan-rule term discard-all-term then discard
Define an action for	1. Next to Then, click <b>Configure</b> .	_
discard-all-term. The action is taken only if a packet	2. From the Designation list, select <b>Discard</b> .	
does not match wan-src-addr-term.	3. Click <b>OK</b> three times.	
Navigate to the <b>Nat</b> level in the configuration	<ol> <li>On the main Configuration page next to Services, click Configure or Edit.</li> </ol>	From the [edit] hierarchy level, enter
hierarchy.	2. Next to Nat, click <b>Configure</b> or <b>Edit</b> .	edit services nat
Define the public-pool	1. Next to Pool, click <b>Add new entry</b> .	Set the address pool name and the range:
address pool name and range.	2. In the Pool name box, type <b>public-pool</b> .	set pool public-pool address-range low 10.148.2.1
5	3. From the Address choice list, select <b>Address range</b> .	high 10.148.2.32
	4. In the High box, type <b>10.148.2.32</b> . In the Low box, <b>10.148.2.1</b> .	

Took	1 34	lab Configuration Editor	OLL Configuration Editor
TASK	J-W	eb Configuration Editor	GLI Configuration Editor
Specify the NAT port pool	1.	Next to Port, click <b>Configure</b> .	Configure the source port translation to be
to be automatically assigned by the router.	2.	From the Port choice list, select <b>Automatic</b> .	automatic:
	3.	Click <b>OK</b> twice.	
Define nat-to-wan-rule and private-public-term.	1.	On the Nat page, next to Rule, click <b>Add</b> new entry.	Set the rule name, match direction, term name, and the term's pool name:
	2.	In the Rule name box, type nat-to-wan-rule.	set rule nat-to-wan-rule match-direction output
	3.	From the Match direction list, select <b>output</b> .	term private-public-term then translated source-pool public-pool
	4.	Next to Term, select Add new entry.	
	5.	In the Term name box, type private-public-term.	
	6.	Next to Then, select <b>Configure</b> .	
	7.	Next to Translated, select Configure.	
	8.	In the Source pool box, type public-pool.	
Set the NAT port translation type for private-public-term.	1.	Next to Translation type, select the check box.	Set the NAT translation type:
	2.	Select Configure.	set rule nat-to-wan-rule match-direction output term private-public-term then translated
	3.	From the Source list, select dynamic.	translation-type source dynamic
	4.	Click <b>OK</b> five times.	

### Table 97: Configuring a Stateful Firewall Filter and NAT (continued)

### Table 98: Applying a Stateful Firewall Filter and NAT to an Interface

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Services</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter
	2.	Next to Services, click <b>Configure</b> or <b>Edit</b> .	
Define wan-service-set and assign the stateful firewall filter rule to-wan-rule to the service set.	1.	Next to Service set, click Add new entry.	Define the service set and assign the rule:
	2.	In the Service set name box, type wan-service-set.	set service-set wan-service-set stateful-firewall-rules to-wan-rule
	3.	From the Stateful firewall rules choice list, select <b>Stateful firewall rules</b> .	
	4.	Next to Stateful firewall rules, click <b>Add new entry</b> .	
	5.	In the Rule name box, type to-wan-rule.	
	6.	Click OK.	

### Table 98: Applying a Stateful Firewall Filter and NAT to an Interface (continued)

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Assign the stateful firewall filter rule from-wan-rule to	1.	Next to Stateful firewall rules, click <b>Add new entry</b> .	Define the service set and assign the rule:
the service set.	2.	In the Rule name box, type from-wan-rule.	set service-set wan-service-set stateful-firewall-rules from-wan-rule
	3.	Click <b>OK</b> .	
Assign the NAT rule nat-to-wan-rule to the	1.	From the Nat rules choice list, select <b>Nat</b> rules.	Assign the rule to the service set:
service set.	2.	Next to Nat rules, click Add new entry.	set service-set wan-service-set nat-rules nat-to-wan-rule
	3.	In the Rule name box, type nat-to-wan-rule.	
	4.	Click <b>OK</b> .	
Define the service set type and virtual interface	1.	From the Service type choice list, select <b>Interface service</b> .	Define the service set type and the service interface:
sp-0/0/0 as the service interface for	2.	Next to Interface service, click <b>Configure</b> .	set service-set wan-service-set interface-service
wan-service-set.	3.	In the Service interface box, type <b>sp-0/0/0</b> .	service-interface sp-0/0/0
(See the interface naming conventions in the J-series Services Router Basic LAN and WAN Access Configuration Guide.)	4.	Click <b>OK</b> .	
Configure the <b>sp–0/0/0</b> service interface.	1.	On the main Configuration page next to Interfaces, click <b>Configure</b> or <b>Edit</b> .	From the [edit] hierarchy level, enter
	2.	Next to Interface, click Add new entry.	set interfaces sp-0/0/0 unit 0 family inet
	3.	In the Interface name box, type <b>sp-0/0/0</b> .	
	4.	Next to Unit, click Add new entry.	
	5.	In the Interface unit number box, type $0.$	
	6.	Next to Inet, select the check box.	
	7.	Click Configure.	
	8.	Click <b>OK</b> .	

Task	J-W	eb Configuration Editor	CLI Configuration Editor
From the Interfaces level of the configuration hierarchy, navigate to the <b>Inet</b> level of the T1 interface—the untrusted interface in this example—and apply wan-service-set to the input and output sides of the t1–0/0/0 interface.	1.	On the main Configuration page next to Interfaces, click <b>Edit</b> .	From the [edit] hierarchy level, apply the service set to the interface:
	2.	Under Interface name, click <b>t1-0/0/0</b> .	set interfaces t1-0/0/0 unit 0 family inet service
	3.	Under Interface unit number, click ${f 0}$ .	input service-set wan-service-set
	4.	Under Family, make sure the Inet check box is selected, and click <b>Configure</b> or <b>Edit</b> .	set interfaces t1-0/0/0 unit 0 family inet service output service-set wan-service-set
	5.	Next to Service, click Configure.	
(See the interface naming conventions in the <i>J-series</i> Services Router Basic LAN and WAN Access Configuration Guide.)	6.	Next to Input, click Configure.	
	7.	Next to Service set, click Add new entry.	
	8.	In the Service set name box, type wan-service-set.	
	9.	Click <b>OK</b> .	
	10.	Next to Output, click <b>Configure</b> .	
	11.	Next to Service set, click Add new entry.	
	12.	In the Service set name box, type wan-service-set.	
	13.	Click <b>OK</b> .	

#### Table 98: Applying a Stateful Firewall Filter and NAT to an Interface (continued)

### **Verifying Stateful Firewall Filter Configuration**

To verify a stateful firewall filter configuration, perform these tasks:

- Displaying Stateful Firewall Filter Configurations on page 217
- Verifying a Stateful Firewall Filter on page 219

### **Displaying Stateful Firewall Filter Configurations**

- **Purpose** Verify the configuration of the stateful firewall filter. You can analyze the flow of the firewall filter terms by displaying the entire configuration.
- Action From the J-Web interface, select

**Configuration > View and Edit > View Configuration Text**. Alternatively, from configuration mode in the CLI, enter the **show services** or **show firewall** command for stateful firewall filters.

The sample output in this section displays the stateful firewall filter and NAT configured in "Configuring a Stateful Firewall Filter with a Configuration Editor" on page 211.

[edit] user@host# **show services** stateful-firewall {

```
rule to-wan-rule {
 match-direction output;
 term app-term {
 from {
 application-sets junos-algs-outbound;
 }
 then {
 accept;
 }
 }
 term accept-all-term {
 then {
 accept;
 }
 }
 }
 rule from-wan-rule {
 match-direction input;
 term wan-src-addr-term {
 from {
 source-address {
 192.168.33.0/24;
 }
 }
 then {
 accept;
 }
 }
 term discard-all-term {
 then {
 discard;
 }
 }
 }
}
nat {
 pool public-pool {
 address-range low 10.148.2.1 high 10.148.2.32;
 port automatic;
 }
 rule nat-to-wan-rule {
 match-direction output;
 term private-public-term {
 then {
 translated {
 source-pool public-pool;
 translation-type source dynamic;
 }
 }
 }
 }
}
service-set wan-service-set {
 stateful-firewall-rules to-wan-rule;
 stateful-firewall-rules from-wan-rule;
 nat-rules nat-to-wan-rule;
```

```
interface-service {
 service-interface sp-0/0/0;
}
```

What It Means Verify that the output shows the intended configuration of the stateful firewall filter.

Verify that the terms are listed in the order in which you want the packets to be tested. You can move terms within a firewall filter by using the **insert** CLI command.

**Related Topics** For more information about the format of a configuration file, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

For information about the insert command, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

# **Verifying a Stateful Firewall Filter**

- **Purpose** Verify the firewall filter configured in "Configuring a Stateful Firewall Filter with a Configuration Editor" on page 211.
  - **Action** To verify that the actions of the firewall filter terms are taken, send packets to and from the untrusted network that match the terms. In addition, verify that actions are *not* taken for packets that do not match.
    - Send packets—associated with the junos-algs-outbound application set—from a host in the trusted network to a host in the untrusted network. Verify that packets received from the host in the untrusted network are responses only to the session originated by the host in the trusted network. To ensure that packets from the host are not accepted because of rule from-wan-rule, do not send packets to the host in the untrusted network with an IP address that matches 192.168.33.0/24.

For example, send a ping request from host **trusted-nw-trusted-host** to host **untrusted-nw-untrusted-host**, and verify that a ping response is returned. Ping requests and responses use ICMP, which belongs to the **junos-algs-outbound** application set.



**NOTE:** To view the configuration of **junos-algs-outbound**, enter the **show groups junos-defaults applications application-set junos-algs-outbound** configuration mode command.

Send packets from a host in the untrusted network to a host in the trusted network. Verify that the host in the trusted network receives packets only from the host in the untrusted network with an IP address that matches 192.168.33.0/24.

For example, send a ping request from host untrusted-nw-trusted-host with an IP address that matches 192.168.33.0/24 to host trusted-nw-trusted-host, and verify that a ping response is returned.

Verify that the ping response displays an IP address from the configured NAT pool.

user@trusted-nw-trusted-host> ping untrusted-nw-untrusted-host
PING untrusted-nw-untrusted-host.acme.net (172.69.13.5): 56 data bytes
64 bytes from 192.169.13.5: icmp_seq=0 ttl=22 time=8.238 ms
64 bytes from 192.169.13.5: icmp_seq=1 ttl=22 time=9.116 ms
64 bytes from 192.169.13.5: icmp_seq=2 ttl=22 time=10.875 ms
...
user@untrusted-nw-trusted-host> ping trusted-nw-trusted-host
PING trusted-nw-trusted-host-ge=000.acme.net (112.148.2.3): 56 data bytes
64 bytes from 10.148.2.3: icmp_seq=0 ttl=253 time=18.248 ms
64 bytes from 10.148.2.3: icmp_seq=1 ttl=253 time=10.906 ms
64 bytes from 10.148.2.3: icmp_seq=2 ttl=253 time=12.845 ms
...
What It Means
Verify the following information:

- A ping request from Host trusted-nw-trusted-host returns a ping response from Host untrusted-nw-untrusted-host.
- A ping request from Host untrusted-nw-trusted-host returns a ping response from Host trusted-nw-trusted-host. Verify that the ping response displays an IP address from the configured NAT pool of 10.148.2.1 through 10.148.2.32.
- **Related Topics** For information about using the J-Web interface to ping a host, see the *J*-series Services Router Administration Guide.

For more information about the ping command, see the *J*-series Services Router Administration Guide or the JUNOS System Basics and Services Command Reference.

# Chapter 13 Configuring Stateless Firewall Filters

A *stateless* firewall filter evaluates the contents of packets transiting the Services Router from a source to a destination, or packets originating from, or destined for, the Routing Engine. Stateless firewall filters applied to the Routing Engine interface protect the processes and resources owned by the Routing Engine. A stateless firewall filter evaluates every packet, including fragmented packets.

A stateless firewall filter, often called a firewall filter or access control list (ACL), statically evaluates packet contents. In contrast, a *stateful* firewall filter uses connection state information derived from past communications and other applications to make dynamic control decisions.

You can use either J-Web Quick Configuration or a configuration editor to configure stateless firewall filters.

This chapter contains the following topics. For more information about stateless firewall filters, see the *JUNOS Policy Framework Configuration Guide*. To configure a *stateful* firewall filter, see "Configuring Stateful Firewall Filters and NAT" on page 205.

If the router is operating in a Common Criteria environment, see the Secure Configuration Guide for Common Criteria and JUNOS-FIPS.

- Before You Begin on page 221
- Configuring a Stateless Firewall Filter with Quick Configuration on page 222
- Configuring a Stateless Firewall Filter with a Configuration Editor on page 238
- Verifying Stateless Firewall Filter Configuration on page 252

# **Before You Begin**

If you do not already have an understanding of firewall filters, read "Stateless Firewall Filters" on page 157.

Unlike a stateful firewall filter, you can configure a stateless firewall filter before configuring the interfaces on which they are applied.



**CAUTION:** If a packet does not match any terms in a firewall filter rule, the packet is discarded. Take care you do not configure a stateless firewall filter that prevents you from accessing the Services Router after you commit the configuration. For

example, if you configure a firewall filter that does not match HTTP or HTTPS packets, you cannot access the router with the J-Web interface.

### **Configuring a Stateless Firewall Filter with Quick Configuration**

The Firewall Filters Quick Configuration pages allow you to configure stateless firewall filters that examine packets traveling to or from a Services Router. You can create new filters or edit existing filters by adding terms to them. Each filter term is defined by a set of match conditions and an associated action. After you define the terms for a filter, you must associate the filter with one or more interfaces on the router.

This section contains the following topics:

- Configuring IPv4 and IPv6 Stateless Firewall Filters on page 222
- Assigning IPv4 and IPv6 Firewall Filters to Interfaces on page 236

# **Configuring IPv4 and IPv6 Stateless Firewall Filters**

Using the Firewall Filters Quick Configuration pages, you can create filters and terms and define match conditions and actions for each filter term. For a description of match conditions, see Table 71 on page 159, and for a description of actions, see Table 73 on page 162.

Figure 18 on page 223 shows the initial Firewall Filters Quick Configuration page that displays existing firewall filters and allows you to add and modify filters.

Figure 19 on page 224 shows the match conditions and actions Quick Configuration page for configuring match conditions and the resulting actions of filter terms.

ondor Configura	tion Diagnose	Manage	Events	Alarm			Lo	gged in as: reg <u>Configuratio</u>	ress Help n > <u>Quick Config</u>	uration > fie	rewall
d Edit 🕨 🕨	Quick Cor	nfigurat	ion								
	Firewall	Filters									
	Firewall F	ilters									
	IPv4 Filter	Summa	ry.					Showing fi	ter 1 to 1 of 1	total. (Paç	je 1
	Fi	lter Nam	e	RRARRER	ererererererererererererererererererer	k k k k k k k k k k k k k k k k k k k	RRRRRRRR			Sectores to	
	X	mfilter									
			Term		Destacal	Source	Source	Destination	Destination	Address	
			Name	Action	Protocol	Address	Port	Address	Port	Address	
			and the second sec	and the second sec	1 mar	10 10 10 0/04					1 A 44
		+X	MyTerm	-	· .	10.10.10.0/24			-	-	
		+X + X	MyLerm Lerm2	×	•	*	·	122.1.1.0/24	•	•	•
	✓ Acce	+X + X	Mylerm Term2	× ×	+ Reject Pa Log Pack	Legend cket : ×	+ Discard Syslog F	122.1.1.0/24 Packet 🖭	<ul> <li>◆</li> <li>↓ Evaluat</li> <li>↓ Count P</li> </ul>	* te Next Te vacket (*)	+
	✓ Acce → Rout PLP Set F Prior	+ X + X ept Packet ing Insta Packet Lo rity 7	Mylerm Term2	× ×	+ Reject Pa Log Pack Logical R	Legend cket : X et : : outer : 4	+ Discard Syslog P Load Ba Packet	122.1.1.0/24	<ul> <li>Evaluat</li> <li>Evaluat</li> <li>Count P</li> <li>Rate Lin</li> <li>Packet</li> </ul>	+ te Next Te Packet (*) mit (Police	• • • •
	Acce Rout Prio Any firewal is immediat of the term the configu Add Ne	A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A	Mylem Term2 t ? mce + iss tch condit firewall fi ilter	× ions that iter is sig	* Reject Pa Log Pack Logical R are colore the term s nificant. Pa	Legend cket ? X et ? outer ? X d red are consid tatement, and th ckets are tested Search	* Discard Syslog F Load Ba Packet ered negative seating against e	122.1.1.0/24 Packet ? Packet ? iance ? teed. If a packer rm in the filter sach term in the	Evaluat     Count P     Rate Lip     Packet t matches a n is evaluated. P	te Next Te Packet (*) mit (Police ? egated con Vote that th ch they are	e) ditione or
	Acce Rout PEP Set 5 Any firewall is immediated of the term the configur Add Ne	+X ept Packet ing Insta Packet Lo rity ? Il term ma tely consists within a ration. w IPv4 F	Mylerm Ierm2 at ? mace ? iss tch conditioned not firewall fil ilter	× ions that to match Iter is sig	* Reject Pa Log Pack Logical R are colore the term s nificant. Pe	Legend cket [* × et * outer [* ] d red are consid tatement, and tatement, and tatement are tested	+ Discard Syslog F Load Ba Packet ered nega he next te d against e	122.1.1.0/24 Packet [?] Packet [?	Evaluat     Count P     Count P     Packet     tmatches a n     s order in whice	te Next Te Packet (r) mit (Police r) egated con vote that th th they are	erm e) ditione e or
	Acce Rout PEP Set 5 Any firewal Any firewal Any firewal Add Ne Name	A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A     A	Mylerm Ierm2 t ? mce * iss tch condit iered not firewall fi ilter	× ions that iter is sig	* Reject Pa Log Pack Logical R are colore the term s milicant. Pe	Legend cket [* × et [* ] outer [* ] d red are consider d red d red	Poiscard     Syslog F Load Ba Packet ered negate next te f against c	122.1.1.0/24       Packet [?]       Packet [?]       Packet [?]       Bance       ?       red, if a packer       rem in the filter       sach term in the       n Name	Evaluat     Count P     Count P     Rate Lis     Packet t matches a n     sorder in whice	te Next Te vacket (*) mit (Police r egated con Vote that th th they are	e) etitio he or liste
	Acce Rout PEP Set I Prior Any firewal is immediat of the term the configu Add Ne Name Location	After     After	Hylerm Ierm2 it ? ince ? isss tch condition firewall fil ilter if Final 4 Filter r Final 4 Filter	× ions that to match Iter is signed P myF	* Reject Pa Log Pack Logical R are colore the term s nificant. Pe	Legend cket : × et : · outer : · d red are consider d red are consider	Poiscard Syslog F Load Ba Packet ered neago ne next te d against c	122.1.1.0/24       Packet ()       Packet () </td <td>Evaluat     Count P     Count P     Rate Lin     Packet     tratches a n     s order in whice     s</td> <td>te Next Te Packet (# mit (Police 7 eqated con Note that th th they are</td> <td>erm e) ditione or</td>	Evaluat     Count P     Count P     Rate Lin     Packet     tratches a n     s order in whice     s	te Next Te Packet (# mit (Police 7 eqated con Note that th th they are	erm e) ditione or
	Acce Rout PF Set I Any firewal of the term the configu Add Ne Name Location	After     After     After	t ? t ? t ? t ? t ? there ? the condition firewall filter ilter r Final 4 Filter r IPv4 pre IPv4	× ions that iter is sig	Reject Pa Log Pack Log Colore the terms inificant. Pe	Legend cket : X et : M outer : X d red are consider d red d r	Poiscard Syslog F Load Ba Packet Packet Pov4 Filter Pv4 Filter terms to 1	Packet 7 Packet 7 Pac	Evaluat     Count P     Count P     Rate Lin     Packet t matches a n     is evaluated. r     sorder in whice	te Next Te Packet @ mit (Police r folice r folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice folice	erm ditione liste
	Acce Rout Prior Set I Any firewal of the term the configu Add Ne Name Location	After     After     After     After	MyTerm Term2 It 9 mice 9 iss tch conditions firewall filter ilter Filter Filter Filter Filter Filter Filter Filter Filter	× ions that to match Iter is sig myF myF	Reject Pa Log Pack Log Pack Logical R are colore the terms s nificant. Pe	Legend cket ? × et ? ? d red are consid tatement, and th ckets are tested Search IF IF	t     t     t     t     t     t     t     t     t     t     t     t     t     t     t     t     t     t     t     t     t     t     t     t     t	122.1.1.0/24       Packet ?       Packet ? <tr< td=""><td>Evaluat     Count P     Rote Lin     packet t matches a n is evaluated. P     order in whice</td><td>te Next Te Packet @ mit (Police r Police dated con Vote that the solution of t</td><td>e)</td></tr<>	Evaluat     Count P     Rote Lin     packet t matches a n is evaluated. P     order in whice	te Next Te Packet @ mit (Police r Police dated con Vote that the solution of t	e)

# Figure 18: Initial Firewall Filters Quick Configuration Page

Monitor Configur	ration Diagnose Manage E	vents Logged in as: regress Help About Lo
		Configuration > Quick Configuration > Enerval1
w and Edit	Quick Configuration	
tory	Firewall Filters	
icue		
	Match Source	Natch Destination Match Source or Destination Match Interface Match Network Action
	Specify the criteria for thi checkbox above the criter this firewall term is match	s firewall term which must be matched. Some options below allow the inverse to be matched. Check the "Excep in that you wish to reverse. Click on the "Action" tab above to define what happens when the firewall criteria fo led.
	Match Source	
	⊖ Source Address	10.10.0/24
		Add Delete
	Source Prefix List	192
		Add Delete
	B Course Port	Event E
	Source Port	http ?
		Add Delete
	OK Cancel	

Figure 19: Match Conditions and Actions Quick Configuration Page

To configure a stateless firewall filter with Quick Configuration:

- In the J-Web interface, select Configuration > Quick Configuration > Firewall Filters.
- 2. Select one of the following options on the Firewall Filters Quick Configuration page:
  - To edit IPv4 firewall filters and terms, select Edit IPv4 Firewall Filters.



**NOTE:** If you have existing IPv4 firewall configurations in both edit firewall filter and edit firewall family inet filter hierarchies, merge the two to one location. The J-Web firewall filter Quick Configuration feature supports configuration in one location only.

- To edit IPv6 firewall filters and terms, select Edit IPv6 Firewall Filters.
- 3. Enter information into the Firewall Filters Quick Configuration pages, as described in Table 99 on page 225.
- 4. Click one of the following buttons on the Firewall Filters Quick Configuration main page:
  - To apply the configuration and stay in the current Firewall Filters Quick Configuration page, click **Apply**.
  - To apply the configuration and return to the previous Quick Configuration page, click **OK**.

- To cancel your entries and return to the previous Quick Configuration page, click **Cancel**.
- 5. Go on to one of the following procedures:
  - If the stateless firewall filter is not already assigned to an interface, see "Assigning IPv4 and IPv6 Firewall Filters to Interfaces" on page 236.
  - To display the configuration, see Displaying Stateless Firewall Filter Configurations on page 252.
  - To verify a stateless firewall filter, see "Verifying Stateless Firewall Filter Configuration" on page 252.

#### Table 99: Firewall Filters Quick Configuration Pages Summary

Field	Function	Your Action
IPv4 Filter Summary		
Action column	Displays up and down arrows and a X, allowing you to delete or change the order of a filter or term. The order of an item is important	To move an item upward, locate the item and click the up arrow from the same row.
	because it determines the order in which corresponding actions are carried out.	To move an item downward, locate the item and click the down arrow from the same row.
		To delete an item, locate the item and click the X from the same row.
Filter Name	Displays the name of the filter and when expanded, lists the terms attached to the filter.	To display the terms added to a filter, click the plus sign next to the filter name. This also displays the match conditions and actions set
	Displays the match conditions and actions that are set for each term.	for the term.
	Allows you to add more terms to a filter or modify filter terms.	To edit a filter, click the filter name. To edit a term, click the name of the term.
Search		
Filter Name	Searches for existing filters by filter name.	To find a specific filter, type the name of the filter in the Filter Name box.
		To list all filters with a common prefix or suffix, use the wildcard character (*) when typing the name of the filter. For example, te* lists all filters with a name starting with the characters <i>te</i> .
Term Name	Searches for existing terms by term name.	To find a specific term, type the name of the term in the Term Name box.
		To list all terms with a common prefix or suffix, use the wildcard character (*) when typing the name of the term. For example, ra* lists all terms with a name starting with the characters ra.

Field	Function	Your Action
Number of Items to Display	Specifies the number of filters or terms to display on one page.	To select the number of items to be displayed on one page, select a number from the list.
Add New IPv4 (or IPv6)	Filter	
Name	Specifies the name for a new filter.	To name a filter, type a string of meaningful characters or integers that allow you to uniquely identify the filter.
Location	<ul> <li>Positions the new filter in one of the following locations:</li> <li>After Final IPv4 Filter—At the end of all filters.</li> <li>After IPv4 Filter—After a specified filter.</li> <li>Before IPv4 Filter—Before a specified filter.</li> </ul>	<ul> <li>To position the new filter:</li> <li>At the end of all filters, select After Final IPv4 Filter.</li> <li>After a specific filter, select After IPv4 Filter then select a name from the filter name list.</li> <li>Before a specific filter, select Before IPv4 Filter then select a name from the filter name list.</li> </ul>
Add	Adds a new filter name. Opens the term summary page for this filter allowing you to add new terms to this filter.	To create a new filter and open the term summary page for this filter, click <b>Add</b> .
Add New IPv4 (or IPv6)	Term	
Name	Defines a term for a specific filter.	To name a term, type a string of meaningful characters or integers that allow you to uniquely identify the term.
Location	<ul> <li>Positions the new term in one of the following locations:</li> <li>After Final IPv4 Term—At the end of all terms.</li> <li>After IPv4 Term—After a specified term.</li> <li>Before IPv4 Term—Before a specified term.</li> <li>Adds a term name for the specific filter.</li> <li>Opens the Filter Term page allowing you to define the match conditions and the action for the specific filter.</li> </ul>	<ul> <li>To position the new term:</li> <li>At the end of all terms, select After Final IPv4 Term.</li> <li>After a specific term, select After IPv4 Term then select a name from the term name list.</li> <li>Before a specific term, select Before IPv4 Term then select a name from the term name list.</li> <li>To add a term name and open the Filter Term page, click Add.</li> </ul>
Match Source		

### Table 99: Firewall Filters Quick Configuration Pages Summary (continued)

Field	Function	Your Action	
Source Address	Specifies IP source addresses to be included in, or excluded from, the match condition.	To specify an IP source address, type an IP address and prefix length.	
	Allows you to remove source IP addresses from the match condition. If you have more than 25 addresses, this field displays a link that allows you to easily scroll through pages, change the order of addresses, and also search for them.	<ul> <li>To include the address in the match condition, click Add.</li> <li>To exclude the address from the match condition, select Except then click Add.</li> <li>To remove an IP source address from the match condition, select it and click Delete.</li> </ul>	
Source Prefix List	Specifies source prefix lists that you have already defined, to be included in the match condition.	To include a predefined source prefix list in the match condition, type the prefix list name and click <b>Add</b> .	
	Allows you to remove a prefix list from the match condition.	To remove a prefix list from the match condition, select it and click <b>Delete</b> .	
	For information about defining prefix lists, see the <i>JUNOS Policy Framework Configuration Guide</i> .		
Source Port	Specifies the source port type to be included in, or excluded from, the match condition. Allows you to remove a source port type from	To specify a known source port type, select the port from the port name list. To specify source port types that do not exist in the port name list, type the port name, number, or range.	
	the match condition.	To include the port in the match condition, click <b>Add</b> .	
	the protocol type being used on the port. Make sure to specify the protocol type (TCP or UDP) match condition in the same term.	• To exclude the port from the match condition, select <b>Except</b> then click <b>Add</b> .	
		To remove a port type from the match condition, select it and click <b>Delete</b> .	
Match Destination			
Destination Address	Specifies destination addresses to be included in, or excluded from, the match condition.	To specify a destination IP address, type an IP address and prefix length.	
	Allows you to remove a destination IP address from the match condition.	To include the address in the match condition, click <b>Add</b> .	
	If you have more than 25 addresses, this field displays a link that allows you to easily scroll	• To exclude the address from the match condition, select <b>Except</b> then click <b>Add</b> .	
	through pages, change the order of addresses, and also search for them.	To remove an IP address from the match condition, select it and click <b>Delete</b> .	

# Table 99: Firewall Filters Quick Configuration Pages Summary (continued)

### Table 99: Firewall Filters Quick Configuration Pages Summary (continued)

Field	Function	Your Action		
Destination Prefix List	Specifies destination prefix lists that you have already defined, to be included in the match condition.	To include a predefined destination prefix list, type the prefix list name and click <b>Add</b> .		
	Allows you to remove a prefix list from the match condition.	To remove a prefix list from the match condition, select it and click <b>Delete</b> .		
	For information about defining prefix lists, see the <i>JUNOS Policy Framework Configuration Guide</i> .			
Destination Port	Specifies destination port types to be included in, or excluded from, the match condition.	To specify a known destination port type, select the port from the port name list. To specify source port types that do not exist in the port		
	Allows you to remove a destination port type from the match condition.	name list, type the port name, number, or range.		
	<b>NOTE:</b> This match condition does not check the protocol type being used on the port. Make	To include the port in the match condition, click <b>Add</b> .		
	sure to specify the protocol type (TCP or UDP) match condition in the same term.	■ To exclude the port from the match condition, select <b>Except</b> then click <b>Add</b> .		
		To remove a destination port type from the match condition, select it and click <b>Delete</b> .		
Match Source or Destinat	ion			
Address	Specifies IP addresses to be included in, or excluded from, the match condition for a	To specify a source or destination IP address, type the IP address and prefix length.		
	Source or destination.	<ul> <li>To include the address in the match condition, click Add.</li> </ul>		
	match condition.	■ To exclude the address from the match condition, select <b>Except</b> then click <b>Add</b> .		
	If you have more than 25 addresses, this field displays a link that allows you to easily scroll through pages, change the order of addresses and also search for them.	To remove an IP address from the match condition, select it and click <b>Delete</b> .		
	<b>NOTE:</b> This address match condition cannot be specified in conjunction with the source address or destination address match conditions in the same term.			
Field	Function	Your Action		
-----------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--	--
Prefix List	Specifies prefix lists that you have already defined, to be included in the match condition for a source or destination.	To include a predefined prefix list in the match condition, type the prefix list name and click <b>Add</b> .		
	Allows you to remove a prefix list from the match condition.	To remove a prefix list from the match condition, select it and click <b>Delete</b> .		
	For information about defining prefix lists, see the <i>JUNOS Policy Framework Configuration Guide</i> .			
	<b>NOTE:</b> This prefix list match condition cannot be specified in conjunction with the source prefix list or destination prefix list match conditions in the same term.			
Port	Specifies a port type to be included in, or excluded from, a match condition for a source or destination.	To specify a known port type in the match condition, select the port from the port name list. To specify port types not included in the port name list, type the port name, number, or		
	Allows you to remove a port from the match condition.	<ul> <li>To include the port in the match condition,</li> </ul>		
	<b>NOTE:</b> This match condition does not check the protocol type being used on the port. Make sure to specify the protocol type (TCP or UDP) match condition in the same term.	<ul> <li>Click Add.</li> <li>To exclude the port from the match condition, select Except then click Add.</li> </ul>		
	Also, this port match condition cannot be specified in conjunction with the source port or destination port match conditions in the same term.	To remove a port from the match condition, select it and click <b>Delete</b> .		
Match Interface				
Interface (See the interface naming	Specifies interfaces to be included in a match condition.	To include an interface in a match condition, either select a name from the interface name list or type the interface name and click <b>Add</b> .		
conventions in the J-series Services Router Basic LAN and WAN Access Configuration Guide.)	Allows you to remove an interface from the match condition.	either select a name from the interface name list or type the interface name and click <b>Add</b> . To remove an interface from the match condition, select it and click <b>Delete</b> .		
Interface Set	Specifies interface sets that you have already defined, to be included in a match condition.	To include a predefined interface set in a match condition, type the interface set name and click <b>Add</b> .		
	Allows you to remove an interface set from the match condition.	To remove an interface set from the match condition, select it and click <b>Delete</b> .		
	For information about defining interface sets, see the <i>JUNOS Policy Framework Configuration Guide</i> .			

Field	Function	Your Action		
Interface Group	Specifies interface groups, that you have already defined, to be included in, or excluded from a match condition	To specify a predefined interface group, type the name of the group.		
	Allows you to remove an interface group from	■ To include the group in the match condition, click <b>Add</b> .		
	the match condition.	■ To exclude the group from the match condition, select <b>Except</b> then click <b>Add</b> .		
	For information about defining interface groups, see the <i>JUNOS Policy Framework Configuration Guide</i> .	To remove an interface group from the match condition, select it and click <b>Delete</b> .		
Match Packet and Networ	k			
First Fragment (IPv4 only)	Matches the first fragment of a fragmented packet.	To match the first fragment, select the check box.		
Is Fragment (IPv4 only)	Matches trailing fragments (all but the first fragment) of a fragmented packet.	To match trailing fragments, select the check box.		
Fragment Flags (IPv4 only)	Specifies fragmentation flags to be included in the match condition.	To specify fragmentation flags, type a text or numeric string defining the flag—for example, more-fragments or 0x2000.		
TCP Established	Matches all TCP packets other than the first packet of a connection.	To match all TCP packets except the first of a connection, select the check box.		
	<b>NOTE:</b> This match condition does not verify that the TCP protocol is used on the port. Make sure to specify the TCP protocol as a match condition in the same term.			
TCP Initial	Matches the first TCP packet of a connection.	To match the first TCP packet of a connection, select the check box.		
	<b>NOTE:</b> This match condition does not verify that the TCP protocol is used on the port. Make sure to specify the TCP protocol as a match condition in the same term.			
TCP Flags	Specifies TCP flags to be included in the match condition.	To specify a TCP flag, type a text or numeric string defining the flag—for example, <b>syn</b> or <b>0x02</b> .		
<b>NOTE:</b> This match condition does not verify that the TCP protocol is used on the port. Make sure to specify the TCP protocol as a match condition in the same term.				

Field	Function	Your Action			
Protocol (IPv4 only)	Specifies IPv4 protocol types to be included in, or excluded from, the match condition.	To specify an IPv4 protocol type, select a protocol name from the list or type a protocol name or number—for example, <b>ospf</b> or <b>89</b> .			
	Allows you to remove an IPv4 protocol type from the match condition.	<ul> <li>To include the protocol in the match condition, click Add.</li> </ul>			
		■ To exclude the protocol from the match condition, select <b>Except</b> then click <b>Add</b> .			
		To remove an IPv4 protocol type from the match condition, select it and click <b>Delete</b> .			
Next Header (IPv6 only)	Specifies IPv6 protocol types to be included in, or excluded from, the match condition.	To specify an IPv6 protocol type, select a protocol name from the list or type the protocol name or number—for example, <b>igmp</b> or <b>2</b> .			
	Allows you to remove an IPv6 protocol type from the match condition.	<ul> <li>To include the protocol in the match condition, click Add.</li> </ul>			
		• To exclude the protocol from the match condition, select <b>Except</b> then click <b>Add</b> .			
		To remove an IPv6 protocol type from the match condition, select it and click <b>Delete</b> .			
ІСМР Туре	Specifies ICMP packet types to be included in, or excluded from, the match condition.	To specify an ICMP packet type, select a packet type from the list or type a packet type name or number—for example, time-exceeded or 11.			
	Allows you to remove an ICMP packet type from the match condition.	<ul> <li>To include the packet type in the match condition, click Add.</li> </ul>			
	<b>NOTE:</b> This protocol does not verify that ICMP is used on the port. Make sure to specify an ICMP type match condition in the same term.	• To exclude the packet type from the match condition, select <b>Except</b> then click <b>Add</b> .			
		To remove an ICMP packet type from the match condtition, select it and click <b>Delete</b> .			
ICMP Code	Specifies the ICMP code to be included in, or excluded from, the match condition.	To specify an ICMP code, select a packet code from the list or type the packet code as text or a number—for example, <b>ip-header-bad</b> or <b>0</b> .			
	Allows you to remove an ICMP code from the match condition.	<ul> <li>To include the ICMP code in the match condition, click Add.</li> </ul>			
	<b>NOTE:</b> The ICMP code is dependent on the ICMP type. Make sure to specify an ICMP type match condition in the same term.	• To exclude the ICMP code from the match condition, select <b>Except</b> then click <b>Add</b> .			
		To remove an ICMP code from the match condition, select it and click <b>Delete</b> .			

Field	Function	Your Action	
Traffic Class (IPv6 only)	Specifies Differentiated Services code points (DSCPs) to be included in, or excluded from, the match condition.	To specify a DSCP, select it from the list or type the DSCP value as a keyword, decimal, or binary string—for example, <b>af11</b> or <b>10</b> .	
	Allows you to remove a DSCP value from the match condition.	<ul> <li>To include the DSCP in the match condition, click Add.</li> <li>To exclude the DSCP from the match</li> </ul>	
	For information about DSCPs, see the <i>J-series</i> Services Router Basic LAN and WAN Access Configuration Guide.	condition, select <b>Except</b> then click <b>Add</b> . To remove a DSCP from the match condition,	
		select it and click <b>Delete</b> .	
Fragment Offset (IPv4 only)	Specifies the fragment offset value to be included in, or excluded from, the match	To specify a fragment offset value, type the fragment offset number or range.	
	condition. The fragment offset value specifies the location of the fragment in the packet. For	■ To include the offset in the match condition, click <b>Add</b> .	
	fragment.	■ To exclude the offset from the match condition, select <b>Except</b> then click <b>Add</b> .	
	Allows you to remove a fragment offset value		
	from the match condition.	To remove a fragment offset value from the match condition, select it and click <b>Delete</b> .	
Precedence (IPv4 only)	Specifies IP precedences to be included in, or excluded from, the match condition.	To specify an IP precedence, select it from the list or type the precedence as a keyword, decimal integer between 0 and 7, or binary	
	Allows you to remove an IP precedence entry	string.	
	nom me match condition.	<ul> <li>To include the precedence in the match condition, click Add.</li> </ul>	
		• To exclude the precedence from the match condition, select <b>Except</b> then click <b>Add</b> .	
		To remove an IP precedence from the match condition, select it and click <b>Delete</b> .	
DSCP (IPv4 only)	Specifies Differentiated Services code points (DSCPs) to be included in, or excluded from, the match condition	To specify a DSCP, select it from the list or type the DSCP value as a keyword, decimal, or binary string—for example, <b>af11</b> or <b>10</b> .	
	Allows you to remove a DSCP entry from the	<ul> <li>To include the DSCP in the match condition click Add</li> </ul>	
	match condition.	To exclude the DSCP from the match condition, select <b>Except</b> then click <b>Add</b> .	
		To remove a DSCP, select it and click <b>Delete</b> .	

Field	Function	Your Action			
TTL (IPv4 only)	Specifies the IPv4 time-to-live (TTL) value to be included in, or excluded from, the match condition. Allows you to remove an IPv4 TTL value from the match condition.	<ul> <li>To specify an IPv4 TTL value, type a number between 1 and 255.</li> <li>To include the TTL in the match condition, click Add.</li> <li>To exclude the TTL from the match condition, select Except then click Add.</li> <li>To remove an IPv4 TTL type from the match condition, select it and click Delete.</li> </ul>			
Packet Length	Specifies the length of received packets, in bytes, to be included in, or excluded from, the match condition. Allows you to remove a packet length value from the match condition.	<ul> <li>To specify a packet length, type a value or range.</li> <li>To include the packet length in the match condition, click Add.</li> <li>To exclude the packet length from the match condition, select Except then click Add.</li> <li>To remove a packet length value from the match condition, select it and click Delete.</li> </ul>			
Forwarding Class	<ul> <li>Specifies forwarding classes to be included in, or excluded from, the match condition.</li> <li>Allows you to a remove forwarding class entry from the match condition.</li> <li>For information about forwarding classes, see the <i>J</i>-series Services Router Basic LAN and WAN Access Configuration Guide.</li> </ul>	<ul> <li>To specify a forwarding class, select it from the list or type it.</li> <li>To include the forwarding class in the match condition, click Add.</li> <li>To exclude the forwarding class from the match condition, select Except then click Add.</li> <li>To remove a forwarding class from the match condition, select it and click Delete.</li> </ul>			
IP Options (IPv4 only)	Specifies IP options to be included in, or excluded from, the match condition. Allows you to remove an IP option from the match condition.	<ul> <li>To specify an IP option, select it from the list or type a text or numeric string identifying the option.</li> <li>To include the IP option in the match condition, click Add.</li> <li>To exclude the IP option from the match condition, select Except then click Add.</li> <li>To remove an IP option from the match condition, select it and click Delete.</li> </ul>			

Field	Function	Your Action		
IPSec ESP SPI (IPv4 only)	Specifies IPSec Encapsulating Security Payload (ESP) security parameter index (SPI) values to be included in or excluded from the match	To specify an ESP SPI value, type a binary, hexadecimal, or decimal SPI value or range.		
	condition.	■ To include the value in the match condition, click <b>Add</b> .		
	Allows you to remove an ESP SPI value from the match condition.	• To exclude the value from the match condition, select <b>Except</b> then click <b>Add</b> .		
		To remove an ESP SPI value from the match condition, select it and click <b>Delete</b> .		
Action				
Nothing	No action is performed. By default, a packet is accepted if it meets the match conditions of the term, and packets that do not match any conditions in the firewall filter are dropped.	To specify no action (or the default action), select <b>Nothing</b> .		
Accept	Accepts a packet that meets the match conditions of the term.	To accept the packet, select <b>Accept</b> .		
Discard	Discards a packet that meets the match conditions of the term.	To discard a packet, select <b>Discard</b> .		
		To name a discard collector, type a filename in		
	Names a discard collector for packets (IPV4 only).	the Accounting box (IPV4 only).		
Reject	Rejects a packet that meets the match conditions of the term and returns a rejection	To reject a packet, select <b>Reject</b> .		
	message.	To specify a message type, select the message from the Reason list.		
	Allows you to specify a message type that denotes the reason the packet was rejected.			
	<b>NOTE:</b> To log and sample rejected packets, specify Log and Sample action modifiers in conjunction with this action.			
Next Term	Evaluates a packet with the next term in the filter if the packet meets the match conditions in this term.	To continue to the next term, select <b>Next Term</b> .		
	This action makes sure that the next term is used for evaluation even when the packet matches the conditions of a term.			
	When this action is not specified, the filter stops evaluating the packet after it matches the conditions of a term, and takes the associated action.			
Routing Instance	Accepts a packet that meets the match conditions, and forwards it to the specified routing instance.	To specify a routing instance, select <b>Routing</b> <b>Instance</b> and type the routing instance name in the box next to Routing Instance.		

Field	Function	Your Action		
Load Balance	Specifies a load-balance group that you have already defined, to be used by packets that meet the match conditions.	To specify a load-balance group, select <b>Load</b> <b>Balance</b> and type the group name in the box next to it.		
	A load-balance group contains interfaces that use the same next-hop group to balance the traffic load.			
	For information about configuring a load-balance group, see the <i>JUNOS Policy</i> Framework Configuration Guide			
Action Modifiers				
Forwarding Class	Classifies the packet as a specific forwarding class.	To specify a forwarding class, select it from the list.		
	For information about forwarding classes, see the J-series Services Router Basic LAN and WAN Access Configuration Guide.			
Count	Counts the packets passing this term.	To count packets passing this term, select <b>Count</b> .		
	Allows you to name a counter, which is specific to this filter. This means that every time a packet transits any interface that uses this filter, it increments the specified counter.	To specify a counter name, type a 24-character string containing letters, numbers, or hyphens.		
Virtual Channel (IPv4 only)	Specifies the virtual channel to be set on a particular logical interface.	To specify the virtual channel, type a string identifying the virtual channel.		
Log	Logs the packet header information in the Routing Engine.	To log packet header information, select <b>Log</b> .		
Syslog	Records packet information in the system log.	To record information in the system log, select <b>Syslog</b> .		
Sample (IPv4 only)	Samples traffic on the interface.	To sample traffic on an interface, select <b>Sample</b> .		
	<b>NOTE:</b> You must enable traffic sampling for this action to work. For more information about traffic sampling and forwarding, see the <i>JUNOS Policy Framework Configuration Guide</i> .			
Loss Priority	Sets the loss priority of the packet. This is the priority of dropping a packet before it is sent, and it affects the scheduling priority of the packet.	To set the loss priority of the packet, select a loss priority from the list.		
	For more information, see the JUNOS Class of Service Configuration Guide.			

## Assigning IPv4 and IPv6 Firewall Filters to Interfaces

For a firewall filter to work, you must assign it to an interface. Use the Firewall Filters Quick Configuration pages to assign IPv4 and IPv6 filters to interfaces. Using these pages you can select a firewall filter to evaluate packets that are received or transmitted on a specific interface.

When assigning firewall filters to interfaces, remember that you can assign only one input and one output firewall filter to each interface. However, you can assign the same filter to multiple interfaces.

Figure 20 on page 236 shows the Firewall Filters Quick Configuration page that displays the Services Router interfaces available for filter assignment and the status of existing filter assignments.

Monitor	Configuration	Diagnose	Manage	Events	Alarms	Logged in as: regress Configuration >	Help Quick Configur.	About ation > Fire	Logou wall Filt
lick Configural	lon -	Quick Conf	igurati	on					
story		Firewall Fi	ilters						
secue								111111111111111111111111111111111111111	
		Logical Interface Name	Lii	nk State	Inj	out Firewall Filters	Output Fi	rewall F	llters
		fe-0/0/0.0	ju j						
		<u>sp-0/0/0.0</u>	s U						
		sp-0/0/0.16	<u>383</u>	•					
		fe-0/0/1.0	U	•					
		dc-6/0.0.32	7 <u>67</u>	•					
		bc-6/0/0:1.0	. D	own					
		bc-6/0/0:2.0		own					
		<u>dl0.0</u>	U	0					
		<u>lo0.0</u>	U						
		ок с	ancel	VlaaA					

#### Figure 20: Firewall Filters Interface Assignment Quick Configuration Page

To assign IPv4 and IPv6 firewall filters to interfaces with Quick Configuration:

- In the J-Web interface, select Configuration > Firewall Filters > Assign Firewall Filters to Interfaces.
- 2. Enter information into the Firewall Filters Quick Configuration pages, as described in Table 100 on page 237.
- 3. Click one of the following buttons on the Firewall Filters Quick Configuration main page:
  - To apply the configuration and stay in current the Firewall Filters Quick Configuration page, click **Apply**.

- To apply the configuration and return to the previous Quick Configuration page, click **OK**.
- To cancel your entries and return to the previous Quick Configuration page, click **Cancel**.
- 4. Go on to one of the following procedures:
  - To display the configuration, see Displaying Stateless Firewall Filter Configurations on page 252.
  - To verify a stateless firewall filter, see "Verifying Stateless Firewall Filter Configuration" on page 252.

#### Table 100: Assigning Firewall Filters Quick Configuration Pages Summary

Field	Function	Your Action		
Firewall Filters				
Logical Interface Name	Displays the logical interfaces on a router.	To apply firewall filters to an interface, click the interface name		
(See the interface naming conventions in the <i>J</i> -series Services Router Basic LAN and WAN Access	Allows you to apply IPv4 and IPv6 firewall filters to packets received on the interface and packets transmitted from the interface.	<ul> <li>To apply an input firewall filter, follow instructions in the input firewall filters section.</li> </ul>		
Configuration Guide.)		<ul> <li>To apply an output firewall filter, follow instructions in the ouput firewall filters section.</li> </ul>		
Link State	Displays the status of the logical interface.	None.		
Input Firewall Filters	Displays the input firewall filter applied on an interface. This filter evaluates all packets received on the interface.	None.		
Output Firewall Filters	Displays the output firewall filter applied on an interface. This filter evaluates all packets transmitted from the interface.	None.		
Input Firewall Filters				
IPv4 Input Filter	Allows you to apply an input firewall filter to an interface. This filter evaluates all packets	To apply an input firewall filter to an interface, select the name of the firewall filter from the		
IPv6 Input Filter	received on the interface.	list.		
<b>Output Firewall Filters</b>				
IPv4 Output Filter	Allows you to apply an output firewall filter to an interface. This filter evaluates all packets	To apply an output firewall filter to an interface, select the name of the firewall filter from the		
IPv6 Output Filter	transmitted on the interface.	list.		

## **Configuring a Stateless Firewall Filter with a Configuration Editor**

The section contains the following topics. For stateless firewall match conditions, actions, and modifiers, see "Stateless Firewall Filter Match Conditions" on page 159 and "Stateless Firewall Filter Actions and Action Modifiers" on page 162.

- Stateless Firewall Filter Strategies on page 238
- Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources on page 238
- Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods on page 241
- Configuring a Routing Engine Firewall Filter to Handle Fragments on page 246
- Applying a Stateless Firewall Filter to an Interface on page 251

## Stateless Firewall Filter Strategies

For best results, use the following sections to plan the purpose and contents of a stateless firewall filter before starting configuration.

#### Strategy for a Typical Stateless Firewall Filter

A primary goal of a typical stateless firewall filter is to protect the Routing Engine processes and resources from malicious or untrusted packets. You can configure a firewall filter like the sample filter **protect-RE** to restrict traffic destined for the Routing Engine based on its source, protocol, and application. In addition, you can limit the traffic rate of packets destined for the Routing Engine to protect against flood, or *denial-of-service* (DoS), attacks.

For details, see "Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources" on page 238 and "Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods" on page 241.

#### **Strategy for Handling Packet Fragments**

You can configure a stateless firewall filter like the sample filter **fragment-filter** to address special circumstances associated with fragmented packets destined for the Routing Engine. Because the Services Router evaluates every packet against a firewall filter (including fragments), you must configure the filter to accommodate fragments that do not contain packet header information. Otherwise, the filter discards all but the first fragment of a fragmented packet.

For details, see "Configuring a Routing Engine Firewall Filter to Handle Fragments" on page 246.

#### Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources

The following example shows how to create a stateless firewall filter, **protect-RE**, that discards all traffic destined for the Routing Engine, except SSH and BGP protocol

packets from specified trusted sources. Table 101 on page 239 lists the terms that are configured in this sample filter.

#### Table 101: Sample Stateless Firewall Filter protect-RE Terms to Allow Packets from Trusted Sources

Term	Purpose
ssh-term	Accepts TCP packets with a source address of <b>192.168.122.0/24</b> and a destination port that specifies SSH.
bgp-term	Accepts TCP packets with a source address of <b>10.2.1.0/24</b> and a destination port that specifies the BGP protocol.
discard-rest-term	For all packets that are not accepted by <b>ssh-term</b> or <b>bgp-term</b> , creates a firewall filter log and system logging records, then discards all packets. To view the log, enter the <b>show firewall log</b> operational mode command. (For more information, see Displaying Stateless Firewall Filter Logs on page 255.)

By applying firewall filter **protect-RE** to the Routing Engine, you specify which protocols and services, or applications, are allowed to reach the Routing Engine, and you ensure the packets are from a trusted source. This protects processes running on the Routing Engine from an external attack.

To use the configuration editor to configure the stateless firewall filter:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 102 on page 239.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to one of the following procedures:
  - To display the configuration, see Displaying Stateless Firewall Filter Configurations on page 252.
  - To apply the firewall filter to the Routing Engine, see "Applying a Stateless Firewall Filter to an Interface" on page 251.
  - To verify the firewall filter, see Verifying a Services, Protocols, and Trusted Sources Firewall Filter on page 257.

#### Table 102: Configuring a Protocols and Services Firewall Filter for the Routing Engine

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Firewall</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit firewall
	2.	Next to Firewall, click <b>Configure</b> or <b>Edit</b> .	

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Define protect-RE and		Next to Filter, click <b>Add new entry</b> .	Set the term name and define the match
ssh-term, and define the protocol, destination port.	2.	In the Filter name box, type protect-RE.	conditions:
and source address match	3.	Next to Term, click Add New Entry.	set family inet filter protect-RE term ssh-term from
conditions.	4.	In the Rule name box, type <b>ssh-term</b> .	source-address 192.168.122.0/24
	5.	Next to From, click Configure.	
	6.	In the Protocol choice list, select <b>Protocol</b> .	
	7.	Next to Protocol, click Add new entry.	
	8.	In the Value keyword list, select <b>tcp</b> .	
	9.	Click <b>OK</b> .	
	10.	In the Destination port choice list, select <b>Destination port</b> .	
	11.	Next to Destination port, click <b>Add new entry</b> .	
	12.	In the Value keyword list, select <b>ssh</b> .	
	13.	Click <b>OK</b> .	
	14.	Next to Source address, click <b>Add new</b> entry.	
	15.	In the Address box, type <b>192.168.122.0/24</b> .	
	16.	Click <b>OK</b> twice.	
Define the actions for ssh-term.	1.	On the Term <b>ssh-term</b> page, next to Then, click <b>Configure</b> .	Set the actions:
	2.	In the Designation list, select Accept.	set family inet filter protect-RE term ssh-term then accept
	3.	Click <b>OK</b> twice.	

## Table 102: Configuring a Protocols and Services Firewall Filter for the Routing Engine (continued)

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Define <b>bgp-term</b> , and define the protocol, destination	1.	On the Filter <b>protect-RE</b> page, next to Term, click <b>Add New Entry</b> .	Set the term name and define the match conditions:
port, and source address match conditions.	2.	In the Rule name box, type bgp-term.	set family inet filter protect-RE term bgp-term from
	3.	Next to From, click Configure.	protocol tcp destination-port bgp
	4.	In the Protocol choice list, select <b>Protocol</b> .	source-address 10.2.1.0/24
	5.	Next to Protocol, click Add new entry.	
	6.	In the Value keyword list, select <b>tcp</b> .	
	7.	Click <b>OK</b> .	
	8.	In the Destination port choice list, select <b>Destination port</b> .	
	9.	Next to Destination port, click <b>Add new</b> entry.	
	10.	In the Value keyword list, select <b>bgp</b> .	
	11.	Click <b>OK</b> .	
	12.	Next to Source address, click <b>Add new</b> entry.	
	13.	In the Address box, type $10.2.1.0/24$ .	
	14.	Click <b>OK</b> twice.	
Define the action for <b>bgp-term</b> .	1.	On the Term <b>bgp-term</b> page, next to Then, click <b>Configure</b> .	Set the action:
	2.	In the Designation list, select Accept.	set family inet filter protect-RE term bgp-term then accept
	3.	Click <b>OK</b> twice.	
Define discard-rest-term and its action.	1.	On the Filter <b>protect-RE</b> page, next to Term, click <b>Add New Entry</b> .	Set the term name and define its actions:
	2.	In the Rule name box, type discard-rest-term.	set family inet filter protect-RE term discard-rest-term then log syslog discard
	3.	Next to Then, click Configure.	
	4.	Next to Log, select the check box.	
	5.	Next to Syslog, select the check box.	
	6.	In the Designation list, select <b>Discard</b> .	
	7.	Click <b>OK</b> four times.	

#### Table 102: Configuring a Protocols and Services Firewall Filter for the Routing Engine (continued)

## **Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods**

The procedure in this section creates a sample stateless firewall filter, **protect-RE**, that limits certain TCP and ICMP traffic destined for the Routing Engine. A router without

this kind of protection is vulnerable to TCP and ICMP flood attacks—also called denial-of-service (DoS) attacks. For example:

- A TCP flood attack of SYN packets initiating connection requests can so overwhelm the Services Router that it can no longer process legitimate connection requests, resulting in denial of service.
- An ICMP flood can overload the Services Router with so many echo requests (ping requests) that it expends all its resources responding and can no longer process valid network traffic, also resulting in denial of service.

Applying a firewall filter like **protect-RE** to the Routing Engine protects against these types of attacks.

For each term in the sample filter, you first create a policer and then incorporate it into the action of the term. For more information about firewall filter policers, see the *JUNOS Policy Framework Configuration Guide*.

If you want to include the terms created in this procedure in the **protect-RE** firewall filter configured in the previous section (see "Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources" on page 238), perform the configuration tasks in this section first, then configure the terms as described in the previous section. This approach ensures that the rate-limiting terms are included as the first two terms in the firewall filter.



**NOTE:** You can move terms within a firewall filter by using the **insert** CLI command. For more information, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

Table 103 on page 242 lists the terms that are configured in this sample filter.

Term	Purpose	Policer	
tcp-connection-term	Polices the following types of TCP packets with a source address of <b>192.168.122.0/24</b> or <b>10.2.1.0/24</b> :	<b>tcp-connection-policer</b> —Limits the traffic rate and burst size of these TCP packets to 500,000 bps and 15,000 bytes. Packets that exceed the traffic rate	
	<ul> <li>Connection request packets (SYN and ACK flag bits equal 1 and 0)</li> </ul>	are discarded.	
	<ul> <li>Connection release packets (FIN flag bit equals 1)</li> </ul>		
	<ul> <li>Connection reset packets (RST flag bit equals 1)</li> </ul>		
icmp-term	Polices the following types of ICMP packets. All are counted in counter icmp-counter.	icmp-policer—Limits the traffic rate and burst size of these ICMP packets to 1,000,000 bps and	
	<ul> <li>Echo request packets</li> </ul>	15,000 bytes. Packets that exceed the traffic rate are discarded.	
	<ul> <li>Echo response packets</li> </ul>		
	<ul> <li>Unreachable packets</li> </ul>		
	<ul> <li>Time-exceeded packets</li> </ul>		

Table 103: Sample Stateless Firewall Filter protect-RE Terms to Protect Against Floods

To use the configuration editor to configure the policers and the stateless firewall filter:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. To configure the firewall filter policers, perform the configuration tasks described in Table 104 on page 243.
- 3. To configure the prefix lists and the firewall filter, perform the configuration tasks described in Table 105 on page 244.
- 4. If you are finished configuring the router, commit the configuration.
- 5. Go on to one of the following procedures:
  - To display the configuration, see Displaying Stateless Firewall Filter Configurations on page 252.
  - To apply the firewall filter to the Routing Engine, see "Applying a Stateless Firewall Filter to an Interface" on page 251.
  - To verify the firewall filter, see Verifying a TCP and ICMP Flood Firewall Filter on page 258.

#### **Table 104: Configuring Policers for TCP and ICMP**

Task	J-M	eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Firewall</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit firewall
	2.	Next to Firewall, click <b>Configure</b> or <b>Edit</b> .	
Define	1.	Next to Policer, click Add new entry.	Set the policer name and its rate limits:
set its rate limits.	2.	In the Policer name box, type tcp-connection-policer.	set policer tcp-connection-policer filter-specific if-exceeding burst-size-limit 15k
The burst size limit can be from 1,500 bytes through	3.	Next to Filter specific, select the check box.	bandwidth-limit 500k
The bandwidth limit can be	4.	Next to If Exceeding, select the check box and click <b>Configure</b> .	
from 32,000 bps through 32,000,000 bps	5.	In the Burst size limit box, type <b>15</b> k.	
Use the following	6.	In the Bandwidth list, select <b>Bandwidth limit</b> .	
abbreviations when specifying these limits:	7.	In the Bandwidth limit box, type <b>500k</b> .	
■ k (1000)	8.	Click <b>OK</b> .	
■ m (1,000,000)			
■ g (1,000,000,000)			

#### Table 104: Configuring Policers for TCP and ICMP (continued)

Task	J-Web	b Configuration Editor	CLI Configuration Editor
Define the policer action for tcp-connection-policer.	1. (	On the Policer <b>tcp-connection-policer</b> page, next to Then, click <b>Configure</b> .	Set the policer action:
	2. 1	Next to Discard, select the check box.	set policer tcp-connection-policer then discard
	3. (	Click <b>OK</b> twice.	
Define <b>icmp-policer</b> and set its rate limits.	1. ( 1	On the Firewall page, next to Policer, click <b>Add new entry</b> .	Set the policer name and its rate limits:
The burst size limit can be	2. I	In the Policer name box, type icmp-policer.	set policer icmp-policer filter-specific if-exceeding burst-size-limit 15k bandwidth-limit 1m
from 1,500 bytes through 100,000,000 bytes.	3. I ł	Next to Filter specific, select the check box.	
The bandwidth limit can be from 32,000 bps through 32,000,000,000 bps.	4. I	Next to If Exceeding, select the check box and click <b>Configure</b> .	
	5. I	In the Burst size limit box, type <b>15k</b> .	
Use the following abbreviations when	6. I I	In the Bandwidth list, select <b>Bandwidth limit</b> .	
specifying these limits:	7. I	In the Bandwidth limit box, type 1m.	
■ k (1000)	8. (	Click <b>OK</b> .	
m (1,000,000)			
■ g(1,000,000,000)			
Define the policer action for icmp-policer.	1. (	On the Policer i <b>cmp-policer</b> page, next to Then, click <b>Configure</b> .	Set the policer action:
	2. 1	Next to Discard, select the check box.	set policer icmp-policer then discard
	3. (	Click <b>OK</b> three times.	

## Table 105: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Policy</b> <b>options</b> level in the	1.	In the J-Web interface, select Configuration > View and Edit > Edit	From the [edit] hierarchy level, enter
configuration nierarchy.	2.	Next to Policy options, click <b>Configure</b> or <b>Edit</b> .	edit policy-options

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Define the prefix list	1.	Next to Prefix list, click Add new entry.	Set the prefix list:
trusted-addresses.	2.	In the Name box, type trusted-addresses.	set prefix-list trusted-addresses
	3.	Next to Prefix list item, click <b>Add new</b> entry.	192.168.122.0/24
	4.	In the Prefix box, type <b>192.168.122.0/24</b> .	set prefix-list trusted-addresses 10.2.1.0/24
	5.	Click <b>OK</b> .	
	6.	Next to Prefix list item, click <b>Add new</b> entry.	
	7.	In the Prefix box, type 10.2.1.0/24.	
	8.	Click <b>OK</b> three times.	
Navigate to the <b>Firewall</b>	On	the main Configuration page next to	From the [edit] hierarchy level, enter
hierarchy.	1 11 (	ewail, eller configure of Lon.	edit firewall
Define protect-RE and	1.	Next to Filter, click Add new entry.	Set the term name and define the source
define the source prefix list	2.	In the Filter name box, type protect-RE.	address match condition:
match condition.	3.	Next to Term, click Add New Entry.	set family inet filter protect-RE
	4.	In the Rule name box, type tcp-connection-term.	source-prefix-list trusted-addresses
	5.	Next to From, click <b>Configure</b> .	
	6.	Next to Source prefix list, click <b>Add new</b> entry.	
	7.	In the Name box, type trusted-addresses.	
	8.	Click <b>OK</b> .	
Define the TCP flags and protocol match conditions	1.	In the TCP flags box, type (syn & !ack)   fin   rst.	Set the TCP flags and protocol and protocol match conditions for the term:
for tcp-connection-term.	2.	In the Protocol choice list, select <b>Protocol</b> .	set family inet filter protect-RE
	3.	Next to Protocol, click Add new entry.	term tcp-connection-term from protocol tcp
	4.	In the Value keyword list, select <b>tcp</b> .	tcp-flags "(syn & lack)   fin   rst"
	5.	Click <b>OK</b> .	
Define the actions for tcp-connection-term.	1.	On the Term <b>tcp-connection-term</b> page, next to Then, click <b>Configure</b> .	Set the actions:
	2.	In the Policer box, type tcp-connection-policer.	set family inet filter protect-RE term tcp-connection-term then policer tcp-connection-policer accept
	3.	In the Designation list, select Accept.	
	4.	Click <b>OK</b> twice.	

## Table 105: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine (continued)

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Define <b>icmp-term</b> , and define the protocol.	1.	On the Filter <b>protect-RE</b> page, next to Term, click <b>Add New Entry</b> .	Set the term name and define the protocol:
	2.	In the Rule name box, type icmp-term.	set family inet filter protect-RE term icmp-term
	3.	Next to From, click Configure.	
	4.	In the Protocol choice list, select <b>Protocol</b> .	
	5.	Next to Protocol, click Add new entry.	
	6.	In the Value keyword list, select <b>icmp</b> .	
	7.	Click <b>OK</b> .	
Define the ICMP type match conditions.	1.	In the Icmp type choice list, select <b>Icmp type</b> .	Set the ICMP type match conditions:
	2.	Next to Icmp type, click Add new entry.	set family inet filter protect-RE term icmp-term
	3.	In the Value keyword list, select echo-request.	unreachable time-exceeded]
	4.	Click <b>OK</b> .	
	5.	Next to Icmp type, click Add new entry.	
	6.	In the Value keyword list, select echo-reply.	
	7.	Click <b>OK</b> .	
	8.	Next to Icmp type, click Add new entry.	
	9.	In the Value keyword list, select unreachable.	
	10.	Click <b>OK</b> .	
	11.	Next to Icmp type, click Add new entry.	
	12.	In the Value keyword list, select <b>time-exceeded</b> .	
	13.	Click <b>OK</b> .	
Define the actions for icmp-term.	1.	On the icmp-term page, next to Then, click Configure.	Set the actions:
	2.	In the Count box, type icmp-counter.	set family inet filter protect-RE term icmp-term
	3.	In the Policer box, type icmp-policer.	accept
	4.	In the Designation list, select Accept.	
	5.	Click <b>OK</b> four times.	

#### Table 105: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine (continued)

## **Configuring a Routing Engine Firewall Filter to Handle Fragments**

The procedure in this section creates a sample stateless firewall filter, **fragment-RE**, that handles fragmented packets destined for the Routing Engine. By applying

fragment-RE to the Routing Engine, you protect against the use of IP fragmentation as a means to disguise TCP packets from a firewall filter.

Table 106 on page 247 lists the terms that are configured in this sample filter.

Table 106: Sample Stateless Firewall Filter fragment-RE Terms

Term	Purpose
small-offset-term	Discards IP packets with a fragment offset of 1 through 5, and adds a record to the system logging facility.
not-fragmented-term	Accepts unfragmented TCP packets with a source address of 10.2.1.0/24 and a destination port that specifies the BGP protocol. A packet is considered unfragmented if its MF flag and its fragment offset in the TCP header equal 0.
first-fragment-term	Accepts the first fragment of a fragmented TCP packet with a source address of 10.2.1.0/24 and a destination port that specifies the BGP protocol.
fragment-term	Accepts all packet fragments with an offset of 6 through 8191.

For example, consider an IP packet that is fragmented into the smallest allowable fragment size of 8 bytes (a 20-byte IP header plus an 8-byte payload). If this IP packet carries a TCP packet, the first fragment (fragment offset of 0) that arrives at the Services Router contains only the TCP source and destination ports (first 4 bytes), and the sequence number (next 4 bytes). The TCP flags, which are contained in the next 8 bytes of the TCP header, arrive in the second fragment (fragment offset of 1). The fragment-RE filter works as follows:

- Term small-offset-term discards small offset packets to ensure that subsequent terms in the firewall filter can be matched against all the headers in the packet.
- Term fragment-term accepts all fragments that were not discarded by small-offset-term. However, only those fragments that are part of a packet containing a first fragment accepted by first-fragment-term are reassembled by the Services Router.

For more information about IP fragment filtering, see RFC 1858, *Security Considerations for IP Fragment Filtering.* 

To use the configuration editor to configure the stateless firewall filter:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. To configure the firewall filter, perform the configuration tasks described in Table 107 on page 248.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to one of the following procedures:
  - To display the configuration, see Displaying Stateless Firewall Filter Configurations on page 252.

- To apply the firewall filter to the Routing Engine, see "Applying a Stateless Firewall Filter to an Interface" on page 251.
- To verify the firewall filter, see Verifying a Firewall Filter That Handles Fragments on page 259.

Table 107: Configuring a	Fragments Firewall	Filter for the Routing Engine
--------------------------	--------------------	-------------------------------

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Firewall</b> level in the configuration	1.	In the J-Web interface, select Configuration > View and Edit > Edit	From the [edit] hierarchy level, enter
hierarchy.		Configuration.	edit firewall
	2.	Next to Firewall, click <b>Configure</b> or <b>Edit</b> .	
Define fragment-RE and	1.	Next to Filter, click Add new entry.	Set the term name and define the fragment
define the fragment offset	2.	In the Filter name box, type fragment-RE.	onset match condition:
match condition.	3.	Next to Term, click Add New Entry.	set family inet filter fragment-RE
The fragment offset can be from 1 through 8191.	4.	In the Rule name box, type small-offset-term.	term small-offset-term from fragment-offset 1-5
Ŭ	5.	Next to From, click Configure.	
	6.	In the Fragment offset choice list, select <b>Fragment offset</b> .	
	7.	Next to Fragment offset, select <b>Add New Entry</b> .	
	8.	In the Range box, type <b>1-5</b> .	
	9.	Click <b>OK</b> twice.	
Define the action for small-offset-term.	1.	On the Term <b>small-offset-term</b> page, next to Then, click <b>Configure</b> .	Set the action:
	2.	Next to Syslog, select the check box.	set family inet filter fragment-RE term small-offset-term then syslog discard
	3.	In the Designation list, select Discard.	
	4.	Click <b>OK</b> twice.	

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Define not-fragmented-term, and define the fragment,	1.	On the Filter fragment-RE page, next to Term, click Add New Entry.	Set the term name and define match conditions:
protocol, destination port, and source address match conditions.	2.	In the Term name box, type not-fragmented-term.	set family inet filter fragment-RE term not-fragmented-term from fragment-flags 0x0
	3.	Next to From, click Configure.	fragment-offset 0 protocol tcp destination-port bgp
	4.	In the Fragment flags box, type <b>0x0</b> .	source-address 10.2.1.0/24
	5.	In the Fragment offset choice list, select <b>Fragment offset</b> .	
	6.	Next to Fragment offset, select <b>Add New Entry</b> .	
	7.	In the Range box, type <b>0</b> .	
	8.	Click <b>OK</b> .	
	9.	In the Protocol choice list, select <b>Protocol</b> .	
	10.	Next to Protocol, click Add new entry.	
	11.	In the Value keyword list, select <b>tcp</b> .	
	12.	Click <b>OK</b> .	
	13.	In the Destination port choice list, select <b>Destination port</b> .	
	14.	Next to Destination port, click <b>Add new</b> entry.	
	15.	In the Value keyword list, select <b>bgp</b> .	
	16.	Click <b>OK</b> .	
	17.	Next to Source address, click <b>Add new</b> entry.	
	18.	In the Address box, type 10.2.1.0/24.	
	19.	Click <b>OK</b> twice.	
Define the action for not-fragmented-term.	1.	On the Term <b>not-fragmented-term</b> page, next to Then, click <b>Configure</b> .	Set the action:
	2.	In the Designation list, select Accept.	set family inet filter fragment-RE term not-fragmented-term then accept
	3.	Click <b>OK</b> twice.	

## Table 107: Configuring a Fragments Firewall Filter for the Routing Engine (continued)

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Define <b>first-fragment-term</b> , and define the fragment,	1.	On the Filter fragment-RE page, next to Term, click Add New Entry.	Set the term name and define match conditions:
protocol, destination port, and source address match conditions.	2.	In the Rule name box, type first-fragment-term.	set family inet filter fragment-RE term first-fragment-term from first-fragment
	3.	Next to From, click Configure.	protocol tcp destination-port bgp
	4.	Next to First fragment, select the check box.	source-address 10.2.1.0/24
	5.	In the Protocol choice list, select <b>Protocol</b> .	
	6.	Next to Protocol, click Add new entry.	
	7.	In the Value keyword list, select <b>tcp</b> .	
	8.	Click <b>OK</b> .	
	9.	In the Destination port choice list, select <b>Destination port</b> .	
	10.	Next to Destination port, click <b>Add new</b> entry.	
	11.	In the Value keyword list, select <b>bgp</b> .	
	12.	Click <b>OK</b> .	
	13.	Next to Source address, click <b>Add new</b> entry.	
	14.	In the Address box, type 10.2.1.0/24.	
	15.	Click <b>OK</b> twice.	
Define the action for first-fragment-term.	1.	On the Term <b>first-fragment-term</b> page, next to Then, click <b>Configure</b> .	Set the action:
	2.	In the Designation list, select Accept.	set family inet filter fragment-RE
	3.	Click <b>OK</b> twice.	
Define <b>fragment-term</b> and define the fragment match	1.	On the Filter fragment-RE page, next to Term, click <b>Add New Entry</b> .	Set the term name and define match conditions:
condition.	2.	In the Rule name box, type fragment-term.	act family inst filter frogment DE
	3.	Next to From, click Configure.	term fragment-term from fragment-offset 6–8191
	4.	In the Fragment offset choice list, select Fragment offset.	
	5.	Next to Fragment offset, select <b>Add New Entry</b> .	
	6.	In the Range box, type <b>6-8191</b> .	
	7.	Click <b>OK</b> twice.	

## Table 107: Configuring a Fragments Firewall Filter for the Routing Engine (continued)

Task	J-M	/eb Configuration Editor	CLI Configuration Editor
Define the action for fragment-term.	1.	On the Term <b>fragment-term</b> page, next to Then, click <b>Configure</b> .	Set the action:
	2.	In the Designation list, select Accept.	set family inet filter fragment-RE term fragment-term then accept
	3.	Click <b>OK</b> four times.	

#### Table 107: Configuring a Fragments Firewall Filter for the Routing Engine (continued)

## Applying a Stateless Firewall Filter to an Interface

You can apply a stateless firewall to the input or output sides, or both, of an interface. To filter packets transiting the router, apply the firewall filter to any non-Routing Engine interface. To filter packets originating from, or destined for, the Routing Engine, apply the firewall filter to the loopback (lo0) interface.

For example, to apply the firewall filter **protect-RE** to the input side of the Routing Engine interface, follow this procedure:

- 1. Perform the configuration tasks described in Table 108 on page 251.
- 2. If you are finished configuring the router, commit the configuration.

#### Table 108: Applying a Firewall Filter to the Routing Engine Interface

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Inet</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, apply the filter to the interface:
(See the interface naming conventions in the <i>J-series</i> <i>Services Router Basic LAN</i> <i>and WAN Access</i> <i>Configuration Guide.</i> )	2.	Next to Interfaces, click Configure or Edit.	set interfaces loO unit O family inet filter input protect-RE
	3.	Under Interface name, click <b>lo0</b> .	
	4.	Under Interface unit number, click ${f 0}.$	
	5.	Under Family, make sure the Inet check box is selected, and click <b>Configure</b> or <b>Edit</b> .	
Apply <b>protect-RE</b> as an input filter to the <b>loO</b> interface.	1.	Next to Filter, click <b>Configure</b> .	-
	2.	In the Input box, type protect-RE.	
	3.	Click <b>OK</b> five times.	

To view the configuration of the Routing Engine interface, enter the **show interfaces IoO** command. For example:

```
user@host# show interfaces lo0
unit 0 {
family inet {
```

```
filter {
 input protect-RE;
 }
 address 127.0.0.1/32;
}
```

## **Verifying Stateless Firewall Filter Configuration**

}

To verify a stateless firewall filter configuration, perform these tasks:

- Displaying Stateless Firewall Filter Configurations on page 252
- Displaying Stateless Firewall Filter Logs on page 255
- Displaying Firewall Filter Statistics on page 256
- Verifying a Services, Protocols, and Trusted Sources Firewall Filter on page 257
- Verifying a TCP and ICMP Flood Firewall Filter on page 258
- Verifying a Firewall Filter That Handles Fragments on page 259

## **Displaying Stateless Firewall Filter Configurations**

**Purpose** Verify the configuration of the firewall filter. You can analyze the flow of the filter terms by displaying the entire configuration.

Action From the J-Web interface, select Configuration > View and Edit > View Configuration Text. Alternatively, from configuration mode in the CLI, enter the show firewall command.

The sample output in this section displays the following firewall filters (in order):

- Stateless protect-RE filter configured in "Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources" on page 238
- Stateless protect-RE filter configured in "Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods" on page 241
- Stateless fragment-RE filter configured in "Configuring a Routing Engine Firewall Filter to Handle Fragments" on page 246

```
[edit]
user@host# show firewall
firewall {
 family inet {
 filter protect-RE {
 term ssh-term {
 from {
 source-address {
 192.168.122.0/24;
 }
 protocol tcp;
 destination-port ssh;
 }
}
```

```
then accept;
 }
 term bgp-term {
 from {
 source-address {
 10.2.1.0/24;
 }
 protocol tcp;
 destination-port bgp;
 }
 then accept;
 }
 term discard-rest-term {
 then {
 log;
 syslog;
 discard;
 }
 }
 }
 }
}
[edit]
user@host# show firewall
firewall {
 policer tcp-connection-policer {
 filter-specific;
 if-exceeding {
 bandwidth-limit 500k;
 burst-size-limit 15k;
 }
 then discard;
 }
 policer icmp-policer {
 filter-specific;
 if-exceeding {
 bandwidth-limit 1m;
 burst-size-limit 15k;
 }
 then discard;
 }
 family inet {
 filter protect-RE {
 term tcp-connection-term {
 from {
 source-prefix-list {
 trusted-addresses;
 }
 protocol tcp;
 tcp-flags "(syn & !ack) | fin | rst";
 }
 then {
 policer tcp-connection-policer;
 accept;
 }
```

```
}
 term icmp-term {
 from {
 protocol icmp;
 icmp-type [echo-request echo-reply unreachable time-exceeded];
 }
 then {
 policer icmp-policer;
 count icmp-counter;
 accept;
 }
 }
 additional terms...
 }
 }
}
[edit]
user@host# show firewall
firewall {
 family inet {
 filter fragment-RE {
 term small-offset-term {
 from {
 fragment-offset 1-5;
 }
 then {
 syslog;
 discard;
 }
 }
 term not-fragmented-term {
 from {
 source-address {
 10.2.1.0/24;
 }
 fragment-offset 0;
 fragment-flags 0x0;
 protocol tcp;
 destination-port bgp;
 }
 then accept;
 }
 term first-fragment-term {
 from {
 source-address {
 10.2.1.0/24;
 }
 first-fragment;
 protocol tcp;
 destination-port bgp;
 }
 then accept;
 }
 term fragment-term {
 from {
```

```
fragment-offset 6-8191;

}

then accept;

}

additional terms ...

}

}
```

What It Means Verify that the output shows the intended configuration of the firewall filter.

Verify that the terms are listed in the order in which you want the packets to be tested. You can move terms within a firewall filter by using the **insert** CLI command.

**Related Topics** For more information about the format of a configuration file, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

For information about the insert command, see the *J*-series Services Router Basic LAN and WAN Access Configuration Guide.

## **Displaying Stateless Firewall Filter Logs**

- **Purpose** Verify that packets are being logged. If you included the log or syslog action in a term, verify that packets matching the term are recorded in the firewall log or your system logging facility.
  - Action From operational mode in the CLI, enter the show firewall log command.

The log of discarded packets generated from the stateless firewall filter configured in "Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources" on page 238 is displayed in the following sample output.

user@host	<pre>&gt; show fir</pre>	ewall 1	og			
Log :						
Time	Filter	Action	Interface	Protocol	Src Addr	Dest Addr
15:11:02	pfe	D	ge-0/0/0.0	ТСР	172.17.28.19	192.168.70.71
15:11:01	pfe	D	ge-0/0/0.0	ТСР	172.17.28.19	192.168.70.71
15:11:01	pfe	D	ge-0/0/0.0	ТСР	172.17.28.19	192.168.70.71
15:11:01	pfe	D	ge-0/0/0.0	ТСР	172.17.28.19	192.168.70.71

**What It Means** Each record of the output contains information about the logged packet. Verify the following information:

- Under Time, the time of day the packet was filtered is shown.
- The **Filter** output is always **pfe**.
- Under Action, the configured action of the term matches the action taken on the packet—A (accept), D (discard), R (reject).
- Under Interface, the inbound (ingress) interface on which the packet arrived is appropriate for the filter.
- Under **Protocol**, the protocol in the IP header of the packet is appropriate for the filter.
- Under Src Addr, the source address in the IP header of the packet is appropriate for the filter.
- Under **Dest Addr**, the destination address in the IP header of the packet is appropriate for the filter.
- **Related Topics** For a complete description of **show firewall log** output, see the *JUNOS Routing Protocols and Policies Command Reference.*

## **Displaying Firewall Filter Statistics**

**Purpose** Verify that packets are being policed and counted.

Action From operational mode in the CLI, enter the show firewall filter *filter-name* command.

The value of the counter, **icmp-counter**, and the number of packets discarded by the policers in the stateless firewall filter configured in "Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods" on page 241 are displayed in the following sample output.

user@host> show firewall filter protect	t-RE	
Filter: protect-RE		
Counters:		
Name	Bytes	Packets
icmp-counter	1040000	5600
Policers:		
Name	Packets	
tcp-connection-policer	643254873	
icmp-policer	7391	

**What It Means** Verify the following information:

- Next to Filter, the name of the firewall filter is correct.
- Under Counters:
  - Under Name, the names of any counters configured in the firewall filter are correct.
  - Under Bytes, the number of bytes that match the filter term containing the count counter-name action are shown.
  - Under Packets, the number of packets that match the filter term containing the count counter-name action are shown.

- Under Policers:
  - Under Name, the names of any policers configured in the firewall filter are correct.
  - Under Packets, the number of packets that match the conditions specified for the policer are shown.
- **Related Topics** For a complete description of the **show** firewall filter command and output, see the *JUNOS Routing Protocols and Policies Command Reference*.

#### Verifying a Services, Protocols, and Trusted Sources Firewall Filter

- **Purpose** Verify the stateless firewall filter configured in "Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources" on page 238.
- **Action** To verify that the actions of the firewall filter terms are taken, send packets to the Services Router that match the terms. In addition, verify that the filter actions are *not* taken for packets that do not match.
  - Use the ssh host-name command from a host at an IP address that matches 192.168.122.0/24 to verify that you can log in to the Services Router using only SSH from a host with this address prefix.
  - Use the show route summary command to verify that the routing table on the Services Router does not contain any entries with a protocol other than Direct, Local, BGP, or Static.

```
% ssh 192.168.249.71
%ssh host
user@host's password:
--- JUNOS 6.4-20040518.0 (JSERIES) #0: 2004-05-18 09:27:50 UTC
user@host>
user@host>
user@host>
user@host> show route summary
Router ID: 192.168.249.71
inet.0: 34 destinations, 34 routes (33 active, 0 holddown, 1 hidden)
Direct: 10 routes, 9 active
Local: 9 routes, 9 active
BGP: 10 routes, 10 active
Static: 5 routes, 5 active
...
What It Means
Verify the following information:
You can successfully log in to the Services Router using SSH.
The show route summary command does not display a protocol other than Direct, Local, BGP, or Static.
```

**Related Topics** For a complete description of **show route summary** output, see the *JUNOS Routing Protocols and Policies Command Reference*.

## Verifying a TCP and ICMP Flood Firewall Filter

- **Purpose** Verify the stateless firewall filter configured in "Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods" on page 241.
  - **Action** To verify that the actions of the firewall filter terms are taken, send packets to the Services Router that match the terms. In addition, verify that the filter actions are *not* taken for packets that do not match.
    - Verify that the Services Router can establish only TCP sessions with a host at an IP address that matches 192.168.122.0/24 or 10.2.1.0/24. For example, log in to the router with the telnet host-name command from another host with one of these address prefixes.
    - Use the ping host-name command to verify that the Services Router responds only to ICMP packets (such as ping requests) that do not exceed the policer traffic rates.
    - Use the ping host-name size bytes command to exceed the policer traffic rates by sending ping requests with large data payloads.

```
user@host> telnet 192.168.249.71
 Trving 192.168.249.71...
 Connected to host.acme.net.
 Escape character is '^]'.
 host (ttyp0)
 login: user
 Password:
 --- JUNOS 6.4-20040521.1 built 2004-05-21 09:38:12 UTC
 user@host>
 user@host> ping 192.168.249.71
 PING host-ge-000.acme.net (192.168.249.71): 56 data bytes
 64 bytes from 192.168.249.71: icmp_seq=0 ttl=253 time=11.946 ms
 64 bytes from 192.168.249.71: icmp_seq=1 ttl=253 time=19.474 ms
 64 bytes from 192.168.249.71: icmp_seq=2 ttl=253 time=14.639 ms
 . . .
 user@host> ping 192.168.249.71 size 20000
 PING host-ge-000.acme.net (192.168.249.71): 20000 data bytes
 ٨C
 --- host-ge-000.acme.net ping statistics ---
 12 packets transmitted, 0 packets received, 100% packet loss
What It Means
 Verify the following information:
 You can successfully log in to the Services Router using Telnet.
```

- The Services Router sends responses to the ping host command.
- The Services Router does not send responses to the **ping host size 20000** command.

**Related Topics** For more information about the ping command, see the *J*-series Services Router Administration Guide or the JUNOS System Basics and Services Command Reference.

For information about using the J-Web interface to ping a host, see the *J-series Services Router Administration Guide*.

For more information about the **telnet** command, see the *J*-series Services Router Administration Guide or the JUNOS System Basics and Services Command Reference.

## **Verifying a Firewall Filter That Handles Fragments**

Direct, Local, BGP, or Static.

- **Purpose** Verify the firewall filter configured in "Configuring a Routing Engine Firewall Filter to Handle Fragments" on page 246.
  - **Action** To verify that the actions of the firewall filter terms are taken, send packets to the Services Router that match the terms. In addition, verify that the filter actions are *not* taken for packets that do not match.
    - Verify that packets with small fragment offsets are recorded in the router's system logging facility.
    - Use the **show route summary** command to verify that the routing table does not contain any entries with a protocol other than **Direct**, **Local**, **BGP**, or **Static**.

```
user@host> show route summary
Router ID: 192.168.249.71
inet.0: 34 destinations, 34 routes (33 active, 0 holddown, 1 hidden)
Direct: 10 routes, 9 active
Local: 9 routes, 9 active
BGP: 10 routes, 10 active
Static: 5 routes, 5 active
...
What It Means Verify that the show route summary command does not display a protocol other than
```

**Related Topics** For a complete description of **show route summary** output, see the *JUNOS Routing Protocols and Policies Command Reference*.

J-series[™] Services Router Advanced WAN Access Configuration Guide

# Part 5 Configuring Class of Service

- Class-of-Service Overview on page 263
- Configuring Class of Service on page 283

J-series[™] Services Router Advanced WAN Access Configuration Guide

## Chapter 14 Class-of-Service Overview

With the class-of-service (CoS) features on a J-series Services Router, you can assign service levels with different delay, jitter (delay variation), and packet loss characteristics to particular applications served by specific traffic flows. CoS is especially useful for networks supporting time-sensitive video and audio applications. To configure CoS features on a Services Router, see "Configuring Class of Service" on page 283.

This chapter contains the following topics. For more information about CoS, see the *JUNOS Class of Service Configuration Guide*.

- CoS Terms on page 263
- Benefits of CoS on page 264
- CoS Across the Network on page 265
- JUNOS CoS Components on page 266
- How CoS Components Work on page 271
- Default CoS Settings on page 272
- Transmission Scheduling on J-series Services Routers on page 280

## **CoS Terms**

Before configuring CoS on a Services Router, become familiar with the terms defined in Table 109 on page 263.

#### Table 109: CoS Terms

Term	Definition
assured forwarding (AF)	CoS packet forwarding class that provides a group of values you can define and includes four subclasses, AF1, AF2, AF3, and AF4, each with three drop probabilities, low, medium, and high.
behavior aggregate (BA) classifier	Feature that can be used to determine the forwarding treatment for each packet. The behavior aggregate classifier maps a code point to a loss priority. The loss priority is used later in the work flow to select one of the two drop profiles used by random early detection (RED).
best-effort (BE)	CoS packet forwarding class that provides no service profile. For the BE forwarding class, loss priority is typically not carried in a code point, and random early detection (RED) drop profiles are more aggressive.

#### Table 109: CoS Terms (continued)

Term	Definition
class of service (CoS)	Method of classifying traffic on a packet-by-packet basis, using information in the type-of-service (TOS) byte to assign traffic flows to different service levels.
Differentiated Services (DiffServ)	Services based on RFC 2474, <i>Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers</i> . The DiffServ method of CoS uses the type-of-service (ToS) byte to identify different packet flows on a packet-by-packet basis. DiffServ adds a Class Selector code point (CSCP) and a DiffServ code point (DSCP).
DiffServ code point (DSCP) values	Values for a 6–bit field defined in IP packet headers that can be used to enforce class-of-service (CoS) distinctions in a Services Router
drop profile	Drop probabilities for different levels of buffer fullness that are used by random early detection (RED) to determine from which Services Router scheduling queue to drop packets.
expedited forwarding (EF)	CoS packet forwarding class that provides end-to-end service with low loss, low latency, low jitter, and assured bandwidth.
multifield (MF) classifier	Firewall filter that scans through a variety of packet fields to determine the forwarding class and loss priority for a packet and polices traffic to a specific bandwidth and burst size. Typically, a classifier performs matching operations on the selected fields against a configured value.
network control (NC)	CoS packet forwarding class that is typically high priority because it supports protocol control.
PLP bit	Packet loss priority bit. Used to identify packets that have experienced congestion or are from a transmission that exceeded a service provider's customer service license agreement. A Services Router can use the PLP bit as part of a congestion control strategy. The bit can be configured on an interface or in a filter.
policer	Feature that limits the amount of traffic passing into or out of an interface. It is an essential component of firewall filters that is designed to thwart denial-of-service (DoS) attacks. A policer applies rate limits on bandwidth and burst size for traffic on a particular Services Router interface.
policing	Applying rate and burst size limits to traffic on an interface.
random early detection (RED)	Gradual drop profile for a given class, used for congestion avoidance. RED attempts to anticipate congestion and reacts by dropping a small percentage of packets from the head of a queue to prevent congestion.
rule	Guide that the Services Router follows when applying services. A rule consists of a match direction and one or more terms.

## **Benefits of CoS**

IP routers normally forward packets independently, without controlling throughput or delay. This type of packet forwarding, known as best-effort service, is as good as your network equipment and links allow. Best-effort service is sufficient for many traditional IP data delivery applications, such as e-mail or Web browsing. However, newer IP applications such as real-time video and audio (or voice) require lower delay, jitter, and packet loss than simple best-effort networks can provide.

CoS features allow a Services Router to improve its processing of critical packets while maintaining best-effort traffic flows, even during periods of congestion. Network
throughput is determined by a combination of available bandwidth and delay. CoS dedicates a guaranteed minimum bandwidth to a particular service class by reducing forwarding queue delays. (The other two elements of overall network delay, serial transmission delays determined by link speeds and propagation delays determined by media type, are not affected by CoS settings.)

Normally, packets are queued for output in their order of arrival, regardless of service class. Queueing delays increase with network congestion and often result in lost packets when queue buffers overflow. CoS packet classification assigns packets to forwarding queues by service class.

Because CoS must be implemented consistently end-to-end through the network, the CoS features on the Services Router are based on IETF Differentiated Services (DiffServ) standards, to interoperate with other vendors' CoS implementations.

#### **CoS Across the Network**

CoS works by examining traffic entering at the edge of your network. The edge routers classify traffic into defined service groups, which allow for the special treatment of traffic across the network. For example, voice traffic can be sent across certain links, and data traffic can use other links. In addition, the data traffic streams can be serviced differently along the network path to ensure that higher-paying customers receive better service. As the traffic leaves the network at the far edge, you can reclassify the traffic.

To support CoS, you must configure each router in the network. Generally, each router examines the packets that enter it to determine their CoS settings. These settings then dictate which packets are first transmitted to the next downstream router. In addition, the routers at the edges of your network might be required to alter the CoS settings of the packets transmitting to the neighboring network.

Figure 21 on page 265 shows an example of CoS operating across an Internet Service Provider (ISP) network.

Figure 21: CoS Across the Network



In the ISP network shown in Figure 21 on page 265, Router A is receiving traffic from your network. As each packet enters, Router A examines the packet's current CoS settings and classifies the traffic into one of the groupings defined by the ISP. This definition allows Router A to prioritize its resources for servicing the traffic streams it is receiving. In addition, Router A might alter the CoS settings (forwarding class and loss priority) of the packets to better match the ISP's traffic groups. When Router B receives the packets, it examines the CoS settings, determines the appropriate traffic group, and processes the packet according to those settings.

Router B then transmits the packets to Router C, which performs the same actions. Router D also examines the packets and determines the appropriate group. Because it sits at the far end of the network, the ISP might decide once again to alter the CoS settings of the packets before Router D transmits them to the neighboring network.

#### **JUNOS CoS Components**

J-series Services Routers support the following CoS components:

- Code-Point Aliases on page 266
- Classifiers on page 266
- Forwarding Classes on page 267
- Loss Priorities on page 267
- Forwarding Policy Options on page 267
- Transmission Queues on page 268
- Schedulers on page 268
- Virtual Channels on page 270
- Policers for Traffic Classes on page 270
- Rewrite Rules on page 271

# **Code-Point Aliases**

A code-point alias assigns a name to a pattern of code-point bits. You can use this name, instead of the bit pattern, when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

# Classifiers

Packet classification refers to the examination of an incoming packet. This function associates the packet with a particular CoS servicing level. In the JUNOS software, classifiers associate incoming packets with a forwarding class and loss priority and, based on the associated forwarding class, assign packets to output queues. Two general types of classifiers are supported—behavior aggregate (BA) classifiers and multifield (MF) classifiers.

#### **Behavior Aggregate Classifiers**

A behavior aggregate (BA) classifier operates on a packet as it enters the router. Using behavior aggregate classifiers the router aggregates different types of traffic into a single forwarding class to receive the same forwarding treatment. The CoS value in the packet header is the single field that determines the CoS settings applied to the packet. Behavior aggregate classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services (DiffServ) code point (DSCP) value, DSCP IPv6 value, IP precedence value, MPLS EXP bits, or IEEE 802.1p value.

The default classifier is based on the IP precedence value. For more information, see "Default Behavior Aggregate Classifiers" on page 277.

#### **Multifield Classifiers**

A multifield (MF) classifier is a second method for classifying traffic flows. Unlike the behavior aggregate classifier, a multifield classifier can examine multiple fields in the packet—for example, the source and destination address of the packet or the source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

# **Forwarding Classes**

Forwarding classes allow you to group packets for transmission. Based on forwarding classes, you assign packets to output queues. The forwarding class plus the loss priority define the per-hop behavior (PHB in DiffServ) of a packet. J-series Services Routers support eight queues (0 through 7). Forwarding classes are mapped one-to-one with these queues. By default, queues 0 through 3 are mapped to forwarding classes—best effort, assured forwarding, expedited forwarding, and network control. Queues 4 through 7 are not mapped to forwarding classes. To use queues 4 through 7, you must create custom forwarding class names and map them to the queues. For more information, see "Forwarding Class Queue Assignments" on page 276.

# **Loss Priorities**

Loss priorities allow you to set the priority of dropping a packet. You can use the loss priority setting to identify packets that have experienced congestion. Typically, you mark packets exceeding some service level with a high loss priority—a greater likelihood of being dropped. You set loss priority by configuring a classifier or a policer. The loss priority is used later in the work flow to select one of the drop profiles used by random early detection (RED).

You can configure the packet loss priority (PLP) bit as part of a congestion control strategy. The PLP bit can be configured on an interface or in a filter. A packet for which the PLP bit is set has an increased probability of being dropped during congestion.

# **Forwarding Policy Options**

Services Routers support CoS-based forwarding (CBF) that enables you to control next-hop selection based on a packet's class of service and, in particular, the value of the IP packet's precedence bits. For example, you can specify a particular interface or next hop to carry high-priority traffic while all best-effort traffic takes some other path. CBF allows path selection based on class. When a routing protocol discovers equal-cost paths, it can pick a path at random or load-balance across the paths through either hash selection or round-robin selection.

Forwarding policy also allows you to create CoS classification overrides. For IPv4 or IPv6 packets, you can override the incoming CoS classification and assign the packets to a forwarding class based on their input interface, input precedence bits, or destination address. When you override the classification of incoming packets, any

mappings you configured for associated precedence bits or incoming interfaces to output transmission queues are ignored.

## **Transmission Queues**

After a packet is sent to the outgoing interface on a router, it is queued for transmission on the physical media. The amount of time a packet is queued on the router is determined by the availability of the outgoing physical media as well as the amount of traffic using the interface.

J-series Services Routers support queues 0 through 7. If you configure more than eight queues on a Services Router, the commit operation fails and the router displays a detailed message stating the total number of queues available.

#### Schedulers

An individual router interface has multiple queues assigned to store packets temporarily before transmission. The router uses a scheduling method, often based on packet type, to determine the order in which the queues are serviced. JUNOS schedulers allow you to define the priority, bandwidth, delay buffer size, rate control status, and RED drop profiles to be applied to a particular queue for packet transmission. For more information, see "Scheduler Settings" on page 277.

On J-series Services Routers, you can configure per-unit scheduling (also called logical interface scheduling). Per-unit scheduling allows you to enable multiple output queues on a logical interface and associate an output scheduler with each queue.

#### **Transmit Rate**

The transmission rate determines the traffic transmission bandwidth for each forwarding class you configure. The rate is specified in bits per second (bps). Each queue is allocated some portion of the bandwidth of the outgoing interface.

This bandwidth amount can be a fixed value, such as 1 megabit per second (Mbps), a percentage of the total available bandwidth, or the rest of the available bandwidth. You can limit the transmission bandwidth to the exact value you configure, or allow it to exceed the configured rate if additional bandwidth is available from other queues. This property helps ensure that each queue receives the amount of bandwidth appropriate to its level of service.

On J-series Services Routers, the minimum transmit rate supported on high-speed interfaces is one-ten thousandth of the speed of that interface. For example, on a Gigabit Ethernet interface with a speed of 1000 Mbps, the minimum transmit rate is 100 Kbps (1000 Mbps x 1/10000). You can configure transmit rates in the range 3200 bps through 160,000,000,000 bps. When the configured rate is less than the minimum transmit rate, the minimum transmit rate is used instead.



**NOTE:** Interfaces with slower interface speeds, like T1, E1, or channelized T1/E1/ISDN PRI, cannot support minimum transmit rates because the minimum transmit rate supported on a Services Router is 3200 bps.

On J-series Services Routers, transmit rate assigns the weighted round-robin (WRR) priority values within a given priority level and not between priorities. For more information, see "Transmission Scheduling on J-series Services Routers" on page 280.

#### **Delay Buffer Size**

You can configure the delay buffer size to control congestion at the output stage. A delay buffer provides packet buffer space to absorb burst traffic up to a specified duration of delay. When the buffer becomes full, packets with 100 percent drop probability are dropped from the head of the buffer.

The system calculates the buffer size for a queue based on the buffer allocation method you specify for it in the scheduler. See "Delay Buffer Size Allocation Methods" on page 341 for different buffer allocation methods and "Specifying Delay Buffer Sizes for Queues" on page 342 for buffer size calculations.

By default, all J-series Services Router interfaces other than channelized T1/E1 interfaces support a delay buffer time of 100,000 microseconds. On channelized T1/E1 interfaces, the default delay buffer time is 500,000 microseconds for clear-channel interfaces, and 1,200,000 microseconds for *N*xDS0 interfaces.

On J-series Services Routers, you can configure larger delay buffers on channelized T1/E1 interfaces. Larger delay buffers help these slower interfaces to avoid congestion and packet dropping when they receive large bursts of traffic. For more information, see "Configuring Large Delay Buffers with a Configuration Editor" on page 340.

## **Scheduling Priority**

Scheduling priority determines the order in which an output interface transmits traffic from the queues, thus ensuring that queues containing important traffic are provided better access to the outgoing interface.

The queues for an interface are divided into sets based on their priority. Each set contains queues of the same priority. The router examines the sets in descending order of priority. If at least one queue in a set has a packet to transmit, the router selects that set. If multiple queues in the set have packets to transmit, the router selects a queue from the set according to the weighted round-robin (WRR) algorithm that operates within the set.

The packets in a queue are transmitted based on the configured scheduling priority, the transmit rate, and the available bandwidth. For more information, see "Transmission Scheduling on J-series Services Routers" on page 280.

# **Shaping Rate**

Shaping rates control the maximum rate of traffic transmitted on an interface. You can configure the shaping rate so that the interface transmits less traffic than it is physically capable of carrying.

On J-series Services Routers, you can configure shaping rates on logical interfaces. By default, output scheduling is not enabled on logical interfaces. Logical interface scheduling (also called per-unit scheduling) allows you to enable multiple output queues on a logical interface and associate an output scheduler and shaping rate with the queues.

By default, the logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. You can specify a peak bandwidth rate in bits per second (bps), either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). The range is from 1000 through 32,000,000,000 bps.

#### **RED Drop Profiles**

A drop profile is a feature of the random early detection (RED) process that allows packets to be dropped before queues are full. Drop profiles are composed of two main values—the queue fullness and the drop probability. The queue fullness represents percentage of memory used to store packets in relation to the total amount that has been allocated for that queue. The drop probability is a percentage value that correlates to the likelihood that an individual packet is dropped from the network. These two variables are combined in a graph-like format.

When a packet reaches the head of the queue, a random number between 0 and 100 is calculated by the router. This random number is plotted against the drop profile having the current queue fullness of that particular queue. When the random number falls above the graph line, the packet is transmitted onto the physical media. When the number falls below the graph line, the packet is dropped from the network.

When you configure the RED drop profile on an interface, the queue no longer drops packets from the tail of the queue (the default). Rather, packets are dropped after they reach the head of the queue.

# **Virtual Channels**

On J-series Services Routers, you can configure virtual channels to limit traffic sent from a corporate headquarters to branch offices. Virtual channels might be required when the headquarters site has an expected aggregate bandwidth higher than that of the individual branch offices. The router at the headquarters site must limit the traffic sent to each branch office router to avoid oversubscribing their links.

You configure virtual channels on a logical interface. Each virtual channel has a set of eight queues with a scheduler and an optional shaper. You can use an output firewall filter to direct traffic to a particular virtual channel. For example, a filter can direct all traffic with a destination address for branch office 1 to virtual channel 1, and all traffic with a destination address for branch office 2 to virtual channel 2.

Although a virtual channel group is assigned to a logical interface, a virtual channel is not the same as a logical interface. The only features supported on a virtual channel are queuing, packet scheduling, and accounting. Rewrite rules and routing protocols apply to the entire logical interface.

#### **Policers for Traffic Classes**

Policers allow you to limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a

different forwarding class, a different loss priority, or both. You define policers with firewall filters that can be associated with input or output interfaces.

## **Rewrite Rules**

A rewrite rule resets the appropriate CoS bits in an outgoing packet. Resetting the bits allows the next downstream router to classify the packet into the appropriate service group. Rewriting or marking outbound packets is useful when the router is at the border of a network and must alter the CoS values to meet the policies of the targeted peer.

#### **How CoS Components Work**

On a Services Router, you configure CoS functions using different components. These components are configured individually or in a combination to define particular CoS services. Figure 22 on page 271 displays the relationship of different CoS components to each other and illustrates the sequence in which they interact. "JUNOS CoS Components" on page 266 defines the components and explains their use.



Figure 22: Packet Flow Through J-series CoS-Configurable Components

Each box in Figure 22 on page 271 represents a CoS component. The solid lines show the direction of packet flow in a router. The upper row indicates an incoming packet, and the lower row an outgoing packet. The dotted lines show the inputs and outputs of particular CoS components. For example, the forwarding class and loss priority are outputs of behavior aggregate classifiers and multifield classifiers and inputs for rewrite markers and schedulers.

Typically, only a combination of some components shown in Figure 22 on page 271 (not all) is used to define a CoS service offering. For example, if a packet's class is determined by a behavior aggregate classifier, it is associated with a forwarding class and loss priority and does not need further classification by the multifield classifier.

## **CoS Process on Incoming Packets**

Classifiers and policers perform the following operations on incoming packets:

- 1. A classifier examines an incoming packet and assigns a forwarding class and loss priority to it.
- 2. Based on the forwarding class, the packet is assigned to an outbound transmission queue.
- 3. Input policers meter traffic to see if traffic flow exceeds its service level. Policers might discard, change the forwarding class and loss priority, or set the PLP bit of a packet. A packet for which the PLP bit is set has an increased probability of being dropped during congestion.

# **CoS Process on Outgoing Packets**

The scheduler map and rewrite rules perform the following operations on outgoing packets:

- 1. Scheduler maps are applied to interfaces and associate the outgoing packets with a scheduler and a forwarding class.
- 2. The scheduler defines how the packet is treated in the output transmission queue based on the configured transmit rate, buffer size, priority, and drop profile.
  - The buffer size defines the period for which the packet is stored during congestion.
  - The scheduling priority and transmit rate determine the order in which the packet is transmitted.
  - The drop profile defines how aggressively to drop packets that are using a particular scheduler.
- 3. Output policers meter traffic and might change the forwarding class and loss priority of a packet if a traffic flow exceeds its service level.
- 4. The rewrite rule writes information to the packet (for example, EXP or DSCP bits) according to the forwarding class and loss priority of the packet.

# **Default CoS Settings**

Even when you do not configure any CoS settings on your routing platform, the software performs some CoS functions to ensure that user traffic and protocol packets are forwarded with minimum delay when the network is experiencing congestion. Some default mappings are automatically applied to each logical interface that you configure. Other default mappings, such as explicit default classifiers and rewrite rules, are in operation only if you explicitly associate them with an interface.

You can display default CoS settings by running the **show class-of-service** operational mode command.

This section contains the following topics:

- Default CoS Values and Aliases on page 273
- Forwarding Class Queue Assignments on page 276
- Scheduler Settings on page 277
- Default Behavior Aggregate Classifiers on page 277
- CoS Value Rewrites on page 279
- Sample Behavior Aggregate Classification on page 279

# **Default CoS Values and Aliases**

Table 110 on page 274 shows the default mappings between the bit values and standard aliases.

# Table 110: Well-Known CoS Aliases and Default CoS Values

CoS Value Type	Alias	CoS Value
DSCP and DSCP IPv6	ef	101110
	af11	001010
	af12	001100
	af13	001110
	af21	010010
	af22	010100
	af23	010110
	af31	011010
	af32	011100
	af33	011110
	af41	100010
	af42	100100
	af43	100110
	be	000000
	cs1	001000
	cs2	010000
	cs3	011000
	cs4	100000
	cs5	101000
	nc1/cs6	110000
	nc2/cs7	111000

MPLS EXPbe000bala001ef010ef1011af11100af12101nc1/cs6110nc2/cs7111EEE 802.1be000fef010ef1011af12101nc1/cs6100ef1011af12101nc1/cs6100fef000fef010ef1011af12101nc1/cs6100fef010fef010af12101nc1/cs6100fef010af11001af12011af14001af15010af14010af14100af15011af16011af17101af18101af19101af10101af11100af12101af12101af14100af15101af16110af17111	CoS Value Type	Alias	CoS Value
be1001ef010ef1011af11100af12101ncl/cs6110nc2/cs7111EEE 802.1be000be1001ef1011af11100af12101ncl/cs6110ncl/cs6110ef1011af11100af12101ncl/cs6110ncl/cs6100ef1001af12101ncl/cs6100ef1010af1100af12010af1100af12011af11100af12101af11100af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af14100af15101af16101af17101af18101af19101af11101af12101af12101af14101af15101af16101af17101af18101 <t< td=""><td>MPLS EXP</td><td>be</td><td>000</td></t<>	MPLS EXP	be	000
ef010ef1011af11100af12101nc1/cs6110nc2/cs7111IEEE 802.1be000be1001ef1010af12101af12101af12101nc2/cs7111IP precedencebe000ef1001af12101nc1/cs6110nc2/cs7111ef1001af11100af12101if1010af11100af12101if11100af12101af12101if12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af14100af15101af16101af17101af18101af19101af11100af12101af12101af14100af15101af16110af17111		bel	001
ef1011if11100if1201nc1/cs6110nc2/cs7111IEE 802.1be000be1010of1011if1010if1100if12101nc2/cs7111IP precedencebe000if1001if1010if1010if100if100if201if101if101if101if101if1100if1100if1100if1100if1100if1100if1100if1100if1100if1100if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1101if1 </td <td></td> <td>ef</td> <td>010</td>		ef	010
af11100af12101nc1/cs6110nc2/cs7111IEEE 802.1be000be1001af1011af1100af12101nc1/cs6110nc2/cs7111IP precedencebe000be1001af11100af12101nc1/cs6110af1011af11001af12011af11100af12101af11100af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af12101af14100af15110af16110af17111		ef1	011
af12101nc1/cs6110nc2/cs7111IEEE 802.1be000be10016ef0106af111006af12101100nc1/cs6110100be10006be10016if11006if1101100if11016if10016if10106if10106if11006if11006if11006if11006if11016if121016if121016if121016if121016if121016if121016if121016if121016if121016if121016if121016if131016if141006if151106if161106if17111if18110if19110if19111if19111if19111if19111if19111if19111if19111if19111if19111if19111 </td <td></td> <td>af11</td> <td>100</td>		af11	100
ncl/cs6110nc2/cs7111IEEE 802.1be000be1001001ef01001ef101101af1210100nc2/cs711100be1000001ef101001nc2/cs711100af11100100ef101101af1201101ef1010100ef101101af1210101af1210101af12101100af12101101nc2/cs7111100		af12	101
nc2/cs7111IEEE 802.1be000be1001ffff010ffaf11100af12af12101nc1/cs6nc1/cs6110nc2/cs7111IP precedencebe000ff1001ffff1010ffaf12101ffaf12101ffff1010ffaf12011ffaf12101ffaf12101ffaf12101ffaf12101ffnc1/cs6110ffnc1/cs6110ffnc2/cs7111		nc1/cs6	110
EEE 802.1be000be1001ef010ef1011af11100af12101nc1/cs6110nc2/cs7111be000be1001ef1010ef1011af11100af12101nc1/cs6110ff1011af12101ff1100af12101af12101af12101af12101af12101nc1/cs6110nc1/cs6110nc1/cs6110		nc2/cs7	111
be1         001           ef         010           ef1         011           af11         100           af12         101           nc1/cs6         110           nc2/cs7         111           be         000           ef1         001           ef1         010           af12         111           hc1/cs6         100           af1         001           ef1         010           af11         100           af12         101           nc2/cs7         111	IEEE 802.1	be	000
ef         010           ef1         011           af11         100           af12         101           nc1/cs6         110           nc2/cs7         111           be         000           fef         010           ef1         010           af12         101           nc2/cs7         111           be1         001           ef1         010           af11         100           af12         101           af12         101           af12         101           af12         101           af12         101           af12         101           nc1/cs6         110		bel	001
ef1         011           af11         100           af12         101           nc1/cs6         110           nc2/cs7         111           P precedence         be         000           be1         001           ef1         011           af11         001           af12         011           if1         010           ef1         010           af11         100           af12         101           af12         101		ef	010
af11         100           af12         101           nc1/cs6         110           nc2/cs7         111           IP precedence         be         000           be1         001           ef         010           ef1         011           af11         100           af12         101           nc2/cs7         111		ef1	011
af12       101         nc1/cs6       110         nc2/cs7       111         IP precedence       be       000         be1       001         ef       010         ef1       011         af11       100         af12       101         nc1/cs6       110         nc1/cs6       110         nc2/cs7       111		af11	100
nc1/cs6         110           nc2/cs7         111           IP precedence         be         000           be1         001           ef         010           ef1         011           af11         100           af12         101           nc1/cs6         110           nc2/cs7         111		af12	101
nc2/cs7         111           IP precedence         be         000           be1         001           ef         010           ef1         011           af11         100           af12         101           nc1/cs6         110           nc2/cs7         111		nc1/cs6	110
IP precedence         be         000           be1         001           ef         010           ef1         011           af11         100           af12         101           nc1/cs6         110           nc2/cs7         111		nc2/cs7	111
be1       001         ef       010         ef1       011         af11       100         af12       101         nc1/cs6       110         nc2/cs7       111	IP precedence	be	000
ef       010         ef1       011         af11       100         af12       101         nc1/cs6       110         nc2/cs7       111		bel	001
ef1       011         af11       100         af12       101         nc1/cs6       110         nc2/cs7       111		ef	010
af11       100         af12       101         nc1/cs6       110         nc2/cs7       111		ef1	011
af12     101       nc1/cs6     110       nc2/cs7     111		af11	100
nc1/cs6 110 nc2/cs7 111		af12	101
nc2/cs7 111		nc1/cs6	110
		nc2/cs7	111

# Table 110: Well-Known CoS Aliases and Default CoS Values (continued)

## **Forwarding Class Queue Assignments**

J-series Services Routers have eight queues built into the hardware. By default, four queues are assigned to four forwarding classes. Table 111 on page 276 shows the four default forwarding classes and queues that Juniper Networks classifiers assign to packets based on the CoS values in arriving packet headers. Queues 4 through 7 have no default assignments to forwarding classes. To use queues 4 through 7, you must create custom forwarding class names and assign them to the queues. For more information about how to assign queues to forwarding classes, see the "Configuring Class of Service" on page 283.

By default, all incoming packets, except the IP protocol control packets, are assigned to the forwarding class associated with queue 0. All IP protocol control packets are assigned to the forwarding class associated with queue 3.

Table 111 on page 276 displays the default assignments of forwarding classes to queues.

Forwarding Queue	Forwarding Class	Forwarding Class Description
Queue 0	best-effort (be)	The Services Router does not apply any special CoS handling to packets with 000000 in the DiffServ field, a backward compatibility feature. These packets are usually dropped under congested network conditions.
Queue 1	expedited-forwarding (ef)	The Services Router delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end-to-end for packets in this service class.
		Routers accept excess traffic in this class, but in contrast to assured forwarding, out-of-profile expedited-forwarding packets can be forwarded out of sequence or dropped.
Queue 2	assured-forwarding (af)	The Services Router offers a high level of assurance that the packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define.
		The router accepts excess traffic, but applies a random early detection (RED) drop profile to determine whether the excess packets are dropped and not forwarded.
		Three drop probabilities (low, medium, and high) are defined for this service class.
Queue 3	network-control (nc)	The Services Router delivers packets in this service class with a low priority. (These packets are not delay sensitive.)
		Typically, these packets represent routing protocol hello or keepalive messages. Because loss of these packets jeopardizes proper network operation, delay is preferable to discard.

**Table 111: Default Forwarding Class Queue Assignments** 

#### Scheduler Settings

Each forwarding class has an associated scheduler priority. Only two forwarding classes, **best-effort** and **network-control** (queue 0 and queue 3), are used in the JUNOS default scheduler configuration.

By default, the **best-effort** forwarding class (queue 0) receives 95 percent, and the **network-control** (queue 3) receives 5 percent of the bandwidth and buffer space for the output link. The default drop profile causes the buffer to fill and then discard all packets until it again has space.

The **expedited-forwarding** and **assured-forwarding** classes have no schedulers, because by default no resources are assigned to queue 1 and queue 2. However, you can manually configure resources for **expedited-forwarding** and **assured-forwarding**.

By default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues. When a forwarding class does not fully use the allocated transmission bandwidth, the remaining bandwidth can be used by other forwarding classes if they receive a larger amount of offered load than the bandwidth allocated. If you do not want a queue to use any leftover bandwidth, you must configure it for strict allocation. For more information, see "Configuring Strict High Priority for Queuing with a Configuration Editor" on page 332.

The router uses the following default scheduler settings. You can modify these settings through configuration. For instructions, see "Configuring Class of Service" on page 283.

```
[edit class-of-service]
schedulers {
 network-control {
 transmit-rate percent 5;
 buffer-size percent 5;
 priority low:
 drop-profile-map loss-priority any protocol any drop-profile terminal;
 best-effort {
 transmit-rate percent 95;
 buffer-size percent 95;
 priority low;
 drop-profile-map loss-priority any protocol any drop-profile terminal;
 }
}
drop-profiles {
 terminal {
 fill-level 100 drop-probability 100;
 }
}
```

#### **Default Behavior Aggregate Classifiers**

Table 112 on page 278 shows the forwarding class and packet loss priority (PLP) that are assigned by default to each well-known DSCP. Although several DSCPs map to the **expedited-forwarding** (ef) and **assured-forwarding** (af) classes, by default no resources are assigned to these forwarding classes. All af classes other than af1x are mapped

to **best-effort**, because RFC 2597, *Assured Forwarding PHB Group*, prohibits a node from aggregating classes. Assignment to **best-effort** implies that the node does not support that class.

You can modify the default settings through configuration. For instructions, see "Configuring Class of Service" on page 283.

Table	112:	Default	Behavior	Aggregate	Classification

DSCP and DSCP IPv6 Alias	Forwarding Class	Packet Loss Priority (PLP)
ef	expedited-forwarding	low
af11	assured-forwarding	low
af12	assured-forwarding	high
af13	assured-forwarding	high
af21	best-effort	low
af22	best-effort	low
af23	best-effort	low
af31	best-effort	low
af32	best-effort	low
af33	best-effort	low
af41	best-effort	low
af42	best-effort	low
af43	best-effort	low
be	best-effort	low
cs1	best-effort	low
cs2	best-effort	low
cs3	best-effort	low
cs4	best-effort	low
cs5	best-effort	low
nc1/cs6	network-control	low
nc2/cs7	network-control	low
other	best-effort	low

# **CoS Value Rewrites**

Typically, a router rewrites CoS values in outgoing packets on the outbound interfaces of an edge router, to meet the policies of the targeted peer. After reading the current forwarding class and loss priority information associated with the packet, the transmitting router locates the chosen CoS value from a table, and writes this CoS value into the packet header.

For instructions for configuring rewrite rules, see "Configuring and Applying Rewrite Rules" on page 312.

## Sample Behavior Aggregate Classification

Table 113 on page 279 shows the router forwarding classes associated with each well-known DSCP code point and the resources assigned to their output queues for a sample DiffServ CoS implementation. This example assigns expedited forwarding to queue 1 and a subset of the assured forwarding classes (af1x) to queue 2, and distributes resources among all four forwarding classes.

Other DiffServ-based implementations are possible. For configuration information, see "Configuring Class of Service" on page 283.

DSCP and DSCP IPv6 Alias	DSCP and DSCP IPv6 Bits	Forwarding Class	PLP	Queue
ef	101110	expedited-forwarding	low	1
af11	001010	assured-forwarding	low	2
af12	001100	assured-forwarding	high	2
af13	001110	assured-forwarding	high	2
af21	010010	best-effort	low	0
af22	010100	best-effort	low	0
af23	010110	best-effort	low	0
af31	011010	best-effort	low	0
af32	011100	best-effort	low	0
af33	011110	best-effort	low	0
af41	100010	best-effort	low	0
af42	100100	best-effort	low	0
af43	100110	best-effort	low	0
be	000000	best-effort	low	0

#### Table 113: Sample Behavior Aggregate Classification Forwarding Classes and Queues

DSCP and DSCP IPv6 Alias	DSCP and DSCP IPv6 Bits	Forwarding Class	PLP	Queue
cs1	0010000	best-effort	low	0
cs2	010000	best-effort	low	0
cs3	011000	best-effort	low	0
cs4	100000	best-effort	low	0
cs5	101000	best-effort	low	0
nc1/cs6	110000	network-control	low	3
nc2/cs7	111000	network-control	low	3
other	_	best-effort	low	0

#### Table 113: Sample Behavior Aggregate Classification Forwarding Classes and Queues (continued)

#### **Transmission Scheduling on J-series Services Routers**

The packets in a queue are transmitted based on their transmission priority, transmit rate, and the available bandwidth.

By default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues. When a forwarding class does not fully use the allocated transmission bandwidth, the remaining bandwidth can be used by other forwarding classes if they receive a larger amount of offered load than the bandwidth allocated. A queue receiving traffic within its bandwidth configuration is considered to have positive bandwidth credit, and a queue receiving traffic in excess of its bandwidth allocation is considered to have negative bandwidth credit.

A queue with positive credit does not need to use leftover bandwidth, because it can use its own allocation. For such queues, packets are transmitted based on the priority of the queue, with packets from higher-priority queues transmitting first. The transmit rate is not considered during transmission. In contrast, a queue with negative credit needs a share of the available leftover bandwidth.

On J-series Services Routers, the leftover bandwidth is allocated to queues with negative credit in proportion to the configured transmit rate of the queues within a given priority set. The queues for an interface are divided into sets based on their priority. For more information, see "Scheduling Priority" on page 269. If no transmit rate is configured, each queue in the set receives an equal percentage of the leftover bandwidth. However, if a transmit rate is configured, each queue in the set receives the configured percentage of the leftover bandwidth.

Table 114 on page 281 shows a sample configuration of priority and transmit rate on six queues. The total available bandwidth on the interface is 100 Mbps.

Queue	Scheduling Priority	Transmit Rate	Incoming Traffic
0	Low	10%	20 Mbps
1	High	20%	20 Mbps
2	High	30%	20 Mbps
3	Low	30%	20 Mbps
4	Medium-high	No transmit rate configured	10 Mbps
5	Medium-high	No transmit rate configured	20 Mbps

#### **Table 114: Sample Transmission Scheduling**

In this example, queues are divided into three sets based on their priority:

- High priority set—Consists of queue 1 and queue 2. Packets use 40 Mbps (20 + 20) of the available bandwidth (100 Mbps) and are transmitted first. Because of positive credit, the configured transmit rate is not considered.
- Medium-high priority set—Consists of queue 4 and queue 5. Packets use 30 Mbps (10 + 20) of the remaining 60 Mbps bandwidth. Because of positive credit, the transmit rate is not considered. If the queues had negative credit, they would receive an equal share of the leftover bandwidth because no transmit rate is configured.
- Low priority set—Consists of queue 0 and queue 3. Packets share the 20 Mbps of leftover bandwidth based on the configured transmit rate. The distribution of bandwidth is in proportion to the assigned percentages. Because the total assigned percentage is 40 (10 + 30), each queue receives a share of bandwidth accordingly. Thus queue 0 receives 5 Mbps (10/40 x 20), and queue 3 receives 15 Mbps (30/40 x 20).

J-series[™] Services Router Advanced WAN Access Configuration Guide

# Chapter 15 Configuring Class of Service

You configure class of service (CoS) when you need to override the default packet forwarding behavior of a Services Router—especially in the three areas identified in Table 115 on page 283.

#### Table 115: Reasons to Configure Class of Service (Cos)

Default Behavior to Override with CoS	CoS Configuration Area
Packet classification—By default, the Services Router does not use behavior aggregate (BA) classifiers to classify packets. Packet classification applies to incoming traffic.	Classifiers
Scheduling queues—By default, the Services Router has only two queues enabled. Scheduling queues apply to outgoing traffic.	Schedulers
Packet headers—By default, the Services Router does not rewrite CoS bits in packet headers. Rewriting packet headers applies to outgoing traffic.	Rewrite rules

You can use either J-Web Quick Configuration or a configuration editor to configure CoS. This chapter contains the following topics. For more information about CoS, see the *JUNOS Class of Service Configuration Guide*.

- Before You Begin on page 283
- Configuring CoS with Quick Configuration on page 284
- Configuring CoS Components with a Configuration Editor on page 305
- Configuring Strict High Priority for Queuing with a Configuration Editor on page 332
- Configuring Large Delay Buffers with a Configuration Editor on page 340
- Verifying a CoS Configuration on page 345

## **Before You Begin**

Before you begin configuring a Services Router for CoS, complete the following tasks:

- If you do not already have a basic understanding of CoS, read "Class-of-Service Overview" on page 263.
- Determine whether the Services Router needs to support different traffic streams, such as voice or video. If so, CoS helps to make sure this traffic receives more than basic best-effort packet delivery service.
- Determine whether the Services Router is directly attached to any applications that send CoS-classified packets. If no sources are enabled for CoS, you must configure and apply rewrite rules on the interfaces facing the sources.
- Determine whether the Services Router must support assured forwarding (AF) classes. Assured forwarding usually requires random early detection (RED) drop profiles to be configured and applied.
- Determine whether the Services Router must support expedited forwarding (EF) classes with a policer. Policers require you to apply a burst size and bandwidth limit to the traffic flow, and set a consequence for packets that exceed these limits—usually a high loss priority, so that packets exceeding the policer limits are discarded first.

# **Configuring CoS with Quick Configuration**

The Class of Service Quick Configuration pages allow you to configure most of the JUNOS CoS components for the IPv4, IPv6, and MPLS traffic on a Services Router. You can configure forwarding classes for transmitting packets, define which packets are placed into each output queue, schedule the transmission service level for each queue, and manage congestion using a random early detection (RED) algorithm. After defining the CoS components you must assign classifiers to the required physical and logical interfaces.

This section contains the following topics:

- Defining CoS Components on page 284
- Assigning CoS Components to Interfaces on page 302

## **Defining CoS Components**

Using the Class of Service Quick Configuration pages, you can configure various CoS components individually or in combination to define particular CoS services. For a description of different CoS components, see "JUNOS CoS Components" on page 266.

Figure 23 on page 285 shows the initial Quick Configuration page for CoS that displays the CoS components.

Monitor	Configuration	Diagnose	Manage	Events	Alarms	Logged in as: reg	ress Help	About Loge
Guick Configura	tien P					Configuration	> Quick Configur	ation > <u>Class of Ser</u>
View and Edit	Þ	Quick Cor	nfigurati	on				
History		Class of	Service					
Rescue		b. Cock Malao	AN					
		Cos value / Define	Allases Class of Se	rvice valu	e aliacec	A CoS value aliaci	s a name you	accion to a
		DiffSer	v Code Poi ence value.	nt (DSCP)	value, a [	SCP IPv6 value, M	IPLS EXP bits, o	or an IPv4
		Forwarding	) Classes					
		Define i number	forwarding '	classes by	y assignin(	) each forwarding c	lass to an inter	mal queue
		Classifiers						
		Define and a lo	classifiers ( oss priority	hat allow based on	you to ass code-poin	ociate incoming pa t.	ckets with a for	warding class
		Rewrite Ru	les					
		Define packets	rewrite rule depending	es that allo ; on the fo	w you to rwarding	edefine the code-p class and loss prior	oint value of o ity.	utgoing
		Schedulers	•					
		Define control control	schedulers parameter status, and	which allo s. Schedul I RED drog	w you to o iers define profiles t	configure transmiss the priority, bandy o be applied to a p	ion scheduling vidth, delay buf articular class o	and rate ifer size, rate of traffic.
		Virtual Cha	nnel Group	\$				
		Define you to s interfac	virtual char setup queu ses.	nnel group eing, pack	s which al et schedu	low you to set up a ing, and accounting	pseudo group g rules to multi	that will allow ple logical
		Assign to li	nterfaces					
		Assign maps to forward	Class of Se physical of ling classes	rvice com or logical in	ponents to nterfaces.	interfaces. Allows Allows assignment	assignment of of classifiers, r	scheduler ewrite rules,

#### Figure 23: Initial Class of Service Quick Configuration Page

To configure CoS components with Quick Configuration:

- 1. In the J-Web interface, select **Configuration > Quick Configuration > Class of Service**.
- 2. On the Class of Service Quick Configuration page, select one of the following options depending on the CoS component that you want to define. Enter information into the pages as described in the respective table:

- To define or edit CoS value aliases, select **CoS Value Aliases** and see "Defining CoS Value Aliases" on page 286.
- To define or edit forwarding classes and assign queues, select **Forwarding Classes** and see "Defining Forwarding Classes" on page 288.
- To define or edit classifiers, select **Classifiers** and see "Defining Classifiers" on page 290.
- To define or edit rewrite rules, select **Rewrite Rules** and see "Defining Rewrite Rules" on page 292.
- To define or edit schedulers, select **Schedulers** and see "Defining Schedulers" on page 294.
- To define or edit virtual channel groups, select **Virtual Channel Groups** and see "Defining Virtual Channel Groups" on page 300.
- 3. Click one of the following buttons after completing configuration on any Quick Configuration page:
  - To apply the configuration and stay in the current Quick Configuration page, click **Apply**.
  - To apply the configuration and return to the previous Quick Configuration page, click **OK**.
  - To cancel your entries and return to the previous Quick Configuration page, click **Cancel**.
- 4. Go on to one of the following procedures:
  - To assign CoS components to interfaces, see "Assigning CoS Components to Interfaces" on page 302.
  - To verify the CoS configuration, see "Verifying a CoS Configuration" on page 345.

# **Defining CoS Value Aliases**

Figure 24 on page 287 shows the initial Quick Configuration page for defining aliases for CoS values, and Table 116 on page 287 describes the related fields. By defining aliases you can assign meaningful names to a particular set of bit values and refer to them when configuring CoS components. For more information about CoS values and aliases, see "Default CoS Values and Aliases" on page 273.

Idek Configuration       Configuration > Quick Configuration         Story       Class of Service         Scue       DSCP       DSCP IPv6       MPLS EXP       IPv4 Precedence         Alias       Default       Configured Value         af11       001010       af12       001100         af12       00100       af12       010100         af21       010010       af22       010100         af22       010100       af22       010100         af22       010100       af22       010100         af21       010010       af22       010100         af22       010100       af22       010100         af21       010010       af22       010100         af22       010100       af22       010100         af21       010100       af22       010100         af21       010100       af22       010100       af22         0       af2       010100       af24       af24         0       af44       af44       af44       af44	Monitor Configuration	Diag	nose	Manage	Events	Alarms	Logged in as: regress	Help	About	Log
Alias       DSCP       DSCP IPv6       MPLS EXP       IPv4 Precedence         Alias       Default       Configured Value         af11       00100       Image: af12       001100         af12       001100       Image: af22       Image: af22       Image: af22         image: af22       01000       Image: af22       Image: af22       Image: af22       Image: af22         image: af22       Image: af22       Image: af22       Image: af22       Image: af22       Image: af22         image: af22       Image: af22       Image: af22       Image: af22       Image: af22       Image: af22       Image: af22         image: af22       Image: af22       Image: af22       Image: af22       Image: af22       Image: af22         image: af22       Image: af22       Image: af22       Image: af22       Image: af22       Image: af22         image: af23       Image: af22       Image: af22       Image: af22       Image: af22       Image: af22         image: af23       Image: af22       Image: af22       Image: af22       Image: af22       Image: af22         image: af24       Image: af22       Image: af22 <t< th=""><th>uick Configuration 🔷 📍</th><th></th><th></th><th></th><th></th><th></th><th>Configuration &gt; Qu</th><th>vick Configurat</th><th>tion &gt; Class</th><th>s of Se</th></t<>	uick Configuration 🔷 📍						Configuration > Qu	vick Configurat	tion > Class	s of Se
Scue         DSCP         DSCP IPv6         MPLS EXP         IPv4 Precedence           Alias         Default         Configured Value           af11         00100         af12         001100           af13         001110         af12         01000           af21         01000         af22         010100           af22         010100         af22         010100           af21         010100         af22         af21         01000           af22         010100         af22         af21         010100           af22         010100         af22         af21         010100           af22         010100         af22         af21         af21         af21           af21         010100         af22         af21	ew and Edit 📃 🕨	Quic	k Conf	iguratio	on					
Scue         DSCP         DSCP IPv6         MPLS EXP.         IPv4 Precedence           Alias         Default Value         Configured Value           af11         001010         Image: Configured Value           af12         001100         Image: Configured Value           af12         001100         Image: Configured Value           af12         001100         Image: Configured Value           af12         010100         Image: Configured Value           af21         010010         Image: Configured Value           af22         010100         Image: Configured Value           Image: Configured Value         Image: Configured Value         Image: Configured Value           Image: Configured Value         Image: Configured Value         Image: Configured Value           Image: Configured Value         Image: Configured Value         Image: Configured Value           Image: Configured Value         Image: Configured Value         Image: Configured Value           Image: Configured Value         Image: Configured Value         Image: Configured Value           Image: Configured Value         Image: Configured Value         Image: Configured Value           Image: Configured Value         Image: Configured Value         Image: Configured Value           Image: Configured Value	story	Clas	s of S	ervice						
Alias Name         Default Value         Configured Value           af11         00100            af12         001100            af13         001110            af21         010010            af22         010100            af22         010100            af22         010100            af22         11000            cs7         111000            nc1         110000	escue	DS	SCP	DSCP I	Pv6	MPLS EXP	IPv4 Precedence			
Image: af11       00100         Image: af12       001100         Image: af13       001110         Image: af12       010010         Image: af22       010100         Image: af23       01110         Image: af24       01110         Image: af24       011100         Image: af24       01000         Image: af24       01000         Image: af24       01000         Image: af25       01000         Image: af26       011000         Image: af26       011000         Image: af26       011000			Alias Name	Defaul Value	t (	Configured V	/alue			
af12       001100         af13       001110         af21       010010         af22       010100         cs7       111000         ef       101110         nc1       110000         nc2       111000			af11	001010						
af13       001110         af21       010010         af22       010100         cs7       111000         ef       101110         nc1       110000         nc2       111000			af12	001100						
af21       010010         af22       010100         cs7       111000         ef       101110         nc1       110000         nc2       111000			af13	001110						
af22       010100         cs7       111000         ef       101110         nc1       110000         nc2       111000		Π	af21	010010						
cs7       111000         ef       101110         nc1       110000         nc2       111000			af22	010100						
ef         101110           nc1         110000           nc2         111000		Π	cs7	111000						
nc1         110000           nc2         111000			ef	101110						
□ nc2 111000			nc1	110000						
Add			nc2	111000						
	I - I - I - I - I	Ad	ld							
			ок с	ancel	Apply					

# Figure 24: CoS Value Aliases Quick Configuration Page

# Table 116: CoS Value Aliases Quick Configuration Pages Summary

Field	Function	Your Action
CoS Value Alias Summary		
DSCP	Allows you to define aliases for DiffServ code point (DSCP) IPv4 values.	To define an alias for a DSCP value, click <b>DSCP</b> .
	You can refer to these aliases when you configure classes and define classifiers.	
DSCP IPv6	Allows you to define aliases for DSCP IPv6 values.	To define an alias for a DSCP IPv6 value, click <b>DSCP IPv6</b> .
	You can refer to these aliases when you configure classes and define classifiers.	
MPLS EXP	Allows you to define aliases for MPLS experimental (EXP) bits.	To define an alias for a set of MPLS EXP bits, click <b>MPLS EXP</b> .
	You can map MPLS EXP bits to the Services Router forwarding classes.	

Table 116: CoS	Value Aliases	<b>Ouick Config</b>	uration Pages	Summary	(continued)
10010 1101 000	Value Allages	Quion voining	aradion r ages	Sammary	(continucu)

Field	Function	Your Action
IPv4 Precedence	Allows you to define aliases for IPv4 precedence values.	To define an alias for an IPv4 precedence value, click <b>IPv4 Precedence</b> .
	Precedence values are modified in the IPv4 type-of-service (TOS) field and mapped to values that correspond to levels of service.	
Alias Name	Displays names given to CoS values—for example, <b>af11</b> or <b>be</b> .	None.
Default Value	Displays the default values mapped to standard aliases. For example, <b>ef</b> (expedited forwarding) is a standard alias for DSCP bits <b>101110</b> .	None.
	You cannot delete default values. The check box next to these values is unavailable.	
Configured Value	Displays the CoS values that you have assigned to specific aliases.	None.
	You can delete a configured alias.	
Add	Opens a page that allows you to define CoS value aliases.	To add a CoS value alias, click <b>Add</b> .
Delete	Allows you to delete a configured CoS value alias.	To delete a CoS value alias, select the check box next to it and click <b>Delete</b> .
	You cannot delete a default alias.	
Add a CoS Value Alias		
CoS Value Alias	Assigns a name to a CoS value. A CoS value can be of different types—DSCP, DSCP IPv6, IP precedence, or MPLS EXP.	To define an alias for a CoS value, type a name—for example, my1.
CoS Value Alias Bits	Specifies the CoS value for which an alias is defined.	To specify a CoS value, type it in an appropriate format:
	Changing this value alters the behavior of all classifiers that refer to this alias.	<ul> <li>For DSCP and DSCP IPv6 CoS values, use the format xxxxx, where x is 1 or 0—for example, 101110.</li> <li>For MPLS EXP and IP precedence CoS uplying the format the format type of the form</li></ul>
		values, use the format xxx, where x is 1 or 0—for example, <b>111</b> .

# **Defining Forwarding Classes**

Figure 25 on page 289 shows the initial Quick Configuration page for defining forwarding classes and assigning them to queues, and Table 117 on page 289 describes the related fields. By assigning a forwarding class to a queue number, you affect the scheduling and marking of a packet as it transits a Services Router. For more

information about forwarding classes and queues, see "JUNOS CoS Components" on page 266.

Monitor C	Configuration Dia	ignose Manage	Events Alarms Logged in as: regress Configuration > Quick	Help About Log Configuration > Class of Se
iew and Edit	Quie	ck Configurati	ion	
story	Cla	ss of Service	•	
		Queue #	Forwarding Class Name	ingening stell bolom
		Queue #	Forwarding Class Name	
		0	best-effort	
		0	expedited-forwarding	
		0 1 2	expedited-forwarding assured-forwarding	
		0 1 2 3	best-effort expedited-forwarding assured-forwarding network-control	

## Figure 25: Forwarding Classes Quick Configuration Page

## Table 117: Forwarding Classes Quick Configuration Pages Summary

Field	Function	Your Action
Forwarding Class Summa	ry	
Queue #	Displays internal queue numbers to which forwarding classes are assigned.	To edit an assigned forwarding class, click the queue number to which the class is assigned.
	By default, if a packet is not classified, it is assigned to the class associated with queue 0.	
	Allows you to edit an assigned forwarding class.	
Forwarding Class Name	Displays the forwarding class names assigned to specific internal queue numbers.	None.
	By default, four forwarding classes are assigned to queue numbers 0 through 3.	
Add	Opens a page that allows you to assign forwarding classes to internal queue numbers.	To add a forwarding class, click <b>Add</b> .
Delete	Deletes an internal queue number and the forwarding class assigned to it.	To delete a queue number, click the check box next to it and click <b>Delete</b> .
Add a Forwarding Class/E	dit Forwarding Class Queue #	

Field	Function	Your Action
Queue #	Specifies the internal queue number to which a forwarding class is assigned.	To specify an internal queue number, type an integer from 0 through 7, as supported by your platform.
Forwarding Class Name	Specifies the forwarding class name assigned to the internal queue number.	To assign a forwarding class name to a queue, type the name—for example, <b>be-class</b> .

#### Table 117: Forwarding Classes Quick Configuration Pages Summary (continued)

## **Defining Classifiers**

Figure 26 on page 290 shows the initial Quick Configuration page for defining classifiers, and Table 118 on page 290 describes the related fields. Classifiers examine the CoS value or alias of an incoming packet and assign it a level of service by setting its forwarding class and loss priority. For more information about classifiers, see "Default Behavior Aggregate Classifiers" on page 277.

#### **Figure 26: Classifiers Quick Configuration Page**

Monitor	Configuration	Diagnose	Manage	Events Alarms Log	gged in as: regress	Help About	Log
lick Configura	ation 🎽				Contiguration > Quick (	configuration > Class	of Se
ew and Edit	•	Quick Con	figurati	on			
tory		Class of \$	Service				
scue					*****	******	
		DSCP	DSCP I	Pv6 MPLS EXP	IPv4 Precedence		
		Class Name	ifier	Incoming Code Point (Alias)	Classify to Forwarding Class	Classify to L Priority	.055
		D ba-sgo	<u>lhafs</u>	010111	best-effort	low	
		Add	Delete		I		

#### Table 118: Classifiers Quick Configuration Page Summary

Field Function		Your Action
<b>Classifier Summary</b>		
DSCP	Allows you to define classifiers for DSCP IPv4 values.	To define a classifier for a DSCP code point value, click <b>DSCP</b> .
DSCP IPv6	Allows you to define classifiers for DSCP IPv6 values.	To define a classifier for a DSCP IPv6 value, click <b>DSCP IPv6</b> .

Field	Function	Your Action
MPLS EXP	Allows you to define classifiers for MPLS experimental (EXP) bits.	To define a classifier for a set of MPLS EXP bits, click <b>MPLS EXP</b> .
IPv4 Precedence	Allows you to define classifiers for IPv4 precedence values.	To define a classifier for an IP precedence value, click <b>IPv4 Precedence</b> .
Classifier Name	Displays the names of classifiers.	To edit a classifier, click its name.
	Allows you to edit a specific classifier.	
Incoming Code Point (Alias)	Displays CoS values and aliases to which forwarding class and loss priority are mapped.	None.
Classify to Forwarding Class	Displays forwarding classes that are assigned to specific CoS values and aliases of a classifier.	None.
Classify to Loss Priority	Displays loss priorities that are assigned to specific CoS values and aliases of a classifier.	None.
Add	Opens a page that allows you to define classifiers.	To add a classifier, click <b>Add</b> .
Delete	Deletes a specified classifier.	To delete a classifier, locate the classifier, select the check box next to it, and click <b>Delete</b> .
Add a Classifier/Edit Clas	ssifier	
Classifier Name	Specifies the name for a classifier.	To name a classifier, type the name—for example, <b>ba-classifier</b> .
Classifier Code Point Mapping	Sets the forwarding classes and the packet loss priorities (PLPs) for specific CoS values and aliases.	None.
Incoming Code Point	Specifies the CoS value in bits and the alias of a classifier for incoming packets.	To specify a CoS value and alias, either select preconfigured ones from the list or type new ones.
		For information about forwarding classes and aliases assigned to well-known DSCPs, see Table 112 on page 278.

# Table 118: Classifiers Quick Configuration Page Summary (continued)

Field	Function	Your Action
Forwarding Class	Assigns the forwarding class to the specified CoS value and alias.	To assign a forwarding class, select either one of following default forwarding classes, or one that you have configured:
		<ul> <li>best-effort—Provides no special CoS handling of packets. Typically, RED drop profile is aggressive and no loss priority is defined.</li> </ul>
		<ul> <li>expedited-forwarding—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service. Packets can be forwarded out of sequence or dropped.</li> </ul>
		<ul> <li>assured-forwarding—Provides high assurance for packets within specified service profile. Excess packets are dropped.</li> </ul>
		<ul> <li>network-control—Packets can be delayed but not dropped.</li> </ul>
Loss Priority	Assigns a loss priority to the specified CoS value and alias.	To assign a loss priority, select one of the following:
		■ <b>low</b> —Packet has a low loss priority.
		■ high—Packet has a high loss priority.
		<ul> <li>medium-low—Packet has a medium-low loss priority.</li> </ul>
		<ul> <li>medium-high-Packet has a medium-high loss priority.</li> </ul>
Add	Assigns a forwarding class and loss priority to the specified CoS value and alias.	To assign a forwarding class and loss priority to a specific CoS value and alias, click <b>Add</b> .
	A classifier examines the incoming packet's header for the specified CoS value and alias and assigns it the forwarding class and loss priority that you have defined.	
Delete	Removes the forwarding class and loss priority assignment from the classifier.	To remove the forwarding class and loss priority assignment, select it and click <b>Delete</b> .

#### Table 118: Classifiers Quick Configuration Page Summary (continued)

# **Defining Rewrite Rules**

Figure 27 on page 293 shows the initial Quick Configuration page for defining rewrite rules, and Table 119 on page 293 describes the related fields. Use the rewrite rules to alter the CoS values in outgoing packets to meet the requirements of the targeted peer. A rewrite rule examines the forwarding class and loss priority of a packet and sets its bits to a corresponding value specified in the rule.

Monitor Configuration	n Diag	nose	Manage	Events	Alarma Log	ged in as: n Configurati	egress on > Quick (	Configuration > Class	Logo s of Serv
View and Edit	Quick Configuration								
History	Class of Service								
Rescue									10101010101641
	D	SCP	DSCP I	Pv6	MPLS EXP	IPv4 Pre	cedence		
		Rewr Nami	ite Rule e	For	warding Class	Loss Priority	Rewrit Point	te Outgoing Cod To	le
		re-ef-	<u>class</u>	exp forw	edited- arding	low	001010	0 (afii)	
		<u>foo</u>		best	-effort	high	101110	) (ef)	
		re he	dage	assu	ured-forwarding	low	101110	) (ef)	
		te-be-	-01435	assu	red-forwarding	high	001010	0 (af11)	
		Id	Delete Cancel	Apply					

# Figure 27: Rewrite Rules Quick Configuration Page

# Table 119: Rewrite Rules Quick Configuration Page Summary

Field	Function	Your Action
<b>Rewrite Rules Summary</b>		
DSCP	Allows you to redefine DSCP IPv4 code point values of outgoing packets.	To redefine a DSCP code point value, click <b>DSCP</b> .
DSCP IPv6	Allows you to redefine DSCP IPv6 code point values.	To redefine a DSCP IPv6 code point value, click <b>DSCP IPv6</b> .
MPLS EXP	Allows you to redefine MPLS experimental (EXP) bits.	To redefine MPLS EXP bits, click <b>MPLS EXP</b> .
IPv4 Precedence	Allows you to redefine IPv4 precedence code point values.	To redefine an IPv4 precedence code point value, click <b>IPv4 Precedence</b> .
Rewrite Rule Name	Displays names of defined rewrite rules.	To edit a rule, click its name.
	Allows you to edit a specific rule.	
Forwarding Class	Displays forwarding classes associated with a specific rewrite rule.	None.
Loss Priority	Displays loss priority values associated with a specific rewrite rule,	None.
Rewrite Outgoing Code Point To	Displays the CoS values and aliases that a specific rewrite rule has set for a specific forwarding class and loss priority.	None.

Field	Function	Your Action
Add	Opens a page that allows you to define a new rewrite rule.	To add a rewrite rule, click <b>Add</b> .
Delete	Removes specified rewrite rules.	To remove a rule, select the check box next to it and click <b>Delete</b> .
Add a Rewrite Rule/Edi	it Rewrite Rule	
Rewrite Rule Name	Specifies a rewrite rule name.	To name a rule, type the name—for example, rewrite-dscps.
Code Point Mapping	Rewrites outgoing CoS values of a packet, based on the forwarding class and loss priority.	To configure the CoS value assignment, follow these steps:
	Allows you to remove a Code Point Mapping entry.	1. From the Forwarding Class list, select a class.
	ç	2. Select a priority from the following:
		■ <b>low</b> —Rewrite rule applies to packets with a low loss priority.
		■ <b>high</b> —Rewrite rule applies to packets with a high loss priority.
		<ul> <li>medium-low—Rewrite rule applies to packets with a medium-low loss priority.</li> </ul>
		<ul> <li>medium-high-Rewrite rule applies to packets with a medium-high loss priority.</li> </ul>
		<ol> <li>For Rewritten Code Point, either select a predefined CoS value and alias or type a new CoS value and alias.</li> </ol>
		For information about predefined CoS values and aliases, see Table 110 on page 274.
		4. Click <b>Add</b> .
		To remove a code point mapping entry, select it and click <b>Delete</b> .

#### Table 119: Rewrite Rules Quick Configuration Page Summary (continued)

#### **Defining Schedulers**

Figure 28 on page 295 shows the initial Quick Configuration page for defining schedulers, scheduler maps, and random early detection (RED) drop profiles. Using schedulers, you can assign attributes to queues and thereby provide congestion control to a particular class of traffic. These attributes include the amount of interface bandwidth, memory buffer size, transmit rate, RED drop profiles and priority.

To configure schedulers using the Quick Configuration pages:

- 1. Create a drop profile by specifying the fill levels and drop probabilities. The drop profile map on the Scheduler page uses this drop profile. For a description of RED drop profile-related fields, see Table 120 on page 295.
- 2. Create a scheduler and specify attributes to it. For a description of scheduler-related fields, see Table 121 on page 297.
- 3. Associate the scheduler to a forwarding class. Because the forwarding class is assigned to a queue number, the queue inherits this scheduler's attributes. For a description of scheduler map-related fields, see Table 122 on page 299.

#### Figure 28: Schedulers Quick Configuration Page

Monitor	Configuration	Diag	nose	Manage	Events	Alarms	Logged in as: regress	Help	About	Logout
Zuick Configure View and Edit	ntion 🕨	Quic	k Con	figuratio	n		Confiduration > Quick	Configuratio	<u>n</u> > <u>cias</u>	or servic
History		Clas	s of S	Service						
Rescue										
		Sc	hedule	rs S	cheduler	Maps	RED Drop Profiles			
			Sched	uler Name	Sched	uler Info	ormation			
		<b>1001</b>		Buffer Schedu Transn Shapin	Buffer Size: 90% Schedule Priority: medium-high Transmit Rate: 20% Shaping Rate: 90%					
			<u>1002</u>		Buffer Schedu Transn Shapin	Size: 819 ule Priorit nit Rate: 3 g Rate: 5	92 microseconds (tempora y: low 20% %	I)		
		Ad	d [	Delete						
		. (	ок 🛛 (	Cancel /	Apply					

#### Table 120: RED Drop Profiles Quick Configuration Page Summary

Field	Function	Your Action			
<b>RED Drop Profiles Summa</b>	RED Drop Profiles Summary				
RED Drop Profile Name	Displays the configured random early detection (RED) drop profile names.	To edit a RED drop profile, click its name.			
	RED attempts to avoid congestion by dropping packets from the head of a queue.				
	Allows you edit a specific drop profile.				
Graph RED Profile	Opens a new window and displays a graph for a specific RED drop profile.	To view the graph for a specific RED drop profile, click <b>Graph</b> .			

# Table 120: RED Drop Profiles Quick Configuration Page Summary (continued)

Field	Function	Your Action
RED Drop Profile Information (Fill Level, Drop Probability)	Displays information about the data point type, the queue buffer fill level, and the drop probability for specific RED drop profiles.	None.
Add	Opens a page that allows you to add a RED drop profile.	To add a RED drop profile, click <b>Add</b> .
Delete	Removes a RED drop profile.	To remove a RED drop profile, select it and click <b>Delete</b> .
Add a RED Drop Profile/E	dit RED Drop Profile	
Graphed RED Profile	Displays a graph of RED drop profiles. Each data point in this graph is defined by a pair of x and y coordinates and represents the relationship between them.	None.
	The x axis represents the queue buffer fill level, which is a percentage value of how full the queue is.	
	The y axis represents the drop probability, which is a percentage value of the chances of a packet being dropped.	
Drop Profile Name	Specifies a name for a drop profile.	To name a drop profile, type the name—for example, be-normal.
	A drop profile consists of pairs of values between 0 and 100, one for queue buffer fill level and one for drop probability, that determine the relationship between a buffer's fullness and the likelihood it will drop packets. The values you assign to each pair must increase relative to the previous pair of values. With a few value pairs the system automatically constructs a drop profile.	
RED Drop Profile Type	Specifies whether a RED drop profile type is interpolated or segmented.	To specify a RED drop profile type, select one of the following:
	For more information about segmented and interpolated drop profiles, see the <i>JUNOS Class</i> of Service Configuration Guide.	<ul> <li>Interpolated—The value pairs are interpolated to produce a smooth profile.</li> <li>Segmented—The value pairs are represented by line fragments, which connect each data point on the graph to produce a segmented profile.</li> </ul>

Field	Function	Your Action
Data Points	<ul> <li>Specifies the points for generating the RED drop profile graph. Each data point is defined by a pair of x and y coordinates and represents the relationship between them.</li> <li>The x axis represents the queue buffer fill level, which is a percentage value of how full the queue is. A value of 100 means the queue is full.</li> <li>The y axis represents the drop probability, which is a percentage value of the chances of a packet being dropped. A value of 0 means that a packet is never dropped, and a value of 100 means that all packets are dropped.</li> </ul>	<ul> <li>To specify x and y coordinates for data points, type a number between 0 and 100 in the following boxes:</li> <li>Fill level—Type the percentage value of queue buffer fullness for the x coordinate—for example, 95.</li> <li>Drop profile—Type the percentage value of drop probability for the y coordinate—for example, 85.</li> </ul>
Add	Adds the specified queue buffer fill level and drop probability as a data point for the graph.	To add the specified fill level and drop probability, click <b>Add</b> .
Delete	Removes a data point.	To remove a data point, select it and click <b>Delete</b> .

# Table 120: RED Drop Profiles Quick Configuration Page Summary (continued)

## Table 121: Schedulers Quick Configuration Page Summary

Field	Function	Your Action
Scheduler Summary		
Scheduler Name	Displays the names of defined schedulers.	To edit a scheduler, click its name.
	Allows you to edit a specific scheduler.	
Scheduler Information	Displays a summary of defined settings for a scheduler, such as bandwidth, delay buffer size, transmit and shaping rates, and RED drop profiles.	None.
Add	Opens a page that allows you to adds a scheduler.	To add a scheduler, click <b>Add</b> .
Delete	Removes a scheduler.	To remove a scheduler, select it and click <b>Delete</b> .
Add a Scheduler/Edit Scl	heduler	
Scheduler Name	Specifies the name for a scheduler.	To name a scheduler, type the name—for example, <b>be-scheduler</b> .

Field	Function	Your Action
Buffer Size	Defines the size of the delay buffer. The delay buffer bandwidth provides packet buffer space to absorb burst traffic up to the specified duration of delay. By default, queues 0 through 7 have the following percentage of the total available buffer space: Queue 0—95 percent Queue 1—0 percent Queue 2—0 percent Queue 3—5 percent Queue 4—0 percent Queue 6—0 percent Queue 7—0 percent Augueue 7—0 percent NOTE: A large buffer size value means a greater possibility for delaying packets in the network. This might not be practical for sensitive traffic such as voice or video.	<ul> <li>To define a delay buffer size for a scheduler, select the appropriate option:</li> <li>To specify no buffer size, select Unconfigured.</li> <li>To specify buffer size as a percentage of the total buffer, select Percent and type an integer from 1 through 100.</li> <li>To specify buffer size as the remaining available buffer, select Remainder.</li> <li>To specify buffer size in microseconds, select Temporal, and type an integer within the range of the buffer size available to you on your platform—for example, 8192.</li> </ul>
Drop Profile Map	Sets the drop profile for a specific packet loss priority (PLP) and protocol type. By default, the drop profile is assigned to packets with low PLP, regardless of protocol type.	<ul> <li>To configure a scheduler drop profile:</li> <li>Select a loss priority from the following: <ul> <li>low—Drop profile applies to packets with a low loss priority.</li> <li>medium-low—Drop profile applies to packets with a medium-low loss priority.</li> <li>high—Drop profile applies to packets with a high loss priority.</li> <li>medium-high—Drop profile applies to packets with a medium-high loss priority.</li> <li>any—Drop profile applies to all packets irrespective of the loss priority.</li> </ul> </li> <li>From the Protocol list, select a profile.</li> <li>Click Add.</li> </ul>

# Table 121: Schedulers Quick Configuration Page Summary (continued)

Field	Function	Your Action
Scheduling Priority	Sets the transmission priority of the scheduler, which determines the order in which an output interface transmits traffic from the queues. You can set scheduling priority at different levels in an order of increasing priority from low to high. A high-priority queue with a high transmission rate might lock out lower-priority traffic.	<ul> <li>To specify a priority, select one of the following:</li> <li>high—Packets in this queue are transmitted first.</li> <li>low—Packets in this queue are transmitted last.</li> <li>medium-high—Packets in this queue are transmitted after high-priority packets.</li> <li>medium-low—Packets in this queue are transmitted before low-priority packets.</li> </ul>
Shaping Rate	Defines the minimum bandwidth allocated to a queue. The default shaping rate is 100 percent, which is the same as no shaping at all.	<ul> <li>To define a shaping rate, select the appropriate option:</li> <li>To specify no shaping rate, select Unconfigured.</li> <li>To specify shaping rate as an absolute number of bits per second, select Absolute Rate and type an integer from 3200 through 3200000000.</li> <li>To specify shaping rate as a percentage, select Percent and type an integer from 0 through 100.</li> </ul>
Transmit Rate	<ul> <li>Defines the transmission rate of a scheduler.</li> <li>The transmit rate determines the traffic bandwidth from each forwarding class you configure.</li> <li>By default, queues 0 through 7 have the following percentage of transmission capacity: <ul> <li>Queue 0—95 percent</li> <li>Queue 1—0 percent</li> <li>Queue 2—0 percent</li> <li>Queue 3—5 percent</li> <li>Queue 4—0 percent</li> <li>Queue 6—0 percent</li> <li>Queue 7—0 percent</li> </ul> </li> </ul>	<ul> <li>To define a transmit rate, select the appropriate option:</li> <li>To not specify transmit rate, select Unconfigured.</li> <li>To specify the remaining transmission capacity, select Remainder Available.</li> <li>To specify a percentage of transmission capacity, select Percent and type an integer from 1 through 100.</li> <li>To enforce the exact transmission rate or percentage you configured, select the Exact Transmit Rate check box.</li> </ul>

# Table 121: Schedulers Quick Configuration Page Summary (continued)

# Table 122: Scheduler Maps Quick Configuration Page Summary

Field	Function	Your Action
Scheduler Maps Summary	,	
Scheduler Map Name	Displays the names of defined scheduler maps. Scheduler maps link schedulers to forwarding classes.	To edit a scheduler map, click its name.
	Allows you to edit a scheduler map.	

Field	Function	Your Action
Scheduler Map Information	For each map, displays the schedulers and the forwarding classes that they are assigned to.	None.
Add	Opens a page that allows you to add a scheduler map.	To add a scheduler map, click <b>Add</b> .
Delete	Removes a scheduler map.	To remove a scheduler map, select it and click <b>Delete</b> .
Add a Scheduler Map/Edit	t Scheduler Map	
Scheduler Map Name	Specifies the name for a scheduler map.	To name a map, type the name—for example, <b>be-scheduler-map</b> .
Scheduler Mapping	Allows you to associate a preconfigured scheduler with a forwarding class.	To associate a scheduler with a forwarding class, locate the forwarding class and select the scheduler in the box next to it.
	Once applied to an interface, the scheduler maps affect the hardware queues, packet schedulers, and RED drop profiles.	

#### Table 122: Scheduler Maps Quick Configuration Page Summary (continued)

## **Defining Virtual Channel Groups**

Figure 29 on page 300 shows the initial Quick Configuration page for defining virtual channel groups, and Table 123 on page 301 describes the related fields. Use virtual channels to avoid oversubscription of links by limiting traffic from a higher aggregated bandwidth to a lower one—for example, to limit traffic from a main office to branch offices. You channelize this traffic by applying queuing, packet scheduling, and accounting rules to logical interfaces.

#### Figure 29: Virtual Channel Group Quick Configuration Page

Monitor Config	uration	Diagnose	Manage	Events	Alerms Logg	ed in as: req	gress Help	About Log
	•					Configuration	> Quick Configu	ration > Class of Se
iew and Edit	• <u>Q</u>	uick Cor	nfiguratio	n				
istory	C	lass of a	Service					
escue								
		Virtua Group	l Channel Name	Vi N	rtual Channel ame	Default	Scheduler Map	Shaping Rate
			branch1-vc	anch1-vc	Default	myMap1	15%	
		wan-ve	wan-vc-group-1		anch2-vc		myMap2	40k bits per second
		Add	Delete					
				Same Las				
Field	Function	Your Action						
-------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------						
Virtual Channel Groups	Summary							
Virtual Channel Group Name	Displays names of defined virtual channel groups.	To edit a virtual channel group, click its name.						
	Allows you to edit a virtual channel group.							
Virtual Channel Name	Displays names of defined virtual channels.	To edit a virtual channel, click its name.						
	Allows you to edit a virtual channel.							
Default	Marks the default virtual channel of a group.	None.						
	One of the virtual channels in a group must be configured as the default channel. Any traffic not explicitly directed to a particular channel is transmitted by this channel.							
Scheduler Map	Displays the scheduler map assigned to a particular virtual channel.	None.						
Shaping Rate	Displays the shaping rate configured for a virtual channel.	None.						
Add	Opens a page that allows you to add a virtual channel group.	To add a virtual channel group, click <b>Add</b> .						
Delete	Removes a specific virtual channel group.	To remove a specific virtual channel group, locate its name, select the check box next to it, and click <b>Delete</b> .						
Add a Virtual Channel Gr	oup/Edit a Virtual Channel Group							
Virtual Channel Group Name	Specifies a name for a virtual channel group.	To name a group, type the name—for example, wan-vc-group.						
Add	Creates a virtual channel group.	To create a virtual channel group, click Add.						
	Opens a page that allows you to add a virtual channel to the specified group.							
Add a Virtual Channel/E	dit Virtual Channel							
Virtual Channel Name	Specifies the name of a virtual channel to be assigned to a virtual channel group.	To name a virtual channel, either select a predefined name from the list or type a new name—for example, <b>branch1–vc</b> .						
Scheduler Map	Specifies a predefined scheduler map to assign to a virtual channel.	To specify a scheduler map, select it from the Scheduler Map list.						
	Scheduler maps associate schedulers with forwarding classes. For information about how to define scheduler maps, see Table 122 on page 299.							

# Table 123: Virtual Channel Group Quick Configuration Page Summary

Field	Function	Your Action
Shaping Rate	Specifies the shaping rate for a virtual channel.	To specify a shaping rate, select one of the following options:
	The shaper limits the maximum bandwidth transmitted by a virtual channel.	<ul> <li>To specify no shaping rate, select</li> <li>Unconfigured.</li> </ul>
	Configuring a shaping rate is optional. If no shaping rate is configured, a virtual channel without a shaper can use the full logical interface bandwidth.	<ul> <li>To configure a shaping rate as an absolute number of bits per second, select Absolute</li> <li>Rate and type a value between 3200 and 320000000000.</li> </ul>
		<ul> <li>To configure a shaping rate as a percentage, select <b>Percent</b> and type a value between 0 and 100.</li> </ul>

#### Table 123: Virtual Channel Group Quick Configuration Page Summary (continued)

### **Assigning CoS Components to Interfaces**

After you have defined CoS components, you must assign them to logical or physical interfaces. The CoS Quick Configuration pages allow you to assign scheduler maps to physical or logical interfaces and to assign forwarding classes, classifiers, rewrite rules, or virtual channel groups to logical interfaces.

Figure 30 on page 303 shows the initial Quick Configuration page for assigning CoS components to interfaces. The page displays the Services Router interfaces available for CoS component assignment and the status of existing CoS components.

Monitor Configurati	on Di	agnose Manage	e Events	Alarms	Logged in as: regress	Help	About Lo	
disk Configuration 🔶	0.0	ck Configura	tion		<u>conigeration</u> > <u>Gold</u>	comiguratio		
ew and Edit -		Class of Service						
scue								
	Cla	ss of Service	Interfac	es				
		Interface Name	Class	of Service	Overview			
		<u>fe-0/0/0</u>	Schedu	ler Map: <b>m</b>	yMap1			
		fe-0/0/0.0	Forward	ding Class:	assured-forwarding			
		fe-0/0/0.1	Forward	ding Class:	best-effort			
		fe-0/0/0.2	Forward	ding Class:	network-control			
		f <u>e-0/0/1</u>	Schedu	ler Map: <b>m</b>	yMap2			
		<u>fe-0/0/1.0</u>	dscp Cl dscp Re	assifier: <b>d</b> write Rule	efault s: re-ef-class			
		<u>fe-0/0/1.1</u>	dscp Re	write Rule	s: foo			
	A	dd Delete			,			
		OK Cancel	Apply					

#### Figure 30: Assignment of CoS Components to Interfaces Quick Configuration Page

To assign CoS components to interfaces with Quick Configuration:

- 1. In the J-Web interface, select Configuration > Quick Configuration > Class of Service > Assign Class of Service Components to Interfaces.
- 2. Enter information into these Quick Configuration pages, as described in Table 124 on page 304.
- 3. Click one of the following buttons after completing configuration on any Quick Configuration main page:
  - To apply the configuration and stay in current the Quick Configuration page, click **Apply**.
  - To apply the configuration and return to the previous Quick Configuration page, click **OK**.
  - To cancel your entries and return to the previous Quick Configuration page, click **Cancel**.
- 4. To verify the CoS configuration, see "Verifying a CoS Configuration" on page 345.

Field	Function	Your Action	
<b>Class of Service Interface</b>	S		
Interface Name	Lists the names of physical and logical interfaces configured on the system.	To edit an interface's CoS assignments, click the interface.	
conventions in the J-series Services Router Basic LAN and WAN Access Configuration Guide.)	Allows you to edit CoS component assignments to physical and logical interfaces.		
Class of Service Overview	Displays the CoS components assigned to a particular interface—for example, information about DSCP classifiers, EXP classifiers, or DSCP rewrite rules.	None.	
Add	Allows you to add a CoS service to a physical interface.	To add a CoS service to a physical interface, click <b>Add</b> .	
Delete	Removes CoS services assigned to a specific interface.	To remove CoS services assigned to a specific interface, locate the interface name, click the check box next to it, and click <b>Delete</b> .	
Add CoS Service to a Phys	sical Interface/Edit CoS Physical Interface		
Physical Interface Name	Specifies the name of a physical interface. Allows you to assign CoS components to a set	To specify an interface for CoS assignment, type its name in the Physical Interface Name box.	
	of interfaces at the same time.	To specify a set of interfaces for CoS assignment, use the wildcard character (*)—for example, ge-0/*/0.	
Scheduler Map	Specifies a predefined scheduler map for the physical interface.	To specify a map for an interface, select it from the Scheduler Map list.	
	A scheduler map enables the physical interface to have more than one set of output queues.		
	<b>NOTE:</b> For 4-port Fast Ethernet ePIMs, if you apply a CoS scheduler map on outgoing (egress) traffic, the router does not divide the bandwidth appropriately among the CoS queues. As a workaround configure enforced CoS shaping on the ports.		
Add	Allows you to add a CoS service to a logical interface on a specified physical interface.	To add a CoS Service to a logical interface, click <b>Add</b> .	
Add CoS Service to a Logi	cal Interface Unit/Edit CoS Logical Interface U	nit	
Logical Interface Unit Name	Specifies the name of a logical interface.	To specify an interface for CoS assignment, type its name in the Logical Interface Unit	
	logical interfaces configured on a physical interface at the same time.	To assign CoS services to all logical interfaces configured on this physical interface, type the wildcard character (*).	

## Table 124: Assigning CoS Components to Interfaces Quick Configuration Summary

Field	Function	Your Action
Scheduler Map	Specifies a predefined scheduler map for this interface.	To assign a scheduler map to the interface, select it from the list.
	<b>NOTE:</b> You can configure either a scheduler map or a virtual channel group on a logical interface, not both.	
Forwarding Class	Assigns a predefined forwarding class to incoming packets on a logical interface.	To assign a forwarding class to the interface, select it.
Virtual Channel Group	Applies a virtual channel group to a logical interface.	To specify a virtual channel group for the interface, select it from the list.
	Applying a virtual channel group creates a set of eight queues for each virtual channel in the group.	
	<b>NOTE:</b> You can configure either a scheduler map or a virtual channel group on a logical interface, not both.	
Classifiers	Allows you to apply classification maps to a logical interface.	To assign a classification map to the interface, select an appropriate classifier for each CoS value type used on the interface.
	Classifiers assign a forwarding class and loss priority to an incoming packet based on its CoS value.	
Rewrite Rules	Allows you to apply rewrite rule configurations to a logical interface.	To apply a rewrite rule configuration to the interface, select a rule for each CoS value type used on the interface.
	Rewrite rules rewrite the CoS values in an outgoing packet based on forwarding class and loss priority.	
	You can choose to apply your own rewrite rule or a default one. The default rewrite assignments are based on the default bit definitions of DSCP, DSCP IPv6, MPLS EXP, and IP precedence.	

### Table 124: Assigning CoS Components to Interfaces Quick Configuration Summary (continued)

## **Configuring CoS Components with a Configuration Editor**

To configure the Services Router as a node in a network supporting CoS, read the section "Before You Begin" on page 283, determine your needs, and select the tasks you need to perform from the following list. For information about using the J-Web and CLI configuration editors, see the *J-series Services Router Basic LAN and WAN Access Configuration Guide*.

- Configuring a Policer for a Firewall Filter on page 306
- Configuring and Applying a Firewall Filter for a Multifield Classifier on page 307

- Assigning Forwarding Classes to Output Queues on page 310
- Configuring and Applying Rewrite Rules on page 312
- Configuring and Applying Behavior Aggregate Classifiers on page 315
- Configuring RED Drop Profiles for Congestion Control on page 319
- Configuring Schedulers on page 321
- Configuring and Applying Scheduler Maps on page 324
- Configuring and Applying Virtual Channels on page 327
- Configuring and Applying Adaptive Shaping for Frame Relay on page 331

# **Configuring a Policer for a Firewall Filter**

You configure a policer to detect packets that exceed the limits established for expedited forwarding. The packets that exceed these limits are given a higher loss priority than packets within the bandwidth and burst size limits.

The following example shows how to configure a policer called **ef-policer** that identifies for likely discard expedited forwarding packets with a burst size greater than 2000 bytes and a bandwidth greater than 10 percent.

For more information about firewall filters, see "Configuring Stateless Firewall Filters" on page 221 and the *JUNOS Policy Framework Configuration Guide*.

To configure an expedited forwarding policer for a firewall filter for the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 125 on page 306.
- 3. Go on to "Configuring and Applying a Firewall Filter for a Multifield Classifier" on page 307.

### Table 125: Configuring a Policer for a Firewall Filter

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Firewall</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the <b>[edit]</b> hierarchy level, enter
	2. Next to Firewall, click <b>Configure</b> or <b>Edit</b> .	edit firewall
Create the policer for expedited forwarding,	1. Click Add new entry next to Policer.	Enter
and give the policer a name—for example, ef-policer.	2. In the Policer name box, type <b>ef-policer</b> .	edit policer ef-policer

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Set the burst limit for the policer—for	1.	Click <b>Configure</b> next to If exceeding.	Enter
example, 2k. Set the bandwidth limit or percentage for	2.	In the Burst size limit box, type a limit for the burst size allowed—for example, <b>2k</b> .	set if-exceeding burst-limit-size 2k
the bandwidth allowed for this type of traffic—for example, use a bandwidth percent of 10	3.	From the Bandwidth list, select <b>bandwidth-percent</b> .	set if-exceeding
	4.	In the Bandwidth percent box, type 10.	bandwidth-percent 10
	5.	Click <b>OK</b> .	
Enter the loss priority for packets exceeding	1.	Click <b>Configure</b> next to Then.	Enter
the limits established by the policer—for example, high.	2.	From the Loss priority list, select <b>high</b> .	set then loss-priority high
-	3.	Click <b>OK</b> .	

#### Table 125: Configuring a Policer for a Firewall Filter (continued)

# **Configuring and Applying a Firewall Filter for a Multifield Classifier**

You configure a multifield (MF) classifier to detect packets of interest to CoS and assign the packet to the proper forwarding class independently of the DiffServ code point (DSCP). To configure a multifield classifier on a customer-facing or host-facing link, configure a firewall filter to classify traffic. Packets are classified as they arrive on an interface.

One common way to detect packets of CoS interest is by source or destination address. The destination address is used in this example, but many other matching criteria for packet detection are available to firewall filters.

This example shows how to configure the firewall filter mf-classifier and apply it to the Services Router's Gigabit Ethernet interface ge-0/0/0. The firewall filter consists of the rules (terms) listed in Table 126 on page 307.

#### **Table 126: Sample mf-classifier Firewall Filter Terms**

Rule (Term)	Purpose	Contents
assured forwarding	Detects packets destined for <b>192.168.44.55</b> , assigns them to an assured forwarding class, and gives them a low likelihood of being dropped.	Match condition: destination address 192.168.44.55
		Forwarding class: af-class
		Loss priority: low
expedited-forwarding	Detects packets destined for <b>192.168.66.77</b> , assigns them to an expedited forwarding class, and subjects them to the EF policer configured in "Configuring a Policer for a	Match condition: destination address 192.168.66.77
	Firewall Filter" on page 306.	Forwarding class: ef-class
		Policer: ef-policer

#### Table 126: Sample mf-classifier Firewall Filter Terms (continued)

Rule (Term)	Purpose	Contents
network control	Detects packets with a network control precedence and forwards them to the network control class.	Match condition: precedence net-control
		Forwarding class: nc-class
best-effort-data	Detects all other packets and assigns them to the best effort class.	Forwarding class: be-class

For more information about firewalls filters see "Configuring Stateless Firewall Filters" on page 221 and the *JUNOS Policy Framework Configuration Guide*.

To configure a firewall filter for a multifield classifier for the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 127 on page 308.
- 3. Go on to "Assigning Forwarding Classes to Output Queues" on page 310.

#### Table 127: Configuring and Applying a Firewall Filter for a Multifield Classifier

Task	J-M	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Firewall</b> level in the configuration hierarchy.		In the J-Web interface, select Configuration > View and Edit > Edit	From the [edit] hierarchy level, enter
		Configuration.	edit firewall
	2.	Next to Firewall, click <b>Configure</b> or <b>Edit</b> .	
Create the multifield classifier filter and	1.	Click Add new entry next to Filter.	Enter
name it—for example, <b>mf-classifier</b> .	2.	In the Filter name box, type mf-classifier.	edit filter mf-classifier
	3.	Select the check box next to Interface specific.	set interface-specific
Create the term for the assured		Click Add new entry next to Term.	Enter
forwarding traffic class, and give it a name—for example, assured-forwarding.	2.	In the Rule name box, type assured-forwarding.	edit term assured-forwarding
Create the match condition for the	1.	Click Configure next to From.	Enter
assured forwarding traffic class. Use the destination address for assured forwarding traffic—for example,	2.	Click <b>Add new entry</b> next to Destination address.	set from destination-address 192.168.44.55
192.168.44.55	3.	In the Address box, type <b>192.168.44.55</b> .	
	4.	Click <b>OK</b> twice.	

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the forwarding class for assured	1. Click <b>Configure</b> next to Then.	Enter
forwarding Diffserv traffic—for example, af-class.	2. In the Forwarding class box, type af-class.	set then forwarding-class af-class
Set the loss priority for the assured	3. From the Loss priority list, select <b>low</b> .	set then loss-priority low
low.	4. Click <b>OK</b> twice.	
Create the term for the expedited	1. Click Add new entry next to Term.	From the [edit firewall filter mf-classifer]
name—for example, expedited-forwarding.	2. In the Rule name box, type expedited-forwarding.	edit term expedited-forwarding
Create the match condition for the	1. Click <b>Configure</b> next to From.	Enter
the destination address for expedited forwarding traffic class. Use the destination address for expedited forwarding traffic—for example,	2. Click <b>Add new entry</b> next to Destination address.	set from destination-address 192.168.66.77
192.168.66.77	<ol> <li>In the Address box, type 192.168.66.77.</li> </ol>	
	4. Click <b>OK</b> twice.	
Create the forwarding class for expedited	1. Click <b>Configure</b> next to Then.	Enter
forwarding DiffServ traffic—for example, ef-class.	2. In the Forwarding class box, type ef-class.	set then forwarding-class ef-class
Apply the policer for the expedited forwarding traffic class. Use the EF	3. From the Policer choice list, select <b>Policer</b> .	set then policer ef-policer
expedited forwarding DiffServ	4. In the Policer box, type <b>ef-policer</b> .	
traffic—ef-policer.	5. Click <b>OK</b> twice.	
(See "Configuring a Policer for a Firewall Filter" on page 306.)		
Create the term for the network control	1. Click Add new entry next to Term.	From the [edit firewall filter mf-classifer]
traffic class, and give it a name—for example, network-control.	2. In the Rule name box, type network-control.	edit term network-control
Create the match condition for the	1. Click <b>Configure</b> next to From.	Enter
network control traffic class.	2. From the Precedence choice list, select <b>Precedence</b> .	set from precedence net-control
	3. Click <b>Add new entry</b> next to Precedence.	
	4. From the Value keyword list, select <b>net-control</b> .	
	5. Click <b>OK</b> twice.	

# Table 127: Configuring and Applying a Firewall Filter for a Multifield Classifier (continued)

#### Table 127: Configuring and Applying a Firewall Filter for a Multifield Classifier (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the forwarding class for the	1. Click <b>Configure</b> next to Then.	Enter
network control traffic class, and give it a name—for example, <b>nc-class</b> .	2. In the Forwarding class box, type nc-class.	set then forwarding-class nc-class
	3. Click <b>OK</b> twice.	
Create the term for the best-effort traffic	1. Click Add new entry next to Term.	From the [edit firewall filter mf-classifer]
class, and give it a name—for example, <b>best-effort-data</b> .	2. In the Rule name box, type	hierarchy level, enter
	best-effort-data.	edit term best-effort-data
Create the forwarding class for the	1. Click <b>Configure</b> next to Then.	Enter
best-effort traffic class, and give it a name—for example, <b>be-class</b> . (Because this is the last term in the filter, it has no	2. In the Forwarding class box, type <b>be-class</b> .	set then forwarding-class be-class
match condition.)	3. Click <b>OK</b> four times.	
Navigate to the <b>Interfaces</b> level in the configuration hierarchy.	On the main Configuration page next to Interfaces, click <b>Configure</b> or <b>Edit</b> .	From the [edit] hierarchy level, enter
	0	edit interfaces
Apply the multifield classifier firewall filter <b>mf-classifier</b> as an input filter on	1. Click the Interface <b>ge-0/0/0</b> and Unit <b>0</b> .	Enter
each customer-facing or host-facing interface that needs the filter—for	2. Click <b>Configure</b> next to Inet.	set ge-0/0/0 unit 0 family inet filter
example, on ge-0/0/0, unit 0.	3. Click <b>Configure</b> next to Filter.	
	4. From the Input choice list, select <b>Input</b> .	
	5. In the Input box, type <b>mf-classifier</b> .	
	6. Click <b>OK</b> .	

# Assigning Forwarding Classes to Output Queues

You must assign the forwarding classes established by the **mf-classifier** multifield classifier to output queues. This example assigns output queues as shown in Table 128 on page 310.

#### Table 128: Sample Output Queue Assignments for mf-classifier Forwarding Queues

mf-classifier Forwarding Class	For Traffic Type	Output Queue
be-class	Best-effort traffic	Queue 0
ef-class	Expedited forwarding traffic	Queue 1
af-class	Assured forwarding traffic	Queue 2
nc-class	Network control traffic	Queue 3

For multifield classifier details, see "Configuring and Applying a Firewall Filter for a Multifield Classifier" on page 307.

To assign forwarding classes to output queues for the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 129 on page 311.
- 3. Go on to "Configuring and Applying Rewrite Rules" on page 312.

### **Table 129: Assigning Forwarding Classes to Output Queues**

Task	J-W	/eb Configuration Editor	CLI Configuration Editor	
Navigate to the <b>Class of</b> <b>service</b> level in the		In the J-Web interface, select <b>Configuration &gt; View and</b> <b>Edit &gt; Edit Configuration</b> .	From the [ <b>edit]</b> hierarchy level, enter	
configuration hierarchy.	2.	Next to Class of service, click <b>Configure</b> or <b>Edit</b> .	edit class-of-service	
Assign best-effort traffic to	1.	Click <b>Configure</b> next to Forwarding classes.	Enter	
queue 0.	2.	Click Add new entry next to Queue.	set forwarding-classes queue	
	3.	In the Queue num box, type <b>0</b> .	0 be-class	
	4.	In the Class name box, type the previously configured name of the best-effort class— <b>be-class</b> .		
	5.	Click <b>OK</b> .		
Assign expedited forwarding	expedited forwarding 1. Click Add new entry next to Queue.		Enter	
traffic to queue 1.	2.	In the Queue num box, type 1.	set forwarding-classes queue 1 ef-class	
	3.	In the Class name box, type the previously configured name of the expedited forwarding class— <b>ef-class</b> .		
4. Click <b>OK</b> .				
Assign assured forwarding	1.	Click Add new entry next to Queue.	Enter	
traffic to queue 2.	2.	In the Queue num box, type 2.	set forwarding-classes queue	
		In the Class name box, type the previously configured name of the assured forwarding class— <b>af-class</b> .	2 af-class	
	4.	Click <b>OK</b> .		
Assign network control traffic	1.	Click Add new entry next to Queue.	Enter	
to queue 3.	2.	In the Queue num box, type <b>3</b> .	set forwarding-classes queue	
		In the Class name box, type the previously configured name of the network control forwarding class— <b>nc-class</b> .	3 nc-class	
	4.	Click <b>OK</b> .		

### **Configuring and Applying Rewrite Rules**

You can configure rewrite rules to replace DiffServ code points (DSCPs) on packets received from the customer or host with the values expected by other routers. You do not have to configure rewrite rules if the received packets already contain valid DSCPs. Rewrite rules apply the forwarding class information and packet loss priority used internally by the Services Router to establish the DSCP on outbound packets. Once configured, you must apply the rewrite rules to the correct interfaces.

The following example shows how to create the rewrite rules **rewrite-dscps** and apply them to the Services Router's Gibabit Ethernet interface ge-0/0/0. The rewrite rules replace the DSCPs on packets in the four forwarding classes, as shown in Table 130 on page 312.

#### Table 130: Sample rewrite-dscps Rewrite Rules to Replace DSCPs

mf-classifier Forwarding Class	For CoS Traffic Type	rewrite-dscps Rewrite Rules
be-class	Best-effort traffic	Low-priority code point: 000000
		High-priority code point: 000001
ef-class	Expedited forwarding traffic	Low-priority code point: 101110
		High-priority code point: 101111
af-class	Assured forwarding traffic	Low-priority code point: 001010
		High-priority code point: 001100
nc-class	Network control traffic	Low-priority code point: 110000
		High-priority code point: 110001

To configure and apply rewrite rules for the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 131 on page 312.
- 3. Go on to "Configuring and Applying Behavior Aggregate Classifiers" on page 315.

#### **Table 131: Configuring and Applying Rewrite Rules**

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Class of service</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>	From the [edit] hierarchy level, enter edit class-of-service
	2. Next to Class of service, click <b>Configure</b> or <b>Edit</b> .	

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure rewrite rules for DiffServ CoS.	<ol> <li>Click Configure next to Rewrite rules.</li> </ol>	Enter
	2. Click Add new entry next to Dscp.	edit rewrite-rules dscp rewrite-dscps
	<ol> <li>In the Name box, type the name of the rewrite rules—for example, rewrite-dscps.</li> </ol>	
Configure best-effort forwarding class rewrite rules.	<ol> <li>Click Add new entry next to Forwarding class.</li> </ol>	Enter
	2. In the Class name box, type the name of the previously configured	set forwarding-class be-class loss-priority low code-point 000000
	best-effort forwarding class— <b>be-class</b> .	set forwarding-class be-class loss-priority
	<ol> <li>Click Add new entry next to Loss priority.</li> </ol>	ngh code-point 000001
	4. From the Loss val list, select <b>low</b> .	
	<ol> <li>In the Code point box, type the value of the low-priority code point for best-effort traffic—for example, 000000.</li> </ol>	
	6. Click <b>OK</b> .	
	7. Click <b>Add new entry</b> next to Loss priority.	
	8. From the Loss val list, select high.	
	<ol> <li>In the Code point box, type the value of the high-priority code point for best-effort traffic—for example, 000001.</li> </ol>	
	10. Click <b>OK</b> twice.	

# Table 131: Configuring and Applying Rewrite Rules (continued)

## Table 131: Configuring and Applying Rewrite Rules (continued)

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Configure expedited forwarding class rewrite rules.	1.	Click <b>Add new entry</b> next to Forwarding class.	Enter
	2.	In the Class name box, type the name of the previously configured expedited forwarding	set forwarding-class ef-class loss-priority low code-point 101110 set forwarding-class ef-class loss-priority
	3.	Class—er-class. Click <b>Add new entry</b> next to Loss	high code-point 101111
	4	From the Loss val list select <b>low</b>	
	5.	In the Code point box, type the value of the low-priority code point for expedited forwarding traffic—for example, <b>101110</b> .	
	6.	Click <b>OK</b> .	
	7.	Click <b>Add new entry</b> next to Loss priority.	
	8.	From the Loss val list, select high.	
	9.	In the Code point box, type the value of the high-priority code point for expedited forwarding traffic—for example, <b>101111</b> .	
	10	Click <b>OK</b> twice.	
Configure assured forwarding class rewrite rules.	1.	Click <b>Add new entry</b> next to Forwarding class.	Enter
	2.	In the Class name box, type the name of the previously configured assured forwarding class—afclass	set forwarding-class af-class loss-priority low code-point 001010
	3.	Click <b>Add new entry</b> next to Loss priority.	set forwarding-class af-class loss-priority high code-point 001100
	4.	From the Loss val list, select low.	
	5.	In the Code point box, type the value of the low-priority code point for assured forwarding traffic—for example, <b>001010</b> .	
	6.	Click <b>OK</b> .	
	7.	Click <b>Add new entry</b> next to Loss priority.	
	8.	From the Loss val list, select <b>high</b> .	
	9.	In the Code point box, type the value of the high-priority code point for assured forwarding traffic—for example, <b>001100</b> .	
	10.	Click <b>OK</b> twice.	

Task	J-N	eb Configuration Editor	CLI Configuration Editor
Configure network control class rewrite rules.	1.	Click <b>Add new entry</b> next to Forwarding class.	Enter
	2.	In the Class name box, type the name of the previously configured	set forwarding-class nc-class loss-priority low code-point 110000
		class—nc-class.	set forwarding-class nc-class loss-priority
	3.	Click <b>Add new entry</b> next to Loss priority.	
	4.	From the Loss val list, select <b>low</b> .	
	5.	In the Code point box, type the value of the low-priority code point for network control traffic—for example, <b>110000</b> .	
	6.	Click <b>OK</b> .	
	7.	Click <b>Add new entry</b> next to Loss priority.	
	8.	From the Loss val list, select high.	
	9.	In the Code point box, type the value of the high-priority code point for network control traffic—for example, <b>110001</b> .	
	10	Click <b>OK</b> four times.	
Apply rewrite rules to an interface.	1.	Click <b>Add new entry</b> next to Interfaces	From the [edit class of service] hierarchy level, enter
(See the interface naming conventions in the <i>J</i> -series Services Router Basic LAN and WAN Access Configuration Guide.)	2.	In the Interface name box, type the name of the interface—for example, ge-0/0/0.	set interfaces ge-0/0/0 unit 0 rewrite-rules dscp rewrite-dscps
	3.	Click Add new entry next to Unit.	
	4.	In the Unit number box, type the logical interface unit number—0.	
	5.	Click <b>Configure</b> next to Rewrite rules.	
	6.	In the Rewrite rules name box, under Dscp, type the name of the previously configured rewrite rules— <b>rewrite-dscps</b> .	
	7.	Click <b>OK</b> .	

### Table 131: Configuring and Applying Rewrite Rules (continued)

# **Configuring and Applying Behavior Aggregate Classifiers**

You configure behavior aggregate classifiers to classify packets that contain valid DSCPs to appropriate queues. Once configured, you must apply the behavior aggregate classifier to the correct interfaces.

The following example shows how to configure the DSCP behavior aggregate classifier **ba-classifier** as the default DSCP map, and apply it to the Services Router's Gigabit Ethernet interface **ge-0/0/0**. The behavior aggregate classifier assigns loss priorities, as shown in Table 132 on page 316, to incoming packets in the four forwarding classes.

#### Table 132: Sample ba-classifier Loss Priority Assignments

mf-classifier Forwarding Class	For CoS Traffic Type	ba-classifier Assignments
be-class	Best-effort traffic	High-priority code point: 000001
ef-class	Expedited forwarding traffic	High-priority code point: 101111
af-class	Assured forwarding traffic	High-priority code point: 001100
nc-class	Network control traffic	High-priority code point: 110001

To configure and apply behavior aggregate classifiers for the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 133 on page 316.
- 3. Go on to "Configuring RED Drop Profiles for Congestion Control" on page 319.

#### Table 133: Configuring and Applying Behavior Aggregate Classifiers

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Class of service</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and	From the [edit] hierarchy level, enter
5		Edit > Edit Configuration.	edit class-of-service
	2.	Next to Class of service, click <b>Configure</b> or <b>Edit</b> .	
Configure behavior aggregate classifiers	1.	Click <b>Configure</b> next to Classifiers.	Enter
for DiffServ CoS.	2.	Click Add new entry next to Dscp.	edit classifiers dscp ba-classifier
	3.	In the Name box, type the name of the behavior aggregate	set import default
		classifier—for example, ba-classifier.	
	4.	In the Import box, type the name of the default DSCP map, <b>default</b> .	

Task	J-M	eb Configuration Editor	CLI Configuration Editor
Configure a best-effort forwarding class classifier.	1.	Click <b>Add new entry</b> next to Forwarding class.	Enter
	2.	In the Class name box, type the name of the previously configured best-effort forwarding class— <b>be-class</b> .	set forwarding-class be-class loss-priority high code-points 000001
	3.	Click <b>Add new entry</b> next to Loss priority.	
	4.	From the Loss val list, select <b>high</b> .	
	5.	Click <b>Add new entry</b> next to Code points.	
	6.	In the Value box, type the value of the high-priority code point for best-effort traffic—for example, <b>00001</b> .	
	7.	Click <b>OK</b> three times.	
Configure an expedited forwarding class classifier.	1.	Click <b>Add new entry</b> next to Forwarding class.	Enter
	2.	In the Class name box, type the name of the previously configured expedited forwarding class—ef-class.	set forwarding-class ef-class loss-priority high code-points 101111
	3.	Click <b>Add new entry</b> next to Loss priority.	
	4.	From the Loss val list, select <b>high</b> .	
	5.	Click <b>Add new entry</b> next to Code points.	
	6.	In the Value box, type the value of the high-priority code point for expedited forwarding traffic—for example, <b>101111</b> .	
	7.	Click <b>OK</b> three times.	

## Table 133: Configuring and Applying Behavior Aggregate Classifiers (continued)

## Table 133: Configuring and Applying Behavior Aggregate Classifiers (continued)

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Configure an assured forwarding class classifier.	1.	Click <b>Add new entry</b> next to Forwarding class.	Enter
	2.	In the Class name box, type the name of the previously configured assured forwarding class— <b>af-class</b> .	set forwarding-class af-class loss-priority high code-points 001100
	3.	Click <b>Add new entry</b> next to Loss priority.	
	4.	From the Loss val list, select <b>high</b> .	
	5.	Click <b>Add new entry</b> next to Code points.	
	6.	In the Value box, type the value of the high-priority code point for assured forwarding traffic—for example, <b>001100</b> .	
	7.	Click <b>OK</b> three times.	
Configure a network control class classifier.	1.	Click <b>Add new entry</b> next to Forwarding class.	Enter
	2.	In the Class name box, type the name of the previously configured network control forwarding class— <b>nc-class</b> .	set forwarding-class nc-class loss-priority high code-points 110001
	3.	Click <b>Add new entry</b> next to Loss priority.	
	4.	From the Loss val list, select <b>high</b> .	
	5.	Click <b>Add new entry</b> next to Code points.	
	6.	In the Value box, type the value of the high-priority code point for network control traffic—for example, <b>110001</b> .	
	7.	Click <b>OK</b> five times.	

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Apply the behavior aggregate classifier to an interface.	1.	Click <b>Add new entry</b> next to Interfaces.	From the <b>[edit class of service]</b> hierarchy level, enter
(See the interface naming conventions in the J-series Services Router Basic LAN and WAN Access Configuration Guide.)	2.	In the Interface name box, type the name of the interface—for example, ge-0/0/0.	set interfaces ge-0/0/0 unit 0 classifiers dscp ba-classifier
	3.	Click Add new entry next to Unit.	
	4.	In the Unit number box, type the logical interface unit number—0.	
	5.	Click <b>Configure</b> next to Classifiers.	
	6.	In the Classifiers box, under Dscp, type the name of the previously configured behavior aggregate classifier— <b>ba-classifier</b> .	
	7.	Click OK.	

#### Table 133: Configuring and Applying Behavior Aggregate Classifiers (continued)

# **Configuring RED Drop Profiles for Congestion Control**

If the Services Router must support assured forwarding, you can control congestion by configuring random early detection (RED) drop profiles. RED drop profiles use drop probabilities for different levels of buffer fullness to determine which scheduling queue on the router is likely to drop assured forwarding packets under congested conditions. The router can drop packets when the queue buffer becomes filled to the configured percentage.

Assured forwarding traffic with the PLP (packet loss priority) bit set is more likely to be discarded than traffic without the PLP bit set. This example shows how to configure a drop probability and a queue fill level for both PLP and non-PLP assured forwarding traffic. It is only one example of how to use RED drop profiles.

The example shows how to configure the RED drop profiles listed in Table 134 on page 319.

Table 134:	Sample	<b>RED Drop</b>	Profiles
------------	--------	-----------------	----------

Drop Profile	Drop Probability	Queue Fill Level
<b>af-normal</b> —For non-PLP (normal) assured forwarding traffic	Between 0 (never dropped) and 100 percent (always dropped)	Between 95 and 100 percent
<b>af-with-plp</b> —For PLP (aggressive packet dropping) assured forwarding traffic	Between 95 and 100 percent (always dropped)	Between 80 and 95 percent

To configure RED drop profiles for assured forwarding congestion control on the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 135 on page 320.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to one of the following tasks:
  - To assign resources, priorities, and profiles to output queues, see "Configuring Schedulers" on page 321.
  - To apply rules to logical interfaces, see "Configuring and Applying Virtual Channels" on page 327.
  - To use adaptive shapers to limit bandwidth for Frame Relay, see "Configuring and Applying Adaptive Shaping for Frame Relay" on page 331.
  - To check the configuration, see "Verifying a CoS Configuration" on page 345.

Task	J-M	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Class of</b> <b>service</b> level in the		In the J-Web interface, select <b>Configuration &gt; View and</b> <b>Edit &gt; Edit Configuration</b> .	From the <b>[edit]</b> hierarchy level, enter
configuration hierarchy.	2.	Next to Class of service, click <b>Configure</b> or <b>Edit</b> .	edit class-of-service
Configure the lower drop	1.	Click Add new entry next to Drop profiles.	Enter
probability for normal, non-PLP traffic.	2.	In the Profile name box, type the name of the drop profile—for example, <b>af-normal</b> .	edit drop-profiles af-normal interpolate
	3.	Click <b>Configure</b> next to Interpolate.	
	4.	Click Add new entry next to Drop probability.	set drop-probability 0
5. 6.		In the Value box, type a number for the first drop point—for example, <b>0</b> .	set drop-probability 100
		Click OK.	
	7.	Click Add new entry next to Drop probability again.	
		In the Value box, type a number for the next drop point—for example, <b>100</b> .	
		Click <b>OK</b> .	
			_
Configure a queue fill level	1.	Click <b>Add new entry</b> next to Fill level.	Enter
probability.	2.	In the Value box, type a number for the first fill level—for example, <b>95</b> .	set fill-level 95
	3.	Click OK.	set fill-level 100
	4.	Click Add new entry next to Fill level.	
	5.	In the Value box, type a number for the next fill level—for example, <b>100</b> .	
	6.	Click <b>OK</b> three times.	

#### Table 135: Configuring RED Drop Profiles for Assured Forwarding Congestion Control

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Configure the higher drop		Click Add new entry next to Drop profiles.	From the [edit class of
probability for PLP traffic.	2.	In the Profile name box, type the name of the drop	service] hierarchy level, enter
	3	Click <b>Configure</b> next to Interpolate	edit drop-profiles af-with-PLP
	Э. 4	Click <b>Add new entry</b> next to Drop probability	interpolate
	г. г	In the Mehrer have a much as fearth a first data a start. For	set drop-probability 95
		example, 95.	set drop-probability 100
	6.	Click OK.	
	7.	Click Add new entry next to Drop probability.	
		In the Value box, type a number for the next drop point—for example, <b>100</b> .	
		Click <b>OK</b> .	
Configure a queue fill level	1.	Click Add new entry next to Fill level.	Enter
for the higher PLP drop probability.	2.	In the Value box, type a number for the first fill level—for example, <b>80</b> .	set fill-level 80
	3.	Click OK.	set fill-level 95
		Click Add new entry next to Fill level.	
	5.	In the Value box, type a number for the next fill level—for example, <b>95</b> .	
	6.	Click OK.	

### Table 135: Configuring RED Drop Profiles for Assured Forwarding Congestion Control (continued)

# **Configuring Schedulers**

You configure schedulers to assign resources, priorities, and drop profiles to output queues. By default, only queues 0 and 3 have resources assigned.

This example creates the schedulers listed in Table 136 on page 321.

Scheduler	For CoS Traffic Type	Assigned Priority	Allocated Portion of Queue Buffer	Assigned Bandwidth (Transmit Rate)
be-scheduler	Best-effort traffic	Low	40 percent	10 percent
ef-scheduler	Expedited forwarding traffic	High	10 percent	10 percent
af-scheduler	Assured forwarding traffic	High	45 percent	45 percent
nc-scheduler	Network control traffic	Low	5 percent	5 percent

#### **Table 136: Sample Schedulers**

To configure schedulers for the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 137 on page 322.
- 3. Go on to "Configuring and Applying Scheduler Maps" on page 324.

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Class of</b> <b>service</b> level in the	<ol> <li>In the J-Web interface, select Configuration &gt; View a Edit &gt; Edit Configuration.</li> </ol>	nd From the [edit] hierarchy level, enter
configuration hierarchy.	2. Next to Class of service, click <b>Configure</b> or <b>Edit</b> .	edit class-of-service
Configure a best-effort	1. Click Add new entry next to Schedulers.	Enter
scheduler.	2. In the Scheduler name box, type the name of the best-e scheduler—for example, <b>be-scheduler</b> .	ffort edit schedulers be-scheduler
Configure a best-effort	1. In the Priority box, type low.	Enter
scheduler priority and buffer size.	2. Click <b>Configure</b> next to Buffer size.	set priority low
	<ol> <li>From the Buffer size choice list, select the basis for the buffer allocation method—for example, Percent.</li> </ol>	e set buffer-size percent 40
	<ol> <li>In the Percent box, type the percentage of the buffer t used by the best-effort scheduler—for example, 40.</li> </ol>	o be
	5. Click <b>OK</b> .	
Configure a best-effort	1. Click <b>Configure</b> next to Transmit rate.	Enter
scheduler transmit rate.	<ol> <li>From the Transmit rate choice list, select the basis for transmit rate method—for example, Percent.</li> </ol>	the set transmit-rate percent 10
	<ol> <li>In the Percent box, type the percentage of the bandw to be used by the best-effort scheduler—for example,</li> </ol>	idth 10.
	4. Click <b>OK</b> twice.	
Configure an expedited	1. Click Add new entry next to Schedulers.	From the [edit class of
forwarding scheduler.	<ol> <li>In the Scheduler name box, type the name of the expect forwarding scheduler—for example, ef-scheduler.</li> </ol>	lited service] hierarchy level, enter
		edit schedulers ef-scheduler

### **Table 137: Configuring Schedulers**

Task	J-W	eb Configuration Editor	<b>CLI Configuration Editor</b>	
Configure an expedited	1.	In the Priority box, type high.	Enter	
forwarding scheduler priority and buffer size.	2.	Click <b>Configure</b> next to Buffer size.	set priority high	
	3.	From the Buffer size choice list, select the basis for the buffer allocation method—for example, <b>Percent</b> .	set buffer-size percent 10	
	4.	In the Percent box, type the percentage of the buffer to be used by the expedited forwarding scheduler—for example, <b>10</b> .		
	5.	Click <b>OK</b> .		
Configure an expedited	1.	Click <b>Configure</b> next to Transmit rate.	Enter	
rate.	2.	From the Transmit rate choice list, select the basis for the transmit rate method—for example, <b>Percent</b> .	set transmit-rate percent 10	
	3.	In the Percent box, type the percentage of the bandwidth to be used by the expedited forwarding scheduler—for example, <b>10</b> .		
	4.	Click <b>OK</b> twice.		
Configure an assured forwarding scheduler.	1.	Click Add new entry next to Schedulers.	From the [edit class of	
	2.	In the Scheduler name box, type the name of the assured forwarding scheduler—for example, <b>af-scheduler</b> .	service] hierarchy level, enter	
			edit schedulers af-scheduler	
Configure an assured	1.	In the Priority box, type high.	Enter	
forwarding scheduler priority and buffer size.	2.	Click <b>Configure</b> next to Buffer size.	set priority high	
	<ol> <li>From the Buffer size choice list, select the basis for buffer allocation method—for example, Percent.</li> </ol>		set buffer-size percent 45	
	4.	In the Percent box, type the percentage of the buffer to be used by the assured forwarding scheduler—for example, 45.		
	5.	Click <b>OK</b> .		
Configure an assured	1.	Click <b>Configure</b> next to Transmit rate.	Enter	
rate.	2.	From the Transmit rate choice list, select the basis for the transmit rate method—for example, <b>Percent</b> .	set transmit-rate percent 45	
	3.	In the Percent box, type the percentage of the bandwidth to be used by the assured forwarding scheduler—for example, <b>45</b> .		
	4.	Click <b>OK</b> .		

# Table 137: Configuring Schedulers (continued)

### Table 137: Configuring Schedulers (continued)

Task	J-W	leb Configuration Editor	CLI Configuration Editor
(Optional) Configure a drop	1.	Click Add new entry next to Drop profile map.	Enter
profile map for assured forwarding low and high	2.	From the Loss priority box, select Low.	set drop-profile-map
priority. (DiffServ can have a	3.	From the Protocol box, select Any.	loss-priority low protocol any
with assured forwarding.)	4.	In the Drop profile box, type the name of the drop profile—for example, <b>af-norma</b> l.	drop-profile af-normal
	5.	Click <b>OK</b> .	loss-priority high protocol any
	6.	Click Add new entry next to Drop profile map.	drop-profile af-with-PLP
	7.	From the Loss priority box, select High.	
	8.	From the Protocol box, select Any.	
	9.	In the Drop profile box, type the name of the drop profile—for example, af-with-PLP.	
	10	Click <b>OK</b> twice.	
Configure a network control		Click Add new entry next to Schedulers.	From the [edit class of
scheduler.	2.	In the Scheduler name box, type the name of the network control scheduler—for example, nc-scheduler.	service] hierarchy level, enter
			edit schedulers nc-scheduler
Configure a network control scheduler priority and buffer size.	1.	In the Priority box, type low.	Enter
	2.	Click <b>Configure</b> next to Buffer size.	set priority low
	3.	From the Buffer size choice list, select the basis for the buffer allocation method—for example, <b>Percent</b> .	set buffer-size percent 5
		In the Percent box, type the percentage of the buffer to be used by the network control scheduler—for example, 5.	
	5.	Click <b>OK</b> .	
Configure a network control	1.	Click <b>Configure</b> next to Transmit rate.	Enter
scheduler transmit rate.	2.	From the Transmit rate choice list, select the basis for the transmit rate method—for example, <b>Percent</b> .	set transmit-rate percent 5
		In the Percent box, type the percentage of the bandwidth to be used by the network control scheduler—for example, 5.	
	4.	Click <b>OK</b> .	

# **Configuring and Applying Scheduler Maps**

You configure a scheduler map to assign a forwarding class to a scheduler, then apply the scheduler map to any interface that must enforce DiffServ CoS.

The following example shows how to create the scheduler map diffserv-cos-map and apply it to the Services Router's Ethernet interface ge-0/0/0. The map associates the

**mf-classifier** forwarding classes configured in "Configuring and Applying a Firewall Filter for a Multifield Classifier" on page 307 to the schedulers configured in "Configuring Schedulers" on page 321, as shown in Table 138 on page 325.

#### Table 138: Sample diffserv-cos-map Scheduler Mapping

mf-classifier Forwarding Class	For CoS Traffic Type	diffserv-cos-map Scheduler
be-class	Best-effort traffic	be-scheduler
ef-class	Expedited forwarding traffic	ef-scheduler
af-class	Assured forwarding traffic	af-scheduler
nc-class	Network control traffic	nc-scheduler

To configure and apply scheduler maps for the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 139 on page 325.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to one of the following tasks:
  - To apply rules to logical interfaces, see "Configuring and Applying Virtual Channels" on page 327.
  - To use adaptive shapers to limit bandwidth for Frame Relay, see "Configuring and Applying Adaptive Shaping for Frame Relay" on page 331.
  - To check the configuration, see "Verifying a CoS Configuration" on page 345.

### Table 139: Configuring Scheduler Maps

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <b>Class of service</b> level in the configuration hierarchy.	1. In the J-Web interface, select Configuration > View and	From the [edit] hierarchy level, enter
	Edit > Edit Configuration.	edit class-of-service
	<ol> <li>Next to Class of service, click Configure or Edit.</li> </ol>	
Configure a scheduler map for DiffServ CoS.	<ol> <li>Click Add new entry next to Scheduler maps.</li> </ol>	Enter
	2. In the Map name box, type the name of the scheduler map—for example, diffserv-cos-map.	edit scheduler-maps diffserv-cos-map

## Table 139: Configuring Scheduler Maps (continued)

Task	J-V	Veb Configuration Editor	CLI Configuration Editor
Configure a best-effort forwarding class and scheduler.	1.	Click <b>Add new entry</b> next to Forwarding class.	Enter
	2.	In the Class name box, type the name of the previously configured best-effort forwarding class— <b>be-class</b> .	set forwarding-class be-class scheduler be-scheduler
	3.	In the Scheduler box, type the name of the previously configured best-effort scheduler— <b>be-scheduler</b> .	
	4.	Click <b>OK</b> .	
Configure an expedited forwarding class and scheduler.	1.	Click <b>Add new entry</b> next to Forwarding class.	Enter
	2.	In the Class name box, type the name of the previously configured expedited forwarding class—ef-class.	set forwarding-class ef-class scheduler ef-scheduler
	3.	In the Scheduler box, type the name of the previously configured expedited forwarding scheduler— <b>ef-scheduler</b> .	
	4.	Click <b>OK</b> .	
Configure an assured forwarding class and scheduler.	1.	Click <b>Add new entry</b> next to Forwarding class.	Enter
	2.	In the Class name box, type the name of the previously configured assured forwarding class— <b>af-class</b> .	set forwarding-class af-class scheduler af-scheduler
	3.	In the Scheduler box, type the name of the previously configured assured forwarding scheduler— <b>af-scheduler</b> .	
	4.	Click <b>OK</b> .	
Configure a network control class and scheduler.	1.	Click <b>Add new entry</b> next to Forwarding class.	Enter
	2.	In the Class name box, type the name of the previously configured network control class—nc-class.	set forwarding-class nc-class scheduler nc-scheduler
	3.	In the Scheduler box, type the name of the previously configured network control scheduler—nc-scheduler.	
	4.	Click <b>OK</b> twice.	

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Apply the scheduler map to an interface.	1.	Click <b>Add new entry</b> next to Interfaces.	From the <b>[edit class of service]</b> hierarchy level, enter
(See the interface naming conventions in the J-series Services Router Basic LAN and WAN Access Configuration Guide.)	2.	In the Interface name box, type the name of the interface—for example, <b>ge-0/0/0</b> .	set interfaces ge-0/0/0 scheduler-map diffserv-cos-map
	3.	Click Add new entry next to Unit.	
	4.	In the Unit number box, type the logical interface unit number—0.	
	5.	In the Scheduler map box, type the name of the previously configured scheduler map—diffserv-cos-map.	
	6.	Click <b>OK</b> .	

#### Table 139: Configuring Scheduler Maps (continued)

## **Configuring and Applying Virtual Channels**

You configure a virtual channel to set up queuing, packet scheduling, and accounting rules to be applied to one or more logical interfaces. You then must apply the virtual channel to a particular logical interface. Virtual channels can be applied in different ways. In the example here, an output firewall filter is used for directing traffic to a particular virtual channel.

The following example shows how to create the virtual channels branch1–vc, branch2–vc, and branch3–vc and apply them in the firewall filter choose-vc to the Services Router's T3 interface t3-1/0/0.

To configure and apply virtual channels for the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 140 on page 328.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to one of the following tasks:
  - To assign resources, priorities, and profiles to output queues, see "Configuring Schedulers" on page 321.
  - To use adaptive shapers to limit bandwidth for Frame Relay, see "Configuring and Applying Adaptive Shaping for Frame Relay" on page 331.
  - To check the configuration, see "Verifying a CoS Configuration" on page 345.

## **Table 140: Configuring and Applying Virtual Channels**

Task	J-M	/eb Configuration Editor	CL	I Configuration Editor	
Navigate to the <b>Class of service</b> level in the configuration hierarchy.	<ol> <li>In the J-Web interface, select Configuration &gt; View and Edit &gt; Edit Configuration.</li> </ol>		Fro ed	From the [edit] hierarchy level, enter edit class-of-service	
	2.	Next to Class of service, click <b>Configure</b> or <b>Edit</b> .			
Define the virtual channels branch1–vc, branch2–vc, branch3–vc, and the default	1.	Click <b>Add new entry</b> next to Virtual channels.	1.	Enter	
virtual channel. You must specify a default virtual channel.	2.	In the Channel name box, type the name of the virtual channel—for	e 2.	set virtual-channels branch1–vc Repeat this statement for	
	3	Click <b>OK</b>		branch2-vc, branch3-vc, and default-vc.	
		Create additional virtual channels for branch2–vc, branch3–vc, and default-vc.			
Define the virtual channel group wan-vc-group to include the four virtua	1.	Click <b>Add new entry</b> next to Virtual channel groups.	1.	Enter	
channels, and assign each virtual channel the scheduler map bestscheduler.		In the Group name box, type the name of the virtual channel group— <b>wan-vc-group</b> .		set virtual-channel-groups wan-vc-group branch1–vc scheduler-map bestscheduler	
		Click <b>Add new entry</b> next to Channel.	2.	Repeat this statement for branch2-vc, branch3-vc, and default-vc	
	4.	In the Channel name box, type the name of the previously configured	3.	Enter	
	5.	In the Scheduler map box, type the name of the previously configured scheduler map— <b>bestscheduler</b> .		set virtual-channel-groups wan-vc-group default–vc default	
	6.	Click <b>OK</b> .			
	7.	Add the virtual channels branch2–vc, branch3–vc, and default-vc. Select the <b>Default</b> box when adding the virtual channel default-vc.			

Task	J-Web Configuration Editor	CLI Configuration Editor	
Specify a shaping rate of 2 Mbps for each virtual channel within the virtual	<ol> <li>Click branch1-vc in the list of virtual channels.</li> </ol>	1. Enter	
channel group.	2. Select the <b>Shaping rate</b> box.	set virtual-channel-groups wan-vc-group branch1-vc shaping-rate	
	3. Click <b>Configure</b> .	2m	
	4. Select <b>Absolute rate</b> from the Rate choice box.	2. Repeat this statement for branch2–vc and branch3–vc.	
	5. In the Absolute rate box, type the shaping rate—2m.		
	<ol> <li>Add the shaping rate for the branch2–vc and branch3–vc virtual channels.</li> </ol>		
	7. Click <b>OK</b> three times.		
Apply the virtual channel group to the logical interface t3-1/0/0.0.	<ol> <li>Click Add new entry next to Interfaces.</li> </ol>	From the [edit class of service] hierarchy level, enter	
(See the interface naming conventions in the I-series Services Router Basic I AN	2. In the Interface name box, type the name of the interface—t3–1/0/0.	set interfaces t3–1/0/0 unit 0	
and WAN Access Configuration Guide.)	3. Click Add new entry next to Unit.		
	<ol> <li>In the Unit number box, type the logical interface unit number—0.</li> </ol>		
	<ol> <li>In the Virtual channel group box, type the name of the previously configured virtual channel group—wan-vc-group.</li> </ol>		
	6. Click <b>OK</b> .		

## Table 140: Configuring and Applying Virtual Channels (continued)

## Table 140: Configuring and Applying Virtual Channels (continued)

Task	J-W	eb Configuration Editor	CL	Configuration Editor
Create the firewall filter <b>choose-vc</b> to select the traffic that is transmitted on a particular virtual channel	1.	On the main Configuration page next to Firewall, click <b>Configure</b> or <b>Edit</b> .	1.	From the [edit] hierarchy level, enter
- F	2.	Click Add new entry next to Filter.	2	Enter
	3.	In the Filter name box, type the name of the firewall filter—choose-vc.	2.	set family inet filter choose-vc term branch1 from destination
	4.	Click Add new entry next to Term.	3	Fnter
	5.	In the Rule name box, type the name of the firewall term— <b>branch1</b> .	5.	set family inet filter choose-vc term branch1 then accept
	6.	Click <b>Configure</b> next to From.	4.	Enter
	7.	Click <b>Add new entry</b> next to Destination address.		set family inet filter choose-vc term
	8.	In the Address box, type the IP address of the destination		branch1-vc
		host—192.168.10.0/24.	5.	Repeat these steps for virtual channels branch2–vc and
	9.	Click OK twice.		branch3-vc.
	10	On the firewall term page, click <b>Configure</b> next to Then.		
	11.	Select <b>Accept</b> from the Designation box.		
	12	In the Virtual channel box, type the name of the previously configured virtual channel—branch1-vc.		
	13	Click OK.		
	14.	Repeat these steps for the virtual channels branch2–vc and branch3–vc.		
Apply the firewall filter <b>choose-vc</b> to output traffic on the <b>t3–1/0/0.0</b>	1.	On the main Configuration page next to Interfaces, click <b>Configure</b>	1.	From the [edit] hierarchy level, enter
interface.		or Edit.		edit interfaces
	2.	Click <b>t3–1/0/0</b> in the list of configured interfaces.	2.	Enter
	3.	Click <b>0</b> in the list of configured logical units for the interface.		set t3–1/0/0 unit 0 family inet filter output choose-vc
	4.	Click <b>Edit</b> next to Inet.		
	5.	Click <b>Configure</b> next to Filter.		
	6.	In the Output box, type the name of the previously configured firewall filter—choose-vc.		
	7.	Click <b>OK</b> .		

# **Configuring and Applying Adaptive Shaping for Frame Relay**

You can use adaptive shaping to limit the bandwidth of traffic flowing on a Frame Relay logical interface. If you configure and apply adaptive shaping, the Services Router checks the backward explicit congestion notification (BECN) bit within the last inbound (ingress) packet received on the interface. If the BECN bit is set, the router limits the transmit bandwidth on the interface to the configured adaptive shaper maximum transmit rate. If the BECN bit is not set, the transmit bandwidth is not limited and is allowed to exceed the adaptive shaper rate.

For more information about adaptive shapers for a Frame Relay interface, see the *JUNOS Class of Service Configuration Guide*.

The following example shows how to create adaptive shaper fr-shaper and apply it to the Services Router's T1 interface t1-0/0/2. The adapter shaper limits the transmit bandwidth on the interface to 64 Kbps.

To configure and apply an adaptive shaper for the Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 141 on page 331.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to one of the following tasks:
  - To assign resources, priorities, and profiles to output queues, see "Configuring Schedulers" on page 321.
  - To apply rules to logical interfaces, see "Configuring and Applying Virtual Channels" on page 327.
  - To check the configuration, see "Verifying a CoS Configuration" on page 345.

### **Table 141: Configuring and Applying an Adaptive Shaper**

Task	J-W	leb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Class of service</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit class-of-service
	2.	Next to Class of service, click <b>Configure</b> or <b>Edit</b> .	

#### Table 141: Configuring and Applying an Adaptive Shaper (continued)

Task	J-M	eb Configuration Editor	CLI Configuration Editor
Define the adaptive shaper name and maximum transmit rate.	1.	Next to Adaptive Shapers, click <b>Add new entry</b> .	Enter
	2.	In the Adaptive shaper name box, type <b>fr-shape</b> r.	set adaptive-shapers fr-shaper trigger becn shaping-rate 64k
	3.	Next to Trigger, click <b>Add new</b> entry.	
	4.	Next to Becn, select the check box.	
	5.	Next to Shaping rate, select the check box and click <b>Configure</b> .	
	6.	From the Rate choice list, select <b>Absolute rate</b> .	
	7.	In the Absolute rate box, type 64k.	
	8.	Click <b>OK</b> three times.	
Apply the adaptive shaper to the logical interface t1-0/0/2.0.	1.	Next to Interfaces, click <b>Add new</b> entry.	Enter
(See the interface naming conventions in the <i>I-series Services Router Basic LAN</i>	2.	In the Interface name box, type the name of the interface—t1-0/0/2.	set interfaces t1-0/0/2 unit 0 adaptive-shaper fr-shaper
and WAN Access Configuration Guide.)	3.	Next to Unit, click Add new entry.	
	4.	In the Unit number box, type the logical interface unit number— <b>0</b> .	
	5.	In the Adaptive shaper box, type the name of the adaptive shaper— <b>fr-shaper</b> .	
	6.	Click <b>OK</b> .	

# **Configuring Strict High Priority for Queuing with a Configuration Editor**

On a Services Router, you can configure one queue per interface to have strict high priority, which causes delay-sensitive traffic, such as voice traffic, to be dequeued and forwarded with minimum delay. Packets that are queued in a strict-priority queue are dequeued before packets in other queues, including high-priority queues.

The strict high-priority queuing feature allows you to configure traffic policing that prevents lower-priority queues from being starved. The strict-priority queue does not cause starvation of other queues because the configured policer allows the queue to exceed the configured bandwidth only when other queues are not congested. If the interface is congested, the software polices strict-priority queues to the configured bandwidth.

To prevent queue starvation of other queues, you must configure an output (egress) policer that defines a limit for the amount of traffic that the queue can service. The software services all traffic in the strict-priority queue that is under the defined limit. When strict-priority traffic exceeds the limit, the policer marks the traffic in excess

of the limit as out-of-profile. If the output port is congested, the software drops out-of-profile traffic.

You can also configure a second policer with an upper limit. When strict-priority traffic exceeds the upper limit, the software drops the traffic in excess of the upper limit, regardless of whether the output port is congested. This upper-limit policer is not a requirement for preventing starvation of the lower-priority queues. The policer for the lower limit, which marks the packets as out-of-profile, is sufficient to prevent starvation of other queues.

The sample strict-high priority queuing configuration does the following:

- 1. Uses a behavior aggregate (BA) classifier to classify traffic based on the IP precedence of the packet. The classifier defines IP precedence value 101 as voice traffic and 000 as data traffic.
- 2. To minimize delay, assigns all delay-sensitive packets to the strict-priority queue.
- 3. Configures two policers on the output interface that identify excess voice traffic belonging to the voice-class forwarding class. If the traffic exceeds 1 Mbps, a policer marks the traffic in excess of 1 Mbps as out-of-profile. If the traffic exceeds 2 Mbps, the second policer discards the traffic in excess of 2 Mbps.

To configure strict-priority queuing and prevent starvation of other queues:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 142 on page 333.
- 3. If you are finished configuring the router, commit the configuration.

#### **Table 142: Configuring Strict-High Priority Queuing and Starvation Prevention**

Task	J-Web Configuration Editor	CLI Configuration Editor			
Configuring a BA Classifier					

Task	J-W	eb Configuration Editor	CLI Configuration Editor		
Use a BA classifier to classify traffic based on the IP precedence of the packet. The classifier defines IP precedence value 101 as voice traffic	1.	In the J-Web interface, select Configuration > View and Edit > Edit	From the [edit] hierarchy level, enter		
		Configuration.	edit Class of service classifiers inet-precedence corp-traffic forwarding-class voice-class loss-priority low		
	2.	Next to Class of service, click <b>Configure</b> or <b>Edit</b> .			
and 000 as data traffic.	3.	Next to Classifiers, click <b>Configure</b> or <b>Edit</b> .	Enter set code-points 101		
	4.	Next to Inet precedence, click <b>Add new</b> entry.	From the [edit] hierarchy level, enter		
	5.	Enter corp-traffic in the Name box.	edit class-of-service classifiers inet-precedence		
	6.	Next to Forwarding class, click <b>Add new entry</b> .	corp-traffic forwarding-class data-class loss-priority high		
	7.	Enter voice-class in the Class name box.	Enter set code-points 000		
	8.	Next to Loss priority, click Add new entry.			
	9.	Enter low in the Loss val box.			
	10.	Next to Code points, click Add new entry.			
	11.	Enter 101 in the Value box.			
	12.	Click <b>OK</b> three times.			
	13.	In the Inet precedence forwarding class page, enter <b>voice-class</b> in the Class name box.			
	14.	Next to Loss priority, click Add new entry.			
	15.	Enter high in the Loss val box.			
	16.	Next to Code points, click Add new entry.			
	17.	Enter 000 in the Value box.			
	18.	Click <b>OK</b> five times.			
Configuring the Forwarding	g Cla	sses			

## Table 142: Configuring Strict-High Priority Queuing and Starvation Prevention (continued)

Task	J-W	eb Configuration Editor	CLI Configuration Editor			
Assign priority queuing to voice and data traffic.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit class-of-service forwarding-classes queue 0			
	2.	Next to Class of service, click Configure or Edit.voice-class enterNext to Forwarding classes, click Configure or Edit.edit class- data-classNext to Queue, click Add new entry.data-class	voice-class			
	3.		edit class-of-service forwarding-classes queue 1			
	4.		data-class			
	5.	Enter <b>0</b> in the Queue num box.				
	6.	Enter voice-class in the Class name box.				
	7.	Click <b>OK</b> to return to the Forwarding Classes page.				
	8.	Next to Queue, click Add new entry.				
	9.	Enter $1$ in the Queue num box.				
	10.	Enter data-class in the Class name box.				
	11.	Click <b>OK</b> three times.				
Configuring the Schedule	r Map	and Schedulers				
Configure the scheduler map and voice scheduler.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit class-of-service scheduler-maps corp-map			
	2.	Next to Class of service, click <b>Configure</b> or <b>Edit</b> .	forwarding-class voice-class			
	3.	Next to Scheduler maps, click <b>Add new</b> entry.	set scheduler voice-sched			
	4.	In the Map name box, type corp-map.				
	5.	Next to Forwarding class, click <b>Add new entry</b> .				
	6.	In the Class name box, type voice-class.				
	7.	In the Scheduler name box, type voice-sched.				
	8.	Click <b>OK</b> three times.				

## Table 142: Configuring Strict-High Priority Queuing and Starvation Prevention (continued)

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Define the voice and data traffic schedulers, and set	1.	In the J-Web interface, select Configuration > View and Edit > Edit	From the [edit] hierarchy level, enter
the priority.		Configuration.	edit class-of-service schedulers voice-sched
	2.	Next to Class of service, click <b>Configure</b> or <b>Edit</b> .	Enter
	3.	Next to Schedulers, click Add new entry.	set priority strict-high
	4.	In the Scheduler name box, type voice-sched.	From the [edit] hierarchy level, enter
	5.	In the Priority box, type strict-high.	edit class-of-service schedulers data-sched
	6.	Click <b>OK</b> .	Enter
	7.	Next to Schedulers, click Add new entry.	Litter
	8.	In the Scheduler name box, type data-sched.	set priority low
	9.	In the Priority box, type low.	
	10.	Click <b>OK</b> twice.	

#### Table 142: Configuring Strict-High Priority Queuing and Starvation Prevention (continued)

Applying the BA Classifier to an Input Interface and Scheduler Map to an Output Interface
Task         J-Web Configuration Editor		eb Configuration Editor	CLI Configuration Editor	
Apply the BA classifier to an input interface—for		In the J-Web interface, select Configuration > View and Edit > Edit	From the [edit] hierarchy level, enter	
example, ge-0/0/0.		Configuration.	edit interfaces ge-0/0/0 unit 0	
Apply the scheduler map	2.	Next to Interfaces, click <b>Configure</b> or <b>Edit</b> .	From the <b>[edit]</b> hierarchy level enter	
to and output	3.	Next to Interface, click Add new entry.		
interface—for example, e1-1/0/0.	4.	In the Interface name box, type ge-0/0/0.	edit class of service classifiers inet-precedence corp-traffic	
(Cas the interface neming	5.	Click <b>OK</b> three times.		
conventions in the <i>J</i> -series	6.	In the Edit Configuration page, next to Class of service, click <b>Configure</b> or <b>Edit</b>	From the [edit] hierarchy level, enter	
and WAN Access	7	Next to Classifiers, click Edit	edit interfaces e1-1/0/0 unit 0	
Configuration Guide.)	1.	Next to least procedures, click <b>Luit</b> .	From the [edit] hierarchy level, enter	
	8.	entry.		
	9.	In the Name box, type corp-traffic.	edit class-of-service scheduler-maps corp-map	
	10.	Click <b>OK</b> three times.		
	11.	In the Edit Configuration page, next to Interfaces, click <b>Configure</b> or <b>Edit</b> .		
	12.	Next to Interface name, type <b>e1-1/0/1</b> .		
	13.	Click <b>OK</b> twice.		
	14.	In the Edit Configuration page, next to Class of service, click <b>Configure</b> or <b>Edit</b> .		
	15.	Next to Scheduler maps, click <b>Add new</b> entry.		
	16.	In the Map name box, type <b>corp-map</b> .		
	17.	Click <b>OK</b> twice.		

### Table 142: Configuring Strict-High Priority Queuing and Starvation Prevention (continued)

**Configuring Two Policers** 

Table 142: Configuring Strict-High Priority Queuing and Starvation Prevention (continued)			
Task	J-Web Configuration Editor	CLI Configuration Editor	

TASK	J-VV	eb Configuration Eultor	CLI Configuration Editor
Configure two policers: one as voice-drop and second as voice-excess.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter
	2.	Next to Firewall, click <b>Configure</b> or <b>Edit</b> .	Pater
	3.	Next to Policer, click Add new entry.	Enter
	4.	In the Policer name box, type voice-drop.	set burst-size-limit 200000 bandwidth-limit 2000000
	5.	Next to If Exceeding, select the check box and click <b>Configure</b> .	Enter
	6.	In the Burst size limit box, type 200000.	set then discard
	7.	In the Bandwidth list, select <b>Bandwidth limit</b> .	From the [edit] hierarchy level, enter
	8.	In the Bandwidth limit box, type 2000000.	edit firewall policer voice-excess if-exceeding
	9.	Click <b>OK</b> .	-
	10.	On the Policer page, next to Then, click <b>Configure</b> .	Enter
	11.	Next to Discard, select the check box.	bandwidth-limit 1000000
	12.	Click Ok twice.	Enter
	13.	In the Firewall Configuration page next to Policer, click <b>Add new entry</b> .	set then out-of-profile
	14.	In the Policer name box, type voice-excess.	
	15.	Next to If Exceeding, select the check box and click <b>Configure</b> .	
	16.	In the Burst size limit box, type 200000.	
	17.	In the Bandwidth list, select <b>Bandwidth limit</b> .	
	18.	In the Bandwidth limit box, type 1000000.	
	19.	Click <b>OK</b> .	
	20.	On the Policer page, next to Then, click <b>Configure</b> .	
	21.	Next to Out of profile, select the check box.	
	22.	Click <b>OK</b> twice.	

Task	J-W	eb Configuration Editor	CLI Configuration Editor
Create a firewall filter voice-term that includes the	1.	In the Firewall Configuration page next to Filter, click <b>Add new entry</b> .	From the [edit] hierarchy level, enter
new policers.	2.	In the Filter name box, type voice-term.	edit firewall filter voice-term term 01 from
First, add the policer	3.	Next to Term click Add new entry.	voice-drop next term
voice-drop to the term.	4.	In the Rule name box, type term 01.	
	5.	Next to Term, click Add new entry.	
	6.	Next to From, click Configure.	
	7.	Next to Forwarding class choice, select <b>forwarding-class</b> .	
	8.	Next to Forwarding class, click <b>Add new entry</b> .	
	9.	In the String box, type voice-class.	
	10.	Click <b>OK</b> twice.	
	11.	In the Term Filter page, next to Then, click <b>Configure</b> .	
	12.	Next to Policer choice, select <b>policer</b> .	
	13.	In the Policer box, type voice-drop.	
	14.	Next to Designation, select <b>Next</b> .	
	15.	In the Next box, select term.	
	16.	Click <b>OK</b> twice.	
Then add the policer voice-excess to the term.	1.	In the Firewall Filter page, next to Term, click <b>Add new entry</b> .	Enter
	2.	In the Rule name box, type term 02.	edit firewall filter voice-term term 02 from
	3.	Next to From, click Configure.	voice-excess accept
	4.	Next to Forwarding class choice, select forwarding-class.	
	5.	Next to Forwarding class, click <b>Add new entry</b> .	
	6.	In the String box, type voice-class.	
	7.	Click <b>OK</b> twice.	
	8.	In the Term Filter page, next to Then, click <b>Configure</b> .	
	9.	Next to Policer choice, select <b>policer</b> .	
	10.	In the Policer box, type voice-excess.	
	11.	Next to Designation, select Accept.	
	12.	Click <b>OK</b> four times.	

### Table 142: Configuring Strict-High Priority Queuing and Starvation Prevention (continued)

Applying the Filter to the Output Interface

Task	J-Web Configuration Editor	CLI Configuration Editor	
Apply filter <b>voice-term</b> to <b>e1-1/0/0</b> using the CLI.		From the [edit] hierarchy level, enter	
,, ,		edit interfaces e1-1/0/1 unit 0 family inet filter output voice-term	
		Enter	
		set family inet address 11.1.1.1/24	

### Table 142: Configuring Strict-High Priority Queuing and Starvation Prevention (continued)

### **Configuring Large Delay Buffers with a Configuration Editor**

Large bursts of traffic from faster interfaces can cause congestion and dropped packets on slower interfaces that have small delay buffers. For example, a J-series Services Router operating at the edge of the network can drop a portion of the burst traffic it receives on a channelized T1/E1 interface from a Fast Ethernet or Gigabit Ethernet interface on a router at the network core.

To ensure that traffic is queued and transmitted properly on slower interfaces, you can configure a buffer size larger than the default maximum. On J-series Services Routers, you can configure large delay buffers on channelized T1/E1 interfaces only.

This section contains the following topics:

- Maximum Delay Buffer Sizes Available to Interfaces on page 340
- Delay Buffer Size Allocation Methods on page 341
- Specifying Delay Buffer Sizes for Queues on page 342
- Configuring a Large Delay Buffer on a Channelized T1 interface on page 343

### Maximum Delay Buffer Sizes Available to Interfaces

When you enable the large delay buffer feature on interfaces, a larger buffer is available for allocation to scheduler queues. The maximum delay buffer size that is available for an interface depends on the maximum available delay buffer time and the speed of the interface.

On channelized T1/E1 interfaces, the maximum delay buffer time varies by the number of DS0 channels configured on the interface as shown in Table 143 on page 341. The default values are as follows:

- Clear-channel interface—The default delay buffer time is 500,000 microseconds (0.5 seconds).
- NxDS0 interface—The default delay buffer time is 1,200,000 microseconds (1.2 seconds).

Channelized (NxDS0) Interfaces	Maximum Available Delay Buffer Time
1xDS0 through 3xDS0	4,000,000 microseconds (4 seconds)
4xDS0 through 7xDS0	2,000,000 microseconds (2 seconds)
8xDS0 through 15xDS0	1,000,000 microseconds (1 second)
16xDS0 through 32xDS0	500,000 microseconds (0.5 second)

#### **Table 143: Maximum Available Delay Buffer Time by Channels**

You can calculate the maximum delay buffer size available for an interface, with the following formula:

interface speed x maximum delay buffer time = maximum available delay buffer size

For example, the following maximum delay buffer sizes are available to 1xDS0 and 2xDS0 interfaces:

 $1 \times DS0 - 64$  kilobits per second x 4 seconds = 256 kilobits (32 kilobytes)

2xDS0—128 kilobits per second x 4 seconds = 512 kilobits (64 kilobytes)

If you configure a delay buffer size larger than the new maximum, the system allows you to commit the configuration but displays a system log warning message and uses the default buffer size setting instead of the configured maximum setting.

### **Delay Buffer Size Allocation Methods**

You can specify delay buffer sizes for each queue using schedulers. The queue buffer can be specified as a period of time (microseconds) or as a percentage of the total buffer or as the remaining buffer. Table 144 on page 341 shows different methods that you can specify for buffer allocation in queues.

#### **Table 144: Delay Buffer Size Allocation Methods**

Buffer Size Allocation Method	Description
Percentage	A percentage of the total buffer.
Temporal	A period of time, value in microseconds. When you configure a temporal buffer, you must also configure a transmit rate. The system calculates the queue buffer size by multiplying the available bandwidth of the interface times the configured temporal value and transmit rate.
	When you specify a temporal method, the drop profile is assigned a static buffer and the system starts dropping packets once the queue buffer size is full. By default, the other buffer types are assigned dynamic buffers that use surplus transmission bandwidth to absorb bursts of traffic.

### Table 144: Delay Buffer Size Allocation Methods (continued)

Buffer Size Allocation Method	Description
Remainder	The remaining buffer available. The remainder is the percentage buffer that is not assigned to other queues. For example, if you assign 40 percent of the delay buffer to queue 0, allow queue 3 to keep the default allotment of 5 percent, and assign the remainder to queue 7, then queue 7 uses approximately 55 percent of the delay buffer.

### **Specifying Delay Buffer Sizes for Queues**

You specify delay buffer sizes for queues using schedulers. The system calculates the buffer size of a queue based on the buffer allocation method you specify for it in the scheduler. See Table 144 on page 341 for different buffer allocation methods and Table 145 on page 342 for buffer size calculations.

### **Table 145: Delay Buffer Allocation Method and Queue Buffer**

Buffer Size Allocation Method	Queue Buffer Calculation	Example
Percentage	available interface bandwidth x configured buffer size percentage x maximum delay buffer time = queue buffer	Suppose you configure a queue on a 1xDS0 interface to use 30 percent of the available delay buffer size. The system uses the maximum available delay buffer time (4 seconds) and allocates the queue 9600 bytes of delay buffer:
		64 Kbps x 0.3 x 4 seconds = 76800 bits = 9600 bytes
Temporal	available interface bandwidth x configured transmit rate percentage x configured temporal buffer size = queue buffer	Suppose you configure a queue on a 1xDS0 interface to use 300,000 microseconds (3 seconds) of delay buffer, and you configure the transmission rate to be 20 percent. The queue receives 4800 bytes of delay buffer:
		64 Kbps x 0.2 x 3 seconds = 38400 bits = 4800 bytes
		When you configure a temporal value that is greater than the maximum available delay buffer time, the system allocates this queue the remaining buffer after other queues are allocated buffer. Suppose you configure a temporal value of 6,000,000 microseconds on a 1xDS0 interface. Because this value is greater than the maximum allowed value of 4,000,000 microseconds, the queue is allocated the remaining delay buffer.

When you specify the buffer size as a percentage, the system ignores the transmit rate and calculates the buffer size based only on the buffer size percentage.

### Configuring a Large Delay Buffer on a Channelized T1 interface

On J-series Services Routers you can configure large delay buffers on channelized T1/E1 interfaces only. To configure large-delay buffer sizes, you must first enable the large buffer feature on the channelized T1/E1 PIM and then configure a buffer size for each queue in the CoS scheduler.

Each channelized T1/E1 interface can be configured as a single clear channel, or for channelized (NxDS0) operation, where N denotes channels 1 to 32 for an E1 interface and channels 1 to 24 for a T1 interface.

In this configuration, you enable the large delay buffer option on a channelized T1 PIM with an interface speed of 1.5 Mbps and a maximum delay buffer time of 500,000 microseconds. Based on the interface speed and the maximum delay buffer time, you can calculate the available delay buffer size for the interface. For more information, see "Maximum Delay Buffer Sizes Available to Interfaces" on page 340.

Next, you specify a queue buffer of 30 percent in a scheduler **be-scheduler** and associate the scheduler to a defined forwarding class **be-class** using a scheduler map **large-buf-sched-map**. Finally, you apply the scheduler map to the channelized T1 interface **t1-3/0/0**. As a result, a buffer of 9600 bytes is assigned to the queue associated with forwarding class **be-class** (see Table 145 on page 342). You can specify a delay buffer size for other queues following the instructions in this example.

To configure large delay buffers for channelized T1/E1 interfaces:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- 2. Perform the configuration tasks described in Table 146 on page 343.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to one of the following tasks:
  - To configure other CoS components, see "Configuring CoS Components with a Configuration Editor" on page 305.
  - From the CLI, enter the **show class of service** command, to check your configuration.

### Table 146: Configuring a Large Delay Buffer

Task	J-W	/eb Configuration Editor	CLI Configuration Editor
Navigate to the <b>Chassis</b> level in the configuration hierarchy.	1.	In the J-Web interface, select Configuration > View and Edit > Edit Configuration.	From the [edit] hierarchy level, enter edit chassis
	2.	Next to Chassis, click <b>Configure</b> or <b>Edit</b> .	

### Table 146: Configuring a Large Delay Buffer (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Enable the large buffer size	1. Next to Fpc, click <b>Add new entry</b> .	Enter
feature on the channelized T1/E1 PIM in slot 3.	2. In the Slot box, type the slot number <b>3</b> .	set foc 3 pic 0 o-pic-large-buffer
	3. Next to Pic, click <b>Add new entry</b> .	
	4. In the Slot box, type <b>0</b> .	
	5. Next to Q pic large buffer, select the chee	ck box.
	6. Click <b>OK</b> .	
Navigate to the <b>Class-of-service</b> level in the configuration hierarchy.	On the main Configuration page next to Classervice, click <b>Configure</b> or <b>Edit</b> .	ss of From the [edit] hierarchy level, enter edit class-of-service
Create be-scheduler and	1. Next to Schedulers, click Add new entr	ry. Enter
specify a buffer size of 30 percent for it.	<ol> <li>In the Scheduler name box, type the na the scheduler—be-scheduler.</li> </ol>	ame of set schedulers be-scheduler buffer-size percent 30
	3. Next to Buffer size, click <b>Configure</b> .	
	4. From the Buffer size choice list, select <b>p</b>	ercent.
	5. In the Percent box, type <b>30</b> .	
	6. Click <b>OK</b> .	
Configure the scheduler map large-buf-scheduler-map	<ol> <li>On the Class of service page, next to Sch maps, click Add new entry.</li> </ol>	neduler From the [edit class-of-service] hierarchy level, enter
For information about classes.	<ol> <li>In the Map name box, type the name o scheduler map—large-buf-sched-map.</li> </ol>	f the set scheduler-maps large-buf-sched-map forwarding-class be-class scheduler
	3. Next to Forwarding class, click Add new	entry. be-scheduler
	<ol> <li>In the Class name box, type the name of forwarding class to be associated with t scheduler—be-class.</li> </ol>	of the he
Output Queues" on page 310.	<ol> <li>In the Scheduler box, type the name of scheduler to be associated with the forw class—be-scheduler.</li> </ol>	the varding
	6. Click <b>OK</b> .	
Apply the scheduler map to the channelized T1	1. On the Class of service page, next to Inte click <b>Add new entry</b> .	erfaces, From the [edit class-of-service] hierarchy level, type
interface. <b>NOTE:</b> For information about configuring channelized T1/E1	<ol> <li>In the Interface name box, type the nar the interface to which the scheduler ma be applied—t1-3/0/0.</li> </ol>	ap is to scheduler-map large-buf-sched-map
	3. Next to Unit, click Add new entry.	
Services Router Basic LAN	4. In the Unit number box, type <b>0</b> .	
and WAN Access Configuration Guide.	<ol> <li>In the Scheduler map box, type the nar the scheduler map—large-buf-sched-map</li> </ol>	ne of D.
	6. Click <b>OK</b> .	

### **Verifying a CoS Configuration**

To verify a CoS configuration, perform the tasks relevant to your CoS configuration from the following:

- Verifying Multicast Session Announcements on page 345
- Verifying a Virtual Channel Configuration on page 345
- Verifying a Virtual Channel Group Configuration on page 346
- Verifying an Adaptive Shaper Configuration on page 346

### **Verifying Multicast Session Announcements**

**Purpose** Verify that the Services Router is listening to the appropriate groups for multicast Session Announcement Protocol (SAP) session announcements.

Action From the CLI, enter the show sap listen command.

user@host> **show sap listen** Group Address Port 224.2.127.254 9875

- **What It Means** The output shows a list of the group addresses and ports that SAP and SDP listen on. Verify the following information:
  - Each group address configured, especially the default 224.2.127.254, is listed.
  - Each port configured, especially the default **9875**, is listed.
- **Related Topics** For a complete description of the **show sap listen** command and output, see the *JUNOS Routing Protocols and Policies Command Reference*.

### **Verifying a Virtual Channel Configuration**

- **Purpose** Verify the virtual channel configuration on a logical interface. Verify the class-of-service (CoS) configuration associated with an interface.
  - Action From the CLI, enter the show class-of-service virtual-channel command.

user@host> **show class-of-service virtual-channel** Virtual channel: vc-1 Index: 1

- What It Means Verify that the name of the configured virtual channel is displayed in the output.
- **Related Topics** For a complete description of the show class-of-service virtual-channel command and output, see the *JUNOS System Basics and Services Command Reference*.

### **Verifying a Virtual Channel Group Configuration**

- Purpose Verify the virtual channel group configuration on a logical interface. Verify the class-of-service (CoS) configuration associated with an interface.
- Action From the CLI, enter the show class-of-service virtual-channel-group command.

user@host> show class-of-service virtual-channel-group Virtual channel group: vc-group, Index: 16321 Virtual channel: vc-1 Scheduler map: sc-map

- What It Means Verify that the name of the configured virtual channel group is displayed in the output.
- For a complete description of the show class-of-service virtual-channel-group command **Related Topics** and output, see the JUNOS System Basics and Services Command Reference.

### **Verifying an Adaptive Shaper Configuration**

- Purpose Verify the adaptive shaper trigger point and its associated transmit rate. Verify the class-of-service (CoS) configuration associated with an interface.
- Action From the CLI, enter the show class-of-service adaptive-shaper and show class-of-service interface t1-0/0/2 commands.

user@host> <b>show clas</b>	s-of-service adaptive-s	haper	
Adaptive shaper: fr-	shaper, Index: 35320		
Trigger type Sh	aping rate		
BECN	64000 bps		
user@host> <b>show clas</b> Physical interface: Queues supported: 8,	<b>s-of-service interface</b> t1-0/0/2, Index: 137 Queues in use: 4	t1-0/0/2	
Scheduler map: <de< td=""><td>efault&gt;, Index: 2</td><td></td><td></td></de<>	efault>, Index: 2		
Logical interface:	t1-0/0/2.0, Index: 69		
Object Adaptive shaper	Name fr.shaper	Туре	Index
Classifier	ipprec-compatibility	ip	11

- What It Means Verify the following information:
  - The trigger type and shaping rate are consistent with the configured adaptive shaper.

- The adaptive shaper applied to the logical interface is displayed under Name.
- **Related Topics** For a complete description of the show class-of-service adaptive-shaper and show class-of-service interface commands and output, see the JUNOS System Basics and Services Command Reference.

# Part 6 Index

■ Index on page 349

J-series[™] Services Router Advanced WAN Access Configuration Guide

# Index

# **Symbols**

#, comments in configuration statements	xx
(), in syntax descriptions	xx
*,G notation, for multicast forwarding states	105
3DES-CBC algorithm	70
< >, in syntax descriptions	xx
[], in configuration statements	xx
{ }, in configuration statements	xx
(pipe), in syntax descriptions	xx

### A

accept, filter action	234
access control lists (ACLs) See stateless firewall filte	rs
ACLs See stateless firewall filters	
action modifiers, stateless firewall filters	
list of	162
setting	235
See also actions	
Action tab, stateless firewall filters	234
actions	
accept, setting	234
count modifier, setting	235
default, routing policy	151
discard, setting	234
final, routing policy	151
forwarding class modifier, setting	235
log modifier, setting	235
loss priority modifier, setting	235
modifiers, list of	162
NAT, list of	167
next term, setting	234
no action, setting	234
reject, setting	234
route list match types	172
routing instance, setting	234
routing policy	153
routing policy, summary of	154
sample modifier, setting	235
stateful firewall filters, list of	157
stateless firewall filters, list of	162
stateless firewall filters, setting actions (Quick	
Configuration)	234
-	

stateless firewall filters, setting modifiers (Qu	ıick
Configuration)	235
syslog modifier, setting	235
virtual channel modifier, setting	235
adaptive shaping	
applying CoS rules to logical interfaces	331
verifying	346
address match conditions	161
address translation See NAT	
addresses	
multicast ranges	104
translating See NAT	
administrative groups, for MPLS path selection	14
administrative scoping	106
Advanced Encryption Standard (AES)	70
AES algorithm	70
AF forwarding class See assured forwarding forwar	ding
class	0
AH (Authentication Header) protocol, IPSec	71
aliases, CoS See CoS value aliases	
AS path, prepending	177
ASs (autonomous systems)	
AS number, in VPNs	40
LSPs through	6
assured forwarding (AF) forwarding class	276
RED drop profiles for	
See also CoS; forwarding classes	
authentication algorithms, IPSec	
Authentication Header (AH) protocol, IPSec	71
Auto-RP	

## В

BA classifiers See classifiers	
bandwidth, for RSVP-signaled LSPs	25
BE forwarding class <i>See</i> best-effort forwarding class	
behavior aggregate classifiers See classifiers	
best-effort (BE) forwarding class	
default assignment	.276
See also CoS; forwarding classes	
typical usage	.264
BGP (Border Gateway Protocol)	
export policy for CLNS	60
for CLNS VPN NLRI	63
injecting OSPF routes into BGP	.174
policy to make routes less preferable	.177

route-flap damping	179
VPNs	39
BGP confederations	
route-flap damping	179
bit-field logical operators, stateless firewall filters	162
bit-field match conditions	161
bit-field synonym match conditions	161
bootstrap router	108
braces, in configuration statements	xx
brackets	
angle, in syntax descriptions	xx
square, in configuration statements	xx
branches	104
See also multicast	
BSR (bootstrap router)	108

# С

Class of Service scheduler maps page	.294
field summary	.299
Class of Service schedulers page	.294
field summary	.297
Class of Service virtual channel groups page	.300
field summary	.301
classifiers	
adding and editing (Ouick Configuration)	.291
applying behavior aggregate classifiers	316
assigning to logical interfaces (Ouick	
Configuration)	305
hehavior aggregate	266
default behavior aggregate classifiers	277
default DSCP CoS classifier for DI Sw	136
defining (Quick Configuration)	290
description	.2.70
multifield classifiers	267
sample behavior addregate classification	.207
sample behavior aggregate classification	.219
	716
assignments	200
sample, for mewall meet	.308
strict high-priority queuing (configuration	774
ealtor)	.334
strict nigh-priority queuing, applying classifier to	) 
interface (configuration editor)	.221
summary (Quick Configuration)	.290
clear services ipsec-vpn certificates service-set	07
command	97
time	740
CLL configuration oditor	.940
CLI COMIGUIATION EURO	57
CLINS	
CoS large delay byffere	.505
CoS, large delay bullers	.340
DLSw (basis)	.332
DLSW (Dasic)	.121
DLSW COS	.134
IPSec tunnels	/5
MPLS traffic engineering	20
multicast network	.110
routing policies	.170
stateful firewall filters	.211
stateless firewall filters	.238
VPNs	34
CLNS (Connectionless Network Service) VPNs	
BGP export policy	60
BGP, to carry CLNS VPN NLRI	63
displaying configurations	63
ES-IS	59
IS-IS.	60
linking hosts	55
overview	56
requirements	57
static routes (without IS-IS)	62
verifying configuration	63
VPN routing instance	58

coloring, link, for MPLS path selection14
comments, in configuration statementsxx
Common Criteria environment, stateless firewall filters
in221
congestion control
with CoS schedulers (Quick Configuration)294
with DiffServ assured forwarding (configuration
editor)
Connectionless Network Service See CLNS
Constrained Shortest Path First See CSPF
conventions
how to use this guidexviii
notice icons xix
text and syntax xix
CoS (class of service)
adaptive shaping for rules 331
aliases See CoS value aliases
assigning components to interfaces (Quick
Configuration) 302
assigning forwarding classes to output
behavior addregate classifiers See classifiers
benefits 264
classifiers Saa classifiers
configuration tacks (configuration oditor) 305
configuration tasks (Quick Configuration) 284
CoS process (ILINOS implementation) 271
Cos process (joinos implementation)271
CoS value allases see CoS value allases
CoS value fewfiles
default DCCD algorities for DLCu
default DSCP classifier for DLSW
default scheduler settings See schedulers
default settings
DI Sus na classification of
DLSW packets, classification of
firewall filter for a multifield classifier
forwarding classes See forwarding classes
interfaces, assigning components to (Quick
Configuration)
JUNOS components
JUNOS implementation
large delay buffers (configuration editor)
overview
See also Class of Service pages
policer for firewall filter
preparation283
Quick Configuration
RED drop profiles See RED drop profiles
rewrite rules See rewrite rules
sample behavior aggregate classification279
scheduler maps See scheduler maps
schedulers See schedulers
slower interfaces, enlarging delay buffers for
(configuration editor)340
starvation prevention for queues (configuration
editor)332

strict high phoney for queuing (configuration	
editor)	332
ToS value for DLSw	136
traffic flow	265
transmission scheduling	280
uses	283
verifying adaptive shaper configuration	346
verifying multicast session announcements	345
verifying virtual channel configuration	345
verifying virtual channel group	
configuration	346
virtual channel groups (Quick	
Configuration)	300
See also virtual channels	
virtual channels for rules See virtual channels	
CoS components	
classifiers	266
code-point alias	266
forwarding classes	267
forwarding policies	267
loss priorities	267
policers	270
RED drop profiles	270
rewrite rules	271
schedulers	268
shaping rate	269
transmission queues	268
virtual channels	270
CoS process	
incoming packets	272
outgoing packets	272
overview (IUNOS implementation)	271
CoS value aliases	
adding (Ouick Configuration)	
default values	
rewrite rules	279
summary (Quick Configuration)	287
CoS values <i>See</i> CoS value aliases	
CoS-based Forwarding (CBF)	267
count_filter action modifier	235
CRLs (certificate revocation lists)	71
CSPF (Constrained Shortest Path First)	
constraints	13
disabling	15 25
link coloring	20
rules	11 13
CSPE algorithm See CSPE	
curly braces in configuration statements	vv
customer edde routers See CE routers	лх
cusionici cuge iouleis see CE Iouleis	vviii
customer support	
customer support	

D
Data Encryption Standard-cipher block chaining
(DES-CBC)70

# data link switching *See* DLSw defaults

uerauits	
behavior aggregate classifiers	278
CoS forwarding class assignments276,	277
DSCP classifier for DLSw	136
junos-algs-outbound group, stateful firewall	
filters	211
routing policy actions	151
delay buffer size	
allocation methods	341
calculation	342
description	269
enlarging	340
enlarging (configuration editor)	343
maximum available	340
denial-of-service attacks preventing	241
dense routing mode, caution for use	105
See also multicast routing modes	105
DEC CPC algorithm	70
designated router, stopping outgoing DIM register	70
messades on	116
messages on	110
destination static NAT	
description	165
example	165
diagnosis	
displaying CLNS VPN configurations	63
displaying stateful firewall filter	
configurations	217
displaying stateless firewall filter	
configurations	252
displaying stateless firewall filter statistics	256
LDP neighbors	25
LDP sessions	26
LDP-signaled LSP	27
RSVP neighbors	28
RSVP sessions	28
RSVP-signaled LSP	29
traffic forwarding over LDP-signaled LSPs	27
verifying adaptive shaper configuration	346
verifying DLSw capabilities	142
verifying DLSw circuit state	143
verifying DLSw circuit state (detail)	143
verifying DLSw Ethernet redundancy interface	
statistics	146
verifying DLSw Ethernet redundancy status	146
verifying DISw LLC2 properties	142
verifying DLSW DEE2 properties	144
verifying DLSw peers (detail)	144
verifying DLSw peers (detail)	145
verifying besw reachability	250
verifying mewan mer nanues nagments	209
verifying if see turner operation.	70 75
venitying WIPLS traine engineering	
verifying multicast IGMP versions	119
verifying multicast SAP and SDP	
configuration	119
verifying multicast session announcements	345

verifying NAT configurations
verifying PIM mode and interface
configuration120
verifying PIM RPF routing table121
verifying PIM RPs120
verifying stateful firewall filters
verifying stateless firewall filter actions
verifying stateless firewall filter DoS
protection258
verifying stateless firewall filter flood
protection258
verifying stateless firewall filters with packet
verifying virtual channel configuration 345, 346
verifying VIItual channel configuration
Differentiated Services See Different
Diffic Hollman exchange IPSec 72
Different (Differentiated Services)
Dilisely (Dilieleninated Services)
queues
behavior addregate classifiers
Denavior aggregate classifiers
Configuration tasks (configuration editor)
interes erability
Interoperability
JUNOS Implementation
policer for lifewall lifter
RED drop promes
rewrite rules
scheduler maps
schedulers
digital contificator
Chaptificates loading on the neuton
CA certificate, loading on the router
captificate authority (CA)
Certificate authority (CA)70
See also CA
Configuring for IPSec tunnels
CRLS
dependence of the second private laws
key pair generating
local cortificate, applying to ap IPSec tuppel
local certificate, applying to an iPSec turner
local certificate, enrorating
local certificate, generating
requesting from a CA
requesting from a CA
revocation of
discard rule Saa discard filter action
discard filter action
automatic stateful firewall filters
automatic, stateless firewall filters
stateless firewall filters (Quick
Configuration) 274
Distance Vector Multicast Routing Protocol 107
Distance vector municast nouting i fotocol

DLSw (data link switching)	
basic configuration (configuration editor)131	1
basic configuration (Quick Configuration)129	9
canureach message128	3
capabilities exchange128	3
circuit establishment128	3
CoS classification of DLSw packets	4
DLSw MIB 125	5
Ethernet redundancy <i>See</i> DLSw Ethernet	
redundancy	
icanreach message 128	2
idle timeout	1
LLC type 2 properties on Ethernet	
interfaces 131	1
IIIterraces	l
interfaces 177	2
Interfaces	2
load balancing See DLSW load balancing	<b>`</b>
local router configuration	2
monitoring capabilities125	)
overview127	(
peers See DLSw peers	_
preparation129	)
promiscuous mode131	1
Quick Configuration129	9
reachability cache, clearing141	1
remote router configuration134	1
sample DLSw network127	7
sample peer router values133	3
SNA forwarding127	7
SSP127	7
stages of operation128	3
ToS precedence for DLSw packets134	1
verifying capabilities142	2
verifying DLSw circuit state143	3
verifying DLSw circuit state (detail)143	3
verifying DLSw peers144	4
verifying DLSw peers (detail)144	4
verifying DLSw reachability145	õ
verifying Ethernet redundancy interface	
statistics146	5
verifying Ethernet redundancy status	5
verifying LLC2 properties	2
DLSw Ethernet redundancy	
configuring 138	R
network topology 137	7
overview 136	, 5
verifying interface statistics	5
verifying status	5
DI Sw load balancing	)
configuring 141	1
network topology	י ר
overview 170	י כ
	ッ つ
field summary	ן 1
neia summary191	L

DLSw peers	
local (configuration editor)	
local (Quick Configuration)	131
remote (configuration editor)	134
remote (Quick Configuration)	131
setting a preference for	
verifving	144
verifying (detail)	144
documentation set	
comments on	xxiii
DoS (denial-of-service) attacks, preventing	241
downstream interfaces	
See also multicast	
DR See designated router	
drop profiles <i>See</i> CoS; RED drop profiles	
DS0 interfaces, maximum delay buffer time	
DSCP IPv6 See CoS: DSCPs	
DSCPs (DiffServ code points)	
default behavior aggregate classifiers	
default classifier for DLSw	
DSCP aliases and values	274
See also CoS	
matching with a filter	232
matching with an IPv4 filter	232
replacing with rewrite rules	
rewrites	
sample behavior aggregate classification	279
DVMRP (Distance Vector Multicast Routing	
Protocol)	
dvnamic LSPs	9
dynamic SAs	
creating (configuration editor)	77
IKE policy (configuration editor)	80
IKE proposal (configuration editor)	
IPSec policy (configuration editor)	82
IPSec proposal (configuration editor)	81
IPSec rules (configuration editor)	83
IPSec services interfaces (configuration	
editor)	
overview	72
service sets (configuration editor)	
See also IPSec service sets	

# Е

# F

Fast Ethernet ports, LLC type 2 properties for DLSw	
See LLC	
filters See firewall filters; stateful firewall filters;	
stateless firewall filters	
firewall filters	
applying CoS rules to logical interfaces	327
in a Common Criteria environment	221
multifield classifier filter terms	307
overview	149
policer for	306
sample classifier terms	308
stateful firewall filters	155
See also stateful firewall filters	
stateless firewall filters	157
See also stateless firewall filters	
term number caution156,	158
verifying fragment handling	259
Firewall Filters configuration pages	
field summary	225
initial page	223
match conditions and actions page	224
Firewall Filters interface assignment pages	
available interfaces and filter status page	236
field summary	237
Firewall/NAT application page	208
field summary	209

Firewall/NAT main page	207
field summary	209
flap damping	179
parameters	179
flooding, preventing	241
flow control, actions in routing policies	154
font conventions	xix
forwarding classes	
adding and editing (Quick Configuration)	
assigning to logical interfaces (Quick	
Configuration)	
assigning to output queues (configuration	
editor)	311
assigning to output queues (Quick	
Configuration)	
default assignments	277
default values	276
defining (Quick Configuration)	
description	
filter action modifier, setting	235
mapping to schedulers (configuration	
editor)	325
matching with a filter	233
policy to group source and destination	
prefixes	176
queue assignments, default	276
sample behavior aggregate classification	279
sample mappings	
summary (Quick Configuration)	
forwarding policy options	267
forwarding states, multicast notation	105
fragment offset, matching with a filter	232
Frame Relay, CoS adaptive shaping for	331
from statement, routing policy match conditions	s152
full-cone NAT	
basic configuration (configuration editor)	190

# G

gateway, IPSec tunnel mode for	73
See also IPSec tunnels	
gateway, local and remote, for IPSec service sets	86
ge-0/0/0, disabling PIM on	113
glossary	
CLNS	55
CoS	263
DLSw	126
firewall filters	149
IPSec	67
MPLS	3
multicast	101
NAT	149
routing policies	149
VPNs	3
groups, default junos-algs-outbound group, for stateful	
firewall filters	211

# Н

handling packet fragments	
high-priority CoS queuing	
host, IPSec transport mode for	73
how to use this guide	xviii

# I

IBM networking See DLSw
icanreach message, DLSw128
ICMP (Internet Control Message Protocol),
policers
ICMP packets, matching with a filter231
idle timeout, DLSw131
IEEE 802.1 CoS value type, aliases and values275
See also CoS
IGMP (Internet Group Management Protocol)
IGMPv1107
IGMPv2107
IGMPv3108
setting the version111
verifying the version119
IGPs (interior gateway protocols)41
VPNs
See also OSPF
IKE (Internet Key Exchange)
description72
dynamic SAs72
IKE policy, configuring80
IKE proposal, configuring78
negotiation phases72
preshared key (configuration editor)81
preshared key (Quick Configuration)75
import routing policy, for Layer 2 VPNs48
import statement, for routing policies152
inbound router, in an LSP7
inet routing table117
ingress router See inbound router; LSPs
injecting routes175
input filters, assigning to interfaces237
interface groups, matching with a filter230
interface service set, for IPSec tunnels
interface sets, matching with a filter229
Intermediate System-to-Intermediate System See IS-IS
internal networks, access with NAT See NAT
Internet Control Message Protocol policers243
Internet Group Management Protocol See IGMP
Internet Key Exchange See IKE
invalid routes, rejecting173
IP addresses, translation with NAT See NAT
IP options, matching with a filter233
IP precedence CoS value type, aliases and values275
See also LOS
IF Security see IPSec

IPSec (IP Security)
AH traffic protection protocol71
authentication algorithms69
authentication methods70
digital certificates authentication70
See also digital certificates
dynamic SAs (configuration editor)77
dynamic SAs for large-scale networks72
encryption algorithms69
ESP SPI values, matching with a filter234
ESP traffic protection protocol71
IKE See IKE
IKE policy (configuration editor)80
IKE proposal (configuration editor)78
IPSec policy (configuration editor)82
IPSec proposal (configuration editor)81
IPSec rules (configuration editor)83
NAT pools (configuration editor)90
overview69
preshared key authentication70
protocol bundle traffic protection72
requirements73
security associations See dynamic SAs; IPSec
security associations
service sets (configuration editor)
See also IPSec service sets
services interfaces (configuration editor)
Services Router as secure gateway or host73
traffic protection protocols71
transport mode73
tunnel mode73
See also IPSec tunnels
verifying tunnels
IPSec security associations
manual SAs76
overview72
See also dynamic SAs; IKE; IPSec tunnels
IPSec service sets
applying rules (configuration editor)
Interface service set (configuration editor)
local gateway (configuration editor)
next-nop services interface (configuration
ealtor)
overview
See also aynamic SAS
digital contificator for
See also digital contificatos
dunamic SAs (confiduration editor)
IKE key (configuration editor)
INE REV (Outer Confiduration)
IKE policy (configuration editor)
IKE proposal (configuration editor) 79
IPSec policy (configuration editor)
IPSec proposal (configuration editor)
IPSec rule (configuration editor)
in See rule (configuration cultor)

IPSec rules (configuration editor)	J-
local endpoint (Quick Configuration)75	
NAT pools (configuration editor)	
private addresses (Quick Configuration)75	
Quick Configuration	
remote endpoint (Quick Configuration)75	
requirements73	
security associations (configuration editor)76	
services interfaces (configuration editor)	
services sets (configuration editor)	
See also IPSec service sets	
verifying98	
VPN policy for digital certificates96	
IPSec Tunnels page	
field summary75	
IPv4 filters	JU
assigning to interfaces (Quick	
Configuration)236	JU
creating and editing (Quick Configuration)222	
See also stateless firewall filters	
IPv6 filters	ju
assigning to interfaces (Quick	
Configuration)236	
creating and editing (Quick Configuration)222	
See also stateless firewall filters	K
IS-IS (Intermediate System-to-Intermediate System)	k
for CLNS route exchange60	k
with CLNS56	

# J

Lseries	
CLNS VPNs.	
CoS	
CoS overview	
DLSw	
firewall filter overview	149
IBM networking	125
IPSec	67
MPLS for VPNs overview	
MPLS traffic engineering	19
multicast	109
multicast overview	101
NAT	185
NAT and stateful firewall filters	205
NAT overview	163
policy framework overview	149
release notes, URL	xvii
routing policies	169
routing policy overview	151
stateful firewall filters	205
stateful firewall filters overview	155
stateless firewall filter overview	157
stateless firewall filters	
VPNs	

I-Web d	configuration	editor
---------	---------------	--------

CLNS	57
CoS	
CoS, large delay buffers	
CoS, strict high priority for queuing	
DLSw (basic)	131
DLSw CoS	134
IPSec tunnels	75
MPLS traffic engineering	20
multicast network	110
NAT	185
routing policies	170
stateful firewall filters	211
stateless firewall filters	238
VPNs	34
JUNOS Internet software	
release notes, URL	xvii
JUNOS software	
CoS components	266
CoS implementation	271
junos-algs-outbound group, for stateful firewa	all
filters	211

# Κ

kee	epalive interval, for LDP-signaled LSPs	22
key	ys	
	preshared	70
	See also preshared keys	
	public, for digital certificates	92
	public-private key pair, generating	94

# L

Label Distribution Protocol See LDP	
label switching	6
label-switched paths See LSPs	
label-switching routers (LSRs)	7
labels, MPLS	8
label operations	8
РНР	9
Layer 2 circuits	
AS number	40
basic, description	
encapsulation	
IGPs	
MPLS	
neighbor address	44
participating interfaces	
signaling protocols	41
task overview	
verifying PE router connections	53
verifying PE router interfaces	53
virtual circuit ID	44

Layer 2 VPNs
AS number40
basic, description32
BGP
encapsulation
export routing policies49
IGPs41
import routing policies48
MPLS
overview17
participating interfaces35
routing instance45
signaling protocols41
task overview
verifying PE router connections53
verifying PE router interfaces53
Layer 3 VPNs
AS number40
basic, description
BGP
IGPs41
overview17
participating interfaces
route target
routing instance
routing policies
signaling protocols
task overview
Verifying PE router connections
and OSEE for VENa
LDP signaled LSPs 21
LDF-Signated LDFS
Incessages
overview 20
requirements 20
verifying LSPs 27
verifying heighbors 25
verifying sessions 26
verifying traffic forwarding 27
LDP neighbors, verifying
LDP-signaled LSP See LDP
leaves
See also multicast
link coloring, for MPLS path selection14
LLC (Logical Link Control) type 2 properties for DLSw
verification142
LLC (Logical Link Control) type 2 properties for DLSw
setting (configuration editor)132
setting (Quick Configuration)131
load balancing, DLSw See DLSw load balancing
local digital certificate See digital certificates
local gateway, for IPSec tunnels86
local router, DLSw132
See also DLSw peers
local tunnel endpoint, IPSec75

logging packet header information		.235
logical interfaces		
adaptive shaping for		.331
adding and editing CoS components (Quick		
Configuration)		.304
assigning CoS components to (Quick		
Configuration)		.302
CoS rules for	327,	331
inside services interface, IPSec		84
outside services interface, IPSec		85
virtual channels for		.327
Logical Link Control (LLC) type 2 properties for I	DLSw	V
See LLC		
longer route list match type		.172
loopback address, for PE routers in VPNs		41
loopback interface, applying stateless firewall filte	ers to	С
(configuration editor)		.251
loose hops, RSVP		12
loss priorities		.267
LSPs (label-switched paths)		
bandwidth		25
description		6
disabling CSPF		25
dynamic LSPs		9
for RSVP in a VPN		38
keepalive interval for LDP link		22
label operations		8
label switching		6
labels		8
LDP		10
LDP-signaled LSPs		21
LSR types		7
overview		19
PHP		9
RSVP		11
RSVP-signaled LSPs		23
static LSPs		9
verifying LDP-signaled LSPs		25
verifying RSVP-signaled LSPs		27
LSRs (label-switching routers)		7

### Μ

Management Information Base See MIB	
management interfaces, disabling PIM on	113
manual SAs	
creating (configuration editor)	76
overview	76
manuals	
comments on	xxiii
mapping, CoS forwarding classes to	
schedulers	294, 325
match conditions	
NAT	167
routing policy	152
routing policy, summary of	152

stateful firewall filters150	6
stateless firewall filters15	9
stateless firewall filters, summary15	9
Match Destination tab, stateless firewall filters22	7
Match Interface tab, stateless firewall filters229	9
Match Network tab, stateless firewall filters230	0
Match Packet and Network tab, stateless firewall	
filters230	0
Match Source or Destination tab, stateless firewall	
filters 228	8
Match Source tab. stateless firewall filters	6
match types 17	2
messages IDP 1	1
MF classifier 30	7
MIR (Management Information Base) DI Sw 12	5
MPLS (Multiprotocol Label Switching)	л
dynamic LSPc	ч 0
label enerations	7
label operations	0
labels	0
labels	8
Layer 2 VPNs and Layer 2 circuits	7
LDP	0
LSP for RSVP in a VPN	8
LSPs	6
LSR types	7
overview	3
PHP	9
RSVP1	1
static I SPs	0
Static Lor S	9
traffic engineering <i>See</i> MPLS traffic engineering	9
traffic engineering <i>See</i> MPLS traffic engineering verifying	9 5
traffic engineering See MPLS traffic engineering verifying	5
traffic engineering <i>See</i> MPLS traffic engineering verifying	9 5 5
traffic engineering See MPLS traffic engineering verifying	9 5 5
traffic engineering See MPLS traffic engineering verifying	9 5 5
traffic engineering See MPLS traffic engineering verifying	9 5 5 0
traffic engineering See MPLS traffic engineering verifying	9 5 5 0 1
traffic engineering See MPLS traffic engineering verifying	9 5 5 0 1 9
traffic engineering See MPLS traffic engineering verifying	5 5 0 1 9 0
traffic engineering See MPLS traffic engineering verifying	5 5 0 1 9 0 0
traffic engineering See MPLS traffic engineering verifying	5 5 0 1 9 0 0 3
traffic engineering <i>See</i> MPLS traffic engineering verifying	5 5 0 1 9 0 3 0
traffic engineering <i>See</i> MPLS traffic engineering verifying	9     5     5       0     1     9       0     3     0
traffic engineering <i>See</i> MPLS traffic engineering verifying	5 5 019003056
traffic engineering <i>See</i> MPLS traffic engineering verifying	5     5       0     1       9     0       0     1       9     0       0     1       9     0       0     1       9     0       0     1       9     0       0     1       9     0       0     1       9     0       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0     1       0
traffic engineering <i>See</i> MPLS traffic engineering verifying	5 5 01900305678
traffic engineering <i>See</i> MPLS traffic engineering verifying	5 5 019003056788
traffic engineering <i>See</i> MPLS traffic engineering verifying	5 5 0190030567880
traffic engineering <i>See</i> MPLS traffic engineering verifying	5 5 0190030567889
traffic engineering <i>See</i> MPLS traffic engineering verifying	5       5       0190030567889       7
traffic engineering <i>See</i> MPLS traffic engineering verifying	5 5 0190030567889 7 °
traffic engineering <i>See</i> MPLS traffic engineering verifying	5 5 0190030567889 78
traffic engineering See MPLS traffic engineering verifying	5 5 0190030567889 78 F
traffic engineering <i>See</i> MPLS traffic engineering verifying	5 5 0190030567889 78 56
traffic engineering See MPLS traffic engineering verifying	9       5       5       0190030567889       78       567
traffic engineering See MPLS traffic engineering verifying	5       5       0       1       9       0       0       3       0       5       6       3       3       3       0       5       6       3       3       3       0       5       6       3       3       3       0       5       6       3       3       3       0       5       6       3       3       3       0       5       6       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3       3

BSR108
downstream interface103
DVMRP107
forwarding state notation105
IGMP See IGMP
IP address ranges104
MSDP108
network elements104
overview101
PGM108
PIM dense mode See PIM
PIM register messages See PIM register messages
PIM source-specific multicast (SSM)107
PIM sparse mode See PIM
preparation109
preventing routing loops105
protocols106
reverse-path forwarding (RPF)105
routing modes See multicast routing modes
S,G notation105
SAP and SDP See SAP; SDP
session announcements110
shortest-path tree (SPT)106
static RP112
See also RP
subnetwork leaves and branches104
upstream interface103
verifying IGMP versions119
verifying PIM mode and interface
configuration120
verifying PIM RPF routing table121
verifying PIM RPs120
verifying SAP and SDP configuration119
multicast routing modes
dense mode105
dense mode, caution for use105
sparse mode105
Multicast Source Discovery Protocol108
multifield classifier
multiple push label operation
Multiprotocol Label Switching See MPLS

### Ν

NAPT (Network Address Port Translation)	
example16	54
overload pool, defining (configuration	
editor)19	93
with dynamic NAT, overview16	54
NAT (Network Address Translation)	
actions16	57
assigning NAT services to interfaces (configuration	
editor)19	98
basic configuration (configuration	
editor)186, 19	90
components16	56

configuring185
destination static NAT processing165
displaying configurations200
interfaces, assigning NAT to (configuration
editor)
match conditions 167
NAPT overload pool defining (configuration
editor) 103
See also NADT
NAT miles without people (appfiguration
NAT rules without pools (configuration
editor)192
overload prefix, defining (configuration
editor)193
oversubscribed pool, defining (configuration
editor)193
overview163
pools See NAT pools
preparation 185
rules for transparent NAT (configuration
editor) 106
comple configuration 200
sample configuration
selective NAT (configuration editor)196
source dynamic NAT with NAPT processing164
source dynamic NAT without NAPT
processing165
source static NAT processing163
stateful firewall filters and See NAT with stateful
firewall filters
IIIewali Iiiteis
transparent, defining rules (configuration
transparent, defining rules (configuration editor)

network interfaces	
assigning CoS components to (Quick	
Configuration)	302
enabling NAT services on	198
enabling PIM on	113
multicast, upstream and downstream	103
verifying PIM on	120
VPN configuration	35
network layer reachability information See NLRI	
network service access points See NSAPs	
networks	32
public-private, access with NAT See NAT	
sample DLSw Ethernet redundancy	
topology	137
sample DLSw load balancing topology	140
sample DLSw topology	127
sample LSP topology	7
sample RSVP topology	13
sample VPN topology	32
trusted	155
untrusted	155
See also VPNs	
next term, filter action	234
next-hop service set, for IPSec tunnels	87
NLRI (network layer reachability information), BGP	
for CLNS	63
for VPNs	16
no filter action	234
notice icons	xix
NSAPs (network service access points)	
overview	56
sample configurations	62
numeric range match conditions	159

# 0

# Ρ

P routers See provider routers
packet encapsulation
Layer 2 circuits
Layer 2 VPNs
packet filters See stateful firewall filters; stateless
firewall filters
packet fragments, matching with a filter230
packet loss priority, setting with a filter235
packets
applying CoS scheduling rules
handling packet fragments238
handling packet fragments (configuration
editor)
ICMP, matching with a filter
TCD matching with a filter
TCP, matching with a filter
PAT (Port Address Translation) See NAPT
nath selection BSVP for MPI S See traffic engineering
database
PF (provider edge) routers 32
description 15
ES-IS for a CLNS island.
route distinguishers
verifying Layer 2 circuit connections
verifying Layer 2 circuit interfaces
verifying Layer 2 VPN connections
verifying Layer 2 VPN interfaces
verifying Layer 3 VPN connections
VPN task overview
VPN topology
See also VPNs
peer routers See DLSw peers
penultimate hop popping (PHP)9
penultimate router, in an LSP7
perfect forward secrecy (PFS), for preshared keys72
PFS (perfect forward secrecy), for preshared keys72
PGM (Pragmatic General Multicast)108
PHP (penultimate nop popping)
physical interfaces
Configuration
configuration)
Configuration) 302
enabling NAT services on 198
PIM (Protocol Independent Multicast)
dense mode 107
disabling on the network management
interface 112
register messages See PIM register messages
RPF routing table group117
source-specific multicast (SSM)107
sparse mode107
static RP router112
supported versions109

verifying the mode	120
verifying the RP	120
PIM register messages	
filtering	114
incoming, rejecting on an RP	115
outgoing, rejecting on a designated router	116
reject policy on designated router	116
reject policy on RP router	115
ping command (NAT configuration)	202
explanation	202
ning command (stateless firewall filter)	258
explanation	258
ning mpls l2circuit interface command	53
ning mpls l2circuit virtual-circuit command	53
ning mpls 12vnn instance	55
ning mpls 12 vpn interface command	53
ning mpls 12 vpn interface command	55
ping mpis 15 vpir command	
ouplanation	220
explanation	220
ping untrusted-nw-untrusted-nost command	220
explanation	220
pinging a VPN connection	52
PKI (public key infrastructure)	
for digital certificate configuration	
overview	70
URLs about	71
policers	
for CoS traffic classes	270
for firewall filter	306
for stateless firewall filters	243
strict high-priority queuing (configuration	
editor)	338
policy framework	149
See also firewall filters; NAT; routing policies	
pools, NAT See NAT pools	
pop label operation	8
Port Address Translation See NAPT	
Pragmatic General Multicast	108
precedence	
matching with a filter	232
ToS value for DLSw	136
prefix-length-range match type	172
preshared keys	
IKE (configuration editor)	
IKE (Ouick Configuration)	75
IKE description	72
overview	70
See also IKE	
PFS for	72
nriority of a packet setting with a filter	12 735
private networks access with NAT See NAT	2,5
promiscuous mode. DI Su	171
propagation suppressing	170
propagation, suppressing	179
Protocol Dullule, IPSec	12
Protocol Independent Multicast See PIM	

### protocols

AH71
Auto-RP108
DVMRP107
ESP71
IGMP See IGMP
IKE See IKE
IPSec See IPSec
IPv4, matching with a filter231
IPv6, matching with a filter231
LDP See LDP
MPLS See MPLS
MSDP108
multicast See multicast
NAT See NAT
PGM108
PIM dense mode See PIM
PIM source-specific multicast (SSM)107
PIM sparse mode <i>See</i> PIM
protocol bundle, IPSec72
RSVP See RSVP
SAP and SDP See SAP; SDP
SSP for DLSw127
provider edge routers See PE routers
provider routers
description15
VPN task overview
VPN topology
See also VPNs
public key infrastructure See PKI
public networks, access with NAT See NAT
public-private key pair, generating for digital
certificates
push label operation8

# Q

queues	268
See also CoS; output queues; queuing	
queuing	
CoS rules	327
starvation prevention (configuration editor)	332
strict high priority (configuration editor)	332
Quick Configuration	
Class of Service initial page	284
Class of Service Interfaces page	302
CoS classifiers page	290
CoS forwarding classes page	288
CoS RED drop profiles page	294
CoS scheduler maps page	294
CoS schedulers page	294
CoS value aliases page	286
DLSw page	130
Firewall Filters initial page	223
Firewall Filters interface assignment page	236

Firewall Filters match conditions and actions	
page	224
Firewall/NAT application page	208
Firewall/NAT main page	207
IPSec Tunnels page	74
rewrite rules page	292
virtual channel groups page	300

# R

random early detection See RED drop profiles	
reachability	16
verifying for DLSw peers	145
See also NLRI	
reachability cache, DLSw, clearing	141
RED (random early detection) drop profiles	
adding and editing (Quick Configuration)	296
defining (configuration editor)	319
defining (Quick Configuration)	294
description	270
sample configurations	319
summary (Quick Configuration)	295
redistributing routes	175
reject, filter action	234
rejecting	
invalid routes	173
unauthorized PIM registration	114
release notes, URL	xvii
remote router, DLSw	134
See also DLSw peers	
remote tunnel endpoint, IPSec	75
request security pki ca-certificate enroll ca-profile	
command	94
request security pki ca-certificate load command	96
request security pki generate-key-pair command	94
request security pki local-certificate enroll ca-profile	
command	95
request security pki local-certificate enroll certificate-id	
command	95
request security pki local-certificate load	
command	95
reservation See RSVP	
Resource Reservation Protocol See RSVP	
reverse-path forwarding See RPF	
rewrite rules	
adding and editing (Quick Configuration)	294
assigning to logical interfaces (configuration	
editor)	312
assigning to logical interfaces (Quick	
Configuration)	305
defining (configuration editor)	312
defining (Quick Configuration)	292
description	271
replacing DSCPs (configuration editor)	312
sample rules	312

summary (Quick Configuration)
when applied279
RIB See routing table
route distinguishers
description16
formats for45
route injection174
route list match types172
route manipulation actions, routing policies154
route redistribution174
route targets, VPN16
in a routing instance46
route-flap damping179
parameters179
routing
configuring VPNs
DLSw See DLSw
filtering and classifying routes149
See also firewall filters; NAT; routing policies
filtering routes with policies169
filtering traffic through a stateful firewall205
filtering traffic through a stateless firewall221
from one source to many destinations109
IBM networking See DLSw
MPLS for VPNs
MPLS traffic engineering19
multicast See multicast
overriding default packet forwarding with
CoS
policies See routing policies
protecting local IP addresses with NAT and stateful
firewall filters205
through IPSec tunnels67
VPNs
with NAT185
Routing Engine
handling packet fragments for (configuration
editor)246
protecting against DoS attacks (configuration
ealtor)
(configuration aditor)
(configuration base See routing table
routing information base see routing table
filter action setting 234
for CLNS static routes (with IS IS)
for CLNS static routes (without IS IS)
VPN confiduration 45
VPN route target 46
VBE instances
VRF table 46
routing policies
actions 153
applying 151
BGP export, for CLNS
configuration tasks

default actions	51
export statement15	52
final actions15	51
forwarding class with source and	
destination 17	76
grouping source and destination prefixes 17	76
	- 1
	51
injecting routes from one protocol into	
another17	74
Layer 2 VPN export policy	19
Layer 2 VPN import policy	18
Laver 3 VPNs	51
making BGP routes less preferable	77
match conditions	50
	) <u>/</u>
overview	51
PIM register messages See PIM register messages	
policy name17	70
preparation16	59
prepending AS paths17	77
reducing update messages with flap	
damning 17	79
rejecting invalid routes	72
registe redistribution	7 A
	74
route-flap damping	(9
terms15	51
terms, creating17	71
VPN configuration	17
routing solutions	
address translation (NAT)18	35
address translation (NAT)	35 33
address translation (NAT)	35 33
address translation (NAT)	35 33 38
address translation (NAT)	35 33 38 38
address translation (NAT)	35 33 38 38
address translation (NAT)	35 33 38 38 38
address translation (NAT)	35 33 38 38 46 77
address translation (NAT)	35 33 38 38 46 77
address translation (NAT)	<ul> <li>35</li> <li>33</li> <li>38</li> <li>38</li> <li>46</li> <li>77</li> <li>19</li> <li>06</li> </ul>
address translation (NAT)	<ul> <li>35</li> <li>33</li> <li>38</li> <li>38</li> <li>46</li> <li>77</li> <li>19</li> <li>06</li> <li>05</li> </ul>
address translation (NAT)	<ul> <li>35</li> <li>33</li> <li>38</li> <li>38</li> <li>46</li> <li>77</li> <li>19</li> <li>06</li> <li>05</li> <li>06</li> </ul>
address translation (NAT)	<ul> <li>35</li> <li>33</li> <li>38</li> <li>38</li> <li>38</li> <li>38</li> <li>46</li> <li>77</li> <li>19</li> <li>06</li> <li>05</li> <li>06</li> <li>19</li> </ul>
address translation (NAT)	35 33 38 38 46 77 19 06 05 06 49
address translation (NAT)	35 33 38 38 46 77 19 06 05 06 49 05
address translation (NAT)	35 33 38 46 77 19 06 05 06 49 05 41
address translation (NAT)	35 33 38 38 46 77 19 06 05 06 49 05 41
address translation (NAT)	35 33 38 38 46 77 19 06 05 06 49 05 41 79
address translation (NAT)	<ul> <li>35</li> <li>33</li> <li>38</li> <li>38</li> <li>38</li> <li>38</li> <li>46</li> <li>77</li> <li>19</li> <li>06</li> <li>05</li> <li>19</li> <li>06</li> <li>05</li> <li>41</li> <li>79</li> <li>72</li> </ul>
address translation (NAT)	35 33 38 46 77 19 06 50 69 05 41 79 72 59
address translation (NAT)	35 338 46 77 905 06 905 19 072 95 5
address translation (NAT)	35       38         46       77         90       55         90       56         97       97         90       55         90       55         90       55         90       55         90       55
address translation (NAT)       18         CoS       263, 28         filtering unwanted services and protocols       22         handling packet fragments       22         handling packet fragments (configuration       24         making BGP routes less preferable       17         MPLS traffic engineering       17         multicast administrative scoping       10         multicast reverse-path forwarding (RPF)       10         multicast shortest-path tree (SPT)       10         policy framework       14         preventing multicast routing loops       10         protecting against DoS attacks       24         reducing update messages with flap       17         damping       17         rejecting invalid routes       15         routing policies       151, 16         stateful firewall filters       15         stateful firewall filters       15	35338       46719       005049       0141         772955       005049       0141       772955       005049
address translation (NAT)       18         CoS       263, 28         filtering unwanted services and protocols       22         handling packet fragments       22         handling packet fragments (configuration       24         making BGP routes less preferable       17         MPLS traffic engineering       17         multicast administrative scoping       10         multicast reverse-path forwarding (RPF)       10         multicast shortest-path tree (SPT)       10         policy framework       14         preventing multicast routing loops       10         protecting against DoS attacks       24         reducing update messages with flap       151, 16         stateful firewall filters       15         stateful firewall filters       15         stateful firewall filters       16         stateless firewall filters       17	353       353         363       467         196       569         197       555         197       555         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000         197       1000
address translation (NAT)       18         CoS       263, 28         filtering unwanted services and protocols       22         handling packet fragments       22         handling packet fragments (configuration       24         making BGP routes less preferable       17         MPLS traffic engineering       17         multicast administrative scoping       10         multicast reverse-path forwarding (RPF)       10         multicast shortest-path tree (SPT)       10         policy framework       14         preventing multicast routing loops       10         protecting against DoS attacks       24         reducing update messages with flap       151, 16         stateful firewall filters       15         stateful firewall filters       15         vPNs       15	35       35         36       46         77       9         90       56         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57         90       57
address translation (NAT)	35       35         36       46         77       9         90       55         91       72         95       55         91       72         95       55         91       72         95       55         91       72         95       55         91       72         95       55         91       72         95       55         91       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75         95       75
address translation (NAT)	35         33         34         46         77         90         90         90         91         77         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90         90
address translation (NAT)       18         CoS       263, 28         filtering unwanted services and protocols       22         handling packet fragments       22         handling packet fragments (configuration       24         making BGP routes less preferable       17         MPLS traffic engineering       17         multicast administrative scoping       10         multicast reverse-path forwarding (RPF)       10         multicast shortest-path tree (SPT)       10         policy framework       14         preventing multicast routing loops       10         protecting against DoS attacks       24         reducing update messages with flap       17         damping       17         rejecting invalid routes       17         routing policies       151, 16         stateful firewall filters       15         stateful firewall filters       15         vPNs       27         routing table       RPF group, for multicast       11         verifying for RPF       12	35       35         38       46         77       96         90       56         91       72         95       55         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100         100       100
address translation (NAT)       18         CoS       263, 28         filtering unwanted services and protocols       22         handling packet fragments       22         handling packet fragments (configuration       24         making BGP routes less preferable       17         MPLS traffic engineering       17         multicast administrative scoping       10         multicast reverse-path forwarding (RPF)       10         multicast shortest-path tree (SPT)       10         policy framework       14         preventing multicast routing loops       10         protecting against DoS attacks       24         reducing update messages with flap       17         damping       17         rejecting invalid routes       17         routing policies       151, 16         stateful firewall filters       15         stateful firewall filters       15         routing table       RPF group, for multicast       11         verifying for RPF       12         verifying LDP-signaled LSPs       27	35338       467790556951         7790556951       7295551         1217
address translation (NAT)       18         CoS       263, 28         filtering unwanted services and protocols       22         handling packet fragments       22         handling packet fragments (configuration       24         making BGP routes less preferable       17         MPLS traffic engineering.       16         multicast administrative scoping.       16         multicast reverse-path forwarding (RPF)       16         multicast shortest-path tree (SPT)       16         policy framework.       14         preventing multicast routing loops.       16         protecting against DoS attacks.       24         reducing update messages with flap       17         damping.       17         rejecting invalid routes.       151, 16         stateful firewall filters.       151, 16         stateful firewall filters.       157, 22         VPNs.       27         routing table       RPF group, for multicast.       11         verifying for RPF.       12         verifying RSVP-signaled LSPs.       27	35338       46779055041       72955511         772955511       721729

RP (rendezvous point)
PIM register messages, incoming, rejecting115
PIM register messages, outgoing, stopping116
reject policy for incoming PIM register
messages115
same reject policy on RP routers in a
network114
static112
verifying120
RP router See RP
RPF (reverse-path forwarding)
description105
routing table group117
verifying the routing table121
RSVP (Resource Reservation Protocol)
and OSPF for VPNs43
bandwidth reservation12
CSPF13
disabling CSPF25
EROs12
fundamentals11
link coloring14
overview20
requirements20
RSVP-signaled LSPs23
verifying LSPs29
verifying neighbors28
verifying sessions28
verifying the routing table on the entry
router
RSVP neighbors, verifying
RSVP-signaled LSP See RSVP
rules, IPSec, applying to service sets

# S

S,G notation, for multicast forwarding states105
SA See dynamic SAs; IPSec security associations
sample configurations
basic source static NAT200
CLNS VPN configuration63
CoS behavior aggregate classification forwarding
classes and queues279
DLSw Ethernet redundancy topology137
DLSw load balancing topology140
DLSw topology127
firewall filter configurations252
stateful firewall filter configuration217
See also networks; topology
sampling traffic on an interface, with a filter235
SAP (Session Announcement Protocol)
description108
session announcements110
verifying119
SCEP request, for IPSec certification authority93

scheduler maps	
adding and editing (Quick Configuration)	00
assigning (configuration editor)	24
assigning to logical interfaces (Quick	
Configuration)3	05
assigning to physical interfaces (Quick	
Configuration)3	04
defining (configuration editor)3	24
defining (Quick Configuration)2	94
strict high-priority queuing (configuration	
editor)	35
strict high-priority queuing, applying scheduler	
map to interface (configuration editor)	37
summary (Quick Configuration)2	99
schedulers	08
adding and editing (Quick Configuration)	97
buffer size	22 60
default settings	77
defining (configuration editor) 3	21
defining (Ouick Configuration)	01 01
description 2	68
mapping to forwarding classes (configuration	00
editor) 3	25
mapping to forwarding classes (Ouick	~~~
Configuration)2	94
RED drop profiles2	70
sample mappings	25
sample schedulers	21
scheduler maps See scheduler maps	
shaping rate2	69
summary (Quick Configuration)2	97
transmission priority2	69
transmit rate2	68
voice and data for strict high-priority queuing	
(configuration editor)	36
voice, for strict high-priority queuing (configuration	
editor)	35
See also transmission scheduling	(0)
scheduling priority2	69
scoping administrative	06
SDP (Session Discovery Protocol)	00
description 1	08
session announcements	10
verifying 1	19
security	. ,
digital certificates	67
IPSec	67
NAT addressing1	85
stateful firewall filters2	05
stateless firewall filters2	21
security association See dynamic SAs; IPSec security	
associations	

service sets
for IPSec tunnels See IPSec service sets
for NAT rules198
for NAT rules, with stateful firewall filters215
for stateful firewall filters215
services interfaces
applying a NAT rule to (configuration
editor)
applying a stateful firewall filter to (configuration
editor)
for IPSec tunnels
Services Router
CLNS VPNs 55
CoS 283
CoS overview 263
DI Sw 125
firewall filter overview 149
host IPSec transport mode 73
IBM networking
IDM TIELWOIKING
MDLS for VDNs overview
MPLS for vrive overview
MPLS traffic engineering
multicast
multicast overview101
NAT
NAT and stateful firewall filters205
NAT overview163
policy framework overview149
routing policies169
routing policy overview151
secure gateway, IPSec tunnel mode73
See also IPSec tunnels
stateful firewall filters205
stateful firewall filters overview155
stateless firewall filter overview157
stateless firewall filters221
VPNs
Session Announcement Protocol See SAP; SDP
sessions
announcements, multicast110
LDP, verifying26
RSVP, verifying28
shaping rate
See also CoS; scheduler maps; schedulers
shortest-path tree106
show class-of-service adaptive-shaper command346
show class-of-service interface command
show class-of-service virtual-channel command
show class-of-service virtual-channel-group
command
show command63
show dlsw peers command. 144
show dlsw peers detail command
explanation
show dlsw reachability command 145
show firewall command

de seu Grande II Gitter annata et DE se anna a d
show lifewall lifter protect-RE command
explanation256
show firewall log command255
explanation255
show igmp interface command119
explanation120
show interfaces command200
show interfaces lo0 command251
show ldp neighbor command25
explanation
show ldp session detail command 26
evplanation 26
show llc2 redundancy interface statistics
show hez redundancy interface statistics
continuation of common d
snow multicast rpi command121
explanation121
show pim interface command120
explanation120
show pim rps command120
explanation121
show route summary command257, 259
explanation257, 259
show route table inet.3 command27, 29
explanation
show rsvp neighbor command 28
explanation 28
show reve session detail command
show isvp session detail command
200
explanation
explanation       29         show sap listen command       119, 345         explanation       119, 345         show services command       200, 217         show services ipsec-vpn ipsec statistics command       98         explanation       98         show services stateful-firewall conversations extensive       202         explanation       202         explanation       202         show statement dlsw circuits detail command       143         show dlsw capabilities command       142         show dlsw circuits command       143         show interfaces fe-3/0/0 command       142
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensivecommand202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command143show unterfaces fe-3/0/0 command142show llc2 redundancy brief command146
explanation
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensivecommand202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show interfaces fe-3/0/0 command142show llc2 redundancy brief command142signaling protocols19overviow10
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensivecommand202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show interfaces fe-3/0/0 command142show llc2 redundancy brief command146signaling protocols19overview10
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensivecommand202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command143show dlsw circuits command143show llc2 redundancy brief command146signaling protocols19overview10VPNs41Source lead DDMDI C set Size of the DD
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensivecommand202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command142show dlsw circuits fe-3/0/0 command142show llc2 redundancy brief command146signaling protocols19overview10VPNs41See also LDP; MPLS traffic engineering; RSVP
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensivecommand202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command143show dlsw circuits command143show llc2 redundancy brief command146signaling protocols19overview10VPNs41See also LDP; MPLS traffic engineering; RSVPSimple Certificate Enrollment Protocol (SCEP) request,
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensivecommand202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command142show interfaces fe-3/0/0 command142show llc2 redundancy brief command146signaling protocols19overview10VPNs41see also LDP; MPLS traffic engineering; RSVPSimple Certificate Enrollment Protocol (SCEP) request, for IPSec certification authority
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensive202command202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command143show dlsw circuits command142show interfaces fe-3/0/0 command142show llc2 redundancy brief command146signaling protocols19overview10VPNs41See also LDP; MPLS traffic engineering; RSVPSimple Certificate Enrollment Protocol (SCEP) request,for IPSec certification authority93Simple Network Management Protocol See SNMP
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensivecommand202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command143show dlsw circuits command142show interfaces fe-3/0/0 command142show llc2 redundancy brief command146signaling protocols19overview10VPNs41See also LDP; MPLS traffic engineering; RSVPSimple Certificate Enrollment Protocol (SCEP) request,for IPSec certification authority93Simple Network Management Protocol See SNMPSNA forwarding See DLSw
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensive202explanation202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command143show dlsw circuits command142show interfaces fe-3/0/0 command142show llc2 redundancy brief command146signaling protocols19overview10VPNs41See also LDP; MPLS traffic engineering; RSVPSimple Certificate Enrollment Protocol (SCEP) request,for IPSec certification authority93Simple Network Management Protocol See SNMPSNA forwarding See DLSwSNMP monitoring, DLSw MIB125
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensivecommand202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command142show dlsw circuits command142show interfaces fe-3/0/0 command142show llc2 redundancy brief command146signaling protocols19overview10VPNs41See also LDP; MPLS traffic engineering; RSVPSimple Certificate Enrollment Protocol (SCEP) request,for IPSec certification authority93Simple Network Management Protocol See SNMPSNA forwarding See DLSwSNMP monitoring, DLSw MIB125source dynamic NAT with NAPT
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensivecommand202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command142show interfaces fe-3/0/0 command142show llc2 redundancy brief command146signaling protocols19overview10VPNs41See also LDP; MPLS traffic engineering; RSVPSimple Certificate Enrollment Protocol (SCEP) request,for IPSec certification authority93Simple Network Management Protocol See SNMPSNA forwarding See DLSwSNMP monitoring, DLSw MIB125source dynamic NAT with NAPT164
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensive202explanation202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command142show interfaces fe-3/0/0 command142show llc2 redundancy brief command146signaling protocols19overview10VPNs41See also LDP; MPLS traffic engineering; RSVPSimple Certificate Enrollment Protocol (SCEP) request,for IPSec certification authority93Simple Network Management Protocol See SNMPSNA forwarding See DLSwSNMP monitoring, DLSw MIB125source dynamic NAT with NAPT164example164
explanation29show sap listen command119, 345explanation119, 345show services command200, 217show services ipsec-vpn ipsec statistics command98explanation98show services stateful-firewall conversations extensive202explanation202explanation202show statement dlsw circuits detail command143show dlsw capabilities command142show dlsw circuits command142show interfaces fe-3/0/0 command142show llc2 redundancy brief command146signaling protocols19overview10VPNs41See also LDP; MPLS traffic engineering; RSVPSimple Certificate Enrollment Protocol (SCEP) request,for IPSec certification authority93Simple Network Management Protocol See SNMPSNA forwarding See DLSwSNMP monitoring, DLSw MIB125source dynamic NAT with NAPT164description164source dynamic NAT without NAPT164

source static NAT
basic configuration (configuration
editor)
description 163
dynamic pool address assignment 187
example 164
sample confiduration 200
varifying 202
verinying
sp-0/0/0
for IPSec tunnels (configuration editor)
no stateful firewall filters211
sparse mode See multicast routing modes
SPT (shortest-path tree)106
ssh command257
explanation257
SSP (Switch-to-Switch Protocol) for DLSw127
starvation prevention, on CoS queues
stateful firewall filters
actions157
applying to an interface (configuration
editor)215
automatic discard rule206
configuration editor211, 212
configuration overview
displaying configurations
do not apply to sp-0/0/0
enabling (Ouick Configuration)
iunos-algs-outbound default group
match conditions
NAT and <i>See</i> NAT with stateful firewall filters
Overview 155
preparation 205
Ouick Confiduration 206
cample rules 211
sample fulles
unitusted network
verifying
verifying configuration
stateless firewall filters
action modifiers (Quick Configuration)235
Action tab (Quick Configuration)
actions and action modifiers
adding (Quick Configuration)226
applying to an interface (configuration
editor)251
assigning to interfaces (Quick
Configuration)236
automatic discard rule158, 222
bit-field logical operators162
chained multiple filters158
destination matching (Quick Configuration)227
displaying configurations252
displaying statistics256
filter actions (Quick Configuration)234
See also actions
handling packet fragments238

handling packet fragments (configuration	246
in a Common Criteria environment	240
in a common criteria environment	.221
Configuration)	237
interface matching (Quick Configuration)	227
IPv4 filters (Quick Confiduration)	.227
IPv4 filters (Quick Configuration)	.222
match conditions	.222
Match Destination tab (Quick	.139
Configuration)	227
Match Interface tab (Quick Configuration)	221
Match Match Network tab (Quick	.22)
Configuration)	230
Match Match Packet and Network tab (Ouick	.290
Configuration)	230
Match Source or Destination tab (Quick	.290
Configuration)	.228
Match Source tab (Quick Configuration)	226
multiple filters chained	158
network matching (Ouick Configuration)	.230
output filters, interface assignment (Ouick	
Configuration)	.237
overview	.157
packet matching (Quick Configuration)	.230
planning	238
policers for	.243
preparation	.221
protecting the Routing Engine against ICMP flood	S
protecting the Routing Engine against ICMP flood (configuration editor)	s .241
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood	s .241 s
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor)	s .241 s .241
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted	s .241 s .241
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor)	s .241 s .241 .238
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted	s .241 s .241 .238
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor)	s .241 s .241 .238 .238
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration	s .241 s .241 .238 .238 .222
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments	s .241 s .241 .238 .238 .222 .247
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and	s .241 s .241 .238 .238 .222 .247
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols	s .241 s .238 .238 .222 .247 .239
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS	s .241 s .238 .238 .238 .222 .247 .239
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks	s .241 s .238 .238 .222 .247 .239 .242
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences	s .241 .238 .238 .228 .228 .247 .247 .239 .242 .158
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences source matching (Quick Configuration)	s .241 s .241 .238 .222 .247 .239 .242 .242 .158 .226
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences source matching (Quick Configuration) source or destination matching (Quick	s .241 s .241 .238 .223 .247 .239 .247 .239 .242 .158 .226
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences source matching (Quick Configuration) source or destination matching (Quick Configuration)	s .241 s .238 .238 .222 .247 .239 .242 .158 .226 .228
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences source matching (Quick Configuration) source or destination matching (Quick Configuration) strict high-priority queuing (configuration	s .241 s .241 .238 .228 .222 .247 .239 .242 .158 .226 .228
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences source matching (Quick Configuration) source or destination matching (Quick Configuration) strict high-priority queuing (configuration editor)	s .241 s .241 .238 .228 .227 .239 .247 .239 .242 .158 .226 .228 .339
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) Quick Configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences source matching (Quick Configuration) surce or destination matching (Quick Configuration) strict high-priority queuing (configuration editor)	s .241 s .241 .238 .223 .238 .223 .247 .239 .242 .158 .226 .228 .228 .228 .228
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences source matching (Quick Configuration) source or destination matching (Quick Configuration) strict high-priority queuing (configuration editor) summary (Quick Configuration) terms, adding (Quick Configuration)	s .241 s .241 .238 .228 .228 .247 .239 .242 .247 .239 .242 .226 .228 .339 .225 .226
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences source matching (Quick Configuration) source or destination matching (Quick Configuration) strict high-priority queuing (configuration editor) summary (Quick Configuration) terms, adding (Quick Configuration) terms, overview	s .241 s .238 .223 .238 .223 .238 .222 .247 .239 .242 .247 .239 .242 .226 .226 .228 .225 .226 .157 .225
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences source matching (Quick Configuration) source or destination matching (Quick Configuration) strict high-priority queuing (configuration editor) summary (Quick Configuration) terms, adding (Quick Configuration) terms, overview typical, planning	s .241 s .238 .238 .238 .228 .238 .222 .247 .239 .242 .247 .239 .242 .225 .226 .228 .225 .226 .225 .226
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) Quick Configuration editor) Quick Configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences source matching (Quick Configuration) source or destination matching (Quick Configuration) strict high-priority queuing (configuration editor) summary (Quick Configuration) terms, adding (Quick Configuration) terms, overview typical, planning verifying actions	s .241 s .238 .238 .222 .247 .239 .242 .247 .239 .242 .247 .239 .242 .247 .239 .242 .247 .238 .225 .226 .157 .238
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration sample terms, to filter fragments sample terms, to filter services and protocols sample terms, to protect against DoS attacks sequences source matching (Quick Configuration) source or destination matching (Quick Configuration) strict high-priority queuing (configuration editor) summary (Quick Configuration) terms, adding (Quick Configuration) terms, overview typical, planning verifying actions verifying configuration	s .241 s .238 .238 .222 .247 .239 .242 .247 .239 .242 .247 .239 .242 .247 .238 .225 .226 .157 .238 .257 .252
protecting the Routing Engine against ICMP flood (configuration editor) protecting the Routing Engine against TCP flood (configuration editor) protecting the Routing Engine against untrusted protocols (configuration editor) protecting the Routing Engine against untrusted services (configuration editor) Quick Configuration	s .241 s .238 .238 .222 .247 .239 .242 .247 .239 .242 .247 .239 .242 .242 .242 .228 .226 .157 .238 .257 .252 .258

static LSPs9
static routes
CLNS VPNs (with IS-IS)58
CLNS VPNs (without IS-IS)62
static RP router112
See also RP
statistics
DLSw Ethernet redundancy interfaces146
IPSec tunnels
stateless firewall filters256
strict high-priority queuing, CoS
applying a scheduler map to interface
(configuration editor)
applying classifier to interface (configuration
editor)
assigning queues
classifying traffic
configuring a scheduler map and schedulers
(configuration editor)
configuring policiers (configuration editor)
creating a stateless firewall filter (configuration
editor)
defining voice and data schedulers (configuration
editor)
overview
strict hops, RSVP12
subnetworks, multicast leaves and branches104
support, technical See technical support
swap and push label operation9
swap label operation8
Switch-to-Switch Protocol (SSP) for DLSw127
syntax conventionsxix
system log, of packet information235
Systems Network Architecture (SNA) forwarding See
DLSw

# Т

TCP packets, matching with a filter230
TCP policers
technical support
contacting JTACxxiii
TED See traffic engineering database
telnet command258
explanation258
terminology
CLNS55
CoS
DLSw126
firewall filters149
IPSec67
MPLS
multicast101
NAT149
routing policies149
VPNs

terms
firewall filter, for multifield classifier
in a routing policy151
in a routing policy, creating171
stateless firewall filters, adding (Quick
Configuration)226
stateless firewall filters, overview157
through route list match type172
time-to-live (TTL) value, matching with an IPv4
filter
to statement, routing policy match conditions152
topology
sample DLSw Ethernet redundancy
topology137
sample DLSw load balancing topology140
sample DLSw topology127
sample LSP network7
sample RSVP-signaled LSP13
sample VPN
ToS (type of service), precedence for DLSw
packets
traceroute source bypass-routing gateway
command27
explanation27
traffic
filtering through a stateful firewall205
filtering through a stateless firewall
protection, IPSec71
sampling on an interface, with a filter235
traffic engineering See MPLS traffic engineering; traffic
engineering database
traffic engineering database
CSPF constraints on path selection13
CSPF rules for path selection
link coloring for CSPF path selection14
transit interfaces
LDP-signaled LSPs for21
RSVP-signaled LSPs for23
transit routers, in an LSP7
transmission priority
See also CoS; scheduler maps; schedulers
transmission scheduling
transmit rate
description268
See also CoS; schedulers; transmission
scheduling
transparent NAT, defining rules196
transport mode, IPSec, for host73
triple Data Encryption Standard-cipher block chaining
(DES-CBC)
trusted networks, firewall filter protection155
TTL (time-to-live) value, matching with an IPv4
filter
tunnel mode, IPSec, for secure gateway73
See also IPSec tunnels

tunnels, through a public network See IPSec tunnels;
VPNs
type of service (ToS), precedence for DLSw
packets134

## U

untrusted networks, firewall filter actions on	155
upstream interfaces	103
See also multicast	
upto route list match type	172
URLs	
digital certificates and PKI	.71
release notes	xvii

# V

traffic forwarding over LDP-signaled LSPs	27
virtual channel groups	346
virtual channels	345
VPNs	52
virtual channel groups	
adding and editing (Quick Configuration)	301
assigning to logical interfaces (Quick	
Configuration)	305
summary (Quick Configuration)	301
verifying	346
virtual channels	
adding and editing (Ouick Configuration)	301
applying CoS rules to logical interfaces.	327
defining groups (Quick Configuration)	300
filter action modifier setting	235
groups See virtual channel groups	
verifying	345
virtual circuit ID for Laver 2 circuits	44
virtual private petworks See VPNs	
voice traffic prioritizing packets for in CoS	
voice trainc, prioritizing packets for, in Cos	330
VDN routing and forwarding (VDE) instances	
VPN fouling and forwarding (VNF) instances	10
VPN routing and forwarding table See VRF table	71
VPNs (virtual private networks)	
AS number	40
basic Layer 2 circuit description	
basic Layer 2 VPN description	
basic Layer 3 VPN description	
BGP	
CLNS See CLNS	
components	15
configuration overview	
configuration task overview	34
IGPs	41
IPSec VPN policy for digital certificates	96
Layer 2 circuit configuration	44
LSP for RSVP	
MPLS	37
overview	3, 14
participating interfaces	35
preparation	34
protocols for	37
route distinguishers	16,45
route target	46
route targets	16
routing information	16
routing instance See routing instance	
routing policies	47
routing requirements	15
sample topology	
signaling protocols	
tunneling process	15
types	17
verifying connectivity	52
voinying connectivity	

VRF instances16
VRF table See VRF table
See also Layer 2 circuits; Layer 2 VPNs: Layer 3
VPNs; MPLS
VRF (VPN routing and forwarding) table46
route targets16
VRF instances16
VRF instances
IPSec dynamic SAs86
overview16

# W

white papers	about digital	certificates71
--------------	---------------	----------------

# X

x and y coordin	ates, CoS	drop pro	files	.29	96
-----------------	-----------	----------	-------	-----	----