




Junos[®] OS for EX Series Ethernet Switches, Release 10.4: Routing Policy and Packet Filtering



Published: 2010-12-06
Revision 1

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

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, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, 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.

Junos® OS for EX Series Ethernet Switches, Release 10.4: Routing Policy and Packet Filtering

Copyright © 2010, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Writing:
Editing:
Illustration:
Cover Design:

Revision History
December 2010—Revision 1

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

END USER LICENSE AGREEMENT

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

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer's principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer's principal office is located outside the Americas) (such applicable entity being referred to herein as "Juniper"), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. "Software" also includes updates, upgrades and new releases of such software. "Embedded Software" means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

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

- a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.
- b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.
- c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.
- d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.
- e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

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

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

Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

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

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

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

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

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

10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.

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

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

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

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

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

Table of Contents

	About This Topic Collection	xi
	How to Use This Guide	xi
	List of EX Series Guides for Junos OS Release 10.4	xi
	Downloading Software	xiii
	Documentation Symbols Key	xiv
	Documentation Feedback	xv
	Requesting Technical Support	xvi
	Self-Help Online Tools and Resources	xvi
	Opening a Case with JTAC	xvi
Part 1	Routing Policy and Packet Filtering (Firewall Filters)	
Chapter 1	Firewall Filters—Overview	3
	Firewall Filters for EX Series Switches Overview	3
	Firewall Filter Types	4
	Firewall Filter Components	5
	Firewall Filter Processing	5
	Understanding Planning of Firewall Filters	7
	Understanding Firewall Filter Processing Points for Bridged and Routed Packets on EX Series Switches	9
	Understanding How Firewall Filters Control Packet Flows	10
	Firewall Filter Match Conditions and Actions for EX Series Switches	11
	Understanding How Firewall Filters Are Evaluated	31
	Understanding Firewall Filter Match Conditions	33
	Filter Match Conditions	33
	Numeric Filter Match Conditions	33
	Interface Filter Match Conditions	34
	IP Address Filter Match Conditions	34
	MAC Address Filter Match Conditions	35
	Bit-Field Filter Match Conditions	35
	Understanding How Firewall Filters Test a Packet's Protocol	36
	Understanding the Use of Policers in Firewall Filters	37
	Understanding Filter-Based Forwarding for EX Series Switches	38
Chapter 2	Examples of Firewall Filters Configuration	39
	Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches	39
	Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches	58
	Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch	62

Chapter 3	Configuring Firewall Filters	67
	Configuring Firewall Filters (CLI Procedure)	67
	Configuring a Firewall Filter	67
	Applying a Firewall Filter to a Port on a Switch	70
	Applying a Firewall Filter to a Management Interface on a Switch	71
	Applying a Firewall Filter to a VLAN on a Network	72
	Applying a Firewall Filter to a Layer 3 (Routed) Interface	73
	Configuring Firewall Filters (J-Web Procedure)	74
	Configuring Policers to Control Traffic Rates (CLI Procedure)	78
	Configuring Policers	79
	Specifying Policers in a Firewall Filter Configuration	80
	Applying a Firewall Filter That Is Configured with a Policer	80
	Assigning Multifield Classifiers in Firewall Filters to Specify Packet-Forwarding Behavior (CLI Procedure)	81
	Configuring Routing Policies (J-Web Procedure)	82
Chapter 4	Verifying Firewall Filter Configuration	89
	Verifying That Firewall Filters Are Operational	89
	Verifying That Policers Are Operational	90
	Monitoring Firewall Filter Traffic	90
	Monitoring Traffic for All Firewall Filters and Policers That Are Configured on the Switch	91
	Monitoring Traffic for a Specific Firewall Filter	91
	Monitoring Traffic for a Specific Policer	91
Chapter 5	Troubleshooting Firewall Filters	93
	Troubleshooting Firewall Filters	93
	Firewall Filter Configuration Returns a No Space Available in TCAM Message	93
Chapter 6	Configuration Statements for Firewall Filters	95
	[edit firewall] Configuration Statement Hierarchy	95
	Firewall Filter Configuration Statements Supported by Junos OS for EX Series Switches	96
	apply-path	99
	as-path	100
	as-path-group	101
	bandwidth-limit	102
	burst-size-limit	103
	community	104
	condition	106
	damping	107
	dynamic-db	108
	family (Firewall Filter)	109
	filter	110
	filter	111
	filter-specific	111
	firewall	112
	from	113
	if-exceeding	114

	interface-specific	115
	policer	116
	policy-statement	117
	prefix-list	119
	routing-instance	120
	term	121
	then	122
	then	123
Chapter 7	Operational Commands for Firewall Filters	125
	clear firewall	126
	clear firewall	127
	show firewall	128
	show firewall	131
	show firewall log	134
	show interfaces filters	136
	show interfaces policers	138
	show policer	140
	show policy	142
	show policy conditions	145
	test policy	147

About This Topic Collection

- How to Use This Guide on page xi
- List of EX Series Guides for Junos OS Release 10.4 on page xi
- Downloading Software on page xiii
- Documentation Symbols Key on page xiv
- Documentation Feedback on page xv
- Requesting Technical Support on page xvi

How to Use This Guide

Complete documentation for the EX Series product family is provided on webpages at http://www.juniper.net/techpubs/en_US/release-independent/information-products/pathway-pages/ex-series/product/index.html. We have selected content from these webpages and created a number of EX Series guides that collect related topics into a book-like format so that the information is easy to print and easy to download to your local computer.

The release notes are at http://www.juniper.net/techpubs/en_US/junos10.4/information-products/topic-collections/release-notes/10.4/junos-release-notes-10.4.pdf.

List of EX Series Guides for Junos OS Release 10.4

Title	Description
<i>Complete Hardware Guide for EX2200 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX2200 Ethernet switches
<i>Complete Hardware Guide for EX3200 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX3200 Ethernet switches
<i>Complete Hardware Guide for EX4200 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX4200 Ethernet switches
<i>Complete Hardware Guide for EX4500 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX4500 Ethernet switches





Title	Description
<i>Complete Hardware Guide for EX8208 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8208 Ethernet switches
<i>Complete Hardware Guide for EX8216 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8216 Ethernet switches
<i>Complete Hardware Guide for the XRE200 External Routing Engine</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for the XRE200 External Routing Engine
<i>Complete Software Guide for Junos® OS for EX Series Ethernet Switches, Release 10.4</i>	Software feature descriptions, configuration examples, and tasks for Junos OS for EX Series switches
Software Topic Collections	Software feature descriptions, configuration examples and tasks, and reference pages for configuration statements and operational commands (This information also appears in the <i>Complete Software Guide for Junos® OS for EX Series Ethernet Switches, Release 10.4.</i>)
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: EX4200 Virtual Chassis</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: EX8200 Virtual Chassis</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Access Control</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Configuration Management</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Class of Service</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Device Security</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Ethernet Switching</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Fibre Channel over Ethernet</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: High Availability</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Interfaces</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Layer 3 Protocols</i>	

Title	Description
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: MPLS</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Multicast</i>	
<i>Junos® OS for EX Series Switches, Release 10.4: Network Management and Monitoring</i>	
<i>Junos® OS for EX Series Switches, Release 10.4: Port Security</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Routing Policy and Packet Filtering</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Software Installation</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: Spanning-Tree Protocols</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: System Monitoring</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: System Services</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: System Setup</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: User and Access Management</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 10.4: User Interfaces</i>	

Downloading Software

You can download Junos OS for EX Series switches from the Download Software area at <http://www.juniper.net/customers/support/>. To download the software, you must have a Juniper Networks user account. For information about obtaining an account, see <http://www.juniper.net/entitlement/setupAccountInfo.do>.

Documentation Symbols Key

Notice Icons		
Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
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
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces important new terms. Identifies book names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	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 <i>domain-name</i>
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> 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 <i>metric</i> >;

Text and Syntax Conventions		
Convention	Description	Examples
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast <i>(string1 string2 string3)</i>
# (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]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> 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 .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. Send e-mail to techpubs-comments@juniper.net with the following:

- Document URL or title
- Page number if applicable
- Software version
- Your name and company

Requesting Technical Support

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf> .
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/> .
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Routing Policy and Packet Filtering (Firewall Filters)

- Firewall Filters—Overview on page 3
- Examples of Firewall Filters Configuration on page 39
- Configuring Firewall Filters on page 67
- Verifying Firewall Filter Configuration on page 89
- Troubleshooting Firewall Filters on page 93
- Configuration Statements for Firewall Filters on page 95
- Operational Commands for Firewall Filters on page 125

CHAPTER 1

Firewall Filters—Overview

- Firewall Filters for EX Series Switches Overview on page 3
- Understanding Planning of Firewall Filters on page 7
- Understanding Firewall Filter Processing Points for Bridged and Routed Packets on EX Series Switches on page 9
- Understanding How Firewall Filters Control Packet Flows on page 10
- Firewall Filter Match Conditions and Actions for EX Series Switches on page 11
- Understanding How Firewall Filters Are Evaluated on page 31
- Understanding Firewall Filter Match Conditions on page 33
- Understanding How Firewall Filters Test a Packet's Protocol on page 36
- Understanding the Use of Policers in Firewall Filters on page 37
- Understanding Filter-Based Forwarding for EX Series Switches on page 38

Firewall Filters for EX Series Switches Overview

Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on a Juniper Networks EX Series Ethernet Switch from a source address to a destination address. You configure firewall filters to determine whether to permit, deny, or forward traffic before it enters or exits a port, VLAN, or Layer 3 (routed) interface to which the firewall filter is applied. An *ingress* firewall filter is a filter that is applied to packets that are entering a network. An *egress* firewall filter is a filter that is applied to packets that are exiting a network. You can configure firewall filters to subject packets to filtering, class-of-service (CoS) marking (grouping similar types of traffic together, and treating each type of traffic as a class with its own level of service priority), and traffic policing (controlling the maximum rate of traffic sent or received on an interface).

This topic describes:

- Firewall Filter Types on page 4
- Firewall Filter Components on page 5
- Firewall Filter Processing on page 5

Firewall Filter Types

The following firewall filter types are supported for EX Series switches:

- Port (Layer 2) firewall filter—Port firewall filters apply to Layer 2 switch ports. You can apply port firewall filters in both ingress and egress directions on a physical port.
- VLAN firewall filter—VLAN firewall filters provide access control for packets that enter a VLAN, are bridged within a VLAN, or leave a VLAN. You can apply VLAN firewall filters in both ingress and egress directions on a VLAN. VLAN firewall filters are applied to all packets that are forwarded to or forwarded from the VLAN.
- Router (Layer 3) firewall filter—You can apply a router firewall filter in both ingress and egress directions on Layer 3 (routed) interfaces and routed VLAN interfaces (RVIs). You can apply a router firewall filter in the ingress direction on the loopback interface (**lo0**) also.



NOTE: You can apply a firewall filter to aggregated Ethernet interfaces and loopback interfaces also. Firewall filters configured on loopback interfaces are applied only to packets that are sent to the Routing Engine CPU for further processing.

On Juniper Networks EX3200, EX4200, and EX8200 Ethernet Switches, you can apply port, VLAN, or router firewall filters to both IPv4 and IPv6 traffic, whereas on Juniper Networks EX4500 Ethernet Switches, you can apply port, VLAN, or router firewall filters to IPv4 traffic only. For information on firewall filters supported on different switches, see “Firewall Filter Match Conditions and Actions for EX Series Switches” on page 11.

You can apply firewall filter match conditions to IPv6 traffic on Layer 3 interfaces, aggregated Ethernet interfaces, and loopback interfaces. To configure port firewall filters and VLAN firewall filters for IPv6 traffic, you must include the match condition **ether-type ipv6** and apply the filter on Layer 2 interfaces or VLANs. When you include the match condition **ether-type ipv6** in a term, you must ensure that other match conditions specified in the term are valid for IPv6 traffic. If the port firewall filter or VLAN firewall filter term contains the match condition **ether-type ipv6**, with no other IPv6 match condition specified, all IPv6 traffic is matched.



NOTE: A term without the match condition **ether-type ipv6** applies only to IPv4 traffic, and a term with that match condition applies only to IPv6 traffic. Hence, to configure port and VLAN firewall filters for both IPv4 and IPv6 traffic, you should configure two different terms, one each for IPv4 and IPv6 traffic.

To apply a firewall filter, you must:

1. Configure the firewall filter.
2. Apply the firewall filter to a port, VLAN, or Layer 3 interface.

Firewall Filter Components

In a firewall filter, you first define the family address type (**ethernet-switching**, **inet**, or **inet6**), and then you define one or more terms that specify the filtering criteria and the action to take if a match occurs.

The maximum number of terms allowed per firewall filter for EX Series switches is:

- 512 for EX2200 switches
- 7168 for EX3200 and EX4200 switches—as allocated by the dynamic allocation of ternary content addressable memory (TCAM) for port, VLAN, and router firewall filters.
- 1536 for EX4500 switches
- 32768 for EX8200 switches



NOTE: The on-demand dynamic allocation of the shared space TCAM in EX8200 switches is achieved by assigning free space blocks to firewall filters. Firewall filters are categorized into two different pools. Port and VLAN filters are pooled together (the memory threshold for this pool is 22K) while router firewall filters are pooled separately (the threshold for this pool is 32K). The assignment happens based on the filter pool type. Free space blocks can be shared only among the firewall filters belonging to the same filter pool type. An error message is generated when you try to configure a firewall filter beyond the TCAM threshold.

Each term consists of the following components:

- Match conditions—Specify the values or fields that the packet must contain. You can define various match conditions, including the IP source address field, IP destination address field, Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port field, IP protocol field, Internet Control Message Protocol (ICMP) packet type, TCP flags, and interfaces.
- Action—Specifies what to do if a packet matches the match conditions. Possible actions are to accept or discard the packet or to send the packet to a specific virtual routing interface. In addition, packets can be counted to collect statistical information. If no action is specified for a term, the default action is to accept the packet.

Firewall Filter Processing

The order of the terms within a firewall filter configuration is important. Packets are tested against each term in the order in which the terms are listed in the firewall filter configuration. When a firewall filter contains multiple terms, the switch takes a top-down approach and compares a packet against the first term in the firewall filter. If the packet matches the first term, the switch executes the action defined by that term to either permit or deny the packet, and no other terms are evaluated. If the switch does not find a match between the packet and first term, it compares the packet to the next term in the firewall filter by using the same match process. If no match occurs between the

packet and the second term, the switch continues to compare the packet to each successive term defined in the firewall filter until a match is found. If a packet does not match any terms in a firewall filter, the default action is to discard the packet.

**Related
Documentation**

- Understanding Planning of Firewall Filters on page 7
- Understanding Firewall Filter Processing Points for Bridged and Routed Packets on EX Series Switches on page 9
- Understanding How Firewall Filters Are Evaluated on page 31
- Understanding Firewall Filter Match Conditions on page 33
- Understanding the Use of Policers in Firewall Filters on page 37
- Understanding Filter-Based Forwarding for EX Series Switches on page 38
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 58

Understanding Planning of Firewall Filters

Before you create a firewall filter and apply 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 goals. You must understand how packets are matched to match conditions, the default and configured actions of the firewall filter, and proper placement of the firewall filter.

You can configure and apply no more than one firewall filter per port, VLAN, or router interface, per direction. The following limits apply for the number of firewall filter terms allowed per filter on various switch models:

- On EX2200 switches, the number of terms per filter cannot exceed 512.
- On EX3200 and EX4200 switches, the number of terms per filter cannot exceed 7168.
- On EX4500 switches, the number of terms per filter cannot exceed 1536.
- On EX8200 switches, the number of terms per filter cannot exceed 32768.

In addition, you should try to be conservative in the number of terms (rules) that you include in each firewall filter because a large number of terms requires longer processing time during a commit and also can make firewall filter testing and troubleshooting more difficult. Similarly, applying firewall filters across many switch and router interfaces can make testing and troubleshooting the rules of those filters difficult.

Before you configure and apply firewall filters, answer the following questions for each of those firewall filters:

1. What is the purpose of the firewall filter?

For example, you can use a firewall filter to limit traffic to source and destination MAC addresses, specific protocols, or certain data rates or to prevent denial of service (DoS) attacks.

2. What are the appropriate match conditions?

- a. Determine the packet header fields that the packet must contain for a match. Possible fields include:

- Layer 2 header fields—Source and destination MAC addresses, dot1q tag, Ethernet type, and VLAN
- Layer 3 header fields—Source and destination IP addresses, protocols, and IP options (IP precedence, IP fragmentation flags, TTL type)
- TCP header fields—Source and destination ports and flags
- ICMP header fields—Packet type and code

- b. Determine the port, VLAN, or router interface on which the packet was received.

3. What are the appropriate actions to take if a match occurs?

Possible actions to take if a match occurs are accept, discard, and forward to a routing instance.

4. What additional action modifiers might be required?

Determine whether additional actions are required if a packet matches a match condition; for example, you can specify an action modifier to count, analyze, or police packets.

5. On what interface should the firewall filter be applied?

Start with the following basic guidelines:

- If all the packets entering a port need to be exposed to filtering, then use port firewall filters.
- If all the packets that are bridged need filtering, then use VLAN firewall filters.
- If all the packets that are routed need filtering, then use router firewall filters.

Before you choose the interface on which to apply a firewall filter, understand how that placement can impact traffic flow to other interfaces. In general, apply a firewall filter that filters on source and destination IP addresses, IP protocols, or protocol information—such as ICMP message types, and TCP and UDP port numbers—nearest to the source devices. However, typically apply a firewall filter that filters only on a source IP address nearest to the destination devices. When applied too close to the source device, a firewall filter that filters only on a source IP address could potentially prevent that source device from accessing other services that are available on the network.



NOTE: Egress firewall filters do not affect the flow of locally generated control packets from the Routing Engine.

6. In which direction should the firewall filter be applied?

You can apply firewall filters to ports on the switch to filter packets that are entering a port. You can apply firewall filters to VLANs, and Layer 3 (routed) interfaces to filter packets that are entering or exiting a VLAN or routed interface. Typically, you configure different sets of actions for traffic entering an interface than you configure for traffic exiting an interface.

Related Documentation

- Firewall Filters for EX Series Switches Overview on page 3
- Understanding the Use of Policers in Firewall Filters on page 37
- Understanding How Firewall Filters Are Evaluated on page 31
- Understanding Filter-Based Forwarding for EX Series Switches on page 38
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 58

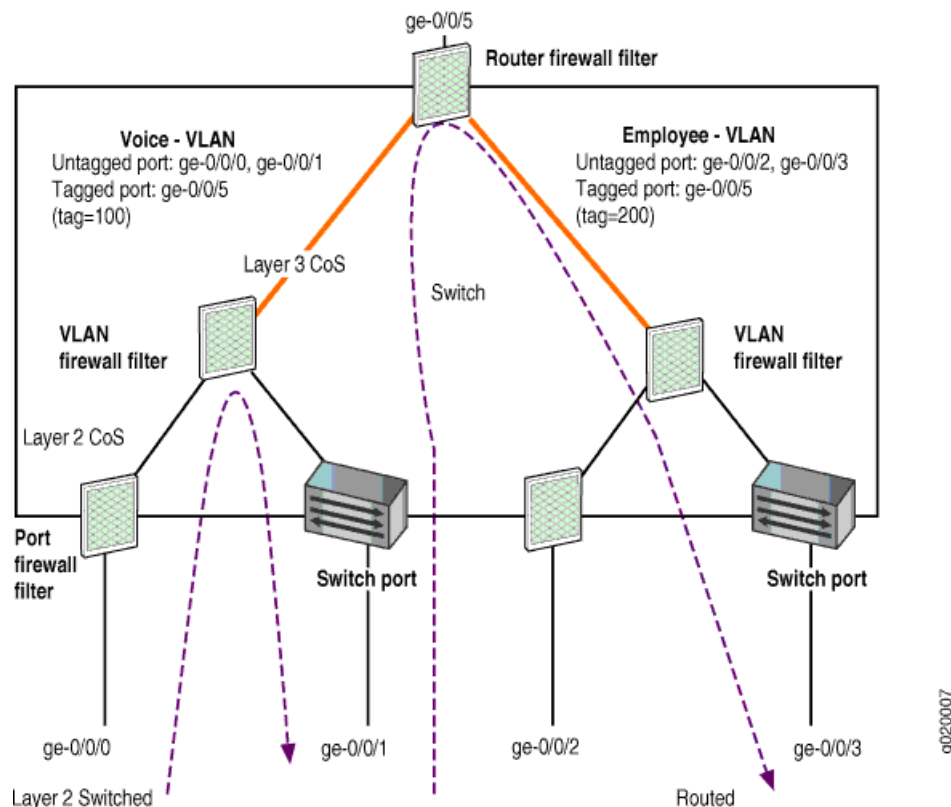
Understanding Firewall Filter Processing Points for Bridged and Routed Packets on EX Series Switches

Juniper Networks EX Series Ethernet Switches are multilayered switches that provide Layer 2 switching and Layer 3 routing. You apply firewall filters at multiple processing points in the packet forwarding path on EX Series switches. At each processing point, the action to be taken on a packet is determined based on the results of the lookup in the switch's forwarding table. A table lookup determines which exit port on the switch to use to forward the packet.

For both bridged unicast packets and routed unicast packets, firewall filters are evaluated and applied hierarchically. First, a packet is checked against the port firewall filter, if present. If the packet is permitted, it is then checked against the VLAN firewall filter, if present. If the packet is permitted, it is then checked against the router firewall filter, if present. The packet must be permitted by the router firewall filter before it is processed.

Figure 1 on page 9 shows the various firewall filter processing points in the packet forwarding path in a multilayered switching platform.

Figure 1: Firewall Filter Processing Points in the Packet Forwarding Path



For a multicast packet that results in replications, an egress firewall filter is applied to each copy of the packet based on its corresponding egress VLAN.

For Layer 2 (bridged) unicast packets, the following firewall filter processing points apply:

- Ingress port firewall filter
- Ingress VLAN firewall filter
- Egress port firewall filter
- Egress VLAN firewall filter

For Layer 3 (routed and multilayer-switched) unicast packets, the following firewall filter processing points apply:

- Ingress port firewall filter
- Ingress VLAN firewall filter (Layer 2 CoS)
- Ingress router firewall filter (Layer 3 CoS)
- Egress router firewall filter
- Egress VLAN firewall filter

**Related
Documentation**

- Firewall Filters for EX Series Switches Overview on page 3
- Understanding How Firewall Filters Control Packet Flows on page 10
- Understanding Bridging and VLANs on EX Series Switches
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39

Understanding How Firewall Filters Control Packet Flows

Juniper Networks EX Series Ethernet Switches support firewall filters that allow you to control flows of data packets and local packets. *Data packets* are chunks of data that transit the switch as they are forwarded from a source to a destination. *Local packets* are chunks of data that are destined for or sent by the switch. Local packets usually contain routing protocol data, data for IP services such as Telnet or SSH, and data for administrative protocols such as the Internet Control Message Protocol (ICMP).

You create firewall filters to protect your switch from excessive traffic transiting the switch to a network destination or destined for the Routing Engine on the switch. Firewall filters that control local packets can also protect your switch from external incidents such as denial-of-service (DoS) attacks.

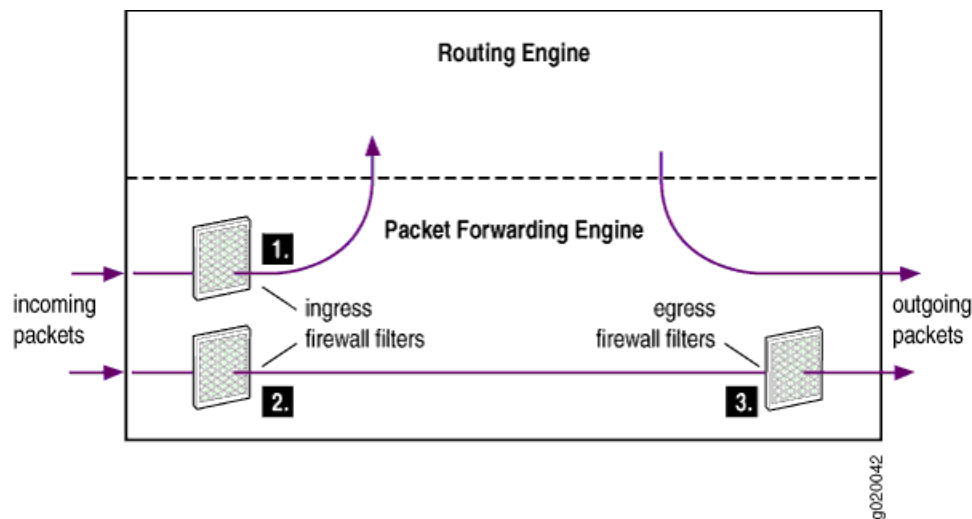
Firewall filters affect packet flows entering in to or exiting from the switch's interfaces:

- Ingress firewall filters affect the flow of data packets that are received by the switch's interfaces. The Packet Forwarding Engine (PFE) handles this flow. When a switch receives a data packet on an interface, the switch determines where to forward the packet by looking in the forwarding table for the best route (Layer 2 switching, Layer 3 routing) to a destination. Data packets are forwarded to their destination through an outgoing interface. Locally destined packets are forwarded to the Routing Engine.

- Egress firewall filters affect the flow of data packets that are transmitted from the switch's interfaces but do not affect the flow of locally generated control packets from the Routing Engine. The Packet Forwarding Engine handles the flow of data packets that are transmitted from the switch, and egress firewall filters are applied here. The Packet Forwarding Engine also handles the flow of control packets from the Routing Engine.

Figure 2 on page 11 illustrates the application of ingress and egress firewall filters to control the flow of packets through the switch.

Figure 2: Application of Firewall Filters to Control Packet Flow



1. Ingress firewall filter applied to control locally destined packets that are received on the switch's interfaces and are destined for the Routing Engine.
2. Ingress firewall filter applied to control incoming packets on the switch's interfaces.
3. Egress firewall filter applied to control packets that are transiting the switch's interfaces.

Related Documentation

- Understanding Firewall Filter Processing Points for Bridged and Routed Packets on EX Series Switches on page 9
- Understanding How Firewall Filters Are Evaluated on page 31

Firewall Filter Match Conditions and Actions for EX Series Switches

Each term in a firewall filter consists of *match conditions* and an *action*. Match conditions are the values or fields that a packet must contain. You can define multiple, single, or no match conditions. If no match conditions are specified for the term, all packets are matched by default. The action is the action that the switch takes if a packet matches the match conditions for the specific term. Action modifiers are optional and specify one or more actions that the switch takes if a packet matches the match conditions for the

specific term. Allowed actions are to accept a packet or discard a packet. In addition, you can specify action modifiers to count, mirror, rate limit, and classify packets.

For each firewall filter, you define the terms that specify the filtering criteria (match conditions) to apply to packets and the action for the switch to take if a match occurs. The string that defines a match condition is called a *match statement*. The following tables list various match conditions, their supported platforms, binding points, and actions.

- Table 1 on page 12 describes the match conditions you can specify when configuring a firewall filter for IPv4 traffic.
- Table 2 on page 22 describes the match conditions you can specify when configuring a firewall filter for IPv6 traffic.
- Table 3 on page 28 shows the actions that you can specify in a term.
- Table 4 on page 29 shows the action modifiers that you can specify in a term.

Table 1: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on EX Series Switches

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
destination-address <i>ip-address</i>	IP destination address field, which is the address of the final destination node.	<ul style="list-style-type: none"> • EX2200—ports, VLANs, and Layer 3 interfaces • EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces • EX4500—ports, VLANs, and Layer 3 interfaces • EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> • EX2200—ports, VLANs, and Layer 3 interfaces • EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces • EX4500—ports, VLANs, and Layer 3 interfaces • EX8200—ports, VLANs, and Layer 3 interfaces
destination-mac-address <i>mac-address</i>	<p>Destination media access control (MAC) address of the packet.</p> <p>You can define a destination MAC address with a prefix, such as from destination-mac-address 00:01:02:03:04:05/24. If no prefix is specified, the default value 48 is used.</p>	<ul style="list-style-type: none"> • EX2200—ports and VLANs • EX3200 and EX4200—ports and VLANs • EX4500—ports and VLANs • EX8200—ports and VLANs 	<ul style="list-style-type: none"> • EX2200—ports and VLANs • EX3200 and EX4200—ports and VLANs • EX4500—ports and VLANs • EX8200—ports and VLANs

Table 1: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
destination-port number	<p>TCP or User Datagram Protocol (UDP) destination port field. Typically, you specify this match in conjunction with the protocol match statement to determine which protocol is used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed):</p> <p>afs (1483), bgp (179), biff (512), bootpc (68), bootps (67),</p> <p>cmd (514), cvspserver (2401),</p> <p>dhcp (67), domain (53),</p> <p>eklogin (2105), ekshell (2106), exec (512),</p> <p>finger (79), ftp (21), ftp-data (20),</p> <p>http (80), https (443),</p> <p>ident (113), imap (143),</p> <p>kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544),</p> <p>ldap (389), login (513),</p> <p>mobileip-agent (434), mobileip-mn (435), msdp (639),</p> <p>netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123),</p> <p>pop3 (110), pptp (1723), printer (515),</p> <p>radacct (1813), radius (1812), rip (520), rkinit (2108),</p> <p>smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514),</p> <p>tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525),</p> <p>who (513),</p> <p>xmcp (177),</p> <p>zephyr-clt (2103), zephyr-hm (2104)</p>	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces

Table 1: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
destination-prefix-list <i>prefix-list</i>	<p>IP destination prefix list field.</p> <p>You can define a list of IP address prefixes under a prefix-list alias for frequent use. You make this definition at the [edit policy-options] hierarchy level.</p>	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces
dot1q-tag number	<p>The tag field in the Ethernet header. The tag values can be 1–4095.</p>	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX4500—ports and VLANs EX8200—ports and VLANs 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX4500—ports and VLANs EX8200—not supported
dot1q-user-priority <i>number</i>	<p>User-priority field of the tagged Ethernet packet. User-priority values can be 0–7.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> background (1)—Background best-effort (0)—Best effort controlled-load (4)—Controlled load excellent-load (3)—Excellent load network-control (7)—Network control reserved traffic standard (2)—Standard or Spare video (5)—Video voice (6)—Voice 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX4500—ports and VLANs EX8200—ports and VLANs 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX4500—ports and VLANs EX8200—ports and VLANs

Table 1: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
dscp number	<p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> • ef (46)—as defined in RFC 2598, <i>An Expedited Forwarding PHB</i>. • af11 (10), af12 (12), af13 (14); af21 (18), af22 (20), af23 (22); af31 (26), af32 (28), af33 (30); af41 (34), af42 (36), af43 (38) <p>These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in RFC 2597, <i>Assured Forwarding PHB</i>.</p>	<ul style="list-style-type: none"> • EX2200—ports, VLANs, and Layer 3 interfaces • EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces • EX4500—ports, VLANs, and Layer 3 interfaces • EX8200—ports and VLANs 	<ul style="list-style-type: none"> • EX2200—ports, VLANs, and Layer 3 interfaces • EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces • EX4500—ports, VLANs, and Layer 3 interfaces • EX8200—ports and VLANs
ether-type [aarp appletalk arp ipv4 ipv6 mpls—multicast mpls-unicast oam ppp pppoe-discovery pppoe-session sna [value]]	<p>Ethernet type field of a packet. The <i>EtherType value</i> specifies what protocol is being transported in the Ethernet frame. In place of the numeric value, you can specify one of the following text synonyms:</p> <ul style="list-style-type: none"> • aarp—EtherType value AARP (0x80F3) • appletalk—EtherType value AppleTalk (0x809B) • arp—EtherType value ARP (0x0806) • ipv4—EtherType value IPv4 (0x0800) • ipv6—EtherType value IPv6 (0x08DD) • mpls multicast—EtherType value MPLS multicast (0x8848) • mpls unicast—EtherType value MPLS unicast (0x8847) • oam—EtherType value OAM (0x88A8) • ppp—EtherType value PPP (0x880B) • pppoe-discovery—EtherType value PPPoE Discovery Stage (0x8863) • pppoe-session—EtherType value PPPoE Session Stage (0x8864) • sna—EtherType value SNA (0x80D5) 	<ul style="list-style-type: none"> • EX2200—ports and VLANs • EX3200 and EX4200—ports and VLANs • EX4500—ports and VLANs • EX8200—ports and VLANs 	<ul style="list-style-type: none"> • EX2200—ports and VLANs • EX3200 and EX4200—ports and VLANs • EX4500—ports and VLANs • EX8200—not supported.

Table 1: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
fragment-flags <i>fragment-flags</i>	IP fragmentation flags, specified in symbolic or hexadecimal formats. You can specify one of the following options: dont-fragment (0x4000), more-fragments (0x2000), or reserved (0x8000)	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—not supported 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—not supported EX4500—not supported EX8200—not supported
icmp-code number	<p>ICMP code field. This value or option provides more specific information than icmp-type. Because the value's meaning depends upon the associated icmp-type, you must specify icmp-type along with icmp-code. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The options are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> parameter-problem—ip-header-bad (0), required-option-missing (1) redirect—redirect-for-host (1), redirect-for-network (0), redirect-for-tos-and-host (3), redirect-for-tos-and-net (2) time-exceeded—ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) unreachable—communication-prohibited-by-filtering (13), destination-host-prohibited (10), destination-host-unknown (7), destination-network-prohibited (9), destination-network-unknown (6), fragmentation-needed (4), host-precedence-violation (14), host-unreachable (1), host-unreachable-for-TOS (12), network-unreachable (0), network-unreachable-for-TOS (11), port-unreachable (3), precedence-cutoff-in-effect (15), protocol-unreachable (2), source-host-isolated (8), source-route-failed (5) 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—VLANs and Layer 3 interfaces EX4500—VLANs and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces

Table 1: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
icmp-type <i>number</i>	ICMP packet type field. Typically, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): echo-reply (0) , echo-request (8) , info-reply (16) , info-request (15) , mask-request (17) , mask-reply (18) , parameter-problem (12) , redirect (5) , router-advertisement (9) , router-solicit (10) , source-quench (4) , time-exceeded (11) , timestamp (13) , timestamp-reply (14) , unreachable (3)	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces
interface <i>interface-name</i>	Interface on which the packet is received. You can specify the wildcard character (*) as part of an interface name.	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces
ip-options	Presence of the options field in the IP header.	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX4500—Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—not supported EX4500—not supported EX8200—not supported
is-fragment	If the packet is a trailing fragment, this match condition does not match the first fragment of a fragmented packet. Use two terms to match both first and trailing fragments.	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—not supported EX4500—not supported EX8200—not supported

Table 1: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
next-header bytes	<p>8-bit protocol field that identifies the type of header immediately following the IPv6 header. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <p>ah (51), dstops (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmpv6 (1), igmp (2), ipip (4), ipv6 (41), no-next-header (59), ospf (89), pim (103), routing (43), rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).</p>	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—Layer 3 interfaces EX4500—not supported EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—Layer 3 interfaces EX4500—not supported EX8200—not supported
precedence precedence	<p>IP precedence. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <p>critical-ecp (5), flash (3), flash-override (4), immediate (2), internet-control (6), net-control (7), priority (1), or routine (0).</p>	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces
protocol list of protocols	<p>IPv4 protocol value. In place of the numeric value, you can specify one of the following text synonyms:</p> <p>egp (8), esp (50), gre (47), icmp (1), igmp (2), ipip (4),</p> <p>ospf (89), pim (103), rsvp (46), tcp (6), udp (17)</p>	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces
source-address ip-address	<p>IP source address field, which is the address of the source node sending the packet. For IPV6, the source-address field is 128 bits in length. The filter description syntax supports the text representations for IPv6 addresses that are described in RFC 2373, <i>IP Version 6 Addressing Architecture</i>.</p>	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces

Table 1: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
source-mac-address <i>mac-address</i>	Source MAC address. You can define a source MAC address with a prefix, such as from destination-mac-address 00:01:02:03:04:05/24 . If no prefix is specified, the default value 48 is used.	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX4500—ports and VLANs EX8200—ports and VLANs 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX4500—ports and VLANs EX8200—ports and VLANs
source-port number	TCP or UDP source-port field. Typically, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. In place of the numeric field, you can specify one of the text synonyms listed under destination-port .	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces
source-prefix-list <i>prefix-list</i>	IP source prefix list field. You can define a list of IP address prefixes under a prefix-list alias for frequent use. You make this definition at the [edit policy-options] hierarchy level.	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces
tcp-established	TCP packets of an established TCP connection. This condition matches packets other than the first packet of a connection. tcp-established is a synonym for the bit names "(ack rst)" . tcp-established does not implicitly check whether the protocol is TCP. To do so, specify the protocol tcp match condition.	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—not supported EX4500—not supported EX8200—not supported

Table 1: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
tcp-flags [<i>flags</i> tcp-initial]	<p>One or more TCP flags:</p> <ul style="list-style-type: none"> bit-name—fin, syn, rst, push, ack, urgent logical operators—& (logical AND), (logical OR), ! (negation) numerical value—0x01 through 0x20 text synonym—tcp-initial <p>To specify multiple flags, use logical operators.</p>	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—not supported EX4500—not supported EX8200—not supported
tcp-initial	<p>Match the first TCP packet of a connection. tcp-initial is a synonym for the bit names "(syn & !ack)".</p> <p>tcp-initial does not implicitly check whether the protocol is TCP. To do so, specify the protocol tcp match condition.</p>	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—not supported EX4500—not supported EX8200—not supported
traffic-class	<p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> ef (46)—as defined in RFC 2598, <i>An Expedited Forwarding PHB</i>. af11 (10), af12 (12), af13 (14); af21 (18), af22 (20), af23 (22); af31 (26), af32 (28), af33 (30); af41 (34), af42 (36), af43 (38) <p>These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in RFC 2597, <i>Assured Forwarding PHB</i>.</p>	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—not supported EX8200—not supported 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX4500—not supported EX8200—not supported

Table 1: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
ttl value	TTL type to match. The value can be 1–255.	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX4500—Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—not supported EX4500—not supported EX8200—not supported
vlan [vlan-name vlan-id]	The VLAN that is associated with the packet.	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX4500—ports and VLANs EX8200—ports and VLANs 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX4500—ports and VLANs EX8200—ports and VLANs

Some of the numeric range and bit-field match conditions allow you to specify a text synonym. For a list of all the synonyms for a match condition, do any of the following:

- If you are using the J-Web Filters Configuration page, select the synonym from the appropriate list.
- If you are using the CLI, type a question mark (?) after the **from** statement.

To specify the bit-field value to match, you must enclose the values in quotation marks (" "). For example, a match occurs if the RST bit in the TCP flags field is set:

```
tcp-flags "rst";
```

For information about logical operators and how to use bit-field logical operations to create expressions that are evaluated for matches, see “Understanding Firewall Filter Match Conditions” on page 33.

On Juniper Networks EX Series Ethernet switches, you can apply a router firewall filter to both IPv4 and IPv6 traffic. You can apply firewall filter match conditions to IPv6 traffic on Layer 3 interfaces, aggregated Ethernet interfaces, and loopback interfaces. Table 2 on page 22 describes the match conditions you can specify when configuring a firewall filter for IPv6 traffic.

Table 2: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
destination-address <i>ip-address</i>	Specifies the 128-bit address that is the final destination node address for the packet. The filter description syntax supports the text representations for IPv6 addresses as described in RFC 2373 , <i>IP Version 6 Addressing Architecture</i> .	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces
destination-mac-address <i>mac-address</i>	<p>Destination media access control (MAC) address of the packet.</p> <p>You can define a destination MAC address with a prefix, such as from destination-mac-address 00:01:02:03:04:05/24. If no prefix is specified, the default value 48 is used.</p>	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX8200—ports and VLANs 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX8200—ports and VLANs

Table 2: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
destination-port number	<p>Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port field. Typically, you specify this match in conjunction with the protocol match statement to determine which protocol is used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed):</p> <p>afs (1483), bgp (179), biff (512), bootpc (68), bootps (67),</p> <p>cmd (514), cvspserver (2401),</p> <p>dhcp (67), domain (53),</p> <p>eklogin (2105), ekshell (2106), exec (512),</p> <p>finger (79), ftp (21), ftp-data (20),</p> <p>http (80), https (443),</p> <p>ident (113), imap (143),</p> <p>kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544),</p> <p>ldap (389), login (513),</p> <p>mobileip-agent (434), mobilip-mn (435), msdp (639),</p> <p>netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123),</p> <p>pop3 (110), pptp (1723), printer (515),</p> <p>radacct (1813), radius (1812), rip (520), rkinit (2108),</p> <p>smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514),</p>	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—VLANs, and Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces

Table 2: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
	tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), xdmcp (177), zephyr-clt (2103), zephyr-hm (2104)		
destination-prefix-list <i>prefix-list</i>	IP destination prefix list field. You can define a list of IP address prefixes under a prefix-list alias for frequent use. You make this definition at the [edit policy-options] hierarchy level.	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces
dot1q-tag <i>number</i>	The tag field in the Ethernet header. The tag values can be 1–4095.	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX8200—ports and VLANs 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX8200—not supported
dot1q-user-priority <i>number</i>	User-priority field of the tagged Ethernet packet. User-priority values can be 0–7. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): <ul style="list-style-type: none"> background (1)—Background best-effort (0)—Best effort controlled-load (4)—Controlled load excellent-load (3)—Excellent load network-control (7)—Network control reserved traffic standard (2)—Standard or Spare video (5)—Video voice (6)—Voice 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX8200—ports and VLANs 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX8200—ports and VLANs

Table 2: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
ether-type (ipv6) value	<p>Ethernet type field of a packet. The EtherType value specifies what protocol is being transported in the Ethernet frame. In place of the numeric value, you can specify the following text synonym:</p> <ul style="list-style-type: none"> ipv6—EtherType value IPv6 (0x08DD) 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX8200—ports and VLANs 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX8200—ports and VLANs.
icmp-code number	<p>ICMP code field. This value or option provides more specific information than icmp-type. Because the value's meaning depends upon the associated icmp-type, you must specify icmp-type along with icmp-code. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The options are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> parameter-problem—ip-header-bad (0), unrecognized-next-header (1), unrecognized-option (2) time-exceeded—ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) destination-unreachable—no-route-to--destination (0), administratively-prohibited (1), address-unreachable (3), port-unreachable (4) 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces
icmp-type number	<p>ICMP packet type field. Typically, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <p>echo-reply (0), echo-request (8), info-reply (16), info-request (15),</p> <p>mask-request (17), mask-reply (18), parameter-problem (12),</p> <p>redirect (5), router-advertisement (9), router-solicit (10), source-quench (4),</p> <p>time-exceeded (11), timestamp (13), timestamp-reply (14), unreachable (3)</p>	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX8200—Layer 3 interfaces

Table 2: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
interface <i>interface-name</i>	Interface on which the packet is received.	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX8200—ports, VLANs, and Layer 3 interfaces
next-header bytes	<p>8-bit protocol field that identifies the type of header immediately following the IPv6 header. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <p>ah (51), dstops (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmpv6 (1), igmp (2), ipip (4), ipv6 (41), no-next-header (59), ospf (89), pim (103), routing (43), rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).</p>	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces
packet-length bytes	<p>Length of the received packet, in bytes.</p> <p>The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.</p>	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX8200—Layer 3 interfaces
source-address <i>ip-address</i>	IP source address field, which is 128 bits in length. The filter description syntax supports the text representations for IPv6 addresses that are described in RFC 2373 , <i>IP Version 6 Addressing Architecture</i> .	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces
source-mac-address <i>mac-address</i>	<p>Source MAC address.</p> <p>You can define a source MAC address with a prefix, such as from destination-mac-address 00:01:02:03:04:05/24. If no prefix is specified, the default value 48 is used.</p>	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX8200—ports and VLANs 	<ul style="list-style-type: none"> EX2200—ports and VLANs EX3200 and EX4200—ports and VLANs EX8200—ports and VLANs

Table 2: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
source-port <i>number</i>	TCP or UDP source-port field. Typically, you specify this match in conjunction with the next-header match statement to determine which next-header is being used on the port. In place of the numeric field, you can specify one of the text synonyms listed under destination-port .	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces
source-prefix-list <i>prefix-list</i>	IP source prefix list field. You can define a list of IP address prefixes under a prefix-list alias for frequent use. You make this definition at the [edit policy-options] hierarchy level.	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—Layer 3 interfaces EX3200 and EX4200—Layer 3 interfaces EX8200—Layer 3 interfaces
tcp-flags (<i>flags</i> <i>tcp-initial</i>)	One or more TCP flags: <ul style="list-style-type: none"> bit-name—fin, syn, rst, push, ack, urgent logical operators—& (logical AND), (logical OR), ! (negation) numerical value—0x01 through 0x20 text synonym—tcp-initial <p>To specify multiple flags, use logical operators.</p>	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—not supported EX8200—not supported
tcp-initial	Match the first TCP packet of a connection. tcp-initial is a synonym for the bit names "(syn & !ack)". tcp-initial does not implicitly check whether the protocol is TCP. To do so, specify the protocol tcp match condition.	<ul style="list-style-type: none"> EX2200—ports, VLANs, and Layer 3 interfaces EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—not supported EX8200—not supported

Table 2: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches (*continued*)

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
traffic-class <i>number</i>	<p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> • ef (46)—as defined in RFC 2598, <i>An Expedited Forwarding PHB</i>. • af11 (10), af12 (12), af13 (14); af21 (18), af22 (20), af23 (22); af31 (26), af32 (28), af33 (30); af41 (34), af42 (36), af43 (38) <p>These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in RFC 2597, <i>Assured Forwarding PHB</i>.</p>	<ul style="list-style-type: none"> • EX2200—not supported • EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces • EX8200—Layer 3 interfaces 	<ul style="list-style-type: none"> • EX2200—not supported • EX3200 and EX4200—ports, VLANs, and Layer 3 interfaces • EX8200—Layer 3 interfaces
vlan (<i>vlan-id</i> <i>vlan-name</i>)	The VLAN that is associated with the packet.	<ul style="list-style-type: none"> • EX2200—ports and VLANs • EX3200 and EX4200—ports and VLANs • EX8200—ports and VLANs 	<ul style="list-style-type: none"> • EX2200—ports and VLANs • EX3200 and EX4200—ports and VLANs • EX8200—ports and VLANs

When you define one or more terms that specify the filtering criteria, you also define the action to take if the packet matches all criteria.

Table 3: Actions for Firewall Filters

Action	Description	Supported Platforms and Direction
accept	Accept a packet.	<ul style="list-style-type: none"> • EX2200—ingress and egress • EX3200 and EX4200—ingress and egress • EX4500—ingress and egress • EX8200—ingress and egress

Table 3: Actions for Firewall Filters (*continued*)

Action	Description	Supported Platforms and Direction
discard	Discard a packet silently without sending an Internet Control Message Protocol (ICMP) message.	<ul style="list-style-type: none"> EX2200—ingress and egress EX3200 and EX4200—ingress and egress EX4500—ingress and egress EX8200—ingress and egress
reject <i>message-type</i>	<p>Discard a packet, and send an ICMPv4 message (type 3) “destination unreachable”. You can log the rejected packets if you configure the syslog action modifier.</p> <p>You can specify one of the following message codes: administratively-prohibited (default), bad-host-tos, bad-network-tos, host-prohibited, host-unknown, host-unreachable, network-prohibited, network-unknown, network-unreachable, port-unreachable, precedence-cutoff, precedence-violation, protocol-unreachable, source-host-isolated, source-route-failed, or tcp-reset.</p> <p>If you specify tcp-reset, a TCP reset is returned if the packet is a TCP packet. Otherwise nothing is returned.</p> <p>If you do not specify a message type, the ICMP notification “destination unreachable” is sent with the default message “communication administratively filtered”.</p> <p>NOTE: reject is not a supported action for IPv6 traffic.</p>	<ul style="list-style-type: none"> EX2200—ingress and egress EX3200 and EX4200—ingress only EX4500—ingress only EX8200—ingress only
routing-instance <i>routing-instance-name</i>	Forward matched packets to a virtual routing instance.	<ul style="list-style-type: none"> EX2200—ingress only EX3200 and EX4200—ingress only EX4500—ingress only EX8200—ingress only
vlan <i>vlan-name</i>	<p>Forward matched packets to a specific VLAN. Ensure that you specify the VLAN name and not the VLAN range because the vlan action does not support the <i>vlan-range</i> option.</p> <p>NOTE: vlan is not a supported action for IPv6 traffic.</p>	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—ingress and egress EX4500—ingress and egress EX8200—ingress and egress

In addition to the actions, you can specify action modifiers.

Table 4: Action Modifiers for Firewall Filters

Action Modifier	Description	Supported Platforms and Direction
analyzer <i>analyzer-name</i>	Mirror port traffic to a specified destination port or VLAN that is connected to a protocol analyzer application. Mirroring copies all packets seen on one switch port to a network monitoring connection on another switch port. The analyzer name must be configured under [edit ethernet-switching-options analyzer] .	<ul style="list-style-type: none"> EX2200—ingress only EX3200 and EX4200—ingress only EX4500—ingress only EX8200—ingress only

Table 4: Action Modifiers for Firewall Filters (*continued*)

Action Modifier	Description	Supported Platforms and Direction
count <i>counter-name</i>	Count the number of packets that pass this filter, term, or policer.	<ul style="list-style-type: none"> EX2200—ingress only <p>NOTE: On EX2200 switches, count is supported for VLAN firewall filters only.</p> <ul style="list-style-type: none"> EX3200 and EX4200—ingress and egress EX4500—ingress only EX8200—ingress only
forwarding-class <i>class</i>	Classify the packet in one of the following forwarding classes: <ul style="list-style-type: none"> assured-forwarding best-effort expedited-forwarding network-control 	<ul style="list-style-type: none"> EX2200—ingress and egress EX3200 and EX4200—ingress and egress EX8200—ingress and egress
interface <i>interface-name</i>	Forward the traffic to the specified interface bypassing the switching lookup.	<ul style="list-style-type: none"> EX2200—not supported EX3200 and EX4200—ingress only EX8200—ingress only
log	Log the packet's header information in the Routing Engine. To view this information, issue the show firewall log command in the CLI. NOTE: log is not a supported action modifier for IPv6 traffic.	<ul style="list-style-type: none"> EX2200—ingress only EX3200 and EX4200—ingress only EX8200—ingress only
loss-priority (<i>high</i> <i>low</i>)	Set the packet loss priority (PLP).	<ul style="list-style-type: none"> EX2200—ingress and egress EX3200 and EX4200—ingress and egress EX8200—ingress and egress
policer <i>policer-name</i>	Apply rate limits to the traffic. You can specify a policer for ingress port, VLAN, and router firewall filters only.	<ul style="list-style-type: none"> EX2200—ingress only EX3200 and EX4200—ingress only EX8200—ingress only
syslog	Log an alert for this packet. You can specify that the log be sent to a server for storage and analysis. NOTE: syslog is not a supported action modifier for IPv6 traffic.	<ul style="list-style-type: none"> EX2200—ingress only EX3200 and EX4200—ingress only EX8200—ingress only



NOTE: On EX Series switches, **accept** and **discard** in the ingress direction, are the only actions supported for firewall filters applied on loopback interfaces.

- Related Documentation**
- Firewall Filter Configuration Statements Supported by Junos OS for EX Series Switches on page 96
 - Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
 - Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 58
 - Understanding Firewall Filter Match Conditions on page 33
 - Understanding How Firewall Filters Are Evaluated on page 31
 - Understanding How Firewall Filters Test a Packet's Protocol on page 36
 - Understanding the Use of Policers in Firewall Filters on page 37
 - Understanding Filter-Based Forwarding for EX Series Switches on page 38

Understanding How Firewall Filters Are Evaluated

A firewall filter consists of one or more terms, and the order of the terms within a firewall filter is important. Before you configure firewall filters, you should understand how Juniper Networks EX Series Ethernet Switches evaluate the terms within a firewall filter and how packets are evaluated against the terms.

When a firewall filter consists of a single term, the filter is evaluated as follows:

- If the packet matches all the conditions, the action in the **then** statement is taken.
- If the packet matches all the conditions, and no action is specified in the **then** statement, the default action **accept** is taken.

When a firewall filter consists of more than one term, the firewall filter is evaluated sequentially:

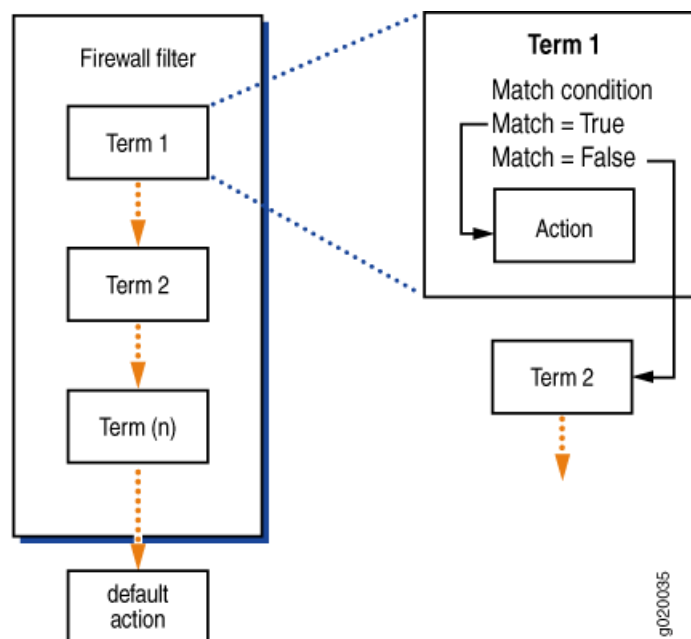
1. The packet is evaluated against the conditions in the **from** statement in the first term.
2. If the packet matches all the conditions in the term, the action in the **then** statement is taken and the evaluation ends. Subsequent terms in the filter are not evaluated.
3. If the packet does not match all the conditions in the term, the packet is evaluated against the conditions in the **from** statement in the second term.

This process continues until either the packet matches the conditions in the **from** statement in one of the subsequent terms or there are no more terms in the filter.

4. If a packet passes through all the terms in the filter without a match, the packet is discarded.

Figure 3 on page 32 shows how an EX Series switch evaluates the terms within a firewall filter.

Figure 3: Evaluation of Terms Within a Firewall Filter



If a term does not contain a **from** statement, the packet is considered to match and the action in the **then** statement of the term is taken.

If a term does not contain a **then** statement, or if an action has not been configured in the **then** statement, and the packet matches the conditions in the **from** statement of the term, the packet is accepted.

Every firewall filter contains an implicit **deny** statement at the end of the filter, which is equivalent to the following explicit filter term:

```
term implicit-rule {
  then discard;
}
```

Consequently, if a packet passes through all the terms in a filter without matching any conditions, the packet is discarded. If you configure a firewall filter that has no terms, all packets that pass through the filter are discarded.



NOTE: Firewall filtering is supported on packets that are at least 48 bytes long.

Related Documentation

- Firewall Filters for EX Series Switches Overview on page 3
- Understanding Firewall Filter Match Conditions on page 33
- Understanding the Use of Policers in Firewall Filters on page 37
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39

Understanding Firewall Filter Match Conditions

Before you define terms for firewall filters, you must understand how the conditions that you specify in a term are handled and how to specify interface filter, numeric filter, address filter, and bit-field filter match conditions to achieve the desired filtering results.

- Filter Match Conditions on page 33
- Numeric Filter Match Conditions on page 33
- Interface Filter Match Conditions on page 34
- IP Address Filter Match Conditions on page 34
- MAC Address Filter Match Conditions on page 35
- Bit-Field Filter Match Conditions on page 35

Filter Match Conditions

In the **from** statement of a firewall filter term, you specify the conditions that the packet must match for the action in the **then** statement to be taken. All conditions in the **from** statement must match for the action to be taken. The order in which you specify match conditions is not important, because a packet must match all the conditions in a term for a match to occur.

If you specify no match conditions in a term, that term matches all packets.

An individual condition in a **from** statement cannot contain a list of values. For example, you cannot specify numeric ranges or multiple source or destination addresses.

Individual conditions in a **from** statement cannot be negated. A negated condition is an explicit mismatch.

Numeric Filter Match Conditions

Numeric filter conditions match packet fields that are identified by a numeric value, such as port and protocol numbers. For numeric filter match conditions, you specify a keyword that identifies the condition and a single value that a field in a packet must match.

You can specify the numeric value in one of the following ways:

- Single number—A match occurs if the value of the field matches the number. For example:
`source-port 25;`
- Text synonym for a single number— A match occurs if the value of the field matches the number that corresponds to the synonym. For example:
`source-port http;`

To specify more than one value in a filter term, you enter each value in its own match statement. For example, a match occurs in the following term if the value of **vlan** field is 10 or 30.

```
[edit firewall family family-name filter filter-name term term-name from]
```

```
vlan 10;  
vlan 30;
```

The following restrictions apply to numeric filter match conditions:

- You cannot specify a range of values.
- You cannot specify a list of comma-separated values.
- You cannot exclude a specific value in a numeric filter match condition. For example, you cannot specify a condition that would match only if the match condition was not equal to a given value.

Interface Filter Match Conditions

Interface filter match conditions can match interface name values in a packet. For interface filter match conditions, you specify the name of the interface, for example:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#  
set interface ge-0/0/1
```

Port and VLAN interfaces do not use logical unit numbers. However, a firewall filter that is applied to a router interface can specify the logical unit number in the interface filter match condition, for example:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#  
set interface ge-0/1/0.0
```

You can include the * wildcard as part of the interface name, for example:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#  
set interface ge-0/*/1user@host# set interface ge-0/1/*user@host# set interface ge-*
```

IP Address Filter Match Conditions

Address filter match conditions can match prefix values in a packet, such as IP source and destination prefixes. For address filter match conditions, you specify a keyword that identifies the field and one prefix of that type that a packet must match.

You specify the address as a single prefix. A match occurs if the value of the field matches the prefix. For example:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#  
set destination-address 10.2.1.0/28;
```

Each prefix contains an implicit 0/0 except statement, which means that any prefix that does not match the prefix that is specified is explicitly considered not to match.

To specify the address prefix, use the notation prefix/prefix-length. If you omit prefix-length, it defaults to /32. For example:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#  
set destination-address 10[edit firewall family family-name filter filter-name term  
term-name from] user@host# showdestination-address {10.0.0.0/32;}
```

To specify more than one IP address in a filter term, you enter each address in its own match statement. For example, a match occurs in the following term if the value of the **source-address** field matches either of the following source-address prefixes:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#  
set source-address 10.0.0.0/8user@host# set source-address 10.1.0.0/16
```

MAC Address Filter Match Conditions

MAC address filter match conditions can match source and destination MAC address values in a packet. For MAC address filter match conditions, you specify a keyword that identifies the field and one value of that type that a packet must match.

You can specify the MAC address as six hexadecimal bytes in the following formats:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#
set destination-mac-address 0011.2233.4455
```

```
[edit firewall family family-name filter filter-name term term-name from]user@host#
set destination-mac-address 00:11:22:33:44:55
```

```
[edit firewall family family-name filter filter-name term term-name from]user@host#
set destination-mac-address 001122334455
```

To specify more than one MAC address in a filter term, you enter each MAC address in its own match statement. For example, a match occurs in the following term if the value of the **source-mac-address** field matches either of the following addresses.

```
[edit firewall family family-name filter filter-name term term-name from]user@host#
set source-mac-address 00:11:22:33:44:55user@host# set source-mac-address 00:11:22:33:20:15
```

Bit-Field Filter Match Conditions

Bit-field filter conditions match packet fields if particular bits in those fields are or are not set. You can match the IP options, TCP flags, and IP fragmentation fields. For bit-field filter match conditions, you specify a keyword that identifies the field and tests to determine that the option is present in the field.

To specify the bit-field value to match, enclose the value in double quotation marks. For example, a match occurs if the **RST** bit in the TCP flags field is set:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#
set tcp-flags "rst"
```

Typically, you specify the bits to be tested by using keywords. Bit-field match keywords always map to a single bit value. You also can specify bit fields as hexadecimal or decimal numbers.

To match multiple bit-field values, use the logical operators, which are described in Table 5 on page 35. The operators are listed in order from highest precedence to lowest precedence. Operations are left-associative.

Table 5: Actions for Firewall Filters

Logical Operators	Description
!	Negation.
&	Logical AND.
	Logical OR.

To negate a match, precede the value with an exclamation point. For example, a match occurs only if the RST bit in the TCP flags field is not set:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#  
set tcp-flags "!rst"
```

In the following example of a logical AND operation, a match occurs if the packet is the initial packet on a TCP session:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#  
set tcp-flags "syn&!ack"
```

In the following example of a logical OR operation, a match occurs if the packet is not the initial packet on a TCP session:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#  
set tcp-flags "syn|ack"
```

For a logical OR operation, you can specify a maximum of two match conditions in a single term. If you need to match more than two bit-field values in a logical OR operation, configure the same match condition in consecutive terms with additional bit-field values. In the following example, the two terms configured match the SYN, ACK, FIN, or RST bit in the TCP flags field:

```
[edit firewall family family-name filter filter-name term term-name1  
from]user@host# set tcp-flags "syn|ack"  
[edit firewall family family-name filter filter-name term term-name2  
from]user@host# set tcp-flags "fin|rst"
```

You can use text synonyms to specify some common bit-field matches. You specify these matches as a single keyword. In the following example of a text synonym, a match occurs if the packet is the initial packet on a TCP session:

```
[edit firewall family family-name filter filter-name term term-name from]user@host#  
set tcp-flags tcp-initial
```

Related Documentation

- Firewall Filters for EX Series Switches Overview on page 3
- Understanding How Firewall Filters Test a Packet's Protocol on page 36
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 58
- Firewall Filter Match Conditions and Actions for EX Series Switches on page 11

Understanding How Firewall Filters Test a Packet's Protocol

When examining match conditions, Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches tests only the field that is specified. The software does not implicitly test the IP header to determine whether a packet is an IP packet. Therefore, in some cases, you must specify **protocol** field match conditions in conjunction with other match conditions to ensure that the filters are performing the expected matches.

If you specify a protocol match condition or a match of the ICMP type or TCP flags field, there is no implied protocol match. For the following match conditions, you must explicitly specify the protocol match condition in the same term:

- **destination-port**—Specify the match **protocol tcp** or **protocol udp**.

- **source-port**—Specify the match **protocol tcp** or **protocol udp**.

If you do not specify the protocol when using the preceding fields, design your filters carefully to ensure that they perform the expected matches. For example, if you specify a match of **destination-port ssh**, the switch deterministically matches any packets that have a value of **22** in the two-byte field that is two bytes beyond the end of the IP header without ever checking the IP protocol field.

**Related
Documentation**

- Firewall Filters for EX Series Switches Overview on page 3
- Understanding Firewall Filter Match Conditions on page 33
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39

Understanding the Use of Policers in Firewall Filters

Policing, or rate limiting, is an important component of firewall filters that lets you control the amount of traffic that enters an interface.

A single firewall filter configured with a policer permits only traffic at specified data rates to provide protection from denial-of-service (DOS) attacks. Traffic that exceeds the rate limits specified by the policer can be discarded. Discard is the only supported policer action. Typically, traffic that exceeds the rate limits specified by the policer is either discarded or marked as lower priority than traffic that meets the rate limits specified by the policer. When necessary, low-priority traffic can be discarded by the switch to prevent congestion.

A policer applies two types of rate limits on traffic:

- **Bandwidth**—The number of bits per second permitted, on average.
- **Maximum burst size**—The maximum size permitted for bursts of data that exceed the given bandwidth limit.

Policing uses an algorithm to enforce a limit on average bandwidth while allowing bursts up to a specified maximum value. You can define specific classes of traffic on an interface and apply a set of rate limits to each class. After you name and configure a policer, it is stored as a template. You can then use a policer in a firewall filter configuration.

Each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. To get filter or term-specific packets counts, you must configure a new policer for each filter or term that requires policing.

**Related
Documentation**

- Firewall Filters for EX Series Switches Overview on page 3
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Firewall Filter Match Conditions and Actions for EX Series Switches on page 11

Understanding Filter-Based Forwarding for EX Series Switches

Administrators of Juniper Networks EX Series Ethernet Switches can use firewall filters in conjunction with virtual routing instances to specify different routes for packets to travel in their networks. To set up this feature, which is called filter-based forwarding, you specify a filter and match criteria and then specify the virtual routing instance to send packets to.

You might want to use filter-based forwarding to route specific types of traffic through a firewall or security device before the traffic continues on its path. You can also use filter-based forwarding to give certain types of traffic preferential treatment or to improve load balancing of switch traffic.

Related Documentation

- [Understanding Virtual Routing Instances on EX Series Switches](#)
- [Firewall Filters for EX Series Switches Overview on page 3](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 58](#)

CHAPTER 2

Examples of Firewall Filters Configuration

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 58
- Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch on page 62

Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches

This example shows how to configure and apply firewall filters to control traffic that is entering or exiting a port on the switch, a VLAN on the network, and a Layer 3 interface on the switch. Firewall filters define the rules that determine whether to forward or deny packets at specific processing points in the packet flow.

- Requirements on page 39
- Overview on page 40
- Configuring an Ingress Port Firewall Filter to Prioritize Voice Traffic and Rate-Limit TCP and ICMP Traffic on page 43
- Configuring a VLAN Ingress Firewall Filter to Prevent Rogue Devices from Disrupting VoIP Traffic on page 48
- Configuring a VLAN Firewall Filter to Count, Monitor, and Analyze Egress Traffic on the Employee VLAN on page 50
- Configuring a VLAN Firewall Filter to Restrict Guest-to-Employee Traffic and Peer-to-Peer Applications on the Guest VLAN on page 52
- Configuring a Router Firewall Filter to Give Priority to Egress Traffic Destined for the Corporate Subnet on page 54
- Verification on page 56

Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.0 or later for EX Series switches.
- Two Juniper Networks EX3200-48T switches: one to be used as an access switch, the other to be used as a distribution switch
- One Juniper Networks EX-UM-4SFP uplink module
- One Juniper Networks J-series router

Before you configure and apply the firewall filters in this example, be sure you have:

- An understanding of firewall filter concepts, policers, and CoS
- Installed the uplink module in the distribution switch. See [Installing an Uplink Module in an EX3200 Switch](#).

Overview

This configuration example show how to configure and apply firewall filters to provide rules to evaluate the contents of packets and determine when to discard, forward, classify, count, and analyze packets that are destined for or originating from the EX Series switches that handle all **voice-vlan**, **employee-vlan**, and **guest-vlan** traffic. Table 6 on page 40 shows the firewall filters that are configured for the EX Series switches in this example.

Table 6: Configuration Components: Firewall Filters


Component	Purpose/Description
Port firewall filter, ingress-port-voip-class-limit-tcp-icmp	<p>This firewall filter performs two functions:</p> <ul style="list-style-type: none"> • Assigns priority queueing to packets with a source MAC address that matches the phone MAC addresses. The forwarding class expedited-forwarding provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service for all voice-vlan traffic. • Performs rate limiting on packets that enter the ports for employee-vlan. The traffic rate for TCP and ICMP packets is limited to 1 Mbps with a burst size up to 30,000 bytes. <p>This firewall filter is applied to port interfaces on the access switch.</p>
VLAN firewall filter, ingress-vlan-rogue-block	<p>Prevents rogue devices from using HTTP sessions to mimic the gatekeeper device that manages call registration, admission, and call status for VoIP calls. Only TCP or UDP ports should be used; and only the gatekeeper uses HTTP. That is, all voice-vlan traffic on TCP ports should be destined for the gatekeeper device. This firewall filter applies to all phones on voice-vlan, including communication between any two phones on the VLAN and all communication between the gatekeeper device and VLAN phones.</p> <p>This firewall filter is applied to VLAN interfaces on the access switch.</p>
VLAN firewall filter, egress-vlan-watch-employee	<p>Accepts employee-vlan traffic destined for the corporate subnet, but does not monitor this traffic. Employee traffic destined for the Web is counted and analyzed.</p> <p>This firewall filter is applied to vlan interfaces on the access switch.</p>

Table 6: Configuration Components: Firewall Filters (*continued*)

Component	Purpose/Description
VLAN firewall filter, ingress-vlan-limit-guest	Prevents guests (non-employees) from talking with employees or employee hosts on employee-vlan . Also prevents guests from using peer-to-peer applications on guest-vlan , but allows guests to access the Web. This firewall filter is applied to VLAN interfaces on the access switch.
Router firewall filter, egress-router-corp-class	Prioritizes employee-vlan traffic, giving highest forwarding-class priority to employee traffic destined for the corporate subnet. This firewall filter is applied to a routed port (Layer 3 uplink module) on the distribution switch.

Figure 4 on page 41 shows the application of port, VLAN, and Layer 3 routed firewall filters on the switch.

Figure 4: Application of Port, VLAN, and Layer 3 Routed Firewall Filters



Network Topology

The topology for this configuration example consists of one EX-3200-48T switch at the access layer, and one EX-3200-48T switch at the distribution layer. The distribution switch's uplink module is configured to support a Layer 3 connection to a J-series router.

The EX Series switches are configured to support VLAN membership. Table 7 on page 42 shows the VLAN configuration components for the VLANs.

Table 7: Configuration Components: VLANs

VLAN Name	VLAN ID	VLAN Subnet and Available IP Addresses	VLAN Description
voice-vlan	10	192.0.2.0/28 192.0.2.1 through 192.0.2.14 192.0.2.15 is subnet's broadcast address	Voice VLAN used for employee VoIP traffic
employee-vlan	20	192.0.2.16/28 192.0.2.17 through 192.0.2.30 192.0.2.31 is subnet's broadcast address	VLAN standalone PCs, PCs connected to the network through the hub in VoIP telephones, wireless access points, and printers. This VLAN completely includes the voice VLAN. Two VLANs (voice-vlan and employee-vlan) must be configured on the ports that connect to the telephones.
guest-vlan	30	192.0.2.32/28 192.0.2.33 through 192.0.2.46 192.0.2.47 is subnet's broadcast address	VLAN for guests' data devices (PCs). The scenario assumes that the corporation has an area open to visitors, either in the lobby or in a conference room, that has a hub to which visitors can plug in their PCs to connect to the Web and to their company's VPN.
camera-vlan	40	192.0.2.48/28 192.0.2.49 through 192.0.2.62 192.0.2.63 is subnet's broadcast address	VLAN for the corporate security cameras.

Ports on the EX Series switches support Power over Ethernet (PoE) to provide both network connectivity and power for VoIP telephones connecting to the ports. Table 8 on page 42 shows the switch ports that are assigned to the VLANs and the IP and MAC addresses for devices connected to the switch ports:

Table 8: Configuration Components: Switch Ports on a 48-Port All-PoE Switch

Switch and Port Number	VLAN Membership	IP and MAC Addresses	Port Devices
ge-0/0/0, ge-0/0/1	voice-vlan, employee-vlan	IP addresses: 192.0.2.1 through 192.0.2.2 MAC addresses: 00.05.85.00.00.01, 00.05.85.00–00.02	Two VoIP telephones, each connected to one PC.

Table 8: Configuration Components: Switch Ports on a 48-Port All-PoE Switch (*continued*)

Switch and Port Number	VLAN Membership	IP and MAC Addresses	Port Devices
ge-0/0/2, ge-0/0/3	employee-vlan	192.0.2.17 through 192.0.2.18	Printer, wireless access points
ge-0/0/4, ge-0/0/5	guest-vlan	192.0.2.34 through 192.0.2.35	Two hubs into which visitors can plug in their PCs. Hubs are located in an area open to visitors, such as a lobby or conference room
ge-0/0/6, ge-0/0/7	camera-vlan	192.0.2.49 through 192.0.2.50	Two security cameras
ge-0/0/9	voice-vlan	IP address: 192.0.2.14 MAC address: 00.05.85.00.00.0E	Gatekeeper device. The gatekeeper manages call registration, admission, and call status for VoIP phones.
ge-0/1/0		IP address: 192.0.2.65	Layer 3 connection to a router; note that this is a port on the switch's uplink module

Configuring an Ingress Port Firewall Filter to Prioritize Voice Traffic and Rate-Limit TCP and ICMP Traffic

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

CLI Quick Configuration

To quickly configure and apply a port firewall filter to prioritize voice traffic and rate-limit packets that are destined for the **employee-vlan** subnet, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall policer tcp-connection-policer if-exceeding burst-size-limit 30k bandwidth-limit 1m
set firewall policer tcp-connection-policer then discard
set firewall policer icmp-connection-policer if-exceeding burst-size-limit 30k bandwidth-limit 1m
set firewall policer icmp-connection-policer then discard
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
from source-mac-address 00.05.85.00.00.01
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
from source-mac-address 00.05.85.00.00.02
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
from protocol udp
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
then forwarding-class expedited-forwarding
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
then loss-priority low
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
network-control from precedence net-control
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
network-control then forwarding-class network-control
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
network-control then loss-priority low
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection from destination-address 192.0.2.16/28
```

```

set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection from protocol tcp
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then policer tcp-connection-policer
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then count tcp-counter
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then forwarding-class best-effort
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then loss-priority high
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection from destination-address 192.0.2.16/28
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection from protocol icmp
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then policer icmp-connection-policer
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then count icmp-counter
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then forwarding-class best-effort
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then loss-priority high
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term best-effort
then forwarding-class best-effort
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term best-effort
then loss-priority high
set interfaces ge-0/0/0 description "voice priority and tcp and icmp traffic rate-limiting filter at
ingress port"
set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
set interfaces ge-0/0/1 description "voice priority and tcp and icmp traffic rate-limiting filter at
ingress port"
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
set class-of-service schedulers voice-high buffer-size percent 15
set class-of-service schedulers voice-high priority high
set class-of-service schedulers net-control buffer-size percent 10
set class-of-service schedulers net-control priority high
set class-of-service schedulers best-effort buffer-size percent 75
set class-of-service schedulers best-effort priority low
set class-of-service scheduler-maps ethernet-diffsrv-cos-map forwarding-class
expedited-forwarding scheduler voice-high
set class-of-service scheduler-maps ethernet-diffsrv-cos-map forwarding-class network-control
scheduler net-control
set class-of-service scheduler-maps ethernet-diffsrv-cos-map forwarding-class best-effort
scheduler best-effort

```

Step-by-Step Procedure To configure and apply a port firewall filter to prioritize voice traffic and rate-limit packets that are destined for the **employee-vlan** subnet:

1. Define the policers **tcp-connection-policer** and **icmp-connection-policer**:

```
[edit]
user@switch# set firewall policer tcp-connection-policer if-exceeding burst-size-limit
30k bandwidth-limit 1m
user@switch# set firewall policer tcp-connection-policer then discard
user@switch# set firewall policer icmp-connection-policer if-exceeding burst-size-limit
30k bandwidth-limit 1m
user@switch# set firewall policer icmp-connection-policer then discard
```

2. Define the firewall filter **ingress-port-voip-class-limit-tcp-icmp**:

```
[edit firewall]
user@switch# set family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp
```

3. Define the term **voip-high**:

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp ]
user@switch# set term voip-high from source-mac-address 00.05.85.00.00.01
user@switch# set term voip-high from source-mac-address 00.05.85.00.00.02
user@switch# set term voip-high from protocol udp
user@switch# set term voip-high then forwarding-class expedited-forwarding
user@switch# set term voip-high then loss-priority low
```

4. Define the term **network-control**:

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp ]
user@switch# set term network-control from precedence net-control
user@switch# set term network-control then forwarding-class network-control
user@switch# set term network-control then loss-priority low
```

5. Define the term **tcp-connection** to configure rate limits for TCP traffic:

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp]
user@switch# set term tcp-connection from destination-address 192.0.2.16/28
user@switch# set term tcp-connection from protocol tcp
user@switch# set term tcp-connection then policer tcp-connection-policer
user@switch# set term tcp-connection then count tcp-counter
user@switch# set term tcp-connection then forwarding-class best-effort
user@switch# set term tcp-connection then loss-priority high
```

6. Define the term **icmp-connection** to configure rate limits for ICMP traffic:

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp]
user@switch# set term icmp-connection from destination-address 192.0.2.16/28
user@switch# set term icmp-connection from protocol icmp
user@switch# set term icmp-connection then policer icmp-policer
user@switch# set term icmp-connection then count icmp-counter
user@switch# set term icmp-connection then forwarding-class best-effort
user@switch# set term icmp-connection then loss-priority high
```

7. Define the term **best-effort** with no match conditions for an implicit match on all packets that did not match any other term in the firewall filter:

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp]
user@switch# set term best-effort then forwarding-class best-effort
user@switch# set term best-effort then loss-priority high
```

8. Apply the firewall filter **ingress-port-voip-class-limit-tcp-icmp** as an input filter to the port interfaces for **employee-vlan** :

```
[edit interfaces]
user@switch# set ge-0/0/0 description "voice priority and tcp and icmp traffic
rate-limiting filter at ingress port"
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
user@switch# set ge-0/0/1 description "voice priority and tcp and icmp traffic
rate-limiting filter at ingress port"
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
```

9. Configure the parameters that are desired for the different schedulers.



NOTE: When you configure parameters for the schedulers, define the numbers to match your network traffic patterns.

```
[edit class-of-service]
user@switch# set schedulers voice-high buffer-size percent 15
user@switch# set schedulers voice-high priority high
user@switch# set schedulers network-control buffer-size percent 10
user@switch# set schedulers network-control priority high
user@switch# set schedulers best-effort buffer-size percent 75
user@switch# set schedulers best-effort priority low
```

10. Assign the forwarding classes to schedulers with a scheduler map:

```
[edit class-of-service]
user@switch# set scheduler-maps ethernet-diffsrv-cos-map
user@switch# set scheduler-maps ethernet-diffsrv-cos-map forwarding-class
expedited-forwarding scheduler voice-high
user@switch# set scheduler-maps ethernet-diffsrv-cos-map forwarding-class
network-control scheduler net-control
user@switch# set scheduler-maps ethernet-diffsrv-cos-map forwarding-class
best-effort scheduler best-effort
```

11. Associate the scheduler map with the outgoing interface:

```
[edit class-of-service]
user@switch# set interfaces ge-0/1/0 scheduler-map ethernet-diffsrv-cos-map
```

Results Display the results of the configuration:

```
user@switch# show
firewall {
  policer tcp-connection-policer {
    if-exceeding {
      bandwidth-limit 1m;
      burst-size-limit 30k;
    }
  }
}
```

```
    then {
        discard;
    }
}
}
policer icmp-connection-policer {
    if-exceeding {
        bandwidth-limit 1m;
        burst-size-limit 30k;
    }
    then {
        discard;
    }
}
}
family ethernet-switching {
    filter ingress-port-voip-class-limit-tcp-icmp {
        term voip-high {
            from {
                destination-mac-address 00.05.85.00.00.01;
                destination-mac-address 00.05.85.00.00.02;
                protocol udp;
            }
            then {
                forwarding-class expedited-forwarding;
                loss-priority low;
            }
        }
        term network-control {
            from {
                precedence net-control ;
            }
            then {
                forwarding-class network-control;
                loss-priority low;
            }
        }
        term tcp-connection {
            from {
                destination-address 192.0.2.16/28;
                protocol tcp;
            }
            then {
                policer tcp-connection-policer;
                count tcp-counter;
                forwarding-class best-effort;
                loss-priority high;
            }
        }
        term icmp-connection
            from {
                protocol icmp;
            }
            then {
                policer icmp-connection-policer;
                count icmp-counter;
                forwarding-class best-effort;
                loss-priority high;
            }
        }
    }
}
```

```

}
}
term best-effort {
    then {
        forwarding-class best-effort;
        loss-priority high;
    }
}
}
}
}
}
interfaces {
    ge-0/0/0 {
        description "voice priority and tcp and icmp traffic rate-limiting filter at ingress port";
        unit 0 {
            family ethernet-switching {
                filter {
                    input ingress-port-voip-class-limit-tcp-icmp;
                }
            }
        }
    }
}
ge-0/0/1 {
    description "voice priority and tcp and icmp traffic rate-limiting filter at ingress port";
    unit 0 {
        family ethernet-switching {
            filter {
                input ingress-port-voip-class-limit-tcp-icmp;
            }
        }
    }
}
}
scheduler-maps {
    ethernet-diffsrv-cos-map {
        forwarding-class expedited-forwarding scheduler voice-high;
        forwarding-class network-control scheduler net-control;
        forwarding-class best-effort scheduler best-effort;
    }
}
interfaces {
    ge-0/1/0 {
        scheduler-map ethernet-diffsrv-cos-map;
    }
}
}

```

Configuring a VLAN Ingress Firewall Filter to Prevent Rogue Devices from Disrupting VoIP Traffic

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

CLI Quick Configuration

To quickly configure a VLAN firewall filter on **voice-vlan** to prevent rogue devices from using HTTP sessions to mimic the gatekeeper device that manages VoIP traffic, copy the following commands and paste them into the switch terminal window:

[edit]


```

set firewall family ethernet-switching filter ingress-vlan-rogue-block term to-gatekeeper from
destination-address 192.0.2.14
set firewall family ethernet-switching filter ingress-vlan-rogue-block term to-gatekeeper from
destination-port 80
set firewall family ethernet-switching filter ingress-vlan-rogue-block term to-gatekeeper then
accept
set firewall family ethernet-switching filter ingress-vlan-rogue-block term from-gatekeeper from
source-address 192.0.2.14
set firewall family ethernet-switching filter ingress-vlan-rogue-block term from-gatekeeper from
source-port 80
set firewall family ethernet-switching filter ingress-vlan-rogue-block term from-gatekeeper then
accept
set firewall family ethernet-switching filter ingress-vlan-rogue-block term not-gatekeeper from
destination-port 80
set firewall family ethernet-switching filter ingress-vlan-rogue-block term not-gatekeeper then
count rogue-counter
set firewall family ethernet-switching filter ingress-vlan-rogue-block term not-gatekeeper then
discard
set vlans voice-vlan description "block rogue devices on voice-vlan"
set vlans voice-vlan filter input ingress-vlan-rogue-block

```

Step-by-Step Procedure

To configure and apply a VLAN firewall filter on **voice-vlan** to prevent rogue devices from using HTTP to mimic the gatekeeper device that manages VoIP traffic:

1. Define the firewall filter **ingress-vlan-rogue-block** to specify filter matching on the traffic you want to permit and restrict:

```

[edit firewall]
user@switch# set family ethernet-switching filter ingress-vlan-rogue-block

```

2. Define the term **to-gatekeeper** to accept packets that match the destination IP address of the gatekeeper:

```

[edit firewall family ethernet-switching filter ingress-vlan-rogue-block]
user@switch# set term to-gatekeeper from destination-address 192.0.2.14
user@switch# set term to-gatekeeper from destination-port 80
user@switch# set term to-gatekeeper then accept

```

3. Define the term **from-gatekeeper** to accept packets that match the source IP address of the gatekeeper:

```

[edit firewall family ethernet-switching filter ingress-vlan-rogue-block]
user@switch# set term from-gatekeeper from source-address 192.0.2.14
user@switch# set term from-gatekeeper from source-port 80
user@switch# set term from-gatekeeper then accept

```

4. Define the term **not-gatekeeper** to ensure all **voice-vlan** traffic on TCP ports is destined for the gatekeeper device:

```

[edit firewall family ethernet-switching filter ingress-vlan-rogue-block]
user@switch# set term not-gatekeeper from destination-port 80
user@switch# set term not-gatekeeper then count rogue-counter
user@switch# set term not-gatekeeper then discard

```

5. Apply the firewall filter **ingress-vlan-rogue-block** as an input filter to the VLAN interface for the VoIP telephones:

```

[edit]
user@switch# set vlans voice-vlan description "block rogue devices on voice-vlan"
user@switch# set vlans voice-vlan filter input ingress-vlan-rogue-block

```

Results Display the results of the configuration:

```
user@switch# show
firewall {
  family ethernet-switching {
    filter ingress-vlan-rogue-block {
      term to-gatekeeper {
        from {
          destination-address 192.0.2.14/32
          destination-port 80;
        }
        then {
          accept;
        }
      }
      term from-gatekeeper {
        from {
          source-address 192.0.2.14/32
          source-port 80;
        }
        then {
          accept;
        }
      }
      term not-gatekeeper {
        from {
          destination-port 80;
        }
        then {
          count rogue-counter;
          discard;
        }
      }
    }
  }
}
vlands {
  voice-vlan {
    description "block rogue devices on voice-vlan";
    filter {
      input ingress-vlan-rogue-block;
    }
  }
}
```

Configuring a VLAN Firewall Filter to Count, Monitor, and Analyze Egress Traffic on the Employee VLAN

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

CLI Quick Configuration

A firewall filter is configured and applied to VLAN interfaces to filter **employee-vlan** egress traffic. Employee traffic destined for the corporate subnet is accepted but not monitored. Employee traffic destined for the Web is counted and analyzed.

To quickly configure and apply a VLAN firewall filter, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-corp
from destination-address 192.0.2.16/28
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-corp
then accept
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-web
from destination-port 80
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-web
then count employee-web-counter
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-web
then analyzer employee-monitor
set vlans employee-vlan description "filter at egress VLAN to count and analyze employee to
Web traffic"
set vlans employee-vlan filter output egress-vlan-watch-employee
```

Step-by-Step Procedure To configure and apply an egress port firewall filter to count and analyze **employee-vlan** traffic that is destined for the Web:

1. Define the firewall filter **egress-vlan-watch-employee**:

```
[edit firewall]
user@switch# set family ethernet-switching filter egress-vlan-watch-employee
```

2. Define the term **employee-to-corp** to accept but not monitor all **employee-vlan** traffic destined for the corporate subnet:

```
[edit firewall family ethernet-switching filter egress-vlan-watch-employee]
user@switch# set term employee-to-corp from destination-address 192.0.2.16/28
user@switch# set term employee-to-corp then accept
```

3. Define the term **employee-to-web** to count and monitor all **employee-vlan** traffic destined for the Web:

```
[edit firewall family ethernet-switching filter egress-vlan-watch-employee]
user@switch# set term employee-to-web from destination-port 80
user@switch# set term employee-to-web then count employee-web-counter
user@switch# set term employee-to-web then analyzer employee-monitor
```



NOTE: See [Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches](#) for information about configuring the **employee-monitor** analyzer.

4. Apply the firewall filter **egress-vlan-watch-employee** as an output filter to the port interfaces for the VoIP telephones:

```
[edit]
user@switch# set vlans employee-vlan description "filter at egress VLAN to count and
analyze employee to Web traffic"
user@switch# set vlans employee-vlan filter output egress-vlan-watch-employee
```

Results Display the results of the configuration:

```
user@switch# show
firewall {
  family ethernet-switching {
```

```
filter egress-vlan-watch-employee {
  term employee-to-corp {
    from {
      destination-address 192.0.2.16/28
    }
    then {
      accept;
    }
  }
  term employee-to-web {
    from {
      destination-port 80;
    }
    then {
      count employee-web-counter;
      analyzer employee-monitor;
    }
  }
}
}
}
vlands {
  employee-vlan {
    description "filter at egress VLAN to count and analyze employee to Web traffic";
    filter {
      output egress-vlan-watch-employee;
    }
  }
}
```

Configuring a VLAN Firewall Filter to Restrict Guest-to-Employee Traffic and Peer-to-Peer Applications on the Guest VLAN

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

CLI Quick Configuration

In the following example, the first filter term permits guests to talk with other guests but not employees on **employee-vlan**. The second filter term allows guests Web access but prevents them from using peer-to-peer applications on **guest-vlan**.


To quickly configure a VLAN firewall filter to restrict guest-to-employee traffic, blocking guests from talking with employees or employee hosts on **employee-vlan** or attempting to use peer-to-peer applications on **guest-vlan**, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family ethernet-switching filter ingress-vlan-limit-guest term guest-to-guest from
destination-address 192.0.2.33/28
set firewall family ethernet-switching filter ingress-vlan-limit-guest term guest-to-guest then
accept
set firewall family ethernet-switching filter ingress-vlan-limit-guest term
no-guest-employee-no-peer-to-peer from destination-mac-address 00.05.85.00.00.DF
set firewall family ethernet-switching filter ingress-vlan-limit-guest term
no-guest-employee-no-peer-to-peer then accept
set vlans guest-vlan description "restrict guest-to-employee traffic and peer-to-peer applications
on guest VLAN"
```

- Step-by-Step Procedure** **set vlans guest-vlan filter input ingress-vlan-limit-guest**
- To configure and apply a VLAN firewall filter to restrict guest-to-employee traffic and peer-to-peer applications on **guest-vlan**:
1. Define the firewall filter **ingress-vlan-limit-guest**:


```
[edit firewall]
set firewall family ethernet-switching filter ingress-vlan-limit-guest
```
 2. Define the term **guest-to-guest** to permit guests on the **guest-vlan** to talk with other guests but not employees on the **employee-vlan**:


```
[edit firewall family ethernet-switching filter ingress-vlan-limit-guest]
user@switch# set term guest-to-guest from destination-address 192.0.2.33/28
user@switch# set term guest-to-guest then accept
```
 3. Define the term **no-guest-employee-no-peer-to-peer** to allow guests on **guest-vlan** Web access but prevent them from using peer-to-peer applications on the **guest-vlan**.

 **NOTE:** The `destination-mac-address` is the default gateway, which for any host in a VLAN is the next-hop router.

```
[edit firewall family ethernet-switching filter ingress-vlan-limit-guest]
user@switch# set term no-guest-employee-no-peer-to-peer from
destination-mac-address 00.05.85.00.00.DF
user@switch# set term no-guest-employee-no-peer-to-peer then accept
```
 4. Apply the firewall filter **ingress-vlan-limit-guest** as an input filter to the interface for **guest-vlan** :


```
[edit]
user@switch# set vlans guest-vlan description "restrict guest-to-employee traffic and
peer-to-peer applications on guest VLAN"
user@switch# set vlans guest-vlan filter input ingress-vlan-limit-guest
```

Results Display the results of the configuration:

```
user@switch# show
firewall {
  family ethernet-switching {
    filter ingress-vlan-limit-guest {
      term guest-to-guest {
        from {
          destination-address 192.0.2.33/28;
        }
        then {
          accept;
        }
      }
      term no-guest-employee-no-peer-to-peer {
        from {
          destination-mac-address 00.05.85.00.00.DF;
        }
      }
    }
  }
}
```

```

        then {
            accept;
        }
    }
}
}
}
vlands {
    guest-vlan {
        description "restrict guest-to-employee traffic and peer-to-peer applications on
        guest VLAN";
        filter {
            input ingress-vlan-limit-guest;
        }
    }
}
}

```

Configuring a Router Firewall Filter to Give Priority to Egress Traffic Destined for the Corporate Subnet

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

CLI Quick Configuration

To quickly configure a firewall filter for a routed port (Layer 3 uplink module) to filter **employee-vlan** traffic, giving highest forwarding-class priority to traffic destined for the corporate subnet, copy the following commands and paste them into the switch terminal window:

```

[edit]
set firewall family inet filter egress-router-corp-class term corp-expedite from destination-address 192.0.2.16/28
set firewall family inet filter egress-router-corp-class term corp-expedite then forwarding-class expedited-forwarding
set firewall family inet filter egress-router-corp-class term corp-expedite then loss-priority low
set firewall family inet filter egress-router-corp-class term not-to-corp then accept
set interfaces ge-0/1/0 description "filter at egress router to expedite destined for corporate network"
set ge-0/1/0 unit 0 family inet address 103.104.105.1
set interfaces ge-0/1/0 unit 0 family inet filter output egress-router-corp-class

```

Step-by-Step Procedure

To configure and apply a firewall filter to a routed port (Layer 3 uplink module) to give highest priority to **employee-vlan** traffic destined for the corporate subnet:

1. Define the firewall filter **egress-router-corp-class**:

```

[edit]
user@switch# set firewall family inet filter egress-router-corp-class

```

2. Define the term **corp-expedite**:

```

[edit firewall]
user@switch# set family inet filter egress-router-corp-class term corp-expedite from destination-address 192.0.2.16/28
user@switch# set family inet filter egress-router-corp-class term corp-expedite then forwarding-class expedited-forwarding
user@switch# set family inet filter egress-router-corp-class term corp-expedite then loss-priority low

```

3. Define the term **not-to-corp**:

```
[edit firewall]
user@switch# set family inet filter egress-router-corp-class term not-to-corp then
accept
```

4. Apply the firewall filter **egress-router-corp-class** as an output filter for the port on the switch's uplink module, which provides a Layer 3 connection to a router:

```
[edit interfaces]
user@switch# set ge-0/1/0 description "filter at egress router to expedite employee
traffic destined for corporate network"
user@switch# set ge-0/1/0 unit 0 family inet address 103.104.105.1
user@switch# set ge-0/1/0 unit 0 family inet filter output egress-router-corp-class
```

Results Display the results of the configuration:

```
user@switch# show
firewall {
  family inet {
    filter egress-router-corp-class {
      term corp-expedite {
        from {
          destination-address 192.0.2.16/28;
        }
        then {
          forwarding-class expedited-forwarding;
          loss-priority low;
        }
      }
      term not-to-corp {
        then {
          accept;
        }
      }
    }
  }
}
interfaces {
  ge-0/1/0 {
    unit 0 {
      description "filter at egress router interface to expedite employee traffic destined
for corporate network";
      family inet {
        source-address 103.104.105.1
        filter {
          output egress-router-corp-class;
        }
      }
    }
  }
}
```

Verification

To confirm that the firewall filters are working properly, perform the following tasks:

- Verifying that Firewall Filters and Policers are Operational on page 56
- Verifying that Schedulers and Scheduler-Maps are Operational on page 56

Verifying that Firewall Filters and Policers are Operational

Purpose Verify the operational state of the firewall filters and policers that are configured on the switch.

Action Use the operational mode command:

```
user@switch> show firewall
Filter: ingress-port-voip-class-limit-tcp-icmp
Counters:
Name                               Packets
icmp-counter                        0
tcp-counter                         0
Policers:
Name                               Packets
icmp-connection-policer            0
tcp-connection-policer              0

Filter: ingress-vlan-rogue-block

Filter: egress-vlan-watch-employee
Counters:
Name                               Packets
employee-web-counter                0
```

Meaning The **show firewall** command displays the names of the firewall filters, policers, and counters that are configured on the switch. The output fields show byte and packet counts for all configured counters and the packet count for all policers.

Verifying that Schedulers and Scheduler-Maps are Operational

Purpose Verify that schedulers and scheduler-maps are operational on the switch.

Action Use the operational mode command:

```
user@switch> show class-of-service scheduler-map

Scheduler map: default, Index: 2

Scheduler: default-be, Forwarding class: best-effort, Index: 20
  Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent,
  Priority: low
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low           non-TCP   1      default-drop-profile
    Low           TCP      1      default-drop-profile
    High          non-TCP   1      default-drop-profile
    High          TCP      1      default-drop-profile

Scheduler: default-nc, Forwarding class: network-control, Index: 22
```



```

Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,
Priority: low
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           non-TCP   1      default-drop-profile
  Low           TCP       1      default-drop-profile
  High          non-TCP   1      default-drop-profile
  High          TCP       1      default-drop-profile
ethernet-diffsrv-cos-map, Index: 21657

```

```

Scheduler: best-effort, Forwarding class: best-effort, Index: 61257
Transmit rate: remainder, Rate Limit: none, Buffer size: 75 percent,
Priority: low
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           non-TCP   1      <default-drop-profile>
  Low           TCP       1      <default-drop-profile>
  High          non-TCP   1      <default-drop-profile>
  High          TCP       1      <default-drop-profile>

```

```

Scheduler: voice-high, Forwarding class: expedited-forwarding, Index: 3123
Transmit rate: remainder, Rate Limit: none, Buffer size: 15 percent,
Priority: high
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           non-TCP   1      <default-drop-profile>
  Low           TCP       1      <default-drop-profile>
  High          non-TCP   1      <default-drop-profile>
  High          TCP       1      <default-drop-profile>

```

```

Scheduler: net-control, Forwarding class: network-control, Index: 2451
Transmit rate: remainder, Rate Limit: none, Buffer size: 10 percent,
Priority: high
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           non-TCP   1      <default-drop-profile>
  Low           TCP       1      <default-drop-profile>
  High          non-TCP   1      <default-drop-profile>
  High          TCP       1      <default-drop-profile>

```

Meaning Displays statistics about the configured schedulers and schedulers-maps.

- Related Documentation**
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches
 - Example: Configuring CoS on EX Series Switches
 - Configuring Firewall Filters (CLI Procedure) on page 67
 - Configuring Firewall Filters (J-Web Procedure) on page 74
 - Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78
 - Firewall Filter Match Conditions and Actions for EX Series Switches on page 11
 - [edit firewall] Configuration Statement Hierarchy on page 95

Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches

Administrators can configure filter-based forwarding on an EX Series switch by using a firewall filter to forward matched traffic to a specific virtual routing instance.

This example describes how to set up filter-based forwarding:

- Requirements on page 58
- Overview and Topology on page 58
- Configuration on page 58
- Verification on page 60

Requirements

This example uses the following software and hardware components:

- One EX Series switch
- Junos OS Release 9.4 or later for EX Series switches

Overview and Topology

In this example, traffic from one application server that is destined for a different application server is matched by a firewall filter based on the IP address. Any matching packets are routed to a particular virtual routing instance that first sends all traffic to a security device, then forwards it to the designated destination address.

Configuration

To configure filter-based forwarding:

CLI Quick Configuration To quickly create and configure filter-based forwarding, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.1/24
set interfaces ge-0/0/3 unit 0 family inet address 10.1.3.1/24
set firewall family inet filter fil term t1 from source-address 1.1.1.1/32
set firewall family inet filter fil term t1 from protocol tcp
set interfaces ge-0/0/0 unit 0 family inet filter input fil
set routing-instances vrf01 instance-type virtual-router
set routing-instances vrf01 interface ge-0/0/1.0
set routing-instances vrf01 interface ge-0/0/3.0
set routing-instances vrf01 routing-options static route 12.34.56.0/24 next-hop 10.1.3.254
set firewall family inet filter fil term t1 then routing-instance vrf01
```

Step-by-Step Procedure

To configure filter-based forwarding:

1. Create interfaces to the application servers:

```
[edit]
user@switch# set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.1/24
user@switch# set interfaces ge-0/0/3 unit 0 family inet address 10.1.3.1/24
```

2. Create a firewall filter that matches the correct source address:

```
[edit]
user@switch# set firewall family inet filter fil term t1 from source-address 1.1.1.1/32
user@switch# set firewall family inet filter fil term t1 from protocol tcp
```

3. Associate the filter with the source application server's interface:

```
[edit]
user@switch# set interfaces ge-0/0/0 unit 0 family inet filter input fil
```

4. Create a virtual router:

```
[edit]
user@switch# set routing-instances vrf01 instance-type virtual-router
```

5. Associate the interfaces with the virtual router:

```
[edit]
user@switch# set routing-instances vrf01 interface ge-0/0/1.0
user@switch# set routing-instances vrf01 interface ge-0/0/3.0
```

6. Configure the routing information for the virtual routing instance:

```
[edit]
user@switch# set routing-instances vrf01 routing-options static route 12.34.56.0/24
next-hop 10.1.3.254
```

7. Set the filter to forward packets to the virtual router you created:

```
[edit]
user@switch# set firewall family inet filter fil term t1 then routing-instance vrf01
```

Results Check the results of the configuration:

```
user@switch> show configuration
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        filter {
          input fil;
        }
        address 10.1.0.1/24;
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.3.1/24;
      }
    }
  }
}
```

```
    }  
  }  
}  
firewall {  
  family inet {  
    filter fil {  
      term t1 {  
        from {  
          source-address {  
            1.1.1.1/32;  
          }  
          protocol tcp;  
        }  
        then {  
          routing-instance vrf01;  
        }  
      }  
    }  
  }  
}  
}  
routing-instances {  
  vrf01 {  
    instance-type virtual-router;  
    interface ge-0/0/1.0;  
    interface ge-0/0/3.0;  
    routing-options {  
      static {  
        route 12.34.56.0/24 next-hop 10.1.3.254;  
      }  
    }  
  }  
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That Filter-Based Forwarding Was Configured on page 60

Verifying That Filter-Based Forwarding Was Configured

Purpose Verify that filter-based forwarding was properly enabled on the switch.

Action 1. Use the **show interfaces filters** command:

```
user@switch> show interfaces filters ge-0/0/0.0
```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/0/0.0	up	down	inet	fil	

2. Use the **show route forwarding-table** command:

```
user@switch> show route forwarding-table
```

```
Routing table: default.inet  
Internet:  
Destination          Type RtRef Next hop          Type Index NhRef Netif
```

```

default      user      1 0:12:f2:21:cf:0    ucst    331    4 me0.0
default      perm      0                    rjct     36     3
0.0.0.0/32   perm      0                    dscd     34     1
10.1.0.0/24  ifdn      0                    rslv     613    1
ge-0/0/0.0
10.1.0.0/32  iddn      0 10.1.0.0           recv     611    1
ge-0/0/0.0
10.1.0.1/32  user      0                    rjct     36     3
10.1.0.1/32  intf      0 10.1.0.1           locl     612    2
10.1.0.1/32  iddn      0 10.1.0.1           locl     612    2
10.1.0.255/32 iddn      0 10.1.0.255         bcst     610    1
ge-0/0/0.0
10.1.1.0/26  ifdn      0                    rslv     583    1 v1an.0
10.1.1.0/32  iddn      0 10.1.1.0           recv     581    1 v1an.0
10.1.1.1/32  user      0                    rjct     36     3
10.1.1.1/32  intf      0 10.1.1.1           locl     582    2
10.1.1.1/32  iddn      0 10.1.1.1           locl     582    2
10.1.1.63/32 iddn      0 10.1.1.63          bcst     580    1 v1an.0
255.255.255.255/32 perm      0                    bcst     32     1

```

Routing table: vrf01.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	559	2	
0.0.0.0/32	perm	0		dscd	545	1	
10.1.3.0/24	ifdn	0		rslv	617	1	
ge-0/0/3.0							
10.1.3.0/32	iddn	0	10.1.3.0	recv	615	1	
ge-0/0/3.0							
10.1.3.1/32	user	0		rjct	559	2	
10.1.3.1/32	intf	0	10.1.3.1	locl	616	2	
10.1.3.1/32	iddn	0	10.1.3.1	locl	616	2	
10.1.3.255/32	iddn	0	10.1.3.255	bcst	614	1	
ge-0/0/3.0							
224.0.0.0/4	perm	0		mdsc	546	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	529	1	
255.255.255.255/32	perm	0		bcst	543	1	

Routing table: default.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	60	1	

Routing table: vrf01.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	600	1	

Meaning The output indicates that the filter was created on the interface and that the virtual routing instance is forwarding matching traffic to the correct IP address.

Related Documentation

- Configuring Firewall Filters (CLI Procedure) on page 67
- Configuring Static Routing (CLI Procedure)
- Configuring Static Routing (J-Web Procedure)
- Understanding Filter-Based Forwarding for EX Series Switches on page 38

Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch

You can configure a firewall filter on a management interface on an EX Series switch to filter ingress or egress traffic on the management interface on the switch. You can use utilities such as SSH or Telnet to connect to the management interface over the network and then use management protocols such as SNMP to gather statistical data from the switch.

This example discusses how to configure a firewall filter on a management interface to filter SSH packets egressing from an EX Series switch:

- Requirements on page 62
- Overview and Topology on page 62
- Configuration on page 63
- Verification on page 64

Requirements

This example uses the following hardware and software components:


- One EX Series switch and one management PC
- Junos OS Release 10.4 or later for EX Series switches

Overview and Topology

In this example, a management PC establishes an SSH connection with the management interface on a switch to remotely manage the switch. The IP address configured for the management interface is 10.204.33.103/20. A firewall filter is configured on the management interface to count the number of packets egressing from a source SSH port on the management interface. When the management PC establishes the SSH session with the management interface, the management interface returns SSH packets to the management PC to confirm that the session is established. These SSH packets are filtered based on the match condition specified in the firewall filter before they are forwarded to the management PC. As these packets are generated from the source SSH port on the management interface, they fulfill the match condition specified for the management interface. The number of matched SSH packets provides a count of the number of packets that have traversed the management interface. A system administrator can use this information to monitor the management traffic and take any action if required.

Figure 5 on page 63 shows the topology for this example in which a management PC establishes an SSH connection with the switch.

Figure 5: SSH Connection From a Management PC to an EX Series Switch



Configuration

To configure a firewall filter on a management interface, perform these tasks:

CLI Quick Configuration To quickly create and configure a firewall filter on the management interface to filter SSH packets egressing from the management interface, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family inet filter mgmt_fil1 term t1 from source-port ssh
set firewall family inet filter mgmt_fil1 term t1 then count c1
set firewall family inet filter mgmt_fil1 term t2 then accept
set interfaces me0 unit 0 family inet filter output mgmt_fil1
```

Step-by-Step Procedure To configure a firewall filter on the management interface to filter SSH packets:

1. Configure the firewall filter that matches SSH packets from the source port:

```
[edit]
user@switch# set firewall family inet filter mgmt_fil1 term t1 from source-port ssh

user@switch# set firewall family inet filter mgmt_fil1 term t1 then count c1

user@switch# set firewall family inet filter mgmt_fil1 term t2 then accept
```

These statements set a counter `c1` to count the number of SSH packets that egress from the source SSH interface on the management interface.

2. Set the firewall filter for the management interface:

```
[edit]
user@switch# set interfaces me0 unit 0 family inet filter output mgmt_fil1
```



NOTE: You can also set the firewall filter for a VME interface.

Results Check the results of the configuration:

```
[edit]
user@switch# show
interfaces {
  me0 {
    unit 0 {
      family inet {
        filter {
          output mgmt_fil1;
        }
      }
    }
  }
}
```

```
        address 10.93.54.6/24;
    }
}
}

firewall {
  family inet {
    filter mgmt_fil1 {
      term t1 {
        from {
          source-port ssh;
        }
        then count c1;
      }
    }
    term t2 {
      then accept;
    }
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Firewall Filter Is Configured on a Management Interface on page 64](#)

Verifying That the Firewall Filter Is Configured on a Management Interface

Purpose Verify that the firewall filter has been enabled on the management interface on the switch.

- Action** 1. Verify that the firewall filter is applied to the management interface:

```
[edit]
user@switch#show interfaces me0
unit 0 {
  family inet {
    filter {
      output mgmt_fil1;
    }
    address 10.204.33.103/20;
  }
}
```

2. Check the counter value that is associated with the firewall filter:

```
user@switch> show firewall

Filter: mgmt_fil1
Counters:
Name                               Bytes      Packets
c1                                  0           0
```

3. From the management PC, establish a secure shell session with the switch:

```
[user@management-pc ~]$ ssh user@10.204.33.103
```

4. Check counter values after SSH packets are generated from the switch in response to the secure shell session request by the management PC:

```
user@switch> show firewall

Filter: mgmt_fil1
Counters:
Name                               Bytes      Packets
c1                                  3533       23
```

Meaning The output indicates that the firewall filter has been applied to the management interface and the counter value indicates that 23 SSH packets were generated from the switch.

- Related Documentation**
- Configuring Firewall Filters (CLI Procedure) on page 67
 - Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39

CHAPTER 3

Configuring Firewall Filters

- Configuring Firewall Filters (CLI Procedure) on page 67
- Configuring Firewall Filters (J-Web Procedure) on page 74
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78
- Assigning Multifield Classifiers in Firewall Filters to Specify Packet-Forwarding Behavior (CLI Procedure) on page 81
- Configuring Routing Policies (J-Web Procedure) on page 82

Configuring Firewall Filters (CLI Procedure)

You configure firewall filters on EX Series switches to control traffic that enters ports on the switch or enters and exits VLANs on the network and Layer 3 (routed) interfaces. To configure a firewall filter you must configure the filter and then apply it to a port, VLAN, or Layer 3 interface.

This topic describes:

- Configuring a Firewall Filter on page 67
- Applying a Firewall Filter to a Port on a Switch on page 70
- Applying a Firewall Filter to a Management Interface on a Switch on page 71
- Applying a Firewall Filter to a VLAN on a Network on page 72
- Applying a Firewall Filter to a Layer 3 (Routed) Interface on page 73

Configuring a Firewall Filter

Before you can apply a firewall filter to a port, VLAN, or Layer 3 interface, you must configure a firewall filter with the required details, such as type of family for the firewall filter, firewall filter name, and match conditions. A match condition in the firewall filter configuration can contain multiple terms that define the criteria for the match condition. For each term, you must specify an action to be performed if a packet matches the conditions in the term. For information on different match conditions and actions, see “Firewall Filter Match Conditions and Actions for EX Series Switches” on page 11.

To configure a firewall filter:

1. Configure the family address type for the firewall filter:

- For a firewall filter that is applied to a port or VLAN, specify the family address type **ethernet-switching** to filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets, for example:

```
[edit firewall]
user@switch# set family ethernet-switching
```

- For a firewall filter that is applied to a Layer 3 (routed) interface:
 - To filter IPv4 packets, specify the family address type **inet**, for example:

```
[edit firewall]
user@switch# set family inet
```

- To filter IPv6 packets, specify the family address type **inet6**, for example:

```
[edit firewall]
user@switch# set family inet6
```



NOTE: You can configure firewall filters for both IPv4 and IPv6 traffic on the same Layer 3 interface.

2. Specify the filter name:

```
[edit firewall family ethernet-switching]
user@switch# set filter ingress-port-filter
```

The filter name can contain letters, numbers, and hyphens (-) and can have a maximum of 64 characters. Each filter name must be unique.

3. If you want to apply a firewall filter to multiple interfaces and name individual firewall counters specific to each interface, configure the **interface-specific** option:

```
[edit firewall family ethernet-switching filter ingress-port-filter]
user@switch# set interface-specific
```

4. Specify a term name:

```
[edit firewall family ethernet-switching filter ingress-port-filter]
user@switch# set term term-one
```

The term name can contain letters, numbers, and hyphens (-) and can have a maximum of 64 characters.

A firewall filter can contain one or more terms. Each term name must be unique within a filter.

**NOTE:**

The maximum number of terms allowed per firewall filter for EX Series switches is:

- 512 for EX2200 switches
- 7168 for EX3200 and EX4200 switches
- 1536 for EX4500 switches
- 32768 for EX8200 switches

If you attempt to configure a firewall filter that exceeds these limits, the switch returns an error message when you commit the configuration.

5. In each firewall filter term, specify the match conditions to use to match components of a packet.

To specify match conditions to match on packets that contain a specific source-address and source-port—for example:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one]
user@switch# set from source-address 192.0.2.14
user@switch# set from source-port 80
```

You can specify one or more match conditions in a single **from** statement. For a match to occur, the packet must match all the conditions in the term.

The **from** statement is optional, but if included in a term, the **from** statement cannot be empty. If you omit the **from** statement, all packets are considered to match.

6. In each firewall filter term, specify the actions to take if the packet matches all the conditions in that term.

You can specify an action and/or action modifiers:

- To specify a filter action, for example, to discard packets that match the conditions of the filter term:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one]
user@switch# set then discard
```

You can specify no more than one action (**accept**, **discard**, or **routing-instance**) per filter term.

- To specify action modifiers, for example, to count and classify packets in a forwarding class:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one]
user@switch# set then count counter-one
user@switch# set then forwarding-class expedited-forwarding
```

You can specify any of the following action modifiers in a **then** statement:

- **analyzer** *analyzer-name*—Mirror port traffic to a specified destination port or VLAN that is connected to a protocol analyzer application. An **analyzer** must be configured under the **ethernet-switching** family address type. See *Configuring Port Mirroring to Analyze Traffic (CLI Procedure)*.
- **count** *counter-name*—Count the number of packets that pass this filter term.



NOTE: We recommend that you configure a counter for each term in a firewall filter, so that you can monitor the number of packets that match the conditions specified in each filter term.

- **forwarding-class** *class*—Classify packets in a forwarding class.
- **loss-priority** *priority*—Set the priority of dropping a packet.
- **policer** *policer-name*—Apply rate-limiting to the traffic.

If you omit the **then** statement or do not specify an action, packets that match all the conditions in the **from** statement are accepted. However, you must always explicitly configure an action and/or action modifier in the **then** statement. You can include no more than one action statement, but you can use any combination of action modifiers. For an action or action modifier to take effect, all conditions in the **from** statement must match.



NOTE: Implicit discard is also applicable to a firewall filter applied to the loopback interface, lo0.

Applying a Firewall Filter to a Port on a Switch

You can apply a firewall filter to a port on a switch to filter ingress or egress traffic on the switch. When you configure the firewall filter, you can specify any match condition, action, and action modifiers specified in “Firewall Filter Match Conditions and Actions for EX Series Switches” on page 11. The action specified in the match condition indicates the action for the matched packets in the ingress or egress traffic.

To apply a firewall filter to a port to filter ingress or egress traffic:



NOTE: For applying a firewall filter to a management interface, see “Applying a Firewall Filter to a Management Interface on a Switch” on page 71

1. Specify the interface name and provide a meaningful description of the firewall filter and the interface to which the filter is applied:

```
[edit interfaces]
user@switch# set ge-0/0/1 description "filter to limit tcp traffic filter at trunk port for
employee-vlan and voice-vlan applied on the interface"
```



NOTE: Providing the description is optional.

- Specify the unit number and family address type for the interface:

```
[edit interfaces]
user@switch# set ge-0/0/1 unit 0 family ethernet-switching
```

For firewall filters that are applied to ports, the family address type must be **ethernet-switching**.

- To apply a firewall filter to filter packets that are entering a port:

```
[edit interfaces]
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input ingress-port-filter
```

To apply a firewall filter to filter packets that are exiting a port:

```
[edit interfaces]
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter output egress-port-filter
```



NOTE: You can apply no more than one firewall filter per port, per direction.

Applying a Firewall Filter to a Management Interface on a Switch

You can configure and apply a firewall filter to a management interface to control traffic that is entering or exiting the interface on a switch. You can use utilities such as SSH or Telnet to connect to the management interface over the network and then use management protocols such as SNMP to gather statistical data from the switch. Similar to configuring a firewall filter on other types of interfaces, you can configure a firewall filter on a management interface using any match condition, action, and action modifier specified in “Firewall Filter Match Conditions and Actions for EX Series Switches” on page 11 except for the following action modifiers:

- **loss-priority**
- **forwarding-class**

You can apply a firewall filter to the management Ethernet interface on any EX Series switch. You can also apply a firewall filter to the virtual management Ethernet (VME) interface on the EX4200 switch. For more information on the management Ethernet interface and the VME interface, see EX Series Switches Interfaces Overview.

To apply a firewall filter on the management interface to filter ingress or egress traffic:

- Specify the interface name and provide a meaningful description of the firewall filter and the interface to which the filter is applied:

```
[edit interfaces]
user@switch# set me0 description "filter to limit tcp traffic filter at management interface"
```



NOTE: Providing the description is optional.

2. Specify the unit number and family address type for the management interface:

```
[edit interfaces]
user@switch# set me0 unit 0 family inet
```



NOTE: For firewall filters that are applied to management interfaces, the family address type can be either `inet` or `inet6`.

3. To apply a firewall filter to filter packets that are entering a management interface:

```
[edit interfaces]
user@switch# set me0 unit 0 family inet filter input ingress-port-filter
```

To apply a firewall filter to filter packets that are exiting a management interface:

```
[edit interfaces]
user@switch# set me0 unit 0 family inet filter output egress-port-filter
```



NOTE: You can apply no more than one firewall filter per management interface, per direction.

Applying a Firewall Filter to a VLAN on a Network

You can apply a firewall filter to a VLAN on a network to filter ingress or egress traffic on the network. To apply a firewall filter to a VLAN, specify the VLAN name and ID, and then apply the firewall filter to the VLAN. When you configure the firewall filter, you can specify any match condition, action, and action modifiers specified in “Firewall Filter Match Conditions and Actions for EX Series Switches” on page 11. The action specified in the match condition indicates the action for the matched packets in the ingress or egress traffic.

To apply a firewall filter to a VLAN:

1. Specify the VLAN name and VLAN ID and provide a meaningful description of the firewall filter and the VLAN to which the filter is applied:

```
[edit vlans]
user@switch# set employee-vlan vlan-id 20 vlan-description "filter to rate limit traffic applied on employee-vlan"
```



NOTE: Providing the description is optional.

2. Apply firewall filters to filter packets that are entering or exiting the VLAN:
 - To apply a firewall filter to filter packets that are entering the VLAN:


```
[edit vlans]
user@switch# set employee-vlan vlan-id 20 filter input ingress-vlan-filter
```

- To apply a firewall filter to filter packets that are exiting the VLAN:

```
[edit vlans]
user@switch# set employee-vlan vlan-id 20 filter output egress-vlan-filter
```



NOTE: You can apply no more than one firewall filter per VLAN, per direction.

Applying a Firewall Filter to a Layer 3 (Routed) Interface

You can apply a firewall filter to a Layer 3 (routed) interface to filter ingress or egress traffic on the switch. When you configure the firewall filter, you can specify any match condition, action, and action modifiers specified in “Firewall Filter Match Conditions and Actions for EX Series Switches” on page 11. The action specified in the match condition indicates the action for the matched packets in the ingress or egress traffic.

To apply a firewall filter to a Layer 3 interface on a switch:

1. Specify the interface name and provide a meaningful description of the firewall filter and the interface to which the filter is applied:

```
[edit interfaces]
user@switch# set ge-0/1/0 description "filter to count and monitor employee-vlan
traffic applied on layer 3 interface"
```



NOTE: Providing the description is optional.

2. Specify the unit number, family address type, and address for the interface:

```
[edit interfaces]
user@switch# set ge-0/1/0 unit 0 family inet address 10.10.10.1/24
```

For firewall filters applied to Layer 3 interfaces, the family address type must be **inet** (for IPv4 traffic) or **inet6** (for IPv6 traffic).

3. You can apply firewall filters to filter packets that are entering or exiting a Layer 3 (routed) interface:

- To apply a firewall filter to filter packets that are entering a Layer 3 interface:

```
[edit interfaces]
user@switch# set ge-0/1/0 unit 0 family inet address 10.10.10.1/24 filter input
ingress-router-filter
```

- To apply a firewall filter to filter packets that are exiting a Layer 3 interface:

```
[edit interfaces]
user@switch# set ge-0/1/0 unit 0 family inet address 10.10.10.1/24 filter output
egress-router-filter
```



NOTE: You can apply no more than one firewall filter per Layer 3 interface, per direction.

**Related
Documentation**

- Configuring Firewall Filters (J-Web Procedure) on page 74
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 58
- Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch on page 62
- Verifying That Firewall Filters Are Operational on page 89
- Monitoring Firewall Filter Traffic on page 90
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78

Configuring Firewall Filters (J-Web Procedure)

You configure firewall filters on EX Series switches to control traffic that enters ports on the switch or enters and exits VLANs on the network and Layer 3 (routed) interfaces. To configure a firewall filter you must configure the filter and then apply it to a port, VLAN, or Layer 3 interface.

To configure firewall filter settings using the J-Web interface:

1. Select **Configure > Security > Filters**.

The Firewall Filter Configuration page displays a list of all configured port/VLAN or router filters and the ports or VLANs associated with a particular filter.



NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one:

- **Add**—Select this option to create a new filter. Enter information as specified in Table 9 on page 75.
- **Edit**—Select this option to edit an existing filter. Enter information as specified in Table 9 on page 75.
- **Delete**—Select this option to delete a filter.

- Term Up—Select this option to move a term up in the filter term list.
- Term Down—Select this option to move a term down in the filter term list.

Table 9: Create a New Filter

Field	Function	Your Action
Filter tab		
Filter type	Specifies the filter type: port/VLAN firewall filter or router firewall filter.	Select the filter type.
Filter name	Specifies the name for the filter.	Enter a name.
Select terms to be part of the filter	Specifies the terms to be associated with the filter. Add new terms or edit existing terms.	Click Add to add new terms. Enter information as specified in Table 10 on page 75 and Table 11 on page 76.
Association tab		
Port Associations	Specifies the ports with which the filter is associated. NOTE: For a port/VLAN filter type, only Ingress direction is supported for port association.	<ol style="list-style-type: none"> 1. Click Add. 2. Select the direction: Ingress or Egress. 3. Select the ports. 4. Click OK.
VLAN Associations	Specifies the VLANs with which the filter is associated. NOTE: Because router firewall filters can be associated with ports only, this section is not displayed for a router firewall filter.	<ol style="list-style-type: none"> 1. Click Add. 2. Select the direction: Ingress or Egress. 3. Select the VLANs. 4. Click OK.

Table 10: Create a New Term

Field	Function	Your Action
Term Name	Specifies the name of the term.	Enter a name.
Protocols	Specifies the protocols to be associated with the term.	<ol style="list-style-type: none"> 1. Click Add. 2. Select the protocols. 3. Click OK.

Table 10: Create a New Term (*continued*)

Field	Function	Your Action
Source	Specifies the source IP address, MAC address, and available ports. NOTE: MAC address is specified only for port/VLAN filters.	To specify the IP address, click Add > IP and enter the IP address. To specify the MAC address, click Add > MAC and enter the MAC address. To specify the ports (interfaces), click Add > Ports and enter the port number. To delete the IP address, MAC address, or port details, select it and click Remove .
Destination	Specifies the destination IP address, MAC address, and available ports. NOTE: MAC address is specified only for port/VLAN filters.	To specify the IP address, click Add > IP and enter the IP address. To specify the MAC address, click Add > MAC and enter the MAC address. To specify the ports (interfaces), click Add > Ports and enter the port number. To delete the IP address, MAC address, or port details, select it and click Remove .
Action	Specifies the packet action for the term.	Select one: <ul style="list-style-type: none"> • Accept • Discard
More	Specifies advanced configuration options for the filter.	Select the match conditions as specified in Table 11 on page 76. Select the packet action for the term as specified in Table 11 on page 76.

Table 11: Advanced Options for Terms

Table	Function	Your Action
ICMP Type	Specifies the ICMP packet type field. Typically, you specify this match condition in conjunction with the protocol match condition to determine which protocol is being used on the port.	Select the option from the list.
ICMP Code	Specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify icmp-type along with icmp-code . The keywords are grouped by the ICMP type with which they are associated.	Select a value from the list.
DSCP	Specifies the Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP.	Select the DSCP number from the list.

Table 11: Advanced Options for Terms (*continued*)

Table	Function	Your Action
Precedence	Specifies IP precedence. NOTE: IP precedence and DSCP number cannot be specified together for the same term.	Select the option from the list.
IP Options	Specifies the presence of the options field in the IP header.	Select the option from the list.
Interface	Specifies the interface on which the packet is received.	Select the interface from the list.
Ether type	Specifies the Ethernet type field of a packet. NOTE: This option is not applicable for a routing filter.	Select a value from the list.
Dot1q user priority	Specifies the user-priority field of the tagged Ethernet packet. User-priority values can be 0–7. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed) <ul style="list-style-type: none"> background (1)—Background best-effort (0)—Best effort controlled-load (4)—Controlled load excellent-load (3)—Excellent load network-control (7)—Network control reserved traffic standard (2)—Standard or Spare video (5)—Video voice (6)—Voice NOTE: This option is not applicable for a routing filter.	Select a value from the list.
VLAN	Specifies the VLAN to be associated with the packet. NOTE: This option is not applicable for a routing filter.	Select the VLAN from the list.
TCP Flags	Specifies one or more TCP flags. NOTE: TCP flags are supported on ingress ports, VLANs, and router interfaces.	Select the option TCP Initial or enter a combination of TCP flags.
Fragmentation Flags	Specifies the IP fragmentation flags. NOTE: Fragmentation flags are supported on ingress ports, VLANs, and router interfaces.	Select either the option is-fragment or enter a combination of fragment action flags.
Dot1q tag	Specifies the value for tag field in the Ethernet header. Values can be from 1 through 4095. NOTE: This option is not applicable for a routing filter.	Enter the value.

Action

Table 11: Advanced Options for Terms (*continued*)

Table	Function	Your Action
Counter name	Specifies the count of the number of packets that pass this filter, term, or policer.	Enter a value.
Forwarding class	Classifies the packet into one of the following forwarding classes: <ul style="list-style-type: none"> assured-forwarding best-effort expedited-forwarding network-control user-defined 	Select the option from the list.
Loss priority	Specifies the packet loss priority. NOTE: Forwarding class and loss priority should be specified together for the same term.	Enter the value.
Analyzer	Specifies whether to perform port-mirroring on packets. Port-mirroring copies all packets entering one switch port to a network monitoring connection on another switch port.	Select the analyzer (port mirroring configuration) from the list.

Related Documentation

- [Configuring Firewall Filters \(CLI Procedure\) on page 67](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39](#)
- [Verifying That Firewall Filters Are Operational on page 89](#)
- [Firewall Filters for EX Series Switches Overview on page 3](#)
- [Firewall Filter Match Conditions and Actions for EX Series Switches on page 11](#)

Configuring Policers to Control Traffic Rates (CLI Procedure)

You can configure policers to rate limit traffic on EX Series switches. After you configure a policer, you can include it in an ingress firewall filter configuration.

When you configure a firewall filter, you can specify a policer action for any term or terms within the filter. All traffic that matches a term that contains a policer action goes through the policer that the term references. Each policer that you configure includes an implicit counter. To get term-specific packet counts, you must configure a new policer for each filter term that requires policing.

The following policer limits apply on the switch:

- A maximum of 512 policers can be configured for port firewall filters.
- A maximum of 512 policers can be configured for VLAN and Layer 3 firewall filters.

If the policer configuration exceeds these limits, the switch returns the following message after the commit operation:

Cannot assign policers: Max policer limit reached

1. Configuring Policers on page 79
2. Specifying Policers in a Firewall Filter Configuration on page 80
3. Applying a Firewall Filter That Is Configured with a Policers on page 80

Configuring Policers

To configure a policer:

1. Specify the name of the policer:

```
[edit firewall]
user@switch# set policer policer-one
```

The policer name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long.

2. Configure rate limiting for the policer:

- a. Specify the bandwidth limit in bits per second (bps) to control the traffic rate on an interface:

```
[edit firewall policer policer-one]
user@switch# set if-exceeding bandwidth-limit 300k
```

The range for the bandwidth limit is 1k through 102.3g bps.

- b. Specify the maximum allowed burst size to control the amount of traffic bursting:

```
[edit firewall policer policer-one]
user@switch# set if-exceeding burst-size-limit 500k
```

To determine the value for the burst-size limit, multiply the bandwidth of the interface on which the filter is applied by the amount of time to allow a burst of traffic at that bandwidth to occur:

burst size = bandwidth * allowable time for burst traffic

The range for the burst-size limit is 1 through 2,147,450,880 bytes.

3. Specify the policer action **discard** to discard packets that exceed the rate limits:

```
[edit firewall policer]
user@switch# set policer-one then discard
```

Discard is the only supported policer action.

Specifying Policers in a Firewall Filter Configuration

To reference a policer for a single firewall, configure a filter term that includes the policer action:

```
[edit firewall family ethernet-switching]
user@switch# set filter limit-hosts term term-one from source-address 192.0.2.16/28
userswitch# set filter limit-hosts term term-one then policer policer-one
```

Applying a Firewall Filter That Is Configured with a Policer

A firewall filter that is configured with one or more policer actions, like any other filter, must be applied to a port, VLAN, or Layer 3 interface. For information about applying firewall filters, see the sections on applying firewall filters in “Configuring Firewall Filters (CLI Procedure)” on page 67.



NOTE: You can include policer actions on ingress firewall filters only.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Configuring Firewall Filters (CLI Procedure) on page 67
- Configuring Firewall Filters (J-Web Procedure) on page 74
- Verifying That Policers Are Operational on page 90
- Understanding the Use of Policers in Firewall Filters on page 37

Assigning Multifield Classifiers in Firewall Filters to Specify Packet-Forwarding Behavior (CLI Procedure)

You can configure firewall filters with multifield classifiers to classify packets transiting a port, VLAN, or Layer 3 interface on an EX Series switch.

You specify multifield classifiers in a firewall filter configuration to set the forwarding class and packet loss priority (PLP) for incoming or outgoing packets. By default, the data traffic that is not classified is assigned to the **best-effort** class associated with queue 0.

You can specify any of the following default forwarding classes:

Forwarding class	Queue
best-effort	0
assured-forwarding	1
expedited-forwarding	5
network-control	7

To assign multifield classifiers in firewall filters:

1. Configure the family name and filter name for the filter at the **[edit firewall]** hierarchy level, for example:

```
[edit firewall]
user@switch# set family ethernet-switching
user@switch# set family ethernet-switching filter ingress-filter
```

2. Configure the terms of the filter, including the **forwarding-class** and **loss-priority** action modifiers as appropriate. When you specify a forwarding class you must also specify the packet loss priority. For example, each of the following terms examines different packet header fields and assigns an appropriate classifier and the packet loss priority:

- The term **voice-traffic** matches packets on the **voice-vlan** and assigns the forwarding class **expedited-forwarding** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term voice-traffic from vlan-id voice-vlan
user@switch# set term voice-traffic then forwarding-class expedited-forwarding
user@switch# set term voice-traffic then loss-priority low
```

- The term **data-traffic** matches packets on **employee-vlan** and assigns the forwarding class **assured-forwarding** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term data-traffic from vlan-id employee-vlan
user@switch# set term data-traffic then forwarding-class assured-forwarding
user@switch# set term data-traffic then loss-priority low
```

- Because loss of network-generated packets can jeopardize proper network operation, delay is preferable to discard of packets. The following term, **network-traffic**, assigns the forwarding class **network-control** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term network-traffic from precedence net-control
user@switch# set term network-traffic then forwarding-class network
user@switch# set term network-traffic then loss-priority low
```

- The last term **accept-traffic** matches any packets that did not match on any of the preceding terms and assigns the forwarding class **best-effort** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term accept-traffic from precedence net-control
user@switch# set term accept-traffic then forwarding-class best-effort
user@switch# set term accept-traffic then loss-priority low
```

3. Apply the filter **ingress-filter** to a port, VLAN or Layer 3 interface. For information about applying the filter, see “Configuring Firewall Filters (CLI Procedure)” on page 67.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Verifying That Firewall Filters Are Operational on page 89
- Monitoring Firewall Filter Traffic on page 90
- Defining CoS Classifiers (CLI Procedure)
- Defining CoS Classifiers (J-Web Procedure)
- Configuring Firewall Filters (CLI Procedure) on page 67
- Configuring Firewall Filters (J-Web Procedure) on page 74

Configuring Routing Policies (J-Web Procedure)

All routing protocols use the Junos OS routing table to store the routes that they learn and to determine which routes are advertised in the protocol packets. Routing policy allows you to control which routes the routing protocols store in and retrieve from the routing table on the routing device.

To configure routing policies for an EX Series switch using the J-Web interface:

1. Select **Configure > Routing > Policies**.



NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one:

- **Global Options**—Configures global options for policies. Enter information into the configuration page as described in Table 12 on page 83.
- **Add**—Configures a new policy. Select **New** and specify a policy name. To add terms, enter information into the configuration page as described in Table 13 on page 84. Select **Clone** to create a copy of an existing policy.
- **Edit**—Edits an existing policy. To modify an existing term, enter information into the configuration page as described in Table 13 on page 84.
- **Term Up**—Moves a term up in the list.
- **Term Down**—Moves a term down in the list.
- **Delete**—Deletes the selected policy.
- **Test Policy**—Tests the policy. Use this option to check whether the policy produces the results that you expect.

Table 12: Policies Global Configuration Parameters

Field	Function	Your Action
Prefix List	Specifies a list of IPv4 address prefixes for use in a routing policy statement.	<p>To add a prefix list:</p> <ol style="list-style-type: none"> 1. Click Add. 2. Enter a name for the prefix list. 3. To add an IP address, click Add. 4. Enter the IP address and the subnet mask and click OK. 5. Click OK. <p>To edit a prefix list, click Edit. Edit the settings and click OK.</p> <p>To delete a prefix list, select it and click Delete.</p>
BGP Community	Specifies a BGP community.	<p>To add a BGP community:</p> <ol style="list-style-type: none"> 1. Click Add. 2. Enter a name for the community. 3. To add a community, click Add. 4. Enter the community ID and click OK. 5. Click OK. <p>To edit a BGP community, click Edit. Edit the settings and click OK.</p> <p>To delete a BGP community, select it and click Delete.</p>

Table 12: Policies Global Configuration Parameters (*continued*)

Field	Function	Your Action
AS Path	Specifies an AS path. This is applicable to BGP only.	<p>To add an AS path:</p> <ol style="list-style-type: none"> 1. Click Add. 2. Enter the AS path name. 3. Enter the regular expression and click OK. 4. Click OK. <p>To edit an AS path, click Edit. Edit the settings and click OK.</p> <p>To delete an AS path, select it and click Delete.</p>

Table 13: Terms Configuration Parameters

Field	Function	Your Action
Term Name	Specifies a term name.	Type or select and edit the name.
Source tab		
Family	Specifies an address family protocol.	Select a value from the list.
Routing Instance	Specifies a routing instance.	Select a value from the list.
RIB	Specifies the name of a routing table.	Select a value from the list
Preference	Specifies the individual preference value for the route.	Type or select and edit the value.
Metric	Specifies a metric value. You can specify up to four metric values.	Type or select and edit the value.
Interface	Specifies a name or IP address of one or more routing device interfaces. Do not use this qualifier with protocols that are not interface-specific, such as internal BGP (IBGP).	<p>To add an interface, select Add > Interface. Select the interface from the list.</p> <p>To add an address, select Add > Address. Select the address from the list.</p> <p>To remove an interface, select it and click Remove.</p>
Prefix List	Specifies a named list of IP addresses. You can specify an exact match with incoming routes.	<p>Click Add. Select the prefix list from the list and click OK.</p> <p>To remove a prefix list, select it and click Remove.</p>
Protocol	Specifies the name of the protocol from which the route was learned or to which the route is being advertised.	<p>Click Add and select the protocol from the list.</p> <p>To remove a protocol, select it and click Remove.</p>

Table 13: Terms Configuration Parameters (*continued*)

Field	Function	Your Action
Policy	Specifies the name of a policy to evaluate as a subroutine.	Click Add . Select the policy from the list. To remove a policy, select it and click Remove .
More	Specifies advanced configuration options for policies.	Click More for advanced configuration.
OSPF Area ID	Specifies the area identifier.	Type the IP address.
BGP Origin	Specifies the origin of the AS path information.	Select a value from the list.
Local Preference	Specifies the BGP local preference.	Type a value.
Route	Specifies the type of route.	Select External . Select the OSPF type from the list.
AS Path	Specifies the name of an AS path regular expression.	Click Add . Select the AS path from the list.
Community	Specifies the name of one or more communities.	Click Add . Select the community from the list.
Destination tab		
Family	Specifies an address family protocol.	Select a value from the list.
Routing Instance	Specifies a routing instance.	Select a value from the list.
RIB	Specifies the name of a routing table.	Select a value from the list.
Preference	Specifies the individual preference value for the route.	Type a value.
Metric	Specifies a metric value.	Type a value.
Interface	Specifies a name or IP address of one or more routing device interfaces. Do not use this qualifier with protocols that are not interface-specific, such as internal BGP (IBGP).	To add an interface, select Add > Interface . Select the interface from the list. To add an address, select Add > Address . Select the address from the list. To delete an interface, select it and click Remove .
Protocol	Specifies the name of the protocol from which the route was learned or to which the route is being advertised.	Click Add and select the protocol from the list. To delete a protocol, select it and click Remove .
Action tab		

Table 13: Terms Configuration Parameters (*continued*)

Field	Function	Your Action
Action	Specifies the action to take if the conditions match.	Select a value from the list.
Default Action	Specifies that any action that is intrinsic to the protocol is overridden. This action is also nonterminating, so that various policy terms can be evaluated before the policy is terminated.	Select a value from the list.
Next	Specifies the default control action if a match occurs, and there are no further terms in the current routing policy.	Select a value from the list.
Priority	Specifies a priority for prefixes included in an OSPF import policy. Prefixes learned through OSPF are installed in the routing table based on the priority assigned to the prefixes.	Select a value from the list.
BGP Origin	Specifies the BGP origin attribute.	Select a value from the list.
AS Path Prepend	Affixes an AS number at the beginning of the AS path. The AS numbers are added after the local AS number has been added to the path. This action adds an AS number to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the affixed AS number is placed within a confederation sequence. Otherwise, the affixed AS number is placed with a nonconfederation sequence.	Enter a value.
AS Path Expand	Extracts the last AS number in the existing AS path and affixes that AS number to the beginning of the AS path n times, where n is a number from 1 through 32. The AS number is added before 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 affixed AS numbers are placed within a confederation sequence. Otherwise, the affixed AS numbers are placed within a nonconfederation sequence. This option is typically used in non-IBGP export policies.	Select the type and type a value.
Load Balance Per Packet	Specifies that all next-hop addresses in the forwarding table must be installed and have the forwarding table perform per-packet load balancing. This policy action allows you to optimize VPLS traffic flows across multiple paths.	Select the check box to enable the option.
Tag	Specifies the tag value. The tag action sets the 32-bit tag field in OSPF external link-state advertisement (LSA) packets.	Select the action and type a value.
Metric	Changes the metric (MED) value by the specified negative or positive offset. This action is useful only in an external BGP (EBGP) export policy.	Select the action and type a value.
Route	Specifies whether the route is external.	Select the External check box to enable the option, and select the OSPF type.
Preference	Specifies the preference value.	Select the preference action and type a value.

Table 13: Terms Configuration Parameters (*continued*)

Field	Function	Your Action
Local Preference	Specifies the BGP local preference attribute.	Select the action and type a value.
Class of Service	<p>Specifies and applies the class-of-service parameters to routes installed into the routing table.</p> <ul style="list-style-type: none"> Source class The value entered here maintains the packet counts for a route passing through your network, based on the source address. Destination class The value entered here maintains packet counts for a route passing through your network, based on the destination address in the packet. Forwarding class 	<p>Type the source class.</p> <p>Type the destination class.</p> <p>Type the forwarding class.</p>

- Related Documentation**
- Configuring BGP Sessions (J-Web Procedure)
 - Configuring an OSPF Network (J-Web Procedure)
 - Configuring a RIP Network (J-Web Procedure)
 - Configuring Static Routing (J-Web Procedure)
 - Layer 3 Protocols Supported on EX Series Switches

Verifying Firewall Filter Configuration

- Verifying That Firewall Filters Are Operational on page 89
- Verifying That Policers Are Operational on page 90
- Monitoring Firewall Filter Traffic on page 90

Verifying That Firewall Filters Are Operational

Purpose After you configure and apply firewall filters to ports, VLANs, or Layer 3 interfaces, you can perform the following task to verify that the firewall filters configured on EX Series switches are working properly.

Action Use the operational mode command to verify that the firewall filters on the switch are working properly:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                                     Bytes      Packets
counter-employee-web                     0           0
Filter: ingress-port-voip-class-limit-tcp-icmp
Counters:
Name                                     Bytes      Packets
icmp-counter                             0           0
Policers:
Name                                     Packets
icmp-connection-policer                  0
tcp-connection-policer                   0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

Meaning The **show firewall** command displays the names of all firewall filters, policers, and counters that are configured on the switch. For each counter that is specified in a filter configuration, the output field shows the byte count and packet count for the term in which the counter is specified. For each policer that is specified in a filter configuration, the output field shows the packet count for packets that exceed the specified rate limits.

- Related Documentation**
- Configuring Firewall Filters (CLI Procedure) on page 67
 - Configuring Firewall Filters (J-Web Procedure) on page 74
 - Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Monitoring Firewall Filter Traffic on page 90

Verifying That Policers Are Operational

Purpose After you configure policers and include them in firewall filter configurations, you can perform the following tasks to verify that the policers configured on EX Series switches are working properly.

Action Use the operational mode command to verify that the policers on the switch are working properly:

```
user@switch> show policer
Filter: egress-vlan-watch-employee
Filter: ingress-port-filter
Filter: ingress-port-voip-class-limit-tcp-icmp
Policers:
Name                                     Packets
icmp-connection-policer                  0
tcp-connection-policer                   0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

Meaning The **show policer** command displays the names of all firewall filters and policers that are configured on the switch. For each policer that is specified in a filter configuration, the output field shows the current packet count for all packets that exceed the specified rate limits.

- Related Documentation**
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78
 - Configuring Firewall Filters (CLI Procedure) on page 67
 - Configuring Firewall Filters (J-Web Procedure) on page 74
 - Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
 - Monitoring Firewall Filter Traffic on page 90

Monitoring Firewall Filter Traffic

You can monitor firewall filter traffic on EX Series switches.

- Monitoring Traffic for All Firewall Filters and Policers That Are Configured on the Switch on page 91
- Monitoring Traffic for a Specific Firewall Filter on page 91
- Monitoring Traffic for a Specific Policer on page 91

Monitoring Traffic for All Firewall Filters and Policers That Are Configured on the Switch

Purpose Perform the following task to monitor the number of packets and bytes that matched the firewall filters and monitor the number of packets that exceeded policer rate limits:

Action Use the operational mode command:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                                     Bytes      Packets
counter-employee-web                    3348        27
Filter: ingress-port-voip-class-limit-tcp-icmp
Counters:
Name                                     Bytes      Packets
icmp-counter                            4100        49
Policers:
Name                                     Packets
icmp-connection-policer                  0
tcp-connection-policer                   0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

Meaning The **show firewall** command displays the names of all firewall filters, policers, and counters that are configured on the switch. The output fields show byte and packet counts for counters and packet count for policers.

Monitoring Traffic for a Specific Firewall Filter

Purpose Perform the following task to monitor the number of packets and bytes that matched a firewall filter and monitor the number of packets that exceeded the policer rate limits.

Action Use the operational mode command:

```
user@switch> show firewall filter ingress-vlan-rogue-block
Filter: ingress-vlan-rogue-block
Counters:
Name                                     Bytes      Packets
rogue-counter                            2308        20
```

Meaning The **show firewall filter *filter-name*** command displays the name of the firewall filter, the packet and byte count for all counters configured with the filter, and the packet count for all policers configured with the filter.

Monitoring Traffic for a Specific Policier

Purpose Perform the following task to monitor the number of packets that exceeded policer rate limits:

Action Use the operational mode command:

```
user@switch> show policer tcp-connection-policer
Filter: ingress-port-voip-class-limit-tcp-icmp
Policers:
```

Name	Packets
tcp-connection-policer	0

Meaning The **show policer *policer-name*** command displays the name of the firewall filter that specifies the policer-action and displays the number of packets that exceeded rate limits for the specified filter.

- Related Documentation**
- Configuring Firewall Filters (CLI Procedure) on page 67
 - Configuring Firewall Filters (J-Web Procedure) on page 74
 - Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78
 - Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
 - Verifying That Firewall Filters Are Operational on page 89

Troubleshooting Firewall Filters

- Troubleshooting Firewall Filters on page 93

Troubleshooting Firewall Filters

1. Firewall Filter Configuration Returns a No Space Available in TCAM Message on page 93

Firewall Filter Configuration Returns a No Space Available in TCAM Message

Problem When a firewall filter configuration exceeds the amount of available TCAM space, the switch returns the following **syslogd** message:

```
No space available in tcam.  
Rules for filter filter-name will not be installed.
```

The switch returns this message during the commit operation if the firewall filter that has been applied to a port, VLAN, or Layer 3 interface exceeds the amount of available TCAM space. However, the commit operation for the firewall filter configuration is completed in the CLI module.

Solution When a firewall filter configuration exceeds the amount of available TCAM table space, you must configure a new firewall filter with fewer filter terms so that the space requirements for the filter do not exceed the available space in the TCAM table.

You can perform either of the following procedures to correct the problem:

To delete the firewall filter and its bind points and apply the new smaller firewall filter to the same bind points:

1. Delete the firewall filter configuration and the bind points to ports, VLANs, or Layer 3 interfaces—for example:

```
[edit]  
user@switch# delete firewall family ethernet-switching filter filter-ingress-vlan  
user@switch# delete vlans voice-vlan description "filter to block rogue devices on  
voice-vlan"  
user@switch# delete vlans voice-vlan filter input mini-filter—ingress-vlan
```

2. Commit the operation:

```
[edit]  
user@switch# commit
```

3. Configure a smaller filter with fewer terms that does not exceed the amount of available TCAM space on the switch—for example:

```
[edit]
user@switch# set firewall family ethernet-switching filter new—filter-ingress-vlan ...
```

4. Apply (bind) the new firewall filter to a port, VLAN , or Layer 3 interface—for example:

```
[edit]
user@switch# set vlans voice-vlan description "filter to block rogue devices on
voice-vlan"
user@switch# set vlans voice-vlan filter input new-filter—ingress-vlan
```

5. Commit the operation:

```
[edit]
user@switch# commit
```

To apply a new firewall filter and overwrite the existing bind points:

1. Configure a firewall filter with fewer terms than the original filter:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-filter-ingress-vlan...
```

2. Apply the firewall filter to the port, VLAN, or Layer 3 interfaces to overwrite the bind points of the original filter—for example:

```
[edit]
user@switch# set vlans voice-vlan description "smaller filter to block rogue devices on
voice-vlan"
user@switch# set vlans voice-vlan filter input new-filter-ingress-vlan
```

3. Commit the operation:

```
[edit]
user@switch# commit
```

Only the original bind points, and not the original firewall filter itself, are deleted.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Verifying That Firewall Filters Are Operational on page 89
- Configuring Firewall Filters (CLI Procedure) on page 67
- Configuring Firewall Filters (J-Web Procedure) on page 74

CHAPTER 6

Configuration Statements for Firewall Filters

- [\[edit firewall\] Configuration Statement Hierarchy on page 95](#)
- [Firewall Filter Configuration Statements Supported by Junos OS for EX Series Switches on page 96](#)

[\[edit firewall\] Configuration Statement Hierarchy](#)

```
firewall {  
  family family-name {  
    filter filter-name {  
      interface-specific;  
      term term-name {  
        from {  
          match-conditions;  
        }  
        then {  
          action;  
          action-modifiers;  
        }  
      }  
    }  
  }  
  policer policer-name {  
    filter-specific;  
    if-exceeding {  
      bandwidth-limit bps;  
      burst-size-limit bytes;  
    }  
    then {  
      policer-action;  
    }  
  }  
}
```

Related Documentation

- [Firewall Filter Configuration Statements Supported by Junos OS for EX Series Switches on page 96](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39](#)

- Configuring Firewall Filters (CLI Procedure) on page 67
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78
- Firewall Filters for EX Series Switches Overview on page 3

Firewall Filter Configuration Statements Supported by Junos OS for EX Series Switches

You configure firewall filters to filter packets based on their components and to perform an action on packets that match the filter.

Table 14 on page 96 lists the options that are supported for the firewall statement in Junos OS for EX Series switches.

Table 14: Supported Options for Firewall Filter Statements

Statement and Option	Description
<code>family <i>family-name</i> {</code> <code>}</code>	<p>The <i>family-name</i> option specifies the version or type of addressing protocol:</p> <ul style="list-style-type: none"> • any—Filter packets based on protocol-independent match conditions. • ethernet-switching—Filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets • inet—Filter IPv4 packets • inet6—Filter IPv6 packets
<code>filter <i>filter-name</i> {</code> <code>}</code>	<p>The <i>filter-name</i> option identifies the filter. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the name in quotation marks (" ").</p>
<code>interface-specific</code>	<p>The interface-specific statement configures unique names for individual firewall counters specific to each interface.</p>
<code>term <i>term-name</i> {</code> <code>}</code>	<p>The <i>term-name</i> option identifies the term. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). Each term name must be unique within a filter.</p>
<code>from {</code> <code> <i>match-conditions</i>;</code> <code>}</code>	<p>The from statement is optional. If you omit it, all packets are considered to match.</p>
<code>then {</code> <code> <i>action</i>;</code> <code> <i>action-modifiers</i>;</code> <code>}</code>	<p>For information about the <i>action</i> and <i>action-modifiers</i> options, see “Firewall Filter Match Conditions and Actions for EX Series Switches” on page 11.</p>
<code>policer <i>policer-name</i> {</code> <code>}</code>	<p>The <i>policer-name</i> option identifies the policer. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the name in quotation marks (" ").</p>

Table 14: Supported Options for Firewall Filter Statements (*continued*)

Statement and Option	Description
<code>filter-specific</code>	The filter-specific statement configures policers and counters for a specific filter name.
<pre>if-exceeding { bandwidth-limit <i>bps</i> burst-size-limit <i>bytes</i> }</pre>	<p>The bandwidth-limit <i>bps</i> option specifies the traffic rate in bits per second (bps).</p> <p>You can specify <i>bps</i> as a decimal value or as a decimal number followed by one of the following abbreviations:</p> <ul style="list-style-type: none"> • k (thousand) • m (million) • g (billion, which is also called a thousand million) <p>Range: 1000 (1k) through 102,300,000,000 (102.3g) bps</p> <p>The burst-size-limit <i>bytes</i> option specifies the maximum allowed burst size to control the amount of traffic bursting. To determine the value for the burst-size limit, you can multiply the bandwidth of the interface on which the filter is applied by the amount of time (in seconds) to allow a burst of traffic at that bandwidth to occur:</p> <p>burst size = bandwidth * allowable time for burst traffic</p> <p>You can specify a decimal value or a decimal number followed by k (thousand) or m (million).</p> <p>Range: 1 through 2,147,450,880 bytes</p>
<pre>then { <i>policer-action</i> }</pre>	Use the <i>policer-action</i> option to specify discard to discard traffic that exceeds the rate limits.

Junos OS for EX Series switches does not support some of the firewall filter statements that are supported by other Junos OS packages. Table 15 on page 98 shows the firewall filter statements that are not supported by Junos OS for EX Series switches.

Table 15: Firewall Filter Statements That Are Not Supported by Junos OS for EX Series Switches

Statements Not Supported	Statement Hierarchy Level
<ul style="list-style-type: none"> • <code>interface-set <i>interface-set-name</i> {</code> <code>}</code> • <code>load-balance-group <i>group-name</i> {</code> <code>}</code> • <code>three-color-policer <i>name</i> {</code> <code>}</code> • <code>logical-interface-policer;</code> • <code>single-rate {</code> <code>}</code> • <code>two-rate {</code> <code>}</code> 	[edit firewall]
<ul style="list-style-type: none"> • <code>prefix-action <i>name</i> {</code> <code>}</code> • <code>prefix-policer {</code> <code>}</code> • <code>service-filter <i>filter-name</i> {</code> <code>}</code> • <code>simple-filter <i>simple-filter-name</i> {</code> <code>}</code> 	[edit firewall family <i>family-name</i>]
<ul style="list-style-type: none"> • <code>accounting-profile <i>name</i>;</code> 	[edit firewall family <i>family-name</i> filter <i>filter-name</i>]
<ul style="list-style-type: none"> • <code>logical-bandwidth-policer;</code> • <code>logical-interface-policer;</code> 	[edit firewall policer <i>policer-name</i>]
<code>bandwidth-percent <i>number</i>;</code>	[edit firewall policer <i>policer-name</i> if-exceeding]

Related Documentation

- Firewall Filter Match Conditions and Actions for EX Series Switches on page 11
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Configuring Firewall Filters (CLI Procedure) on page 67
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78
- Firewall Filters for EX Series Switches Overview on page 3

apply-path

Syntax	<code>apply-path path;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> policy-options prefix-list <i>name</i>], [edit policy-options prefix-list <i>name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Expand a prefix list to include all prefixes pointed to by a defined path.
Options	path —String of elements composed of identifiers or configuration keywords that points to a set of prefixes. You can include wildcards (enclosed in angle brackets) to match more than one identifier. You cannot add a path element, including wildcards, after a leaf statement. Path elements, including wildcards, can only be used after a container statement.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring Prefix Lists

as-path

Syntax	<code>as-path name regular-expression;</code>
Hierarchy Level	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
Description	Define an autonomous system (AS) path regular expression for use in a routing policy match condition.
Options	name —Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). regular-expression —One or more regular expressions used to match the AS path.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring AS Path Regular Expressions to Use as Routing Policy Match ConditionsConfiguring Routing Policies and Policy Objects in the Dynamic Databasedynamic-db on page 108

as-path-group

Syntax	<pre>as-path-group <i>group-name</i> { as-path <i>name</i> <i>regular-expression</i>; }</pre>
Hierarchy Level	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for dynamic database configuration introduced in Junos OS Release 9.5. Support for dynamic database configuration introduced in Junos OS Release 9.5 for EX Series switches.
Description	Define a group containing multiple AS path regular expressions for use in a routing policy match condition.
Options	<p><i>group-name</i>—Name that identifies the AS path group. One or more AS path regular expressions must be listed below the as-path-group hierarchy.</p> <p><i>name</i>—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p><i>regular-expression</i>—One or more regular expressions used to match the AS path.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring AS Path Regular Expressions to Use as Routing Policy Match Conditions Configuring Routing Policies and Policy Objects in the Dynamic Database dynamic-db on page 108

bandwidth-limit

Syntax	<code>bandwidth-limit <i>bps</i>;</code>
Hierarchy Level	<code>[edit firewall policer <i>policer-name</i> if-exceeding]</code> <code>[edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i> if-exceeding]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Logical systems support introduced in Junos OS Release 9.3.
Description	Specify the traffic rate in bits per second.
Options	<p><i>bps</i> —Traffic rate to be specified in bits per second. Specify <i>bps</i> as a decimal value or as a decimal number followed by one of the following abbreviations:</p> <ul style="list-style-type: none">• k (thousand)• m (million)• g (billion, which is also called a thousand million) <p>Range:</p> <ul style="list-style-type: none">• 1000 (1k) through 102,300,000,000 (102.3g) bps (EX Series switches)• 8000 (8k) through 40,000,000,000 (40g) bps (routers)
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39• Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78• Understanding the Use of Policers in Firewall Filters on page 37• Rate Limiting• Single-Rate Two-Color Policer Overview• Configuring a Single-Rate Two-Color Policer

burst-size-limit

Syntax	<code>burst-size-limit <i>bytes</i>;</code>
Hierarchy Level	<code>[edit firewall policer <i>policer-name</i> if-exceeding]</code> <code>[edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i> if-exceeding]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Logical systems support introduced in Junos OS Release 9.3.
Description	Specify the maximum allowed burst size to control the amount of traffic bursting.
Options	<i>bytes</i> —Decimal value or a decimal number followed by k (thousand) or m (million). Range: <ul style="list-style-type: none"> • 1 through 2,147,450,880 bytes (EX Series switches) • 1500 through 1,00,000,000,000 bytes (routers)
Required Privilege Level	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39 • Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78 • Understanding the Use of Policers in Firewall Filters on page 37 • Rate Limiting • Single-Rate Two-Color Policer Overview • Configuring a Single-Rate Two-Color Policer

community

Syntax	<pre>community <i>name</i> { invert-match; members [<i>community-ids</i>]; }</pre>
Hierarchy Level	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
Description	Define a community or extended community for use in a routing policy match condition.
Options	<p><i>name</i>—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters. To include spaces in the name, enclose it in quotation marks (" ").</p> <p><i>invert-match</i>—Invert the results of the community expression matching.</p> <p><i>members community-ids</i>—One or more community members. If you specify more than one member, you must enclose all members in brackets.</p> <p>The format for <i>community-ids</i> is:</p> <p><i>as-number:community-value</i></p> <p><i>as-number</i> is the AS number and can be a value in the range from 0 through 65,535. <i>community-value</i> is the community identifier and can be a number in the range from 0 through 65,535.</p> <p>You also can specify <i>community-ids</i> for communities as one of the following well-known community names, which are defined in RFC 1997, <i>BGP Communities Attribute</i>:</p> <ul style="list-style-type: none">• <i>no-export</i>—Routes containing this community name are not advertised outside a BGP confederation boundary.• <i>no-advertise</i>—Routes containing this community name are not advertised to other BGP peers.• <i>no-export-subconfed</i>—Routes containing this community name are not advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation. <p>You can explicitly exclude BGP community information with a static route using the <i>none</i> option. Include <i>none</i> when configuring an individual route in the <i>route</i> portion of the <i>static</i> statement to override a <i>community</i> option specified in the <i>defaults</i> portion of the statement.</p>

The format for extended **community-ids** is the following:

type:administrator:assigned-number

type is the type of extended community and can be either a **bandwidth**, **target**, **origin**, **domain-id**, **src-as**, or **rt-import** community or a 16-bit number that identifies a specific BGP extended community. The **target** community identifies the destination to which the route is going. The **origin** community identifies where the route originated. The **domain-id** community identifies the OSPF domain from which the route originated. The **src-as** community identifies the autonomous system from which the route originated. The **rt-import** community identifies the route to install in the routing table.



NOTE: For **src-as**, you can specify only an AS number and not an IP address. For **rt-import**, you can specify only an IP address and not an AS number.

administrator is the administrator. It is either an AS number or an IPv4 address prefix, depending on the type of extended community.

assigned-number identifies the local provider.

The format for linking a bandwidth with an AS number is:

bandwidth:as-number:bandwidth

as-number specifies the AS number and **bandwidth** specifies the bandwidth in bytes per second.



NOTE: In Junos OS Release 9.1 and later, you can specify 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*, as well as the 2-byte AS numbers that are supported in earlier releases of the Junos OS. In plain-number format, you can configure a value in the range from 1 through 4,294,967,295. To configure a target or origin extended community that includes a 4-byte AS number in the plain-number format, append the letter “L” to the end of number. For example, a target community with the 4-byte AS number 334,324 and an assigned number of 132 is represented as **target:334324L:132**.

In Junos OS Release 9.2 and later, you can also use AS-dot notation when defining a 4-byte AS number for the target and origin extended communities. Specify two integers joined by a period: *16-bit high-order value in decimal.16-bit low-order value in decimal*. For example, the 4-byte AS number represented in plain-number format as 65546 is represented in AS-dot notation as 1.10.



For more information about configuring AS numbers, see the *Junos OS Routing Protocols Configuration Guide*.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Overview of BGP Communities and Extended Communities as Routing Policy Match Conditions• Defining BGP Communities and Extended Communities for Use in Routing Policy Match Conditions• Configuring Routing Policies and Policy Objects in the Dynamic Database• dynamic-db on page 108

condition

Syntax	<pre>condition <i>condition-name</i> { if-route-exists <i>address</i> table <i>table-name</i>; }</pre>
Hierarchy Level	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
Description	Define a policy condition based on the existence of routes in specific tables for use in BGP export policies.
Options	if-route-exists <i>address</i> —Specify the address of the route in question. table <i>table-name</i> —Specify a routing table.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Routing Policy Match Conditions Based on Routing Table Entries• Configuring Routing Policies and Policy Objects in the Dynamic Database• dynamic-db on page 108

damping

Syntax	<pre>damping <i>name</i> { disable; half-life <i>minutes</i>; max-suppress <i>minutes</i>; reuse <i>number</i>; suppress <i>number</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Define route flap damping properties to set on BGP routes.
Options	<p>disable—Disable damping on a per-prefix basis. Any damping state that is present in the routing table for a prefix is deleted if damping is disabled.</p> <p>half-life <i>minutes</i>—Decay half-life. <i>minutes</i> is the interval after which the accumulated figure-of-merit value is reduced by half if the route remains stable. Range: 1 through 45 Default: 15 minutes</p> <p> NOTE: For the half-life, configure a value that is less than the max-suppress.</p> <p>max-suppress <i>minutes</i>—Maximum hold-down time. <i>minutes</i> is the maximum time that a route can be suppressed no matter how unstable it has been. Range: 1 through 720 Default: 60 minutes</p> <p> NOTE: For the max-suppress, configure a value that is greater than the half-life.</p> <p><i>name</i>—Name that identifies the set of damping parameters. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p>reuse <i>number</i>—Reuse threshold. <i>number</i> is the figure-of-merit value below which a suppressed route can be used again. Range: 1 through 20,000 Default: 750 (unitless)</p>

suppress *number*—Cutoff (suppression) threshold. *number* is the figure-of-merit value above which a route is suppressed for use or inclusion in advertisements.

Range: 1 through 20,000

Default: 3000 (unitless)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BGP Flap Damping Parameters

dynamic-db

Syntax dynamic-db;

Hierarchy Level [edit logical-systems *logical-system-name* policy-options as-path *path-name*],
[edit logical-systems *logical-system-name* policy-options as-path-group *group-name*],
[edit logical-systems *logical-system-name* policy-options community *community-name*],
[edit logical-systems *logical-system-name* policy-options condition *condition-name*],
[edit logical-systems *logical-system-name* policy-options policy-statement *policy-statement-name*],
[edit logical-systems *logical-system-name* policy-options prefix-list *prefix-list-name*],
[edit policy-options as-path *path-name*],
[edit policy-options as-path-group *group-name*],
[edit policy-options community *community-name*],
[edit policy-options condition *condition-name*],
[edit policy-options policy-statement *policy-statement-name*],
[edit policy-options prefix-list *prefix-list-name*]

Release Information Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description Define routing policies and policy objects that reference policies configured in the dynamic database at the **[edit dynamic]** hierarchy level.

Required Privilege Level routing—To view this statement in the configuration.
routing-control-level—To add this statement to the configuration.

Related Documentation

- Configuring Routing Policies Based on Dynamic Database Configuration

family (Firewall Filter)

Syntax	<pre> family <i>family-name</i> { filter <i>filter-name</i> { interface-specific; term <i>term-name</i> { from { <i>match-conditions</i>; } then { <i>action</i>; <i>action-modifiers</i>; } } } } </pre>
Hierarchy Level	[edit firewall]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Option interface-specific introduced in Junos OS Release 9.5 for EX Series switches.
Description	Configure a firewall filter for IP version 4 or IP version 6.
Options	<p><i>family-name</i>—Version or type of addressing protocol:</p> <ul style="list-style-type: none"> • any—Filter packets based on protocol-independent match conditions. • ethernet-switching—Filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets. • inet—Filter IPv4 packets. • inet6—Filter IPv6 packets. <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions for EX Series Switches on page 11 • Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39 • Configuring Firewall Filters (CLI Procedure) on page 67 • Configuring Firewall Filters (J-Web Procedure) on page 74 • Firewall Filters for EX Series Switches Overview on page 3

filter

Syntax `filter filter-name {
 interface-specific;
 term term-name {
 from {
 match-conditions;
 }
 then {
 action;
 action-modifiers;
 }
 }
 }`

Hierarchy Level [edit firewall family *family-name*]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Option **interface-specific** introduced in Junos OS Release 9.5 for EX Series switches.

Description Configure firewall filters.

Options *filter-name*—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.

The remaining statements are explained separately.

Required Privilege Level firewall—To view this statement in the configuration.
 firewall-control—To add this statement to the configuration.

Related Documentation

- Firewall Filter Match Conditions and Actions for EX Series Switches on page 11
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Configuring Firewall Filters (CLI Procedure) on page 67
- Configuring Firewall Filters (J-Web Procedure) on page 74
- Firewall Filters for EX Series Switches Overview on page 3

filter

Syntax	<code>filter (input output) <i>filter-name</i>;</code>
Hierarchy Level	<code>[edit vlans <i>vlan-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Apply a firewall filter to traffic coming into or exiting from the VLAN.
Default	All incoming traffic is accepted unmodified to the VLAN, and all outgoing traffic is sent unmodified from the VLAN.
Options	<i>filter-name</i> —Name of a firewall filter defined in a filter statement. <ul style="list-style-type: none">• input—Apply a firewall filter to VLAN ingress traffic.• output—Apply a firewall filter to VLAN egress traffic.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39• Configuring Firewall Filters (CLI Procedure) on page 67• Configuring Firewall Filters (J-Web Procedure) on page 74• Firewall Filters for EX Series Switches Overview on page 3

filter-specific

Syntax	<code>filter-specific;</code>
Hierarchy Level	<code>[edit firewall policer <i>policer-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.5 for EX Series switches.
Description	Configure a policer to act as a filter-specific policer.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39• Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78• Understanding the Use of Policers in Firewall Filters on page 37

firewall

```
Syntax  firewall {  
        family family-name {  
            filter filter-name {  
                interface-specific;  
                term term-name {  
                    from {  
                        match-conditions;  
                    }  
                    then {  
                        action;  
                        action-modifiers;  
                    }  
                }  
            }  
        }  
        policer policer-name {  
            filter-specific;  
            if-exceeding {  
                bandwidth-limit bps;  
                burst-size-limit bytes;  
            }  
            then {  
                policer-action;  
            }  
        }  
    }
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Options **interface-specific** and **filter-specific** introduced in Junos OS Release 9.5 for EX Series switches.

Description Configure firewall filters and policers.

The remaining statements are explained separately.

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Firewall Filter Match Conditions and Actions for EX Series Switches on page 11
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39
- Configuring Firewall Filters (CLI Procedure) on page 67
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78
- Firewall Filters for EX Series Switches Overview on page 3

from

Syntax	<pre>from { match-conditions; }</pre>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Match packet fields to values specified in a match condition. If the from statement is not included in a firewall filter configuration, all packets are considered to match and the actions and action modifiers in the then statement are taken.
Options	match-conditions —Conditions that define the values or fields that the incoming or outgoing packets must contain for a match. You can specify one or more match conditions. If you specify more than one, they all must match for a match to occur and for the action in the then statement to be taken.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Firewall Filter Match Conditions and Actions for EX Series Switches on page 11• Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39• Configuring Firewall Filters (CLI Procedure) on page 67• Configuring Firewall Filters (J-Web Procedure) on page 74• Understanding Firewall Filter Match Conditions on page 33

if-exceeding

Syntax	<pre>if-exceeding { bandwidth-limit <i>bps</i>; bandwidth-percent <i>percent</i> burst-size-limit <i>bytes</i>; }</pre>
Hierarchy Level	<pre>[edit firewall policer <i>policer-name</i>] [edit logical-systems logical-system-name firewall policer <i>policer-name</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Logical systems support introduced in Junos OS Release 9.3.</p>
Description	<p>Configure policer rate limits.</p> <p>The bandwidth-percent statement is supported on routers only.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39• Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78• Understanding the Use of Policers in Firewall Filters on page 37• Rate Limiting• Single-Rate Two-Color Policer Overview• Configuring a Single-Rate Two-Color Policer

interface-specific

Syntax	interface-specific;
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5 for EX Series switches.
Description	Configure firewall counters that are interface-specific.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Firewall Filter Match Conditions and Actions for EX Series Switches on page 11• Configuring Firewall Filters (CLI Procedure) on page 67• Configuring Firewall Filters (J-Web Procedure) on page 74• Firewall Filters for EX Series Switches Overview on page 3

policer

Syntax	<pre>policer <i>policer-name</i> { filter-specific; if-exceeding { bandwidth-limit <i>bps</i>; bandwidth-percent <i>percent</i> burst-size-limit <i>bytes</i>; } then { <i>policer-action</i>; } }</pre>
Hierarchy Level	[edit firewall] [edit logical-systems <i>logical-system-name</i> firewall]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Logical systems support introduced in Junos OS Release 9.3.
Description	Configure policer rate limits and actions. To activate a policer, you must include the policer action modifier in the then statement in a firewall filter term. Each policer that you configure includes an implicit counter. To ensure term-specific packet counts, you configure a policer for each term in the filter that requires policing.
Options	<i>policer-name</i> —Name that identifies the policer. The name can contain letters, numbers, hyphens (-), and can be up to 64 characters long. The remaining statements are explained separately.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39• Example: Configuring CoS on EX Series Switches• Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78• Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure)• Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)• Configuring Firewall Filters (CLI Procedure) on page 67• Configuring Firewall Filters (J-Web Procedure) on page 74• Understanding the Use of Policers in Firewall Filters on page 37• Single-Rate Two-Color Policer Overview• Configuring a Single-Rate Two-Color Policer

policy-statement

Syntax	<pre> policy-statement <i>policy-name</i> { term <i>term-name</i> { from { family <i>family-name</i>; match-conditions; policy <i>subroutine-policy-name</i>; prefix-list <i>prefix-list-name</i>; prefix-list-filter <i>prefix-list-name</i> match-type <actions>; route-filter <i>destination-prefix</i> match-type <actions>; source-address-filter <i>source-prefix</i> match-type <actions>; } to { match-conditions; policy <i>subroutine-policy-name</i>; } then <i>actions</i>; } } </pre>
Hierarchy Level	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>inet-mdt option introduced in Junos OS Release 10.0R2.</p>
Description	Define a routing policy, including subroutine policies.
Options	<p>actions—(Optional) One or more actions to take if the conditions match. The actions are described in Configuring Flow Control Actions.</p> <p>family <i>family-name</i>—(Optional) Specify an address family protocol. Specify inet for IPv4. Specify inet6 for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify iso. For IPv4 multicast VPN traffic, specify inet-mvpn. For IPv6 multicast VPN traffic, specify inet6-mvpn. For multicast-distribution-tree (MDT) IPv4 traffic, specify inet-mdt.</p>



NOTE: When **family** is not specified, the routing device uses the default IPv4 setting.

from—(Optional) Match a route based on its source address.

match-conditions—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in Configuring Match Conditions in Routing Policy Terms.

policy-subroutine-policy-name—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). For information about how to configure subroutines, see Configuring Subroutines in Routing Policy Match Conditions.

policy-name—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").

prefix-list prefix-list-name —Name of a list of IPv4 or IPv6 prefixes.

prefix-list-filter prefix-list-name—Name of a prefix list to evaluate using qualifiers; **match-type** is the type of match (see Configuring Prefix List Filters), and **actions** is the action to take if the prefixes match.

route-filter destination-prefix match-type <actions>—(Optional) List of routes on which to perform an immediate match; **destination-prefix** is the IPv4 or IPv6 route prefix to match, **match-type** is the type of match (see Configuring Route Lists), and **actions** is the action to take if the **destination-prefix** matches.

source-address-filter source-prefix match-type <actions>—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. **source-prefix** is the IPv4 or IPv6 route prefix to match, **match-type** is the type of match (see Configuring Route Lists), and **actions** is the action to take if the **source-prefix** matches.

term term-name—Name that identifies the term.

to—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

then—(Optional) Actions to take on matching routes. The actions are described in Configuring Flow Control Actions and Configuring Actions That Manipulate Route Characteristics.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Defining Routing Policies• Configuring Routing Policies and Policy Objects in the Dynamic Database• dynamic-db on page 108
------------------------------	---

prefix-list

Syntax	<pre>prefix-list name { ip-addresses; apply-path path; }</pre>
Hierarchy Level	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches. Support for the vpls protocol family introduced in Junos OS Release 10.2.
Description	<p>Define a list of IPv4 or IPv6 address prefixes for use in a routing policy statement or firewall filter statement.</p> <p>You can configure up to 85,325 prefixes in each prefix list. To configure more than 85,325 prefixes, configure multiple prefix lists and apply them to multiple firewall filter terms.</p>
Options	<p>name—Name that identifies the list of IPv4 or IPv6 address prefixes.</p> <p>ip-addresses—List of IPv4 or IPv6 address prefixes, one IP address per line in the configuration.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Prefix Lists for Use in Routing Policy Match Conditions Configuring Routing Policies and Policy Objects in the Dynamic Database dynamic-db on page 108

routing-instance

Syntax	<code>routing-instance <i>routing-instance-name</i>;</code>
Hierarchy Level	[edit firewall family inet filter <i>filter-name</i> term <i>term-name</i> then]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Specify a specific virtual routing instance to which the switch sends matched packets.
Options	<i>routing-instance-name</i> —Name of a virtual routing instance.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 58• Configuring Virtual Routing Instances (CLI Procedure)• Understanding Filter-Based Forwarding for EX Series Switches on page 38

term

Syntax	<pre>term <i>term-name</i> { from { <i>match-conditions</i>; } then { <i>action</i>; <i>action-modifiers</i>; } }</pre>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Define a firewall filter term.
Options	<p><i>term-name</i>—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions for EX Series Switches on page 11 • Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39 • Configuring Firewall Filters (CLI Procedure) on page 67 • Configuring Firewall Filters (J-Web Procedure) on page 74 • Firewall Filters for EX Series Switches Overview on page 3

then

Syntax	<pre>then { action; action-modifiers; }</pre>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure a filter action.
Options	<p>action—Actions to accept, discard, or forward packets that match all match conditions specified in a filter term.</p> <p>action-modifiers—Additional actions to analyze, classify, count, or police packets that match all conditions specified in a filter term.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Firewall Filter Match Conditions and Actions for EX Series Switches on page 11• Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39• Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 58• Configuring Firewall Filters (CLI Procedure) on page 67• Configuring Firewall Filters (J-Web Procedure) on page 74• Understanding Firewall Filter Match Conditions on page 33


then

Syntax	<pre>then { <i>policer-action</i>; }</pre>
Hierarchy Level	<pre>[edit firewall policer <i>policer-name</i>] [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Configure a policer action.
Options	<p><i>policer-action</i>—Actions to take are:</p> <ul style="list-style-type: none"> • discard—Discard traffic that exceeds the rate limits defined by the policer. • forwarding-class <i>class-name</i>—For routers only, classify traffic that exceeds the rate limits defined by the policer. • loss-priority—Set the loss priority for traffic that exceeds the rate limits defined by the policer.
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall -control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39 • Configuring Policers to Control Traffic Rates (CLI Procedure) on page 78 • Configuring Firewall Filters (CLI Procedure) on page 67 • Configuring Firewall Filters (J-Web Procedure) on page 74 • Understanding the Use of Policers in Firewall Filters on page 37 • Example: Configuring CoS for a PBB Network on MX Series Routers • Single-Rate Two-Color Policer Overview • Configuring a Single-Rate Two-Color Policer

CHAPTER 7

Operational Commands for Firewall Filters

clear firewall

Syntax	clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i> logical-system <i>logical-system-name</i>)
Syntax (EX Series Switch)	clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i>)
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3.
Description	Clear statistics about configured firewall filters.
<div> NOTE: The clear firewall command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for GRES.</div>	
Options	<p>all—Clear the packet and byte counts for all filters.</p> <p>counter <i>counter-name</i>—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.</p> <p>filter <i>filter-name</i>—Clear the packet and byte counts for the specified firewall filter.</p> <p>logical-system <i>logical-system-name</i>—Clear the packet and byte counts for the specified logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show firewall on page 128
List of Sample Output	clear firewall all on page 126
Output Fields	When you enter this command, you are provided feedback on the status of your request.
clear firewall all	user@host> clear firewall all

clear firewall

Syntax	clear firewall <all> <counter <i>counter-name</i> > <filter <i>filter-name</i> >
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Clear statistics about configured firewall filters.
Options	<p>none—Clear the packet and byte counts for all firewall filter counters and clear the packet counts for all policer counters.</p> <p>all—(Optional) Clear the packet and byte counts for all firewall filter counters and clear the packet counts for all policer counters.</p> <p>counter <i>counter-name</i> —(Optional) Clear the packet and byte counts for the specified firewall filter counter.</p> <p>filter <i>filter-name</i> —(Optional) Clear the packet and byte counts for the specified firewall filter.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39 • Verifying That Firewall Filters Are Operational on page 89 • Verifying That Policers Are Operational on page 90 • Firewall Filters for EX Series Switches Overview on page 3 • Understanding the Use of Policers in Firewall Filters on page 37
clear firewall (all)	user@host> clear firewall all
clear firewall (counter counter-name)	user@host> clear firewall counter port-filter-counter
clear firewall (filter filter-name)	user@host> clear firewall filter ingress-port-filter

show firewall

Syntax	show firewall <filter <i>filter-name</i> > <counter <i>counter-name</i> > <log> <logical-system (all <i>logical-system-name</i>)> <terse>
Syntax (EX Series Switch)	show firewall <filter <i>filter-name</i> > <counter <i>counter-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3. terse option introduced in Junos OS Release 9.4.
Description	Display statistics about configured firewall filters.
Options	none—(Optional) Display statistics about configured firewall filters. filter <i>filter-name</i> —(Optional) Name of a configured filter. counter <i>counter-name</i> —(Optional) Name of a filter counter. logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular system. log—(Optional) Display log entries for firewall filters. terse—(Optional) Display firewall filter names only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear firewall on page 126
List of Sample Output	show firewall filter on page 130 show firewall filter (Dynamic Input Filter) on page 130 show firewall (Logical Systems) on page 130
Output Fields	Table 16 on page 129 lists the output fields for the show firewall command. Output fields are listed in the approximate order in which they appear.

Table 16: show firewall Output Fields

Field Name	Field Description
Filter	<p>Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level.</p> <p>When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either -i for an input filter or -o for an output filter.</p> <p>When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either -in for an input filter or -out for an output filter. When a logical system-specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, __ls1/filter1).</p>
Counters	<p>Display filter counter information:</p> <ul style="list-style-type: none"> • Name—Name of a filter counter that has been configured with the counter firewall filter action. • Bytes—Number of bytes that match the filter term under which the counter action is specified. • Packets—Number of packets that matched the filter term under which the counter action is specified.
Policers	<p>Display policer information:</p> <ul style="list-style-type: none"> • Name—Name of policer. • Packets—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.

```
show firewall filter user@host> show firewall filter test
Filter: test
Counters:
Name                               Bytes          Packets
Counter-1                          0              0
Counter-2                          0              0
Policers:
Name                               Packets
Policer-1                          0

show firewall filter user@host> show firewall filter dfwd-ge-5/0/0.1-in
(Dynamic Input Filter) Filter: dfwd-ge-5/0/0.1-in
Counters:
Name                               Bytes          Packets
c1-ge-5/0/0.1-in                   0              0

show firewall (Logical user@host>show firewall
Systems)
Filter: __lr1/test
Counters:
Name                               Bytes          Packets
icmp                               420            5
Filter: __default_bpdu_filter__
Filter: __lr1/inet_filter1
Counters:
Name                               Bytes          Packets
inet_tcp_count                     0              0
inet_udp_count                     0              0
Filter: __lr1/inet_filter2
Counters:
Name                               Bytes          Packets
inet_icmp_count                    0              0
inet_pim_count                     0              0
Filter: __lr2/inet_filter1
Counters:
Name                               Bytes          Packets
inet_tcp_count                     0              0
inet_udp_count                     0              0
```

show firewall

Syntax	<pre>show firewall <counter <i>counter-name</i>> <filter <i>filter-name</i>> log (detail interface <i>interface-name</i>) terse</pre>	
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.	
Description	Display statistics about configured firewall filters.	
Options	<p>none—Display statistics about all configured firewall filters, counters, and policers.</p> <p>counter <i>counter-name</i>—(Optional) Display statistics about a particular firewall filter counter.</p> <p>filter <i>filter-name</i>—(Optional) Display statistics about a particular firewall filter.</p> <p>log (detail interface <i>interface-name</i>)—(Optional) Display detailed log entries of firewall activity or log information about a specific interface.</p> <p>terse—(Optional) Display firewall filter names only.</p>	
Required Privilege Level	view	
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39 • Verifying That Firewall Filters Are Operational on page 89 • Verifying That Policers Are Operational on page 90 • Firewall Filters for EX Series Switches Overview on page 3 • Understanding the Use of Policers in Firewall Filters on page 37 	
List of Sample Output	<pre>show firewall on page 132 show firewall (filter filter-name) on page 132 show firewall (counter counter-name) on page 132 show firewall log on page 132</pre>	
Output Fields	Table 17 on page 131 lists the output fields for the show firewall command. Output fields are listed in the approximate order in which they appear.	

Table 17: show firewall Output Fields

Field Name	Field Description	Level of Output
Filter	Name of the filter that is configured with the filter statement at the [edit firewall] hierarchy level.	All levels

Table 17: show firewall Output Fields (*continued*)

Field Name	Field Description	Level of Output
Counters	Display filter counter information: <ul style="list-style-type: none"> Name—Name of a filter counter that has been configured with the counter firewall filter action Bytes—Number of bytes that match the filter term where the counter action was specified. Packets—Number of packets that matched the filter term where the counter action was specified. 	All levels
Policers	Display policer information: <ul style="list-style-type: none"> Name—Name of policer. Packets—Number of packets that matched the filter term where the policer action was specified. This is the number of packets that exceed the rate limits that the policer specifies. 	All levels

show firewall	<pre> user@host> show firewall Filter: egress-vlan-filter Counters: Name Bytes Packets employee-web-counter 0 0 Filter: ingress-port-filter Counters: Name Bytes Packets ingress-port-counter 0 0 Filter: ingress-port-voip-class-filter Counters: Name Bytes Packets icmp-counter 0 0 Policers: Name Packets icmp-connection-policer 0 tcp-connection-policer 0 </pre>
show firewall (filter filter-name)	<pre> user@host> show firewall filter egress-vlan-filter Filter: egress-vlan-filter Counters: Name Bytes Packets employee-web-counter 0 0 </pre>
show firewall (counter counter-name)	<pre> user@host> show firewall counter icmp-counter Filter: ingress-port-voip-class-filter Counters: Name Bytes Packets icmp-counter 0 0 </pre>
show firewall log	<pre> user@host> show firewall log Log : Time Filter Action Interface Protocol Src Addr 08:00:53 pfe R ge-1/0/1.0 ICMP 192.168.3.5 192.168.3.4 </pre>

08:00:52	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				
08:00:51	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				
08:00:50	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				
08:00:49	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				
08:00:48	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				
08:00:47	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				

show firewall log

Syntax	show firewall log <detail> <interface <i>interface-name</i> > <logical-system (<i>logical-system-name</i> all)>
Syntax (EX Series Switch)	show firewall log <detail> <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3.
Description	Display log information about firewall filters.
Options	none—Display log information about firewall filters. detail—(Optional) Display detailed information. interface <i>interface-name</i> —(Optional) Display log information about a specific interface. logical-system (<i>logical-system-name</i> all)—(Optional) Perform this operation on all logical systems or on a particular system.
Required Privilege Level	view
List of Sample Output	show firewall log on page 135 show firewall log detail on page 135
Output Fields	Table 18 on page 134 lists the output fields for the show firewall log command. Output fields are listed in the approximate order in which they appear.

Table 18: show firewall log Output Fields

Field Name	Field Description
Time of Log	Time that the event occurred.
Filter	<p>Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level.</p> <ul style="list-style-type: none"> A hyphen (-) indicates that the packet was handled by the Packet Forwarding Engine. A space (no hyphen) indicates the packet was handled by the Routing Engine. The notation pfe indicates packets logged by the Packet Forwarding Engine hardware filters.

Table 18: show firewall log Output Fields (*continued*)

Field Name	Field Description
Filter Action	Filter action: <ul style="list-style-type: none"> • A—Accept • D—Discard • R—Reject
Name of Interface	Ingress interface for the packet.
Name of protocol	Packet's protocol name: egp, gre, icmp, ipip, ospf, pim, rsvp, tcp, or udp.
Packet length	Length of the packet.
Source address	Packet's source address.
Destination address	Packet's destination address and port.

show firewall log

```
user@host>show firewall log
Time      Filter  Action Interface    Protocol  Src Addr    Dest Addr
13:10:12  pfe      D      rlsq0.902      ICMP      180.1.177.2  180.1.177.1
13:10:11  pfe      D      rlsq0.902      ICMP      180.1.177.2  180.1.177.1
```

show firewall log detail

```
user@host> show firewall log detail
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
....
```

show interfaces filters

Syntax	show interfaces filters <i><interface-name></i>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display firewall filters that are configured on each interface in a system.
Options	none—Display firewall filter information about all interfaces. <i>interface-name</i> —(Optional) Display firewall filter information about a particular interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show interfaces policers on page 138 • show firewall on page 131
List of Sample Output	show interfaces filters on page 136 show interfaces filters <interface-name> on page 137
Output Fields	Table 19 on page 136 lists the output fields for the show interfaces filters command. Output fields are listed in the approximate order in which they appear.

Table 19: show interfaces filters Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the physical interface.	All levels
Admin	Interface state: up or down.	All levels
Link	Link state: up or down.	All levels
Proto	Protocol that is configured on the interface.	All levels
Input Filter	Name of the firewall filter to be evaluated when packets are received on the interface.	All levels
Output Filter	Name of the firewall filter to be evaluated when packets are transmitted on the interface.	All levels

```

show interfaces filters  user@host> show interfaces filters
Interface      Admin Link Proto Input Filter      Output Filter
ge-0/0/0       up   down
ge-0/0/0.0     up   down eth-switch unknown
ge-0/0/1       up   down
ge-0/0/1.0     up   down eth-switch unknown
ge-0/0/2       up   down
ge-0/0/3       up   down

```


ge-0/0/4	up	down
ge-0/0/5	up	down
ge-0/0/6	up	down
ge-0/0/7	up	down
ge-0/0/8	up	down
ge-0/0/9	up	down
ge-0/0/10	up	down
ge-0/0/10.0	up	down

```
show interfaces filters <interface-name>
user@host> show interfaces filters ge-0/0/0
Interface      Admin Link Proto Input Filter      Output Filter
ge-0/0/0       up    down
ge-0/0/0.0     up    down eth-switch unknown
```

show interfaces policers

Syntax	show interfaces policers <i><interface-name></i>
Release Information	Command introduced before Junos OS Release 9.0 for EX Series switches.
Description	Display all policers that are configured on each interface in a system.
Options	none—Display policer information about all interfaces. <i>interface-name</i> —(Optional) display firewall filters information about a particular interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show interfaces filters on page 136 • show policer on page 140
List of Sample Output	show interfaces policers on page 138 show interfaces policers on page 139 show interfaces policers (interface-name) on page 139
Output Fields	Table 20 on page 138 lists the output fields for the show interfaces policers command. Output fields are listed in the approximate order in which they appear.

Table 20: show interfaces policers Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Admin	Interface state: up or down.	All levels
Link	Link state: up or down.	All levels
Proto	Protocol configured on the interface.	All levels
Input Policer	Policer to be evaluated when packets are received on the interface. It has the format <i>interface-name-in-policer</i> .	All levels
Output Policer	Policer to be evaluated when packets are transmitted on the interface. It has the format <i>interface-name-out-policer</i> .	All levels

```

show interfaces  user@host> show interfaces policers
policers         Interface  Admin Link Proto Input Policer      Output Policer
                  ge-0/0/0    up   down
                  ge-0/0/0.0  up   down
                                eth-switch

                  Interface  Admin Link Proto Input Policer      Output Policer

```

```

ge-0/0/1      up    down
ge-0/0/1.0    up    down
eth-switch

Interface      Admin Link Proto Input Policer      Output Policer
ge-0/0/2      up    down
ge-0/0/3      up    down
ge-0/0/4      up    down
ge-0/0/5      up    down
ge-0/0/6      up    down
ge-0/0/7      up    down
ge-0/0/8      up    down
ge-0/0/9      up    down
ge-0/0/10     up    down
ge-0/0/10.0   up    down
eth-switch

show interfaces user@host> show interfaces policers
policers      Interface      Admin Link Proto Input Policer      Output Policer
ge-0/0/0      up    down
ge-0/0/0.0    up    down
eth-switch

Interface      Admin Link Proto Input Policer      Output Policer
ge-0/0/1      up    down
ge-0/0/1.0    up    down
eth-switch

Interface      Admin Link Proto Input Policer      Output Policer
ge-0/0/2      up    down
ge-0/0/3      up    down
ge-0/0/4      up    down
ge-0/0/5      up    down
ge-0/0/6      up    down
ge-0/0/7      up    down
ge-0/0/8      up    down
ge-0/0/9      up    down
ge-0/0/10     up    down
ge-0/0/10.0   up    down
eth-switch

show interfaces user@host> show interfaces policers ge-0/0/1
policers (    Interface      Admin Link Proto Input Policer      Output Policer
interface-name) ge-0/0/0      up    down
ge-0/0/0.0    up    down
eth-switch

```

show policer

Syntax	show policer <i><policer-name></i>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display statistics about configured policers.
Options	<p>none—Display the count of policed packets for all configured policers in the system.</p> <p><i>policer-name</i>—(Optional) Display the count of policed packets for the specified policer.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 39 • Verifying That Firewall Filters Are Operational on page 89 • Verifying That Policers Are Operational on page 90 • Firewall Filters for EX Series Switches Overview on page 3 • Understanding the Use of Policers in Firewall Filters on page 37
List of Sample Output	<p>show policer on page 140</p> <p>show policer (policer-name) on page 141</p>
Output Fields	Table 21 on page 140 lists the output fields for the show policer command. Output fields are listed in the approximate order in which they appear.

Table 21: show policer Output Fields

Field Name	Field Description	Level of Output
Filter	Name of filter that is configured with the filter statement at the [edit firewall] hierarchy level.	All levels
Policers	Display policer information: <ul style="list-style-type: none"> • Filter—Name of filter that specifies the policer action. • Name—Name of policer. • Packets—Number of packets that matched the filter term where the policer action is specified. This is the number of packets that exceed the rate limits that the policer specifies. 	All levels

```

show policer      user@host> show policer
                    Filter: egress-vlan-filter
                    Filter: ingress-port-filter
                    Policers:
                    Name                               Packets
  
```

```
icmp-connection-policer          0
tcp-connection-policer          0
Filter: ingress-vlan-rogue-block
```

```
show policer (policer-name) user@host> show policer tcp-connection-policer
Filter: ingress-port-filter
Policers:
Name                                     Packets
tcp-connection-policer                  0
```

show policy

Syntax	show policy <logical-system (all <i>logical-system-name</i>)> < <i>policy-name</i> >
Syntax (EX Series Switch)	show policy < <i>policy-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about configured routing policies.
Options	<p>none—List the names of all configured routing policies.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>policy-name</i>—(Optional) Show the contents of the specified policy.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show policy damping
List of Sample Output	<p>show policy on page 142</p> <p>show policy <i>policy-name</i> on page 144</p> <p>show policy (Multicast Scoping) on page 144</p>
Output Fields	Table 22 on page 142 lists the output fields for the show policy command. Output fields are listed in the approximate order in which they appear.

Table 22: show policy Output Fields

Field Name	Field Description
<i>policy-name</i>	Name of the policy listed.
<i>term</i>	Policy term listed.
<i>from</i>	Match condition for the policy.
<i>then</i>	Action for the policy.

show policy

```

user@host> show policy
Configured policies:
__vrf-export-red-internal__
__vrf-import-red-internal__

```

```
red-export  
all_routes
```

```
show policy user@host> show policy test-statics
policy-name Policy test-statics:
              from
                3.0.0.0/8 accept
                3.1.0.0/16 accept
              then reject
```

```
show policy (Multicast user@host> show policy test-statics
  Scoping) Policy test-statics:
              from
                multicast-scoping == 8
```


show policy conditions

Syntax	<pre>show policy conditions <condition-name> <detail> <dynamic> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switch)	<pre>show policy conditions <condition-name> <detail> <dynamic></pre>
Release Information	<p>Command introduced in Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Display all the configured conditions as well as the routing tables with which the configuration manager is interacting. If the detail keyword is included, the output also displays dependent routes for each condition.</p>
Options	<p>none—Display all configured conditions and associated routing tables.</p> <p>condition-name—(Optional) Display information about the specified condition only.</p> <p>detail—(Optional) Display the specified level of output.</p> <p>dynamic—(Optional) Display information about the conditions in the dynamic database.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show policy conditions detail on page 146
Output Fields	<p>Table 23 on page 145 lists the output fields for the show policy conditions command. Output fields are listed in the approximate order in which they appear.</p>

Table 23: show policy conditions Output Fields

Field Name	Field Description	Level of Output
Condition	Name of configured condition.	All levels
event	Condition type. If the if-route-exists option is configured, the event type is: Existence of a route in a specific routing table.	All levels
Dependent routes	List of routes dependent on the condition, along with the latest generation number.	detail
Condition tables	List of routing tables associated with the condition, along with the latest generation number and number of dependencies.	All levels

Table 23: show policy conditions Output Fields (*continued*)

Field Name	Field Description	Level of Output
If-route-exists conditions	List of conditions configured to look for a route in the specified table.	All levels

```
show policy conditions detail
user@host> show policy conditions detail
Configured conditions:
Condition cond1, event: Existence of a route in a specific routing table
Dependent routes:
  4.4.4.4/32, generation 3
  6.6.6.6/32, generation 3
  10.10.10.10/32, generation 3

Condition cond2, event: Existence of a route in a specific routing table
Dependent routes:
  None

Condition tables:
Table inet.0, generation 4, dependencies 3, If-route-exists conditions: cond1
cond2
```

test policy

Syntax	<code>test policy <i>policy-name</i> <i>prefix</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Test a policy configuration to determine which prefixes match routes in the routing table.
Options	<i>policy-name</i> —Name of a policy. <i>prefix</i> —Destination prefix to match.
Additional Information	All prefixes in the default unicast routing table (inet.0) that match prefixes that are the same as or longer than the specific prefix are processed by the from clause in the specified policy. All prefixes accepted by the policy are displayed. The test policy command evaluates a policy differently from the Border Gateway Protocol (BGP) import process. When testing a policy that contains an interface match condition in the from clause, the test policy command uses the match condition. In contrast, BGP does not use the interface match condition when evaluating the policy against routes learned from internal BGP (IBGP) or external BGP (EBGP) multihop peers.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show policy damping
List of Sample Output	test policy on page 147
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.
test policy	<pre> user@host> test policy test-statics 3.0.0.1/8 inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden) Prefixes passing policy: 3.0.0.0/8 *[BGP/170] 16:22:46, localpref 100, from 10.255.255.41 AS Path: 50888 I > to 10.11.4.32 via en0.2, label-switched-path 12 3.3.3.1/32 *[IS-IS/18] 2d 00:21:46, metric 0, tag 2 > to 10.0.4.7 via fxp0.0 3.3.3.2/32 *[IS-IS/18] 2d 00:21:46, metric 0, tag 2 > to 10.0.4.7 via fxp0.0 3.3.3.3/32 *[IS-IS/18] 2d 00:21:46, metric 0, tag 2 > to 10.0.4.7 via fxp0.0 3.3.3.4/32 *[IS-IS/18] 2d 00:21:46, metric 0, tag 2 > to 10.0.4.7 via fxp0.0 Policy test-statics: 5 prefixes accepted, 0 prefixes rejected </pre>

