



---

# Security Feature Guide for the NFX250 Network Services Platform

Release  
15.1



---

Modified: 2016-04-14

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

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.

*Security Feature Guide for the NFX250 Network Services Platform*  
15.1  
Copyright © 2016, Juniper Networks, Inc.  
All rights reserved.

The information in this document is current as of the date on the title page.

#### YEAR 2000 NOTICE

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

#### END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

# Table of Contents

	About the Documentation . . . . .	xiii
	Documentation and Release Notes . . . . .	xiii
	Using the Examples in This Manual . . . . .	xiii
	Merging a Full Example . . . . .	xiv
	Merging a Snippet . . . . .	xiv
	Documentation Conventions . . . . .	xv
	Documentation Feedback . . . . .	xvii
	Requesting Technical Support . . . . .	xvii
	Self-Help Online Tools and Resources . . . . .	xvii
	Opening a Case with JTAC . . . . .	xviii
<b>Part 1</b>	<b>Firewall Filters</b>	
<b>Chapter 1</b>	<b>Configuring Firewall Filters . . . . .</b>	<b>3</b>
	Overview of Firewall Filters . . . . .	3
	Firewall Filter Types . . . . .	4
	Firewall Filter Components . . . . .	5
	Firewall Filter Processing . . . . .	5
	How Many Filters Are Supported? . . . . .	5
	Understanding How Firewall Filters Are Evaluated . . . . .	6
	Understanding How Firewall Filters Control Packet Flows . . . . .	7
	Understanding Firewall Filter Match Conditions . . . . .	9
	Filter Match Conditions . . . . .	9
	Numeric Filter Match Conditions . . . . .	9
	Interface Filter Match Conditions . . . . .	10
	IP Address Filter Match Conditions . . . . .	10
	MAC Address Filter Match Conditions . . . . .	11
	Bit-Field Filter Match Conditions . . . . .	11
	Understanding How a Firewall Filter Tests a Protocol . . . . .	12
	Understanding Firewall Filter Planning . . . . .	13
	Planning the Number of Firewall Filters to Create . . . . .	14
	Understanding How Many Firewall Filters Are Supported . . . . .	15
	Egress Filters . . . . .	16
	Avoid Configuring too Many Filters . . . . .	16
	Configuring TCAM Error Messages . . . . .	17
	Policers can Limit Egress Filters . . . . .	17
	Planning for Filter-Specific Policers . . . . .	18
	Planning for Filter-Based Forwarding . . . . .	18

	Understanding Firewall Filter Processing Points for Bridged and Routed Packets . . . . .	19
	Configuring Firewall Filters . . . . .	20
	Configuring a Firewall Filter . . . . .	20
	Applying a Firewall Filter to a Port . . . . .	22
	Applying a Firewall Filter to a VLAN . . . . .	22
	Applying a Firewall Filter to a Layer 3 (Routed) Interface . . . . .	22
	Applying Firewall Filters to Interfaces . . . . .	23
	Monitoring Firewall Filter Traffic . . . . .	24
	Monitoring Traffic for All Firewall Filters and Policers That Are Configured . . . . .	24
	Monitoring Traffic for a Specific Firewall Filter . . . . .	25
	Monitoring Traffic for a Specific Policer . . . . .	25
	Verifying That Firewall Filters Are Operational . . . . .	25
	Troubleshooting Firewall Filters . . . . .	26
	Troubleshooting QFX10000 Switches . . . . .	26
	Do Not Combine Match Conditions for Different Layers . . . . .	26
	Layer 2 Packets Cannot be Discarded with Firewall Filters . . . . .	26
	Troubleshooting Other Switches . . . . .	27
	Firewall Filter Configuration Returns a No Space Available in TCAM Message . . . . .	27
	Filter Counts Previously Dropped Packet . . . . .	29
	Matching Packets Not Counted . . . . .	29
	Counter Reset When Editing Filter . . . . .	30
	Cannot Include loss-priority and policer Actions in Same Term . . . . .	30
	Cannot Egress Filter Certain Traffic Originating on QFX Switch . . . . .	30
	Firewall Filter Match Condition Not Working with Q-in-Q Tunneling . . . . .	31
	Egress Firewall Filters with Private VLANs . . . . .	31
	Egress Filtering of L2PT Traffic Not Supported . . . . .	32
	Cannot Drop BGP Packets in Certain Circumstances . . . . .	32
	Invalid Statistics for Policer . . . . .	32
	Policers can Limit Egress Filters . . . . .	32
<b>Part 2</b>	<b>Policers</b>	
<b>Chapter 2</b>	<b>Configuring Policers . . . . .</b>	<b>37</b>
	Overview of Policers . . . . .	37
	Policer Overview . . . . .	38
	Policer Types . . . . .	38
	Policer Actions . . . . .	39
	Policer Colors . . . . .	40
	Filter-Specific Policers . . . . .	40
	Suggested Naming Convention for Policers . . . . .	41
	Policer Counters . . . . .	41
	Policer Algorithms . . . . .	41
	How Many Policers Are Supported? . . . . .	42
	Policers Can Limit Egress Firewall Filters . . . . .	42
	Understanding Policers with Link Aggregation Groups . . . . .	43
	Understanding Color-Blind Mode for Single-Rate Tricolor Marking . . . . .	43

Understanding Color-Aware Mode for Single-Rate Tricolor Marking . . . . .	44
Summary of PLP Changes . . . . .	44
Effect on Green Packets (Low PLP) . . . . .	45
Effect on Yellow Packets (Medium PLP) . . . . .	45
Effect on Red Packets (High PLP) . . . . .	45
Understanding Color-Blind Mode for Two-Rate Tricolor Marking . . . . .	46
Understanding Color-Aware Mode for Two-Rate Tricolor Marking . . . . .	46
Summary of PLP Changes . . . . .	46
Effect on Green Packets (Low PLP) . . . . .	47
Effect on Yellow Packets (Medium PLP) . . . . .	47
Effect on Red Packets (High PLP) . . . . .	48
Example: Using Two-Color Policers and Prefix Lists . . . . .	48
Example: Using Policers to Manage Oversubscription . . . . .	51
Assigning Forwarding Classes and Loss Priority . . . . .	53
Configuring Color-Blind Egress Policers for Medium-Low PLP . . . . .	54
Configuring Two-Color and Three-Color Policers to Control Traffic Rates . . . . .	55
Configuring Two-Color Policers . . . . .	56
Configuring Three-Color Policers . . . . .	56
Specifying Policers in a Firewall Filter Configuration . . . . .	57
Applying a Firewall Filter That Includes a Policer . . . . .	57
Verifying That Two-Color Policers Are Operational . . . . .	57
Verifying That Three-Color Policers Are Operational . . . . .	58
Troubleshooting Policer Configuration . . . . .	58
Incomplete Count of Packet Drops . . . . .	59
Counter Reset When Editing Filter . . . . .	59
Invalid Statistics for Policer . . . . .	59
Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured . . . . .	60
Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured . . . . .	61
Policers Can Limit Egress Filters . . . . .	61

## Part 3

### Chapter 3

## Configuring Device Security

<b>Device Security . . . . .</b>	<b>65</b>
Understanding Storm Control . . . . .	65
Example: Configuring Storm Control to Prevent Network Outages . . . . .	66
Verifying That the Port Error Disable Setting Is Working Correctly . . . . .	68
Understanding Unicast RPF . . . . .	69
Unicast RPF for Switches Overview . . . . .	70
Unicast RPF Implementation . . . . .	71
Unicast RPF Packet Filtering . . . . .	71
Bootstrap Protocol (BOOTP) and DHCP Requests . . . . .	71
Default Route Handling . . . . .	71
When to Enable Unicast RPF . . . . .	71
When Not to Enable Unicast RPF . . . . .	72

	Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches . . . . .	73
	Configuring Unicast RPF (CLI Procedure) . . . . .	74
	Disabling Unicast RPF (CLI Procedure) . . . . .	75
	Verifying Unicast RPF Status . . . . .	76
	Understanding Unknown Unicast Forwarding . . . . .	79
	Configuring Unknown Unicast Forwarding (CLI Procedure) . . . . .	79
	Configuring Unknown Unicast Forwarding on EX4300 Switches . . . . .	80
	Configuring Unknown Unicast Forwarding on EX9200 Switches . . . . .	80
<b>Part 4</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 4</b>	<b>Configuration Statements (Firewall Filters) . . . . .</b>	<b>85</b>
	family . . . . .	86
	filter . . . . .	87
	filter (Layer 2 and Layer 3 Interfaces) . . . . .	88
	filter (VLANs) . . . . .	89
	firewall . . . . .	90
	from . . . . .	91
	input (Forwarding Table) . . . . .	92
	interface-specific . . . . .	92
	output (Forwarding Table) . . . . .	93
	term . . . . .	94
	then (Filters) . . . . .	95
<b>Chapter 5</b>	<b>Configuration Statements (Policers) . . . . .</b>	<b>97</b>
	action . . . . .	98
	bandwidth-limit . . . . .	98
	burst-size-limit . . . . .	99
	color-aware . . . . .	100
	color-blind . . . . .	101
	committed-burst-size . . . . .	102
	committed-information-rate . . . . .	103
	excess-burst-size . . . . .	104
	filter-specific . . . . .	105
	firewall . . . . .	106
	if-exceeding . . . . .	107
	loss-priority high then discard (Three-Color Policer) . . . . .	108
	peak-burst-size . . . . .	109
	peak-information-rate . . . . .	110
	policer . . . . .	111
	single-rate . . . . .	112
	then (Policers) . . . . .	113
	three-color-policer . . . . .	114
	two-rate . . . . .	115
<b>Chapter 6</b>	<b>Operational Commands (Firewall Filters) . . . . .</b>	<b>117</b>
	clear firewall . . . . .	118
	show firewall . . . . .	119
	show firewall policer . . . . .	123

show interfaces filters . . . . .	125
show pfe filter hw summary . . . . .	127





# List of Figures

<b>Part 1</b>	<b>Firewall Filters</b>	
<b>Chapter 1</b>	<b>Configuring Firewall Filters .....</b>	<b>3</b>
	Figure 1: Evaluation of Terms Within a Firewall Filter .....	7
	Figure 2: Application of Firewall Filters to Control Packet Flow .....	8
<b>Part 2</b>	<b>Policers</b>	
<b>Chapter 2</b>	<b>Configuring Policers .....</b>	<b>37</b>
	Figure 3: Flow of Tricolor Marking Policer Operation .....	38
<b>Part 3</b>	<b>Configuring Device Security</b>	
<b>Chapter 3</b>	<b>Device Security .....</b>	<b>65</b>
	Figure 4: Symmetrically Routed Interfaces .....	72
	Figure 5: Asymmetrically Routed Interfaces .....	73



# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xiii</b>
	Table 1: Notice Icons . . . . .	xv
	Table 2: Text and Syntax Conventions . . . . .	xv
<b>Part 1</b>	<b>Firewall Filters</b>	
<b>Chapter 1</b>	<b>Configuring Firewall Filters</b> . . . . .	<b>3</b>
	Table 3: Supported Firewall Filter Numbers for Specific Switches . . . . .	5
	Table 4: Actions for Firewall Filters . . . . .	12
	Table 5: Supported Firewall Filter Numbers . . . . .	15
<b>Part 2</b>	<b>Policers</b>	
<b>Chapter 2</b>	<b>Configuring Policers</b> . . . . .	<b>37</b>
	Table 6: Policer Actions . . . . .	39
	Table 7: Color-Blind Mode TCM Color-to-PLP Mapping . . . . .	44
	Table 8: Color-Aware Mode Single-Rate PLP Mapping . . . . .	44
	Table 9: Color-Blind Mode TCM Color-to-PLP Mapping . . . . .	46
	Table 10: Color-Aware Mode Two-Rate PLP Mapping . . . . .	46
	Table 11: Servers Connected to Switch . . . . .	51
	Table 12: Unicast Forwarding Classes . . . . .	53
<b>Part 4</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 6</b>	<b>Operational Commands (Firewall Filters)</b> . . . . .	<b>117</b>
	Table 13: show firewall Output Fields . . . . .	119
	Table 14: show firewall policer Output Fields . . . . .	123
	Table 15: show interfaces filters Output Fields . . . . .	125
	Table 16: show pfe filter hw summary Output Fields . . . . .	127



# About the Documentation

- Documentation and Release Notes on page xiii
- Using the Examples in This Manual on page xiii
- Documentation Conventions on page xv
- Documentation Feedback on page xvii
- Requesting Technical Support on page xvii

## Documentation and Release Notes

---

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

## Using the Examples in This Manual

---

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

## Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

## Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

## Documentation Conventions

Table 1 on page xv defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b> No alarms currently active

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<i>Italic text like this</i>	<ul style="list-style-type: none"><li>Introduces or emphasizes important new terms.</li><li>Identifies guide names.</li><li>Identifies RFC and Internet draft titles.</li></ul>	<ul style="list-style-type: none"><li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li><li><i>Junos OS CLI User Guide</i></li><li>RFC 1997, <i>BGP Communities Attribute</i></li></ul>
<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>
<b>Text like this</b>	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"><li>To configure a stub area, include the <b>stub</b> statement at the [edit protocols ospf area area-id] hierarchy level.</li><li>The console port is labeled <b>CONSOLE</b>.</li></ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub</b> <default-metric <i>metric</i> >;
(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.	<b>broadcast   multicast</b>  ( <i>string1</i>   <i>string2</i>   <i>string3</i> )
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [ community-ids ]</b>
Indentation and braces ( { } )	Identifies 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.	
<b>GUI Conventions</b>		
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"><li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li><li>To cancel the configuration, click <b>Cancel</b>.</li></ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .



## Documentation Feedback

---

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software version (if applicable).

## 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 Partner Support Service 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: <http://kb.juniper.net/InfoCenter/>

- 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

# Firewall Filters

- [Configuring Firewall Filters on page 3](#)



## CHAPTER 1

# Configuring Firewall Filters

- [Overview of Firewall Filters on page 3](#)
- [Understanding How Firewall Filters Are Evaluated on page 6](#)
- [Understanding How Firewall Filters Control Packet Flows on page 7](#)
- [Understanding Firewall Filter Match Conditions on page 9](#)
- [Understanding How a Firewall Filter Tests a Protocol on page 12](#)
- [Understanding Firewall Filter Planning on page 13](#)
- [Planning the Number of Firewall Filters to Create on page 14](#)
- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 19](#)
- [Configuring Firewall Filters on page 20](#)
- [Applying Firewall Filters to Interfaces on page 23](#)
- [Monitoring Firewall Filter Traffic on page 24](#)
- [Verifying That Firewall Filters Are Operational on page 25](#)
- [Troubleshooting Firewall Filters on page 26](#)

## Overview of Firewall Filters

---

**Supported Platforms** [EX4600, QFabric System, QFX Series standalone switches](#)

Firewall filters provide rules that define whether to accept or discard packets that are transiting an interface. If a packet is accepted, you can configure additional actions to perform on the packet, such as 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). You configure firewall filters to determine whether to accept or discard a packet before it enters or exits any of these:

- Port
- VLAN
- Layer 3 (routed) interface
- Routed VLAN interface (RVI)

An *ingress* firewall filter is applied to packets that are entering an interface or VLAN, and an *egress* firewall filter is applied to packets that are exiting an interface or VLAN.



**NOTE:** Firewall filters are sometimes called *access control lists (ACLs)*.

- [Firewall Filter Types on page 4](#)
- [Firewall Filter Components on page 5](#)
- [Firewall Filter Processing on page 5](#)
- [How Many Filters Are Supported? on page 5](#)

## Firewall Filter Types

The following firewall filter types are supported:

- Port (Layer 2) firewall filter—Port firewall filters apply to Layer 2 traffic transiting system ports.
- VLAN firewall filter—VLAN firewall filters provide access control for packets that enter a VLAN, are bridged within a VLAN, or leave a VLAN.
- Router (Layer 3) firewall filter—You can apply a router firewall filter in both ingress and egress directions on IPv4 or IPv6 Layer 3 (routed) interfaces, routed VLAN interfaces (RVI) and a loopback interface, which filters traffic sent to the switch itself or generated by the switch. (You apply a filter to a loopback interface in the input direction to protect the switch from unwanted traffic. You also might want to apply a filter to a loopback interface in the output direction so that you can set the forwarding class and DSCP bit value for packets that originate on the switch itself. This feature gives you very fine control over the classification of CPU generated packets. For example, you might want to assign different DSCP values and forwarding classes to traffic generated by different routing protocols so the traffic for those protocols can be treated in a differentiated manner by other devices. You can apply a filter to a loopback interface in the output direction starting with Junos OS 13.2X51-D15.)



**NOTE:** You can apply a firewall filter to a management interface (for example, `me0`) on a QFX and EX4600 standalone switch. You cannot apply a firewall filter to a management interface on a QFX3000-G or QFX3000-M system.

- MPLS filter—You can apply a firewall filter to an MPLS interface

To apply a firewall filter:

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



**NOTE:** You can apply only one firewall filter to a port, VLAN, or interface for a given direction. For example, for interface `ge-0/0/6.0`, you can apply one filter for the ingress direction and one for the egress direction.

## Firewall Filter Components

In a firewall filter, you first define the family address type (ethernet-switching, inet (for IPv4), inet6 (for IPv6), or mpls) and then define one or more terms that specify the filtering criteria and the action to take if a match occurs.

Each term consists of the following components:

- **Match conditions**—Specify values that a packet must contain to be considered a match. You can specify values for most fields in the IP, TCP, UDP, or ICMP headers. You can also match on interface names.
- **Action**—Specifies what to do if a packet matches the match conditions. A filter can accept, discard, or reject a matching packet and then perform additional actions, such as counting, classifying, and policing. If no action is specified for a term, the default is to accept the matching packet.

## Firewall Filter Processing

If there are multiple terms in a filter, the order of the terms is important. If a packet matches the first term, the switch executes the action defined by that term, and no other terms are evaluated. If the switch does not find a match between the packet and the first term, it compares the packet to the next term. If no match occurs between the packet and the second term, the system continues to compare the packet to each successive term in the filter until a match is found. If the packet does not match any terms in the filter, the switch discards the packet by default.

## How Many Filters Are Supported?

QFX10000 switches support 8K firewall filters and 64K firewall filter terms.

QFX3500, QFX3600, QFX5100, and EX4600 switches, QFabric Node devices, and VCF members support the maximum numbers of firewall filter terms per type of attachment point shown in [Table 3 on page 5](#).

**Table 3: Supported Firewall Filter Numbers for Specific Switches**

Filter Type	QFX3500, QFX3600	QFX5100, EX4600
Ingress	768	1536
Egress	1024	1024

These totals are applied in aggregate. For example, on the QFX3500 and QFX3600 you can apply a total of 768 terms in all your port filters, Layer 3 filters, and VLAN filters that are applied in the input direction and 1024 terms in port filters, Layer 3 filters, and VLAN filters that are applied in the output direction. The actual number of filters that these

switches will support depends on how the filters are stored in ternary content addressable memory (TCAM). See [“Planning the Number of Firewall Filters to Create” on page 14](#) for detailed information about this topic.

**Related  
Documentation**

- [Understanding Firewall Filter Planning on page 13](#)
- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 19](#)
- [Understanding How Firewall Filters Are Evaluated on page 6](#)
- [Understanding Firewall Filter Match Conditions on page 9](#)
- [Overview of Policers on page 37](#)
- [Configuring Firewall Filters on page 20](#)

---

## Understanding How Firewall Filters Are Evaluated

---

**Supported Platforms**    EX4600, OCX1100, QFabric System, QFX Series standalone switches

A firewall filter consists of one or more terms, and the order of the terms within a filter is important. Before you configure firewall filters, you should understand how switches evaluate the terms within a 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.
- If the packet does not match all the conditions, the switch discards it.

When a firewall filter consists of more than one term, the 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 the packet matches all 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 switch discards it.



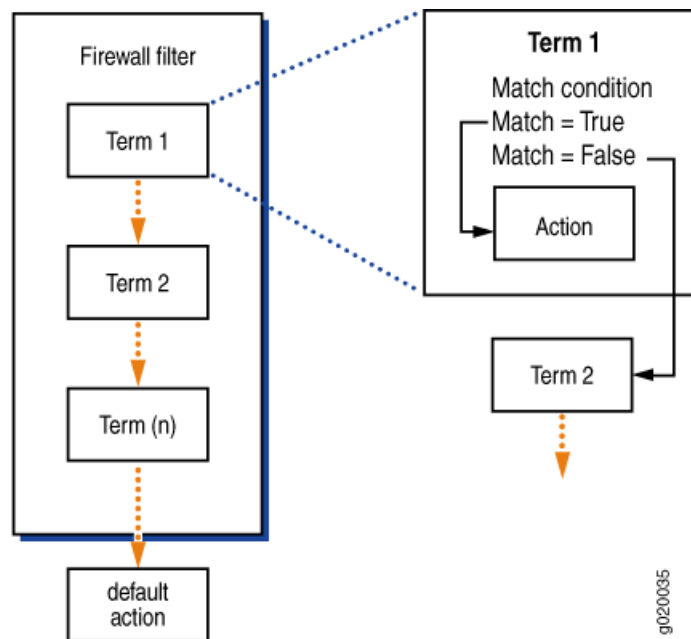
**NOTE:** The order of conditions in a **from** statement is not important because a packet must match all the conditions to be considered a match.

---

[Figure 1 on page 7](#) shows how switches evaluate the terms within a firewall filter.



Figure 1: Evaluation of Terms Within a Firewall Filter



If you do not include a **from** statement in a term, all packets will match the term and be processed by the **then** statement. If a term does not contain a **then** statement or if an action has not been configured in the **then** statement, the term accepts any matching packets.

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, a packet that does not match any of the terms in a firewall filter is discarded. If you configure a 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 64 bytes long.

#### Related Documentation

- [Understanding Firewall Filter Match Conditions on page 9](#)
- [Overview of Policers on page 37](#)
- [Configuring Firewall Filters on page 20](#)

## Understanding How Firewall Filters Control Packet Flows

**Supported Platforms** EX4600, OCX1100, QFabric System, QFX Series standalone switches

A switch supports firewall filters that allow you to control flows of data packets and local packets. *Data packets* transit a switch as they are forwarded from a source to a destination. *Local packets* are destined for or sent by a Routing Engine (they do not transit a switch). Local packets usually contain routing protocol data, data for IP services such as Telnet or SSH, or data for administrative protocols such as the Internet Control Message Protocol (ICMP).

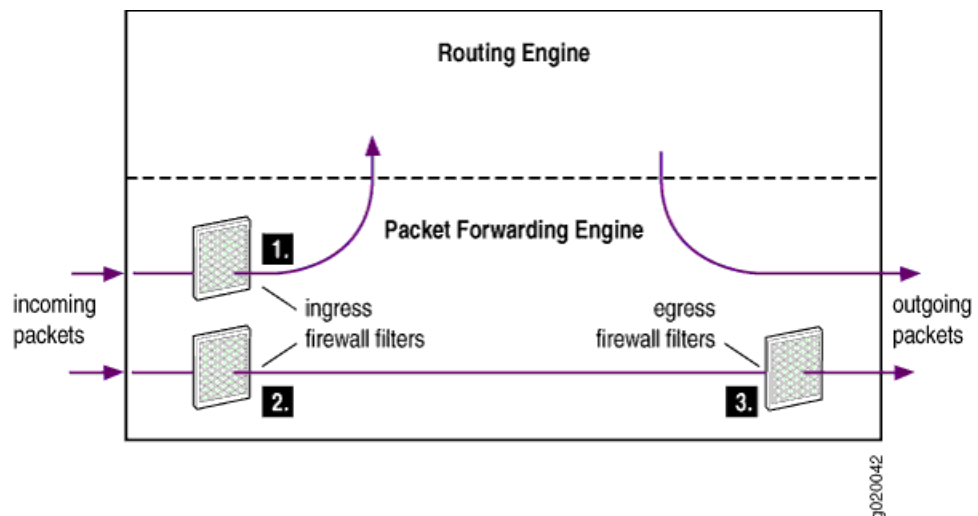
Firewall filters affect packet flows entering into or exiting from a switch as follows:

- Ingress firewall filters affect the flow of data packets that are received on switch interfaces. When a switch receives a data packet, the Packet Forwarding Engine in the system that contains the ingress interface determines where to forward the packet by looking in its Layer 2 or Layer 3 forwarding table for the best route to the destination. Data packets are forwarded to an egress interface. Locally destined packets are forwarded to the Routing Engine.
- Egress firewall filters affect data packets that are transiting a switch but do not affect packets sent by the Routing Engine. These filters are applied by the Packet Forwarding Engine in the system that contains the egress interface.

Figure 2 on page 8 illustrates the application of ingress and egress firewall filters to control the flow of packets through a switch:

1. Ingress firewall filter applied to locally destined packets that are received on switch interfaces and are destined for the Routing Engine.
2. Ingress firewall filter applied to data packets that are received on switch interfaces and will transit the switch.
3. Egress firewall filter applied to data packets that are transiting the switch.

Figure 2: Application of Firewall Filters to Control Packet Flow



#### Related Documentation

- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 19](#)

- [Understanding How Firewall Filters Are Evaluated on page 6](#)
- [Configuring Firewall Filters on page 20](#)

## Understanding Firewall Filter Match Conditions

**Supported Platforms** [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

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

- [Filter Match Conditions on page 9](#)
- [Numeric Filter Match Conditions on page 9](#)
- [Interface Filter Match Conditions on page 10](#)
- [IP Address Filter Match Conditions on page 10](#)
- [MAC Address Filter Match Conditions on page 11](#)
- [Bit-Field Filter Match Conditions on page 11](#)

### 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 must match for the action to be implemented. 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 multiple values for the same condition, a match on any one of those values matches that condition. For example, if you specify multiple IP source addresses using the **source-address** statement, a packet that contains any one of those IP source addresses matches the condition. In some cases you can specify multiple values for the same condition by enclosing the possible values in square brackets, as in:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set protocol (icmp | udp)
```

In other cases you must enter multiple statements, as in:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-address 10.1.1.1
user@switch# set source-address 10.1.1.2
```

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



**NOTE:** Unlike traditional Junos OS firewall filters, you cannot use **except** in a condition statement to negate the condition.

### Numeric Filter Match Conditions

You can specify numeric filter match conditions that are identified by a numeric value, such as port and protocol numbers. For numeric filter match conditions, you specify the

condition and a single value that a field in a packet must contain to be considered a 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, to match Telnet traffic:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port 23
```

- 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, to match Telnet traffic:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port telnet
```

- To specify multiple values for the same match condition in a filter term, enter each value in its own match statement. For example, a match occurs in the following term if the value of the source port in the packet is 22 or 23.

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port 22
user@switch# set source-port 23
```

## Interface Filter Match Conditions

You can specify an interface filter match condition to match an interface on which a packet is received or transmitted. For example, if you apply a filter to a VLAN you might want the filter to match on some interfaces that participate in the VLAN and not match on other interfaces in the VLAN. When you specify the name of the interface, you must include a logical unit.

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set interface ge-0/0/6.0
```

In this example, the final character (0) specifies the logical unit. 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@switch# set interface ge-0/*/6.0
user@switch# set interface ge-0/1/*0
user@switch# set interface ge-0/0/6.*
```

Note that you must specify a value or a wildcard for the logical unit.

## IP Address Filter Match Conditions

You can specify an address filter match condition to match an IP source or destination address or prefix in a packet. Specify the address or prefix type and the address or prefix itself. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-address 10.2.1.0/24;
```

If you omit the prefix length, it defaults to /32. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-address 10
[edit firewall family family-name filter filter-name term term-name from]
user@switch# show
destination-address {
```

```
10.0.0.0/32;
}
```

To specify more than one IP address or prefix in a filter term, enter each address or prefix in its own match statement. For example, a match occurs in the following term if the source address of a packet matches either of the following prefixes:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-address 10.1.0.0/16
user@switch# set source-address 10.2.0.0/16
```

## MAC Address Filter Match Conditions

You can specify a MAC address filter match condition to match a source or destination MAC address. You specify the address type and value that a packet must contain to be considered a match.

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

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

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

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

Regardless of the formats you use, the system resolves the address to the standard format, in this case 00:11:22:33:44:55.

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

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

## Bit-Field Filter Match Conditions

You can specify bit-field filter match conditions to match particular bits within certain fields in Ethernet frames and IP, TCP, UDP, and ICMP headers. You usually specify the field and the bit within the field that must be set in a packet to be considered a match.

In most cases you can use a keyword to specify the bit you want to match on. For example, to match on a TCP SYN packet you can enter **syn**, as in:

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

You can also enter **0x02** because the SYN bit is the third least-significant bit of the 8-bit tcp-flags field:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags 0x02
```

To match multiple bit-field values, use the logical operators, which are described in [Table 4 on page 12](#). The operators are listed in order from highest precedence to lowest precedence. Operations are evaluated from left to right.

Table 4: Actions for Firewall Filters

Logical Operators	Description
!	Negation
&	Logical AND
	Logical OR

If you use a logical operator, enclose the values in quotation marks and do not include any spaces. For example, the following statement matches the second packet of a TCP handshake:

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

To negate a match, precede the value with an exclamation point. For example, the following statement matches only the initial packet of a TCP handshake:

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

You can use text synonyms to specify some common bit-field matches. For example, the following statement also matches the initial packet of a TCP handshake:

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

#### Related Documentation

- [Understanding How a Firewall Filter Tests a Protocol on page 12](#)
- [Firewall Filter Match Conditions and Actions](#)
- [Configuring Firewall Filters on page 20](#)

## Understanding How a Firewall Filter Tests a Protocol

**Supported Platforms** EX4600, OCX1100, QFabric System, QFX Series standalone switches

When examining match conditions in a firewall filter, a switch tests only the fields that you specify. It does not implicitly test any fields that you do not explicitly configure. For example, if you specify a match condition of **source-port ssh**, there is no implied test to determine if the protocol is TCP. In this case, the switch considers any packet that has a value of **22** (decimal) in the 2-byte field that follows a *presumed* IP header to be a match. To ensure that the term matches on TCP packets, you also specify an **ip-protocol tcp** match condition.

For the following match conditions, you should explicitly specify the protocol match condition in the same term:

- **destination-port**—Specify protocol **tcp** or protocol **udp**.
- **icmp-code**—Specify protocol **icmp** and **icmp-type**.
- **icmp-type**—Specify protocol **icmp** or protocol **icmp6**.

- **source-port**—Specify protocol **tcp** or protocol **udp**.
- **tcp-flags**—Specify protocol **tcp**.

**Related  
Documentation**

- [Understanding Firewall Filter Match Conditions on page 9](#)
- [Configuring Firewall Filters on page 20](#)

---

## Understanding Firewall Filter Planning

---

**Supported Platforms**    [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

Before you create a firewall filter and apply it, determine what you want the filter to accomplish and how to use its match conditions and actions to achieve your goals. It is important that you understand how packets are matched, the default and configured actions of the firewall filter, and where to apply the firewall filter.

You can apply no more than one firewall filter per port, VLAN, or router interface per direction (input and output). For example, for a given port you can apply at most one filter in the input direction and one filter in the output direction. 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 operation and can make testing and troubleshooting more difficult.

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

1. What is the purpose of the filter?

For example, the system can drop packets based on header information, rate-limit traffic, classify packets into forwarding classes, log and count packets, or prevent denial-of-service attacks.

2. What are the appropriate match conditions? 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, 802.1Q tag, Ethernet type, or VLAN.
- Layer 3 header fields—Source and destination IP addresses, protocols, and IP options (IP precedence, IP fragmentation flags, or TTL type).
- TCP header fields—Source and destination ports and flags.
- ICMP header fields—Packet type and code.

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

The system can accept, discard, or reject packets.

4. What additional action modifiers might be required?

For example, you can configure the system to mirror (copy) packets to a specified port, count matching packets, apply traffic management, or police packets.

5. On what port, router interface, or VLAN should the firewall filter be applied?

Start with the following basic guidelines:

- If packets entering or leaving a Layer 2 interface (port) need to be filtered, apply the filter at the **[edit family ethernet switching filter]** hierarchy level. This is a port filter.
- If packets entering or leaving any port in a specific VLAN need to be filtered, use a VLAN filter.
- If packets entering or leaving a Layer 3 (routed) interface or routed VLAN interface (RVI) need to be filtered, use a router firewall filter. Apply the filter to the interface at the **[edit family inet]** hierarchy level. You can also apply a router firewall filter on a loopback interface.

Before you choose the interface or VLAN on which to apply a firewall filter, understand how that placement can affect traffic flow to other interfaces. In general, apply a filter close to the source device if the filter matches on source or destination IP addresses, IP protocols, or protocol information—such as ICMP message types, and TCP or UDP port numbers. However, you should apply a filter close to the destination device if the filter matches *only* on a source IP address. When you apply a filter too close to the source device, the filter could 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 typically configure different actions for traffic entering an interface than you configure for traffic exiting an interface.

7. How many filters should I create?

See “[Planning the Number of Firewall Filters to Create](#)” on page 14 for information about how many firewall filters you can apply.

**Related  
Documentation**

- [Overview of Policers on page 37](#)
- [Understanding How Firewall Filters Are Evaluated on page 6](#)
- [Configuring Firewall Filters on page 20](#)

---

## Planning the Number of Firewall Filters to Create

---

**Supported Platforms**    [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

- [Understanding How Many Firewall Filters Are Supported on page 15](#)
- [Egress Filters on page 16](#)
- [Avoid Configuring too Many Filters on page 16](#)
- [Configuring TCAM Error Messages on page 17](#)



- [Policers can Limit Egress Filters on page 17](#)
- [Planning for Filter-Specific Policers on page 18](#)
- [Planning for Filter-Based Forwarding on page 18](#)

## Understanding How Many Firewall Filters Are Supported

**Supported Platforms** [OCX1100, QFabric System, QFX Series standalone switches](#)

QFX3500, QFX3600, QFX5100, and EX4600 switches, QFabric Node devices, and VCF members support the maximum numbers of firewall filter terms per type of attachment point shown in [Table 5 on page 15](#).

**Table 5: Supported Firewall Filter Numbers**

Filter Type	QFX3500, QFX3600	QFX5100, EX4600
Ingress	768	1536
Egress	1024	1024

These totals are applied in aggregate. For example, on the QFX3500 and QFX3600 you can apply a total of 768 terms in all your port filters, Layer 3 filters, and VLAN filters that are applied in the input direction and 1024 terms in port filters, Layer 3 filters, and VLAN filters that are applied in the output direction.



**NOTE:** If you want to create more than 512 egress VLAN filters, your first VLAN ID should be 6 and the subsequent VLAN IDs should increase by 1. For example, to create 1024 egress VLAN filters, the first VLAN ID would be 6, the second ID would be 7, and the sequence would continue through VLAN ID 1029. Similarly, if you want to create fewer than 512 egress VLAN filters but want the total number of terms in those filters to exceed 512, you should number your VLAN IDs in the same manner. If you do not use this approach to create your VLAN IDs, the total number of allowed terms or filters will be less than 1024 and might be 512.

The ternary content addressable memory (TCAM) for firewall filters is divided into slices that accommodate 256 terms, and all the terms in a memory slice must be in filters of the same type and applied in the same direction. A memory slice is reserved as soon as you commit a filter. For example, if you create a port filter and apply it in the input direction, a memory slice is reserved that will only store ingress port filters. If you create and apply only one ingress port filter and that filter has only one term, the rest of this slice is unused and is unavailable for other filter types.

Continuing with the above example for QFX3500 and QFX3600 switches, assume that you create and apply 256 ingress port filters with one term each so that one memory slice is filled. This leaves two more memory slices available for ingress filters. (Remember that the maximum number of ingress terms is 768.) If you then create and apply an ingress Layer 3 filter with one term, another memory slice is reserved for ingress Layer 3

filters. As before, the rest of the slice is unused and is unavailable for different filter types. At this point there is one memory slice available for any ingress filter type.

Now assume that you create and apply a VLAN ingress filter. The final memory slice is reserved for VLAN ingress filters. Memory allocation for ingress filters (once again assuming one term per filter) is as follows:

- Slice 1: Filled with 256 ingress port filters. You cannot commit any more ingress port filters.
- Slice 2: Contains one ingress Layer 3 filter with one term. You can commit 255 more terms in ingress Layer 3 filters.
- Slice 3: Contains one ingress VLAN filter with one term. You can commit 255 more terms in ingress VLAN filters.

Here is another example for QFX3500 and QFX3600 switches. Assume that you create 257 ingress port filters with one term per filter—that is, you create one more term than a single memory slice can accommodate. When you apply the filters and commit the configuration, the filter memory allocation is:

- Slice 1: Filled with 256 ingress port filters. You cannot apply any more ingress port filters.
- Slice 2: Contains one ingress port filter. You can apply 255 more terms in ingress port filters.
- Slice 3: This slice is unassigned. You can create and apply 256 terms in ingress filters of any type (port, Layer 3, or VLAN), but all the filters must be of the same type.

## Egress Filters

All of the preceding principles also apply to egress filters, but four memory slices are used because IPv4 Layer 3 filters and IPv6 Layer 3 filters are stored in separate slices. The memory slices for egress filters are the same size as those for ingress filters, so the maximum number of egress filter terms is therefore 1024.

## Avoid Configuring too Many Filters

If you violate any of these restrictions and commit a configuration that is not in compliance, Junos OS rejects the excessive filters. For example, if you configure 300 ingress port filters and 300 ingress Layer 3 filters and try to commit the configuration, Junos OS does the following (again assuming one term per filter):

- Accepts the 300 ingress port filters (storing them in two memory slices).
- Accepts the first 256 ingress Layer 3 filters it processes (storing them in the third memory slice).
- Rejects the remaining 44 ingress Layer 3 filters.



**NOTE:** In this situation, be sure to delete excessive filters (for example, the remaining 44 ingress Layer 3 filters) from the configuration before you reboot the device. If you reboot a device that has a noncompliant configuration, you cannot predict which filters are installed after the reboot. Using the example above, the 44 ingress Layer 3 filters that were originally rejected might be installed, and 44 of the port filters that were originally accepted might be rejected.

---

## Configuring TCAM Error Messages

You can configure your switch to display error messages if a filter cannot be installed because there isn't enough TCAM space available. To have TCAM error messages sent to a syslog file, enter

```
set system syslog file filename pfe emergency
```

To have TCAM error messages sent to the console, enter

```
set system syslog console pfe emergency
```

To have TCAM error messages sent to an SSH terminal session, enter

```
set system syslog user user-login pfe emergency
```

## Policers can Limit Egress Filters

On some switches, the number of egress policers that you configure can affect the total number of allowed egress firewall filters. (This issue does not affect QFX10000 switches.) On the affected switches, every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:

- Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
- Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem from occurring by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

## Planning for Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in ternary content addressable memory (TCAM). If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented above to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

## Planning for Filter-Based Forwarding

You 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—called filter-based forwarding—you specify a filter and match criteria and then specify the virtual routing instance to send packets to. Filters used in this way also consume memory in an additional TCAM. See *Understanding FIP Snooping, FBF, and MVR Filter Scalability* for more information. The section *FBF Filter VFP TCAM Consumption* in this topic specifically addresses the number of supported filters when using filter-based forwarding.



**WARNING:** Filter-based forwarding does not work with IPv6 interfaces on some Juniper switches.

**Related  
Documentation**

- [Understanding How Firewall Filters Are Evaluated on page 6](#)
- [Understanding Firewall Filter Planning on page 13](#)
- [Configuring Firewall Filters on page 20](#)
- [Understanding Filter-Based Forwarding](#)

---

## Understanding Firewall Filter Processing Points for Bridged and Routed Packets

---

**Supported Platforms**    [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

You apply firewall filters at multiple processing points in the forwarding path. At each processing point, the action to be taken on a packet is determined by the configuration of the filter and the results of the lookup in the forwarding or routing table.

For both bridged (Layer 2) unicast packets and routed (Layer 3) unicast packets, firewall filters are applied in the prescribed order shown below (assuming that each filter is present and a packet is accepted by each one).

Bridged packets:

1. Ingress port filter
2. Ingress VLAN filter
3. Egress VLAN filter
4. Egress port filter

Routed packets:

1. Ingress port firewall filter
2. Ingress VLAN firewall filter (Layer 2 CoS)
3. Ingress router firewall filter (Layer 3 CoS)
4. Egress router firewall filter
5. Egress VLAN firewall filter
6. Egress port filter



**NOTE:** MAC learning occurs before filters are applied, so switches learn the MAC addresses of packets that are dropped by ingress filters.

**Related  
Documentation**

- [Overview of Firewall Filters on page 3](#)

- [Understanding How Firewall Filters Control Packet Flows on page 7](#)
- [Configuring Firewall Filters on page 20](#)

## Configuring Firewall Filters

---

**Supported Platforms** [EX4600, QFabric System, QFX Series standalone switches](#)

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

- [Configuring a Firewall Filter on page 20](#)
- [Applying a Firewall Filter to a Port on page 22](#)
- [Applying a Firewall Filter to a VLAN on page 22](#)
- [Applying a Firewall Filter to a Layer 3 \(Routed\) Interface on page 22](#)

### Configuring a Firewall Filter

To configure a firewall filter:

1. Configure the family address type, filter name, term name, and at least one match condition—for example, match on packets that contain a specific source address:

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

For a firewall filter that is applied to a port or VLAN, specify the family address type **ethernet-switching**. For a firewall filter that is applied to a Layer 3 (routed) interface, specify the family address type **inet**.

The filter and term names can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. Each filter name must be unique. A filter can contain one or more terms, and each term name must be unique within a filter.

2. Configure additional match conditions. For example, match on packets that contain a specific source port:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one from]
user@switch# set 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, it cannot be empty. If you omit the **from** statement, all packets are considered to match.

3. If you want to apply a firewall filter to multiple interfaces and be able to see 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. 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 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 then]
user@switch# set discard
```

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

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

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

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

- **analyzer *analyzer-name***—Mirror port traffic to a specified analyzer, which you must configure at the **[ethernet-switching-options]** level.
- **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.



**NOTE:** On QFX3500 and QFX3600 switches, filters automatically count packets that have been dropped on ingress because of cyclic redundancy check (CRC) errors.

- **forwarding-class *class***—Assign packets to a forwarding class.
- **log**—Log the packet header information in the Routing Engine.
- **loss-priority *priority***—Set the priority of dropping a packet.
- **policer *policer-name***—Apply rate-limiting to the traffic.
- **syslog**—Log an alert for this packet.

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 should always explicitly configure an action 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

To apply a firewall filter to an ingress port:

1. Provide a meaningful description of the firewall filter in the configuration of the port to which the filter will be applied:

```
[edit]
user@switch# set interfaces ge-0/0/6 description "filter to limit tcp traffic at trunk port for employee-vlan"
```

2. Apply the filter to the interface, specifying the unit number, family address type, the direction of the filter (for packets entering the port), and the filter name:

```
[edit]
user@switch# set ge-0/0/6 unit 0 family ethernet-switching filter input ingress-port-filter
```

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



**NOTE:** You can apply only one filter to a port for a given direction (ingress or egress).

## Applying a Firewall Filter to a VLAN

To apply a firewall filter to a VLAN:

1. Provide a meaningful description of the firewall filter in the configuration of the VLAN to which the filter will be applied:

```
[edit]
user@switch# set vlans employee-vlan vlan-id 20 description "filter to block rogue devices on employee-vlan"
```

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

- To apply a filter to match packets that are entering the VLAN:

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

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

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



**NOTE:** You can apply only one filter to a VLAN for a given direction (ingress or egress).

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

To apply a firewall filter to a Layer 3 routed interface:

1. Provide a meaningful description of the firewall filter in the configuration of the interface to which the filter will be applied:



```
[edit]
```

```
user@switch# set interfaces ge-0/1/6 description "filter to count and monitor traffic on layer 3 interface"
```

2. You can apply firewall filters to filter packets that enter or exit a Layer 3 routed interface:

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

```
[edit]
```

```
user@switch# set interfaces ge-0/1/6 unit 0 family inet filter input ingress-router-filter
```

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

```
[edit]
```

```
user@switch# set interfaces ge-0/1/6 unit 0 family inet filter output egress-router-filter
```

For firewall filters applied to Layer 3 routed interfaces, the family address type must be **inet**.



**NOTE:** You can apply only one filter to an interface for a given direction (ingress or egress).

#### Related Documentation

- [Overview of Firewall Filters on page 3](#)
- [Firewall Filter Match Conditions and Actions](#)
- [Verifying That Firewall Filters Are Operational on page 25](#)
- [Monitoring Firewall Filter Traffic on page 24](#)
- [Configuring Port Mirroring](#)

## Applying Firewall Filters to Interfaces

**Supported Platforms** [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

For a firewall filter to work, you must apply it to at least one interface. To do this, include the **filter** statement when configuring a logical interface at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
```

```
user@switch# set interface-name unit logical-unit-number family family-name filter (input | output) filter-name
```

In the **input** statement, specify a firewall filter to be evaluated when packets are received on the interface. Input filters applied to a loopback interface affect only traffic destined for the Routing Engine.

In the **output** statement, specify a filter to be evaluated when packets exit the interface.



**NOTE:** When you create a loopback interface, it is important to apply an ingress filter to it so the Routing Engine is protected. We recommend that when you apply a filter to the loopback interface lo0, you include the `apply-groups` statement. Doing so ensures that the filter is automatically inherited on every loopback interface, including lo0 and other loopback interfaces.

**Related Documentation**

- [Configuring Firewall Filters on page 20](#)

## Monitoring Firewall Filter Traffic

**Supported Platforms** EX4600, OCX1100, QFabric System, QFX Series standalone switches

You can use operational mode commands to monitor firewall filter traffic.

- [Monitoring Traffic for All Firewall Filters and Policers That Are Configured on page 24](#)
- [Monitoring Traffic for a Specific Firewall Filter on page 25](#)
- [Monitoring Traffic for a Specific Policier on page 25](#)

### Monitoring Traffic for All Firewall Filters and Policers That Are Configured

**Purpose** 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 `show firewall` operational mode command:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                               Bytes          Packets
counter-employee-web               3348            27
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                               Bytes          Packets
icmp-counter                       560             10
Policers:
Name                               Packets
icmp-connection-policer            10
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, counters, and policers that are configured. 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.

## Monitoring Traffic for a Specific Firewall Filter

**Purpose** Monitor the number of packets and bytes that matched a firewall filter and monitor the number of packets that exceeded policer rate limits.

**Action** Use the **show firewall filter *filter-name*** operational mode command:

```
user@switch> show firewall filter ingress-port-limit-tcp-icmp
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                                     Bytes      Packets
icmp-counter                             560         10
```

**Meaning** The **show firewall filter *filter-name*** command limits the display information to the counters and policers that are defined for the specified filter.

## Monitoring Traffic for a Specific Policer

**Purpose** Monitor the number of packets that exceeded the rate limits of a policer:

**Action** Use the **show firewall policer *policer-name*** operational mode command:

```
user@switch> show firewall policer icmp-connection-policer
Filter: ingress-port-limit-tcp-icmp
Policers:
Name                                     Packets
icmp-connection-policer                  10
```

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

**Related Documentation**

- [Configuring Firewall Filters on page 20](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55](#)
- [Verifying That Firewall Filters Are Operational on page 25](#)

## Verifying That Firewall Filters Are Operational

**Supported Platforms** EX4600, OCX1100, QFabric System, QFX Series standalone switches

**Purpose** Verify that firewall filters are working properly.

**Action** Use the **show firewall** operational mode command to verify that the firewall filters are working properly:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                                     Bytes      Packets
counter-employee-web                     0           0
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                                     Bytes      Packets
```

icmp-counter	560	10
Policers:		
Name	Packets	
icmp-connection-policer	10	
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, counters, and policers that are configured. 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. In the above example, the **icmp-counter** in the filter **ingress-port-limit-tcp-icmp** shows that the filter matched 10 packets. 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. The policer **icmp-connection-policer** shows that 10 ICMP packets were policed.

**Related Documentation**

- [Configuring Firewall Filters on page 20](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55](#)
- [Monitoring Firewall Filter Traffic on page 24](#)

## Troubleshooting Firewall Filters

**Supported Platforms** [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

Use the following information to troubleshoot your firewall filter configuration.

- [Troubleshooting QFX10000 Switches on page 26](#)
- [Troubleshooting Other Switches on page 27](#)

## Troubleshooting QFX10000 Switches

This section describes issues specific to QFX10000 switches:

- [Do Not Combine Match Conditions for Different Layers on page 26](#)
- [Layer 2 Packets Cannot be Discarded with Firewall Filters on page 26](#)

### Do Not Combine Match Conditions for Different Layers

On QFX10000 switches, do not combine match conditions for Layer 2 and any other layer in a **family ethernet-switching** filter. (For example, do not include conditions that match MAC addresses and IP addresses in the same filter.) If you do so, the filter will commit successfully but will not work. You will also see the following log message: **L2 filter *filter-name* doesn't support mixed L2 and L3/L4 match conditions. Please re-config.**

### Layer 2 Packets Cannot be Discarded with Firewall Filters

**Problem Description:** Layer 2 (L2) control packets such as Link Layer Discovery Protocol (LLDP) and bridge protocol data unit (BPDU) cannot be discarded with firewall filters.

**Solution** Configure distributed denial-of-service (DDoS) protection on the L2 control packet and set the aggregate policer bandwidth and burst values to the minimum value of 1. For example,

```
[edit system ddos-protection protocols protocol name]  
user@host# set aggregate bandwidth 1
```

```
[edit system ddos-protection protocols protocol name]  
user@host# set aggregate burst 1
```

## Troubleshooting Other Switches

This section describes issues specific to QFX switches other than QFX10000 switches. This information also applies to OCX1100 switches and EX4600 switches.

- [Firewall Filter Configuration Returns a No Space Available in TCAM Message on page 27](#)
- [Filter Counts Previously Dropped Packet on page 29](#)
- [Matching Packets Not Counted on page 29](#)
- [Counter Reset When Editing Filter on page 30](#)
- [Cannot Include loss-priority and policer Actions in Same Term on page 30](#)
- [Cannot Egress Filter Certain Traffic Originating on QFX Switch on page 30](#)
- [Firewall Filter Match Condition Not Working with Q-in-Q Tunneling on page 31](#)
- [Egress Firewall Filters with Private VLANs on page 31](#)
- [Egress Filtering of L2PT Traffic Not Supported on page 32](#)
- [Cannot Drop BGP Packets in Certain Circumstances on page 32](#)
- [Invalid Statistics for Policer on page 32](#)
- [Policers can Limit Egress Filters on page 32](#)

### Firewall Filter Configuration Returns a No Space Available in TCAM Message

**Problem** **Description:** When a firewall filter configuration exceeds the amount of available Ternary Content Addressable Memory (TCAM) space, the system returns the following **syslogd** message:

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

A 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 space available in the TCAM table. The filter is not applied, but 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 filter and its binding and apply the new smaller firewall filter to the same binding:

1. Delete the filter and its binding to ports, VLANs, or Layer 3 interfaces. For example:

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

2. Commit the changes:

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

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

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

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

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

5. Commit the changes:

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

To apply a new firewall filter and overwrite the existing binding but not delete the original filter:

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

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

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

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

Because you can apply no more than one firewall filter per VLAN per direction, the binding of the original firewall filter to the VLAN is overwritten with the new firewall filter **new-ingress-vlan-rogue-block**.

3. Commit the changes:

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



**NOTE:** The original filter is not deleted and is still available in the configuration.

---

### Filter Counts Previously Dropped Packet

---

- Problem**    **Description:** If you configure two or more filters in the same direction for a physical interface and one of the filters includes a counter, the counter will be incorrect if the following circumstances apply:
- You configure the filter that is applied to packets first to discard certain packets. For example, imagine that you have a VLAN filter that accepts packets sent to 10.10.1.0/24 addresses and implicitly discards packets sent to any other addresses. You apply the filter to the **admin** VLAN in the output direction, and interface xe-0/0/1 is a member of that VLAN.
  - You configure a subsequent filter to accept and count packets that are dropped by the first filter. In this example, you have a port filter that accepts and counts packets sent to 192.168.1.0/24 addresses that is also applied to xe-0/0/1 in the output direction.

The egress VLAN filter is applied first and correctly discards packets sent to 192.168.1.0/24 addresses. The egress port filter is applied next and counts the discarded packets as matched packets. The packets are not forwarded, but the counter displayed by the egress port filter is incorrect.

Remember that the order in which filters are applied depends on the direction in which they are applied, as indicated here:

Ingress filters:

1. Port (Layer 2) filter
2. VLAN filter
3. Router (Layer 3) filter

Egress filters:

1. Router (Layer 3) filter
2. VLAN filter
3. Port (Layer 2) filter

**Solution**    This is expected behavior.

### Matching Packets Not Counted

---

- Problem**    **Description:** If you configure two egress filters with counters for a physical interface and a packet matches both of the filters, only one of the counters includes that packet. For example:
- You configure an egress port filter with a counter for interface xe-0/0/1.
  - You configure an egress VLAN filter with a counter for the **admin** VLAN, and interface xe-0/0/1 is a member of that VLAN.

- A packet matches both filters.

In this case, the packet is counted by only one of the counters even though it matched both filters.

**Solution** This is expected behavior.

---

### Counter Reset When Editing Filter

---

**Supported Platforms** EX4600, OCX1100, QFabric System, QFX Series standalone switches

**Problem** **Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

**Solution** This is expected behavior.

---

### Cannot Include loss-priority and policer Actions in Same Term

---

**Problem** **Description:** You cannot include both of the following actions in the same firewall filter term in a QFX Series switch:

- **loss-priority**
- **policer**

If you do so, you see the following error message when you attempt to commit the configuration: "cannot support policer action if loss-priority is configured."

**Solution** This is expected behavior.

---

### Cannot Egress Filter Certain Traffic Originating on QFX Switch

---

**Problem** **Description:** On a QFX Series switch, you cannot filter certain traffic with a firewall filter applied in the output direction if the traffic originates on the QFX switch. This limitation applies to control traffic for protocols such as ICMP (ping), STP, LACP, and so on.

**Solution** This is expected behavior.



### Firewall Filter Match Condition Not Working with Q-in-Q Tunneling

---

**Supported Platforms** OCX1100, QFabric System, QFX Series standalone switches

**Problem** **Description:** If you create a firewall filter that includes a match condition of **dot1q-tag** or **dot1q-user-priority** and apply the filter on input to a trunk port that participates in a service VLAN, the match condition does not work if the Q-in-Q EtherType is not 0x8100. (When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network and therefore participate in service VLANs.)

**Solution** This is expected behavior. To set the Q-in-Q EtherType to 0x8100, enter the **set dot1q-tunneling ethertype 0x8100** statement at the **[edit ethernet-switching-options]** hierarchy level. You must also configure the other end of the link to use the same EtherType.

### Egress Firewall Filters with Private VLANs

---

**Problem** **Description:** If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port that carries an isolated VLAN to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a secondary VLAN trunk port that carries a community VLAN to a PVLAN trunk port
- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).

- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).
- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

**Solution** These are expected behaviors. They occur only if you apply a firewall filter to a private VLAN in the output direction and do not occur if you apply a firewall filter to a private VLAN in the input direction.

---

### Egress Filtering of L2PT Traffic Not Supported

---

**Supported Platforms** OCX1100, QFabric System, QFX Series standalone switches

**Problem** **Description:** Egress filtering of L2PT traffic is not supported on the QFX3500 switch. That is, if you configure L2PT to tunnel a protocol on an interface, you cannot also use a firewall filter to filter traffic for that protocol on that interface in the output direction. If you commit a configuration for this purpose, the firewall filter is not applied to the L2PT-tunneled traffic.

**Solution** This is expected behavior.

---

### Cannot Drop BGP Packets in Certain Circumstances

---

**Problem** **Description:** BGP packets with a time-to-live (TTL) value greater than 1 cannot be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface. BGP packets with TTL value of 1 or 0 can be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface.

**Solution** This is expected behavior.

---

### Invalid Statistics for Policer

---

**Supported Platforms** EX4600, OCX1100, QFabric System, QFX Series standalone switches

**Problem** **Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

**Solution** This is expected behavior.

---

### Policers can Limit Egress Filters

---

**Problem** **Description:** On some switches, the number of egress policers that you configure can affect the total number of allowed egress firewall filters. (This issue does not affect QFX10000 switches.) On the affected switches, every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including

counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

**Solution** You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.



## PART 2

# Policers

- [Configuring Policers on page 37](#)



## CHAPTER 2

# Configuring Policers

- [Overview of Policers on page 37](#)
- [Understanding Policers with Link Aggregation Groups on page 43](#)
- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 43](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 44](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 46](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 46](#)
- [Example: Using Two-Color Policers and Prefix Lists on page 48](#)
- [Example: Using Policers to Manage Oversubscription on page 51](#)
- [Assigning Forwarding Classes and Loss Priority on page 53](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 54](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55](#)
- [Verifying That Two-Color Policers Are Operational on page 57](#)
- [Verifying That Three-Color Policers Are Operational on page 58](#)
- [Troubleshooting Policer Configuration on page 58](#)

## Overview of Policers

---

**Supported Platforms** [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing (or rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service.

- [Policer Overview on page 38](#)
- [Policer Types on page 38](#)
- [Policer Actions on page 39](#)
- [Policer Colors on page 40](#)
- [Filter-Specific Policers on page 40](#)
- [Suggested Naming Convention for Policers on page 41](#)
- [Policer Counters on page 41](#)


- [Policer Algorithms on page 41](#)
- [How Many Policers Are Supported? on page 42](#)
- [Policers Can Limit Egress Firewall Filters on page 42](#)

## Policer Overview

You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets.

Policers provide two functions: metering and marking. A policer meters (measures) each packet against traffic rates and burst sizes that you configure. It then passes the packet and the metering result to the marker, which assigns a packet loss priority that corresponds to the metering result. [Figure 3 on page 38](#) illustrates this process.

**Figure 3: Flow of Tricolor Marking Policer Operation**



After you name and configure a policer, you can use it by specifying it as an action in one or more firewall filters.

## Policer Types

A switch supports three types of policers:

- **Single-rate two-color marker**—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit with a specified PLP or simply discard them.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-color policer is most useful for metering traffic at the port (physical interface) level.

- **Single-rate three-color marker**—This type of policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on one rate—the configured committed information rate (CIR) as well as the committed burst size (CBS) and the excess burst size (EBS). The CIR specifies the average rate at which bits are admitted to the switch. The CBS



specifies the usual burst size in bytes and the EBS specifies the maximum burst size in bytes. The EBS must be greater than or equal to the CBS, and neither can be 0.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A single-rate three-color marker (TCM) is most useful when a service is structured according to packet length and not peak arrival rate.

- Two-rate three-color marker—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior classification system for a Differentiated Services environment. This type of policer meters traffic based on two rates—the CIR and peak information rate (PIR) along with their associated burst sizes, the CBS and peak burst size (PBS). The PIR specifies the maximum rate at which bits are admitted to the network and must be greater than or equal to the CIR.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

See [Table 6 on page 39](#) for information about how metering results are applied for each of these policer types.

## Policer Actions

Policer actions are implicit or explicit and vary by policer type. *Implicit* means that Junos OS assigns the loss priority automatically. [Table 6 on page 39](#) describes the policer actions.

**Table 6: Policer Actions**

Policer	Marking	Implicit Action	Configurable Action
Single-rate two-color	Green (conforming)	Assign low loss priority	None
	Red (nonconforming)	None	Discard
Single-rate three-color	Green (conforming)	Assign low loss priority	None
	Yellow (above the CIR and CBS)	Assign medium-high loss priority	None
	Red (above the EBS)	Assign high loss priority	Discard

Table 6: Policer Actions (*continued*)

Policer	Marking	Implicit Action	Configurable Action
Two-rate three-color	Green (conforming)	Assign low loss priority	None
	Yellow (above the CIR and CBS)	Assign medium-high loss priority	None
	Red (above the PIR and PBS)	Assign high loss priority	Discard



**NOTE:** If you specify a policer in an egress firewall filter, the only supported action is **discard**.

## Policer Colors

Single-rate and two-rate three-color policers can operate in two modes:

- **Color-blind**—In color-blind mode, the three-color policer assumes that all packets examined have not been previously marked or metered. In other words, the three-color policer is “blind” to any previous coloring a packet might have had.
- **Color-aware**—In color-aware mode, the three-color policer assumes that all packets examined have been previously marked or metered. In other words, the three-color policer is “aware” of the previous coloring a packet might have had. In color-aware mode, the three-color policer can increase the PLP of a packet but cannot decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high but cannot reduce the PLP level to low.

## Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this on some QFX switches, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in TCAM. If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps. (This behavior does not occur in QFX10000 switches.)

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 14](#) to

organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

## Suggested Naming Convention for Policers

We recommend that you use the naming convention ***policertypeTCM#-color type*** when configuring three-color policers and ***policer#*** when configuring two-color policers. TCM stands for three-color marker. Because policers can be numerous and must be applied correctly to work, a simple naming convention makes it easier to apply the policers properly. For example, the first single-rate, color-aware three-color policer configured would be named ***srTCM1-ca***. The second two-rate, color-blind three-color configured would be named ***trTCM2-cb***. The elements of this naming convention are explained below:

- sr (single-rate)
- tr (two-rate)
- TCM (tricolor marking)
- 1 or 2 (number of marker)
- ca (color-aware)
- cb (color-blind)

## Policer Counters

On some QFX switches, 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. If you use the same policer in multiple terms—either within the same filter or in different filters—the implicit counter counts all the packets that are policed in all of these terms and provides the total amount. (This does not apply to QFX10000 switches.) If you want to obtain separate packet counts for each term on an affected switch, use these options:

- Configure a unique policer for each term.
- Configure only one policer, but use a unique, explicit counter in each term.

## Policer Algorithms

Policing uses the *token-bucket algorithm*, which enforces a limit on average bandwidth while allowing bursts up to a specified maximum value. It offers more flexibility than the *leaky bucket algorithm* in allowing a certain amount of bursty traffic before it starts discarding packets.



**NOTE:** In an environment of light bursty traffic, QFX5200 might not replicate all multicast packets to two or more downstream interfaces. This occurs only at a line rate burst—if traffic is consistent, the issue does not occur. In addition, the issue occurs only when packet size increases beyond 6k in a one gigabit traffic flow.

## How Many Policers Are Supported?

QFX10000 switches support 8K policers (all policer types). QFX5100 and QFX5200 switches support 1535 ingress policers and 1024 egress policers (assuming one policer per firewall filter term).

QFX3500 and QFX3600 standalone switches and QFabric Node devices support the following numbers of policers (assuming one policer per firewall filter term):

- Two-color policers used in ingress firewall filters: 767
- Three-color policers used in ingress firewall filters: 767
- Two-color policers used in egress firewall filters: 1022
- Three-color policers used in egress firewall filters: 512

## Policers Can Limit Egress Firewall Filters

On some switches, the number of egress policers that you configure can affect the total number of allowed egress firewall filters. (This issue does not affect QFX10000 switches.) On the affected switches, every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In

this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

#### Related Documentation

- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 43](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 46](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 44](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 46](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55](#)

## Understanding Policers with Link Aggregation Groups

**Supported Platforms** [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

If you apply a policer to a link aggregation group (LAG) on a standalone switch or QFabric node, the policer applies to all the interfaces in the LAG in aggregate. For example, if you configure a policer to rate-limit at 1 Gbps and apply the policer (by using a firewall filter) to a LAG that has two member interfaces on a single switch or node, the total allowed throughput for both members is 1 Gbps.

If you apply a policer to a LAG that has members on different nodes in a QFabric network Node group or redundant server Node group, the configured rate applies to the interface on each node. For example, if you configure a policer to rate-limit at 1 Gbps and apply the policer to a LAG that has one member on server node A and one member on server node B, the allowed throughput for each member is 1 Gbps, for a total allowed throughput of 2 Gbps.

#### Related Documentation

- [Overview of Policers on page 37](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55](#)

## Understanding Color-Blind Mode for Single-Rate Tricolor Marking

**Supported Platforms** [EX4600, OCX1100, QFX Series standalone switches](#)

With the color-blind mode of single-rate tricolor marking, all packets are evaluated against the CBS. If a packet exceeds the CBS, it is evaluated against the EBS. In color-blind mode, the policer supports three loss priorities only: low, medium-high, and high.

Packets that exceed the CBS but are below the EBS are marked yellow (medium-high). Packets that exceed the EBS are marked red (high), as shown in [Table 7 on page 44](#).

Table 7: Color-Blind Mode TCM Color-to-PLP Mapping

Color	PLP	Meaning
Green	low	Conforming.
Yellow	medium-high	Packet exceeds the CIR and CBS but does not exceed the EBS.
Red	high	Packet exceeds the EBS.

**Related Documentation**

- [Overview of Policers on page 37](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 54](#)

## Understanding Color-Aware Mode for Single-Rate Tricolor Marking

**Supported Platforms** [EX4600, OCX1100, QFX Series standalone switches](#)

In color-aware mode, the treatment the packet receives depends on its classification. Marking can increase a preassigned PLP but cannot decrease it.

### Summary of PLP Changes

[Table 8 on page 44](#) shows how a packet's incoming priority can be modified with single-rate marking.

Table 8: Color-Aware Mode Single-Rate PLP Mapping

Incoming PLP	Packet Metered Against	Possible Cases	Outgoing PLP
low	CIR, CBS, and EBS	Conforming	low
		Packet exceeds the CIR and CBS but does not exceed the EBS.	medium-high
		Packet exceeds the EBS.	high
medium-low	EBS only	Packet does not exceed the EBS.	medium-low
		Packet exceeds the EBS.	high
medium-high	EBS only	Packet does not exceed the EBS.	medium-high
		Packet exceeds the EBS.	high
high	Not metered by the policer.	All cases.	high

The following sections describe single-rate color-aware PLP mapping in more detail.

### Effect on Green Packets (Low PLP)

---

Packets belonging to the green class have already been marked by a classifier with low PLP. The marking policer can leave the PLP unchanged or increase it to medium-high or high, so these packets are therefore metered against both the CBS and the EBS. For example, if a behavior aggregate or multifield classifier marks a packet with low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CIR, packets remain marked as low PLP.
- If bursts exceed the CBS but not the EBS, some of the packets are marked as medium-high PLP, and some of the packets remain marked as low PLP.
- If bursts exceed the EBS, some of the packets are marked as high PLP, and some of the packets remain marked as low PLP.

### Effect on Yellow Packets (Medium PLP)

---

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the PLP unchanged or increase it to high, so these packets are therefore metered against the EBS only. For example, if a behavior aggregate or multifield classifier marks a packet with medium-low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CBS, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the CBS but less than the EBS, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the EBS, some of the packets are marked as high PLP and some remain marked as medium-low PLP.

If a BA or multifield classifier marks a packet with medium-high PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CBS, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the CBS but less than the EBS, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the EBS, some of the packets are marked as high PLP and some remain marked as medium-high PLP.

### Effect on Red Packets (High PLP)

---

Packets belonging to the red class have already been marked by a classifier with high PLP. Because the policer cannot decrease the PLP, it does not change it, and these packets are not metered against the CBS or the EBS.

- Related Documentation**
- [Overview of Policers on page 37](#)
  - [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 54](#)

## Understanding Color-Blind Mode for Two-Rate Tricolor Marking

**Supported Platforms** [EX4600, OCX1100, QFX Series standalone switches](#)

With the color-blind mode of two-rate tricolor marking, all packets are evaluated against the committed information rate (CIR). If a packet exceeds the CIR, it is evaluated against the peak information rate (PIR). Packets that exceed the CIR but are below the PIR are marked yellow (medium-high). Packets that exceed the PIR are marked red (high).

**Table 9: Color-Blind Mode TCM Color-to-PLP Mapping**

Color	PLP	Meaning
Green	<b>low</b>	Packet does not exceed the CIR.
Yellow	<b>medium-high</b>	Packet exceeds the CIR but does not exceed the PIR.
Red	<b>high</b>	Packet exceeds the PIR.

- Related Documentation**
- [Overview of Policers on page 37](#)
  - [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 54](#)

## Understanding Color-Aware Mode for Two-Rate Tricolor Marking

**Supported Platforms** [EX4600, OCX1100, QFX Series standalone switches](#)

In color-aware mode, the treatment the packet receives depends on its classification. Marking can increase the preassigned PLP but cannot decrease it.

### Summary of PLP Changes

[Table 10 on page 46](#) shows how a packet's incoming priority can be modified with two-rate marking.

**Table 10: Color-Aware Mode Two-Rate PLP Mapping**

Incoming PLP	Packet Metered Against	Possible Cases	Outgoing PLP
<b>low</b>	CIR and PIR	Packet does not exceed the CIR.	<b>low</b>
		Packet exceeds the CIR but not the PIR.	<b>medium-high</b>
		Packet exceeds the PIR.	<b>high</b>



Table 10: Color-Aware Mode Two-Rate PLP Mapping (*continued*)

Incoming PLP	Packet Metered Against	Possible Cases	Outgoing PLP
medium-low	PIR only	Packet does not exceed the PIR.	medium-low
		Packet exceeds the PIR.	high
medium-high	PIR only	Packet does not exceed the PIR.	medium-high
		Packet exceeds the PIR.	high
high	Not metered by the policer.	All cases.	high

The following sections describe color-aware two-rate PLP mapping in more detail.

### Effect on Green Packets (Low PLP)

Packets belonging to the green class have already been marked by a classifier with low PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to medium-high or high. These packets are therefore metered against both the CIR and the PIR. For example, if a behavior aggregate or multifield classifier marks a packet with low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as low PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, some of the packets are marked as medium-high PLP and some of the packets remain marked as low PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as low PLP.

### Effect on Yellow Packets (Medium PLP)

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the PLP unchanged or increase it to high. These packets are therefore metered against the PIR only. For example, if a behavior aggregate (BA) or multifield classifier marks a packet with medium-low PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as medium-low PLP.

If a BA or multifield classifier marks a packet with medium-high PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as medium-high PLP.

### Effect on Red Packets (High PLP)

Packets belonging to the red class have already been marked by a classifier with high PLP. Because the policer cannot decrease the PLP, it does not change it, and these packets are not metered against the CIR or the PIR.

#### Related Documentation

- [Overview of Policers on page 37](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 54](#)

---

## Example: Using Two-Color Policers and Prefix Lists

**Supported Platforms**    [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

If you provide specific amounts of bandwidth to internal or external customers, you can use policing to make sure that customers do not consume more bandwidth than they should receive. For example, you might connect many customers to one 10-Gbps interface and want to ensure that none of them congest the interface by using more bandwidth than they have been allotted.

You could accomplish this by creating a two-color policer similar to the following for each customer:

```
firewall {
  policer Limit-Customer-1 {
    if-exceeding {
      bandwidth-limit 100m;
      burst-size-limit 150m;
    }
    then discard;
  }
}
```

Creating a policer for each customer is clearly not a scalable solution, however. As an alternative, you can create prefix lists that group classes of customers and then create policers for each prefix list. For example, you could create prefix lists such as **Class-A-Customer-Prefixes**, **Class-B-Customer-Prefixes**, and **Class-C-Customer-Prefixes** (at the **[edit policy-options]** hierarchy level) and create the following corresponding policers:

```
firewall {
  policer Class-A {
```

```

    if-exceeding {
        bandwidth-limit 100m;
        burst-size-limit 150m;
    }
    then discard;
}
policer Class-B {
    if-exceeding {
        bandwidth-limit 75m;
        burst-size-limit 100m;
    }
    then discard;
}
policer Class-C {
    if-exceeding {
        bandwidth-limit 50m;
        burst-size-limit 75m;
    }
    then discard;
}
}

```

You must create filter terms that specify the prefix lists in their **from** statements and the corresponding policers in their **then** statements similar to the following:

```

firewall
family inet {
    filter Class-A-Customers {
        term term-1 {
            from {
                destination-prefix-list {
                    Class-A-Customer-Prefixes;
                }
            }
            then policer Class-A;
        }
    }
    filter Class-B-Customers {
        term term-1 {
            from {
                destination-prefix-list {
                    Class-B-Customer-Prefixes;
                }
            }
            then policer Class-B;
        }
    }
    filter Class-C-Customers {
        term term-1 {
            from {
                destination-prefix-list {
                    Class-C-Customer-Prefixes;
                }
            }
            then policer Class-C;
        }
    }
}

```

```
}
}
```

Here are the steps to create this firewall configuration:

1. Create the first policer:

```
[edit firewall]
user@switch# set policer Class-A if-exceeding bandwidth-limit 100m burst-size-limit 150m
user@switch# set policer Class-A then discard
```

2. Create the second policer:

```
[edit firewall]
user@switch# set policer Class-B if-exceeding bandwidth-limit 75m burst-size-limit 100m
user@switch# set policer Class-B then discard
```

3. Create the third policer:

```
[edit firewall]
user@switch# set policer Class-C if-exceeding bandwidth-limit 50m burst-size-limit 75m
user@switch# set policer Class-C then discard
```

4. Create a filter for class A customers:

```
[edit firewall]
user@switch# edit family inet filter Class-A-Customers
```

5. Configure the filter to send packets matching the **Class-A-Customer-Prefixes** prefix list to the **Class-A** policer:

```
[edit firewall family inet filter Class-A-Customers]
user@switch# set term term-1 from source-prefix-list Class-A-Customers
user@switch# set term term-1 then policer Class-A
```

6. Create a filter for class B customers:

```
[edit firewall]
user@switch# edit family inet filter Class-B-Customers
```

7. Configure the filter to send packets matching the **Class-B-Customer-Prefixes** prefix list to the **Class-B** policer:

```
[edit firewall family inet filter Class-B-Customers]
user@switch# set term term-1 from source-prefix-list Class-B-Customers
user@switch# set term term-1 then policer Class-B
```

8. Create a filter for class C customers:

```
[edit firewall]
user@switch# edit family inet filter Class-C-Customers
```

9. Configure the filter to send packets matching the **Class-C-Customer-Prefixes** prefix list to the **Class-C** policer:

```
[edit firewall family inet filter Class-C-Customers]
user@switch# set term term-1 from source-prefix-list Class-C-Customers
user@switch# set term term-1 then policer Class-C
```

10. Apply the filters you created to the appropriate interfaces in the output direction.



**NOTE:** Note that the implicit deny statement in this filter will block traffic from any source that does not match one of the prefix lists. If you want the filter to allow this traffic, you must include an explicit term for this purpose.

- Related Documentation**
- [Overview of Policers on page 37](#)
  - [prefix-list](#)

## Example: Using Policers to Manage Oversubscription

**Supported Platforms** [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

You might want to use a policer when an interface is oversubscribed and you want to control what will happen if congestion occurs. For example, you might have servers connected to a switch as listed in [Table 11 on page 51](#).

**Table 11: Servers Connected to Switch**

Server Type	Connection	IP Address
Network application server	1-gigabit interface	10.0.0.1
Authentication server	1-gigabit interface	10.0.0.2
Database server	10-gigabit interface	10.0.0.3

In this example, users access services provided by the network application server, which requests information from the database server as appropriate. When it receives a request from a user, the network application server first contacts the authentication server to verify the user's credentials. When a user is authenticated and the network application server provides the requested service, all the packets sent from the database server to the application server must transit the 1-Gigabit Ethernet interface connected to the application server twice—once on ingress to the application server and again on egress to the user.

The sequence of events for a user session is as follows:

1. A user connects to the application server and requests a service.
2. The application server requests the user's credentials and relays them to the authentication server.
3. If the authentication server verifies the credentials, the application server initiates the requested service.
4. The application server requests the files necessary to meet the user's request from the database server.
5. The database server sends the requested files to the application server.
6. The application server includes the requested files in its response to the user.

Traffic from the database server to the application server might congest the 1-gigabit interface to which that the application server is connected. This congestion might prevent the server from responding to requests from users and creating new sessions for them. You can use policing to make sure that this does not occur.

To create this firewall configuration, perform the following steps on the database server:

1. Create a policer to drop traffic from the database server to the application server if it exceeds certain limits:

```
[edit firewall]
user@switch# set policer Database-Egress-Policer if-exceeding bandwidth-limit 400
burst-size-limit 500m
user@switch# set policer Database-Egress-Policer then discard
```

2. Create a filter to examine traffic from the database server to the application server:

```
[edit firewall]
user@switch# edit family inet filter Database-Egress-Filter
```

3. Configure the filter to apply the policer to traffic egressing the database server and destined for the application server:

```
[edit firewall family inet filter Database-Egress-Filter]
user@switch# set term term-1 from destination-address 10.0.0.1
user@switch# set term term-1 then policer Database-Egress-Policer
```

4. If required, configure a term to allow traffic from the database server to other destinations (otherwise the traffic will be dropped by the implicit deny statement):

```
[edit firewall family inet filter Database-Egress-Filter]
user@switch# set term term-2 then accept
```

Note that omitting a **from** statement causes the term to match all packets, which is the desired behavior.

5. Install the egress filter as an output filter on the database server interface that is connected the application server:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet filter output Database-Egress-Filter
```

Here is how the final configuration would appear:

```
firewall {
  policer Database-Egress-Policer {
    if-exceeding {
      bandwidth-limit 400;
      burst-size-limit 500m;
    }
    then discard;
  }
  family inet {
    filter Database-Egress-Filter {
      term term-1 {
        from {
          destination-address {
            10.0.0.1/24;
          }
        }
        then policer Database-Egress-Policer;
      }
      term term-2 { # If required, include this term so that traffic from the database server
                    # to other destinations is allowed.
        then accept;
      }
    }
  }
}
```

```
}
]
```

#### Related Documentation

- [Overview of Policers on page 37](#)

## Assigning Forwarding Classes and Loss Priority

**Supported Platforms** [EX4600, OCX1100, QFX Series standalone switches](#)

You can configure firewall filters to assign packet loss priority (PLP) and forwarding classes so that if congestion occurs, the marked packets can be dropped according to the priority you set. The valid match conditions are one or more of the six packet header fields: destination address, source address, IP protocol, source port, destination port, and DSCP. In other words, you can set the forwarding class and the PLP for each packet entering or an interface with a specific destination address, source address, IP protocol, source port, destination port, or DSCP.



**NOTE:** Junos OS assigns forwarding classes and PLP on ingress only. Do not use a filter that assigns forwarding classes or PLP as an egress filter.

When tricolor marking is enabled, a switch supports four PLP designations: **low**, **medium-low**, **medium-high**, and **high**. You can also specify any of the forwarding classes listed in [Table 12 on page 53](#)

**Table 12: Unicast Forwarding Classes**

Unicast Forwarding Class	For CoS Traffic Type
<b>be</b>	Best-effort traffic
<b>no-loss</b>	Guaranteed delivery for TCP traffic
<b>fcoe</b>	Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic
<b>nc</b>	Network-control traffic

To assign forwarding classes in firewall filters:

1. Configure the family address type and filter name:
 

```
[edit]
user@switch# edit firewall family ethernet-switching filter ingress-filter
```
2. Configure the terms of the filter as appropriate, including the **forwarding-class** and **loss-priority** action modifiers. For example, each of the following terms in the filter examines various packet header fields and assigns the appropriate forwarding class and packet loss priority:
  - The term **corp-traffic** matches all IPv4 packets with a **10.1.1.0/24** source address and assigns the packets to forwarding class **no-loss** with a loss priority of **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term corp-traffic from source-address 10.1.1.0/24;
user@switch# set term corp-traffic then forwarding-class no-loss
user@switch# set term corp-traffic then loss-priority low
```

- The term **data-traffic** matches all IPv4 packets with a 10.1.2.0/24 source address and assigns the packets to forwarding class **be** (best effort) with a loss priority of **medium-high**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term data-traffic from source-address 10.1.2.0/24;
user@switch# set term data-traffic then forwarding-class be
user@switch# set term data-traffic then loss-priority medium-high
```

- Because the loss of network-generated packets can jeopardize proper network operation, the delay of these packets is preferable to discarding these packets. The term **network-traffic** assigns the packets with an IP precedence of **net-control** to forwarding class **nc** (network control) with a loss priority of **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 nc
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 packets to forwarding class **be** with a loss priority of **high**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term accept-traffic then forwarding-class be
user@switch# set term accept-traffic then loss-priority high
```

3. Apply the filter **ingress-filter** to a port, VLAN, or Layer 3 interface. For information about applying the filter, see [“Configuring Firewall Filters” on page 20](#). (Assigning forwarding classes and PLP is supported only on ingress filters.)

#### Related Documentation

- [Configuring Firewall Filters on page 20](#)
- [Verifying That Firewall Filters Are Operational on page 25](#)
- [Monitoring Firewall Filter Traffic on page 24](#)
- [Overview of Policers on page 37](#)
- [Understanding CoS Classifiers](#)
- [Understanding CoS Forwarding Classes](#)

---

## Configuring Color-Blind Egress Policers for Medium-Low PLP

**Supported Platforms**    EX4600, OCX1100, QFabric System, QFX Series standalone switches



If you use color-blind mode and want to configure an egress policer that marks packets to have medium-low PLP, you must configure a single-rate two-color policer at the **[edit firewall policer *policer-name*]** hierarchy level, because color-blind mode does not support medium-low priority. For example:

1. Specify the name of the policer, the bandwidth limit in bits per second (bps) to control the traffic rate on an interface, and the maximum allowed burst size to control the amount of traffic bursting:

```
[edit]
user@switch# set firewall policer policer-name if-exceeding bandwidth-limit bytes
burst-size-limit bytes
```

2. Specify medium-low loss priority for matching packets:

```
[edit]
user@switch# set firewall policer policer-name then loss-priority medium-low;
```

3. Apply the filter to a port, VLAN, or Layer 3 interface.

#### Related Documentation

- [Overview of Policers on page 37](#)
- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 43](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 46](#)
- [Configuring Firewall Filters on page 20](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55](#)

## Configuring Two-Color and Three-Color Policers to Control Traffic Rates

**Supported Platforms** [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

You can rate-limit traffic by configuring a policer and specifying it as an action modifier for a term in a firewall filter. By default, if you specify the same policer in multiple terms, Junos OS creates a separate policer instance for each term and applies rate limiting separately for each instance. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, each policer instance enforces a 1-Gbps limit. In this case, the total bandwidth allowed by the filter is 3 Gbps.

You can also configure a policer to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps.



**NOTE:** You can include two-color policer actions on ingress firewall filters only. You can include three-color policer actions on ingress and egress filters.

1. [Configuring Two-Color Policers on page 56](#)
2. [Configuring Three-Color Policers on page 56](#)

3. [Specifying Policers in a Firewall Filter Configuration on page 57](#)
4. [Applying a Firewall Filter That Includes a Policer on page 57](#)

## Configuring Two-Color Policers

To configure a two-color policer:

1. Specify the name of the policer, the bandwidth limit to control the traffic rate on an interface, and the maximum allowed burst size to control the amount of traffic bursting:

```
[edit firewall]
user@switch# set policer policer-name <filter-specific> if-exceeding bandwidth-limit bps
burst-size-limit bytes
```

The policer name can contain letters, numbers, and hyphens (-) and can have as many as 64 characters.

The range for the bandwidth limit is 32000 (32k) through 102,300,000,000 (102300m) bps.

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 and divide the result by 8:

**maximum burst size = (interface bandwidth) X (allowable time for burst) / (8 bits/byte)**

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

2. Specify the policer action to discard or assign a loss priority to packets that exceed the rate limits:

```
[edit firewall policer policer-name]
user@switch# set then (discard | loss-priority low | loss-priority high)
```

## Configuring Three-Color Policers

To configure a three-color policer:

1. Specify the name of the policer and (optionally) whether to automatically discard packets with high loss priority (PLP):

```
[edit firewall]
user@switch# set three-color-policer policer-name
user@switch# set three-color-policer policer-name action loss-priority high then discard
```

2. Specify whether the three-color policer should be single-rate or two-rate and whether it should be color-aware or color-blind:

```
[edit firewall three-color-policer policer-name]
user@switch# set (single-rate | two-rate) (color-aware | color-blind)
```

3. For single-rate three-color policers, configure the CIR, CBS, and EBS:

```
[edit firewall three-color-policer policer-name single-rate]
user@switch# set committed-information-rate bps
user@switch# set committed-burst-size bytes
user@switch# set excess-burst-size bytes
```

4. For two-rate three-color policers, configure the CIR, CBS, PIR, and PBS:

```
[edit firewall three-color-policer policer-name single-rate]
user@switch# set committed-information-rate bps
user@switch# set committed-burst-size bytes
```

```

user@switch# set peak-information-rate bps
user@switch# set peak-burst-size bytes

```

## Specifying Policers in a Firewall Filter Configuration

To use a two-color policer, configure a filter term that includes the action **policer**:

```

[edit firewall family family-name]
user@switch# set filter filter-name term name then name

```

For example, the following commands apply a two-color policer to all packets sent from 192.0.2.0/24.

```

[edit firewall family family-name]
user@switch# set filter limit-hosts term term1 from source-address 192.0.2.0/24
user@switch# set filter limit-hosts term term1 then policer policer1

```

To use a three-color policer, configure a filter term that includes the action **three-color-policer**:

```

[edit firewall family name]
user@switch# set filter name term name from match-condition
user@switch# set filter name term name then three-color-policer (single-rate | two-rate) name

```

For example, the following commands apply a single-rate three-color policer to all packets received or sent by interface **ge-0/0/6** (depending on whether the filter is an ingress or egress filter).

```

[edit firewall family name]
user@switch# set filter srTCM term term-one from interface ge-0/0/6
user@switch# set filter srTCM term term-one then three-color-policer single-rate srTCM1-ca

```

You must specify whether the three-color policer is single-rate or two-rate, and this must match the policer itself. Otherwise, the configuration listing includes an error message indicating that the three-color policer you referenced in the filter does not exist.

## Applying a Firewall Filter That Includes a Policer

A firewall filter that includes one or more policer action modifiers must be applied to a port, VLAN, or Layer 3 interface like any other filter. For information about applying firewall filters, see “Configuring Firewall Filters” on page 20.



**NOTE:** You can include two-color policer actions on ingress firewall filters only. You can include three-color policer actions on ingress and egress filters.

### Related Documentation

- [Configuring Firewall Filters on page 20](#)
- [Overview of Policers on page 37](#)
- [Verifying That Two-Color Policers Are Operational on page 57](#)
- [Verifying That Three-Color Policers Are Operational on page 58](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 54](#)

## Verifying That Two-Color Policers Are Operational

**Supported Platforms** EX4600, OCX1100, QFabric System, QFX Series standalone switches

**Purpose** Verify that two-color policers in firewall filter configurations are working properly.

**Action** Use the **show firewall policer** operational mode command to verify that the policers are working properly:

```
user@switch> show firewall policer
Filter: egress-vlan-watch-employee
Filter: ingress-port-filter
Filter: ingress-port-limit-tcp-icmp
Policers:
Name                                     Packets
icmp-connection-policer                  10
tcp-connection-policer                   539
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

**Meaning** The **show firewall policer** command displays the names of all firewall filters and policers that are configured. 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 Two-Color and Three-Color Policers to Control Traffic Rates on page 55](#)
- [Configuring Firewall Filters on page 20](#)
- [Monitoring Firewall Filter Traffic on page 24](#)

---

## Verifying That Three-Color Policers Are Operational

---

**Supported Platforms** [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

**Purpose** Verify that three-color policers in firewall filter configurations are working properly.

**Action** Use the following operational mode commands to verify that a three-color policer is working properly:

- **show class-of-service forwarding-table classifiers**
- **show interfaces *interface-name* extensive**
- **show interfaces queue *interface-name***

**Related Documentation**

- [Overview of Policers on page 37](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55](#)

---

## Troubleshooting Policer Configuration

---

**Supported Platforms** [EX4600, OCX1100, OCX1100, QFabric System, QFX Series standalone switches](#)

- [Incomplete Count of Packet Drops on page 59](#)
- [Counter Reset When Editing Filter on page 59](#)
- [Invalid Statistics for Policer on page 59](#)

- [Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 60](#)
- [Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 61](#)
- [Policers Can Limit Egress Filters on page 61](#)

## Incomplete Count of Packet Drops

### Supported Platforms

**Problem** **Description:** Under certain circumstances, Junos OS might display a misleading number of packets dropped by an ingress policer.

If packets are dropped because of ingress admission control, policer statistics might not show the number of packet drops you would expect by calculating the difference between ingress and egress packet counts. This might happen if you apply an ingress policer to multiple interfaces, and the aggregate ingress rate of those interfaces exceeds the line rate of a common egress interface. In this case, packets might be dropped from the ingress buffer. These drops are not included in the count of packets dropped by the policer, which causes policer statistics to underreport the total number of drops.

**Solution** This is expected behavior.

## Counter Reset When Editing Filter

**Supported Platforms** [EX4600](#), [OCX1100](#), [QFabric System](#), [QFX Series standalone switches](#)

**Problem** **Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

**Solution** This is expected behavior.

## Invalid Statistics for Policer

**Supported Platforms** [EX4600](#), [OCX1100](#), [QFabric System](#), [QFX Series standalone switches](#)

**Problem** **Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

**Solution** This is expected behavior.

## Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

**Problem Description:** If you configure a policer to rate-limit throughput and apply it on egress to multiple interfaces on a QFX3500 switch or Node, the measured aggregate policed rate might be twice the configured rate, depending on which interfaces you apply the policer to. The doubling of the policed rate occurs if you apply a policer to multiple interfaces and *both* of the following are true:

- There is at least one policed interface in the range xe-0/0/0 to xe-0/0/23 or the range xe-0/1/1 to xe-0/1/7.
- There is at least one policed interface in the range xe-0/0/24 to xe-0/0/47 or the range xe-0/1/8 to xe-0/1/15.

For example, if you configure a policer to rate-limit traffic at 1 Gbps and apply the policer (by using a firewall filter) to xe-0/0/0 and xe-0/0/24 in the output direction, each interface is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps. The same behavior occurs if you apply the policer to xe-0/1/1 and xe-0/0/24—each interface is rate-limited at 1 Gbps.

If you apply the same policer on egress to multiple interfaces in these groups, each *group* is rate-limited at 1 Gbps. For example, if you apply the policer to xe-0/0/0 through xe-0/0/4 (five interfaces) and xe-0/0/24 through xe-0/0/33 (ten interfaces), each group is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps.

Here is another example: If you apply the policer to xe-0/0/0 through xe-0/0/4 and xe-0/1/1 through xe-0/1/5 (a total of ten interfaces), that group is rate-limited at 1 Gbps in aggregate. If you also apply the policer to xe-0/0/24, that one interface is rate-limited at 1 Gbps while the other ten are still rate-limited at 1 Gbps in aggregate.

Interfaces xe-0/1/1 through xe-0/1/15 are physically located on the QSFP+ uplink ports, according to the following scheme:

- xe-0/1/1 through xe-0/1/3 are on Q0.
- xe-0/1/4 through xe-0/1/7 are on Q1.
- xe-0/1/8 through xe-0/1/11 are on Q2.
- xe-0/1/12 through xe-0/1/15 are on Q3.

The doubling of the policed rate occurs only if the policer is applied in the output direction. If you configure a policer as described above but apply it in the input direction, the total allowed throughput for all interfaces is 1 Gbps.

**Solution** This is expected behavior.

## Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

**Problem**    **Description:** You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in ternary content addressable memory (TCAM). If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

**Solution**    To prevent this unexpected behavior, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 14](#) to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

## Policers Can Limit Egress Filters

**Problem**    **Description:** On some switches, the number of egress policers that you configure can affect the total number of allowed egress firewall filters. (This issue does not affect QFX10000 switches.) On the affected switches, every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.

- Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

**Solution** You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.



## PART 3

# Configuring Device Security

- [Device Security on page 65](#)



## CHAPTER 3

# Device Security

- [Understanding Storm Control on page 65](#)
- [Example: Configuring Storm Control to Prevent Network Outages on page 66](#)
- [Verifying That the Port Error Disable Setting Is Working Correctly on page 68](#)
- [Understanding Unicast RPF on page 69](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 74](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 75](#)
- [Verifying Unicast RPF Status on page 76](#)
- [Understanding Unknown Unicast Forwarding on page 79](#)
- [Configuring Unknown Unicast Forwarding \(CLI Procedure\) on page 79](#)

## Understanding Storm Control

---

**Supported Platforms** [EX4600, MX480, MX80, QFabric System, QFX Series standalone switches](#)

A traffic storm occurs when broadcast packets prompt receiving devices to broadcast packets in response. This prompts further responses, creating a snowball effect. The switch is flooded with packets, which creates unnecessary traffic that leads to poor performance or even a complete loss of service by some clients. Storm control causes a device to monitor traffic levels and take a specified action when a specified traffic level—called the *storm control level*—is exceeded, thus preventing packets from proliferating and degrading service. You can configure devices to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when the storm control level is exceeded.

Storm control is enabled by default on ELS platforms and disabled by default on non-ELS platforms. If storm control is enabled, the default level is 80 percent of the available bandwidth for ingress traffic. You can change the storm control level by configuring it as a specific bandwidth value. (The **level** configuration statement, which allows you to configure the storm control level as a percentage of the combined broadcast and unknown unicast streams, is deprecated and might be removed from future releases. We recommend that you phase out its use and replace it with the **bandwidth** statement.)



---

**NOTE:** Storm control is not enabled by default on MX platforms.

---



**NOTE:** When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually enforced. For example, if you configure a bandwidth limit of 150 Kbps, storm control enforces a bandwidth limit of 128 Kbps.



**NOTE:** On an FCoE-FC gateway, storm control must be disabled on all Ethernet interfaces that belong to an FCoE VLAN to prevent FCoE traffic from being dropped. Configuring storm control on an Ethernet interface that is included in an FCoE-FC gateway may have undesirable effects, including FCoE packet loss. After disabling storm control on all interfaces, enable storm control on any interfaces that are not part of an FCoE-FC gateway on which you want to use storm control. However, on an FCoE transit switch, you can enable storm control on interfaces that carry FCoE traffic.



**CAUTION:** The Junos OS allows you to configure a storm control value that exceeds the bandwidth of the interface. If you configure an interface this way, storm control does not drop broadcast or unknown unicast packets even if they consume all the available bandwidth.

To recognize a storm, you must be able to identify when traffic has reached an abnormal level. Suspect a storm when operations begin timing out and network response times slow down. Users might be unable to access expected services. Monitor the percentage of broadcast and unknown unicast traffic in the network when it is operating normally. This data can then be used as a benchmark to determine when traffic levels are too high. You can then configure storm control to set the level at which you want to drop broadcast and unknown unicast traffic.

**Related  
Documentation**

- *action-shutdown*
- *port-error-disable*
- *storm-control*

---

## Example: Configuring Storm Control to Prevent Network Outages

**Supported Platforms**    [EX4600, QFX Series standalone switches](#)

Storm control enables you to prevent network outages caused by broadcast storms on the LAN. You can configure storm control on to rate-limit broadcast traffic, multicast traffic, and unknown unicast traffic at a specified level and to have packets dropped when the specified traffic level is exceeded, thereby preventing packets from proliferating and degrading the LAN.



**NOTE:** This example uses a Junos OS release that supports the Enhanced Layer 2 Software (ELS) configuration style.

- [Requirements on page 67](#)
- [Overview and Topology on page 67](#)
- [Configuration on page 68](#)

## Requirements

This example uses the following hardware and software components:

- One QFX Series switch running Junos OS with ELS
- Junos OS Release 13.2 or later

## Overview and Topology

A storm is generated when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own messages on the network. This, in turn, prompts further responses, creating a snowball effect and resulting in a broadcast storm that can cause network outages.

You can use storm control to prevent broadcast storms by specifying the amount, also known as the *storm control level*, of broadcast traffic, multicast traffic, and unknown unicast traffic to be allowed on an interface. You specify the storm control level as the traffic rate in kilobits per second (Kbps) of the combined applicable traffic streams or as the percentage of available bandwidth used by the combined applicable traffic streams. On ELS systems, storm control is enabled by default on all interfaces at a level of 80 percent of the available bandwidth.

Storm control monitors the level of applicable incoming traffic and compares it with the level that you specify. If the combined level of the applicable traffic exceeds the specified level, the switch drops packets for the controlled traffic types. As an alternative to having the switch drop packets, you can configure storm control to shut down interfaces or temporarily disable interfaces (see the **action-shutdown** statement or the **recovery-timeout** statement) when the storm control level is exceeded.



**NOTE:** If you configure storm control on an aggregated Ethernet interface, the storm-control level is applied to each member interface individually. For example, if the aggregated interface has two members and you configure a storm-control level of 20 kbps, Junos will not detect a storm if one or both of the member interfaces receives traffic at 15 kbps because in neither of these cases does an individual member receive traffic at a rate greater than the configured storm-control level. In this example, Junos detects a storm only if at least one member interface receives traffic at greater than 20 Kbps.

The topology used in this example consists of one switch connected to various network devices. This example shows how to configure the storm control level on interface xe-0/0/0 by setting the level to a traffic rate of 15,000 Kbps, based on the traffic rate of the combined applicable traffic streams. If the combined traffic exceeds this level, the switch drops packets for the controlled traffic types to prevent a network outage.

## Configuration

**CLI Quick Configuration** To quickly configure storm control based on the traffic rate in kilobits per second of the combined traffic streams, copy the following command and paste it into the switch terminal window:

```
[edit]
set forwarding-options storm-control-profiles sc-profile all bandwidth-level 15000
set interfaces xe-0/0/0 unit 0 family ethernet-switching storm-control sc-profile
```

**Step-by-Step Procedure** To configure storm control:

1. Configure a storm control profile, **sc-profile**, and specify the traffic rate in kilobits per second of the combined traffic streams:

```
[edit]
user@switch> set forwarding-options storm-control-profiles sc-profile all bandwidth-level 15000
```

2. Bind the storm control profile, **sc**, to a logical interface:

```
[edit]
user@switch> set interfaces xe-0/0/0 unit 0 family ethernet-switching storm-control sc-profile
```

**Results** Display the results of the configuration:

```
[edit forwarding-options]
user@switch> show storm-control-profiles sc-profile
all {
  bandwidth 15000;
}

[edit]
user@switch> show interfaces xe-0/0/0
unit 0 {
  family ethernet-switching {
    vlan {
      members default;
    }
    storm-control sc-profile;
  }
}
```

**Related Documentation**

- [Understanding Storm Control on page 65](#)
- [Configuring Autorecovery for MAC Limited or Storm Control Interfaces \(CLI Procedure\)](#)

---

## Verifying That the Port Error Disable Setting Is Working Correctly

**Supported Platforms** EX4600, QFabric System, QFX Series standalone switches

**Purpose** Verify that the port error disable setting is working as expected for MAC limited and storm control interfaces.

**Action** Display information about interfaces:

```
user@switch> show ethernet-switching interfaces
```

Interface	State	VLAN members	Blocking
xe-2:0/0/0.0	up	T1122	unblocked
xe-2:0/0/1.0	down	default	MAC limit exceeded
xe-2:0/0/2.0	down	default	Storm control in effect
xe-2:0/0/3.0	down	default	unblocked
xe-2:0/0/4.0	down	default	unblocked
xe-2:0/0/5.0	down	default	unblocked
xe-2:0/0/6.0	down	default	unblocked

**Meaning** For interfaces disabled by port security features, the sample output from the **show ethernet-switching interfaces** command specifies the reasons that the interfaces are disabled:

- **MAC limit exceeded**—The interface is temporarily disabled because of a *mac-limit* error. The disabled interface is automatically restored to service when the disable-timeout expires.
- **MAC move limit exceeded**—The interface is temporarily disabled because of a *mac-move-limit* error. The disabled interface is automatically restored to service when the disable-timeout expires.
- **Storm control in effect**—The interface is temporarily disabled because of a storm control error. The disabled interface is automatically restored to service when the disable-timeout expires.

**Related Documentation**

- *Understanding MAC Limiting and MAC Move Limiting for Port Security*
- *port-error-disable*

## Understanding Unicast RPF

**Supported Platforms** EX Series, OCX1100, QFX Series standalone switches

Unicast reverse-path forwarding (RPF) helps protect the switch against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled. It also helps ensure that traffic arriving on ingress interfaces comes from a network source that the receiving interface can reach.

When you enable unicast RPF, by default the switch forwards a packet only if the receiving interface is the best return path to the packet's unicast source address. This is known as strict mode unicast RPF. You can also enable loose mode, which means that the system checks to see if the packet has a source address with a corresponding prefix in the routing table but does not check whether the receiving interface is the best return path to the packet's unicast source address.



**NOTE:** On Juniper Networks EX3200, EX4200, and EX4300 Ethernet Switches, the switch applies unicast RPF *globally* to all interfaces when unicast RPF is configured on any interface. For additional information, see [“Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches” on page 73](#).

This topic covers:

- [Unicast RPF for Switches Overview on page 70](#)
- [Unicast RPF Implementation on page 71](#)
- [When to Enable Unicast RPF on page 71](#)
- [When Not to Enable Unicast RPF on page 72](#)
- [Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches on page 73](#)

## Unicast RPF for Switches Overview

Unicast RPF functions as an ingress filter that reduces the forwarding of IP packets that might be spoofing an address. By default, unicast RPF is disabled on the switch interfaces.

The type of unicast RPF provided on the switches—that is, strict mode unicast RPF is especially useful on untrusted interfaces. An untrusted interface is an interface where untrusted users or processes can place packets on the network segment.

The switch supports only the active paths method of determining the best return path back to a unicast source address. The active paths method looks up the best reverse path entry in the forwarding table. It does not consider alternate routes specified using routing-protocol-specific methods when determining the best return path.

If the forwarding table lists the receiving interface as the interface to use to forward the packet back to its unicast source, it is the best return path interface.

Use strict mode unicast RPF only on symmetrically routed interfaces. (For information about symmetrically routed interfaces, see [“When to Enable Unicast RPF” on page 71](#).)



For more information about strict unicast RPF, see RFC 3704, *Ingress Filtering for Multihomed Networks* at <http://www.ietf.org/rfc/rfc3704.txt>.

## Unicast RPF Implementation

This section includes:

- [Unicast RPF Packet Filtering on page 71](#)
- [Bootstrap Protocol \(BOOTP\) and DHCP Requests on page 71](#)
- [Default Route Handling on page 71](#)

---

### Unicast RPF Packet Filtering

When you enable unicast RPF on the switch, the switch handles traffic in the following manner:

- If the switch receives a packet on the interface that is the best return path to the unicast source address of that packet, the switch forwards the packet.
- If the best return path from the switch to the packet's unicast source address is not the receiving interface, the switch discards the packet.
- If the switch receives a packet that has a source IP address that does not have a routing entry in the forwarding table, the switch discards the packet.

---

### Bootstrap Protocol (BOOTP) and DHCP Requests

Bootstrap protocol (BOOTP) and DHCP request packets are sent with a broadcast MAC address and therefore the switch does not perform unicast RPF checks on them. The switch forwards all BOOTP packets and DHCP request packets without performing unicast RPF checks.

---

### Default Route Handling

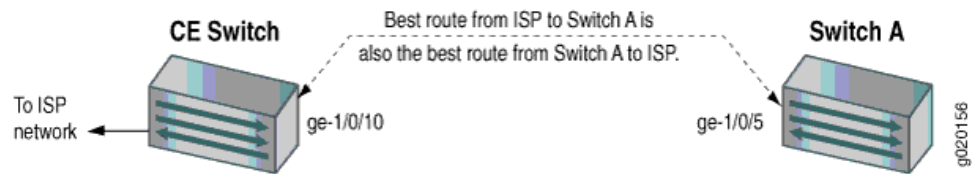
If the best return path to the source is the default route (**0.0.0.0**) and the default route points to **reject**, the switch discards the packets. If the default route points to a valid network interface, the switch performs a normal unicast RPF check on the packets.

## When to Enable Unicast RPF

Enable unicast RPF when you want to ensure that traffic arriving on a network interface comes from a source that resides on a network that that interface can reach. You can enable unicast RPF on untrusted interfaces to filter spoofed packets. For example, a common application for unicast RPF is to help defend an enterprise network from DoS/DDoS attacks coming from the Internet.

Enable unicast RPF only on symmetrically routed interfaces. A symmetrically routed interface uses the same route in both directions between the source and the destination, as shown in [Figure 4 on page 72](#). Symmetrical routing means that if an interface receives a packet, the switch uses the same interface to send a reply to the packet source (the receiving interface matches the forwarding-table entry for the best return path to the source).

Figure 4: Symmetrically Routed Interfaces



Enabling unicast RPF on asymmetrically routed interfaces (where different interfaces receive a packet and reply to its source) results in packets from legitimate sources being filtered (discarded) because the best return path is not the same interface that received the packet.

The following switch interfaces are most likely to be symmetrically routed and thus are candidates for unicast RPF enabling:

- The service provider edge to a customer
- The customer edge to a service provider
- A single access point out of the network (usually on the network perimeter)
- A terminal network that has only one link



**NOTE:** Because unicast RPF is enabled globally on EX3200, EX4200, and EX4300 switches, ensure that *all* interfaces are symmetrically routed before you enable unicast RPF on these switches. Enabling unicast RPF on asymmetrically routed interfaces results in packets from legitimate sources being filtered.



**TIP:** Enabling unicast RPF as close as possible to the traffic source stops spoofed traffic before it can proliferate or reach interfaces that do not have unicast RPF enabled.

## When Not to Enable Unicast RPF


Typically, you will not enable unicast RPF if:

- Switch interfaces are multihomed.
- Switch interfaces are trusted interfaces.
- BGP is carrying prefixes and some of those prefixes are not advertised or are not accepted by the ISP under its policy. (The effect in this case is the same as filtering an interface by using an incomplete access list.)
- Switch interfaces face the network core. Core-facing interfaces are usually asymmetrically routed.

An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, as shown in [Figure 5 on page 73](#). This means

that if an interface receives a packet, that interface does not match the forwarding table entry as the best return path back to the source. If the receiving interface is not the best return path to the source of a packet, unicast RPF causes the switch to discard the packet even though it comes from a valid source.

Figure 5: Asymmetrically Routed Interfaces



**NOTE:** Do not enable unicast RPF on EX3200, EX4200, and EX4300 switches if any switch interfaces are asymmetrically routed, because unicast RPF is enabled globally on all interfaces of these switches. All switch interfaces must be symmetrically routed for you to enable unicast RPF without the risk of the switch discarding traffic that you want to forward.

### Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches

On EX3200, EX4200, and EX4300 switches, the switch implements unicast RPF on a global basis. You cannot enable unicast RPF on a per-interface basis. Unicast RPF is globally disabled by default.

- When you enable unicast RPF on any interface, it is automatically enabled on all switch interfaces, including link aggregation groups (LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs).
- When you disable unicast RPF on the interface (or interfaces) on which you enabled unicast RPF, it is automatically disabled on all switch interfaces.



**NOTE:** You must explicitly disable unicast RPF on every interface on which it was explicitly enabled or unicast RPF remains enabled on all switch interfaces.

QFX switches, OCX switches, and EX3200 and EX4200 switches do not perform unicast RPF filtering on equal-cost multipath (ECMP) traffic. The unicast RPF check examines only one best return path to the packet source, but ECMP traffic employs an address block consisting of multiple paths. Using unicast RPF to filter ECMP traffic on these switches can result in the switch discarding packets that you want to forward because the unicast RPF filter does not examine the entire ECMP address block.

#### Related Documentation

- [Example: Configuring Unicast RPF on an EX Series Switch](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 74](#)

- [Disabling Unicast RPF \(CLI Procedure\) on page 75](#)

## Configuring Unicast RPF (CLI Procedure)

---

**Supported Platforms**   [EX Series, OCX1100, QFX Series standalone switches](#)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. When you enable unicast RPF, by default the switch forwards a packet only if the receiving interface is the best return path to the packet's unicast source address. This is known as strict mode unicast RPF. You can also enable loose mode, which means that the system checks to see if the packet has a source address with a corresponding prefix in the routing table but does not check whether the receiving interface is the best return path to the packet's unicast source address.



**NOTE:** On EX3200, EX4200, and EX4300 switches, you can enable unicast RPF only globally—that is, on all switch interfaces. You cannot enable unicast RPF on a per-interface basis.

Before you begin:

- On an EX8200, EX6200, QFX Series switch, or OCX Series switch, ensure that the selected switch interface is symmetrically routed before you enable unicast RPF. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.
- On an EX3200, EX4200, or EX4300 switch, ensure that *all* switch interfaces are symmetrically routed before you enable unicast RPF on an interface. When you enable unicast RPF on any interface, it is enabled globally on all switch interfaces. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.

To enable unicast RPF, configure it explicitly on a selected customer-edge interface:

```
[edit interfaces]
user@switch# set interface-name unit 0 family inet rpf-check
```

To enable unicast RPF loose mode, enter:

```
[edit interfaces]
user@switch# set interface-name unit 0 family inet rpf-check mode loose
```



**BEST PRACTICE:** On EX3200, EX4200, and EX4300 switches, unicast RPF is enabled globally on *all* switch interfaces, regardless of whether you configure it explicitly on only one interface or only on some interfaces.

On EX3200, EX4200, and EX4300 switches, we recommend that you enable unicast RPF explicitly on either all interfaces or only one interface. To avoid possible confusion, do not enable it on only some interfaces:

- Enabling unicast RPF explicitly on only one interface makes it easier if you choose to disable it in the future because you must explicitly disable unicast RPF on every interface on which you explicitly enabled it. If you explicitly enable unicast RPF on two interfaces and you disable it on only one interface, unicast RPF is still implicitly enabled globally on the switch. The drawback of this approach is that the switch displays the flag that indicates that unicast RPF is enabled only on interfaces on which unicast RPF is explicitly enabled, so even though unicast RPF is enabled on all interfaces, this status is not displayed.
- Enabling unicast RPF explicitly on all interfaces makes it easier to know whether unicast RPF is enabled on the switch because every interface shows the correct status. (Only interfaces on which you explicitly enable unicast RPF display the flag that indicates that unicast RPF is enabled.) The drawback of this approach is that if you want to disable unicast RPF, you must explicitly disable it on every interface. If unicast RPF is enabled on any interface, it is implicitly enabled on all interfaces.

#### Related Documentation

- [Example: Configuring Unicast RPF on an EX Series Switch](#)
- [Verifying Unicast RPF Status on page 76](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 75](#)
- [Troubleshooting Unicast RPF](#)
- [Understanding Unicast RPF on page 69](#)

## Disabling Unicast RPF (CLI Procedure)

**Supported Platforms** [EX Series, OCX1100, QFX Series standalone switches](#)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Unicast RPF filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. If the network configuration changes so that an interface that has unicast RPF enabled becomes a trusted interface or becomes asymmetrically routed (the interface that receives a packet is not the best return path to the packet's source), disable unicast RPF.

To disable unicast RPF on an EX3200, EX4200, or EX4300 switch, you must delete it from every interface on which you explicitly configured it. If you do not disable unicast RPF on every interface on which you explicitly enabled it, it remains implicitly enabled on all interfaces. If you attempt to delete unicast RPF from an interface on which it was not explicitly enabled, the **warning: statement not found** message appears. If you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces of the EX3200, EX4200, or EX4300 switch.

On EX8200, EX6200, QFX Series switches, and OCX Series switches, the switch does not apply unicast RPF to an interface unless you explicitly enable that interface for unicast RPF.

To disable unicast RPF, delete its configuration from the interface:

[edit interfaces]

user@switch# **delete ge-1/0/10 unit 0 family inet rpf-check**



**NOTE:** On EX3200, EX4200, and EX4300 switches, if you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces.

---

**Related  
Documentation**

- [Example: Configuring Unicast RPF on an EX Series Switch](#)
- [Verifying Unicast RPF Status on page 76](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 74](#)
- [Understanding Unicast RPF on page 69](#)

---

## Verifying Unicast RPF Status

---

**Supported Platforms** [EX Series, OCX1100, QFX Series standalone switches](#)

**Purpose** Verify that unicast reverse-path forwarding (RPF) is enabled and is working on the interface.

**Action** Use one of the **show interfaces *interface-name*** commands with either the **extensive** or **detail** options to verify that unicast RPF is enabled and working on the switch. The following example displays output from the **show interfaces ge- extensive** command.

```
user@switch> show interfaces ge-1/0/10 extensive
Physical interface: ge-1/0/10, Enabled, Physical link is Down
Interface index: 139, SNMP ifIndex: 58, Generation: 140
```

```

Link-level type: Ethernet, MTU: 1514, Speed: Auto, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
Auto-negotiation: Enabled, Remote fault: Online
Device flags   : Present Running
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues    : 8 supported, 8 maximum usable queues
Hold-times    : Up 0 ms, Down 0 ms
Current address: 00:19:e2:50:95:ab, Hardware address: 00:19:e2:50:95:ab
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
  Input bytes :                0                0 bps
  Output bytes :                0                0 bps
  Input packets:                0                0 pps
  Output packets:                0                0 pps
IPv6 transit statistics:
  Input bytes :                0
  Output bytes :                0
  Input packets:                0
  Output packets:                0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort                0                0                0

  1 assured-forw                0                0                0

  5 expedited-fo                0                0                0

  7 network-cont                0                0                0

Active alarms : LINK
Active defects : LINK
MAC statistics:
  Total octets                Receive      Transmit
  Total packets                0            0
  Unicast packets              0            0
  Broadcast packets            0            0
  Multicast packets            0            0
  CRC/Align errors             0            0
  FIFO errors                  0            0
  MAC control frames           0            0
  MAC pause frames             0            0
  Oversized frames             0
  Jabber frames                0
  Fragment frames              0
  VLAN tagged frames           0
  Code violations              0
Filter statistics:
  Input packet count           0
  Input packet rejects         0
  Input DA rejects             0

```

```

Input SA rejects                                0
Output packet count                            0
Output packet pad count                        0
Output packet error count                      0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Incomplete
Packet Forwarding Engine configuration:
Destination slot: 1

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes :                                0
Output bytes :                               0
Input packets:                               0
Output packets:                              0
IPv6 transit statistics:
Input bytes :                                0
Output bytes :                               0
Input packets:                               0
Output packets:                              0
Local statistics:
Input bytes :                                0
Output bytes :                               0
Input packets:                               0
Output packets:                              0
Transit statistics:
Input bytes :                                0                0 bps
Output bytes :                               0                0 bps
Input packets:                               0                0 pps
Output packets:                              0                0 pps
IPv6 transit statistics:
Input bytes :                                0
Output bytes :                               0
Input packets:                               0
Output packets:                              0
Protocol inet, Generation: 144, Route table: 0
Flags: uRPF
Addresses, Flags: Is-Preferred Is-Primary

```

**Meaning** The `show interfaces ge-1/0/10 extensive` command (and the `show interfaces ge-1/0/10 detail` command) displays in-depth information about the interface. The **Flags:** output field near the bottom of the display reports the unicast RPF status. If unicast RPF has not been enabled, the **uRPF** flag is not displayed.

On EX3200, EX4200, and EX4300 switches, unicast RPF is implicitly enabled on *all* switch interfaces, including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs) when you enable unicast RPF on a single interface. However, the unicast RPF status is shown as enabled only on interfaces for which you have explicitly configured unicast RPF. Thus, the **uRPF** flag is not displayed on interfaces for which you have not explicitly configured unicast RPF even though unicast RPF is implicitly enabled on all interfaces on EX3200 and EX4200 switches.

**Related Documentation**

- *show interfaces xe-*



- *Example: Configuring Unicast RPF on an EX Series Switch*
- [Configuring Unicast RPF \(CLI Procedure\) on page 74](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 75](#)
- *Troubleshooting Unicast RPF*

## Understanding Unknown Unicast Forwarding

**Supported Platforms** [EX Series, OCX1100, QFX Series standalone switches](#)

Unknown unicast traffic consists of unicast packets with unknown destination MAC addresses. By default, the switch floods these unicast packets that traverse a VLAN to all interfaces that are members of that VLAN. Forwarding this type of traffic can create unnecessary traffic that leads to poor network performance or even a complete loss of network service. This flooding of packets is known as a traffic storm.

To prevent a traffic storm, you can disable the flooding of unknown unicast packets to all VLAN interfaces by configuring specific VLANs or all VLANs to forward all unknown unicast traffic traversing them to a specific interface. You can configure multiple VLANs to forward unknown unicast packets to the same interface or configure different interfaces for different VLANs. This channels the unknown unicast traffic traversing VLANs to specific interfaces instead of flooding all interfaces.

**Related Documentation**

- *Configuring Unknown Unicast Forwarding (CLI Procedure)*
- [Configuring Unknown Unicast Forwarding \(CLI Procedure\) on page 79](#)
- *Understanding Storm Control on EX Series Switches*
- *Understanding Storm Control for Managing Traffic Levels on Switching Devices*
- *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*
- *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*

## Configuring Unknown Unicast Forwarding (CLI Procedure)

**Supported Platforms** [EX Series, OCX1100, QFX Series standalone switches](#)



**NOTE:** This task uses Junos OS for EX Series switches or QFX Series with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see *Configuring Unknown Unicast Forwarding (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*

Unknown unicast traffic consists of packets with unknown destination MAC addresses. By default, the switch floods these packets that traverse a VLAN to all interfaces associated with that VLAN. This flooding of packets is known as a traffic storm and can negatively impact network performance.

To prevent flooding unknown unicast traffic across the switch, configure unknown unicast forwarding to direct all unknown unicast packets within a VLAN to a specific interface. You can configure each VLAN to divert unknown unicast traffic to a different interface or use the same interface for multiple VLANs.

- [Configuring Unknown Unicast Forwarding on EX4300 Switches on page 80](#)
- [Configuring Unknown Unicast Forwarding on EX9200 Switches on page 80](#)

## Configuring Unknown Unicast Forwarding on EX4300 Switches

To configure unknown unicast forwarding options on EX4300 switches:

- Configure unknown unicast forwarding for a specific VLAN and specify the interface to which all unknown unicast traffic will be forwarded:

```
[edit switch-options]
user@switch# set unknown-unicast-forwarding vlan vlan-name interface interface-name
```

- Configure unknown unicast forwarding for all VLANs and specify the interface to which all unknown unicast traffic will be forwarded:

```
[edit switch-options]
user@switch# set unknown-unicast-forwarding vlan all interface interface-name
```

## Configuring Unknown Unicast Forwarding on EX9200 Switches

To configure unknown unicast forwarding on EX9200 switches, you must configure a flood filter and apply it to VLANs for which you want to configure unknown unicast forwarding. Flood filters are firewall filters that are applied only to broadcast, unknown unicast, and multicast (BUM) traffic. If a flood filter is configured, only traffic packets that are of the packet type **unknown-unicast** are forwarded to the interface on which unicast forwarding is configured. A next-hop group redirects the packets according to the action specified in the flood filter.

To configure the next-hop group that receives Layer 2 packets and then configure the interface to which these packets are forwarded:

1. Configure the **next-hop-group** action for the Layer 2 interface expected to receive unknown unicast packets:

```
[edit forwarding-options]
user@switch# set next-hop-group next-hop-group-name group-type layer-2
[edit forwarding-options]
user@switch# set next-hop-group next-hop-group-name interface interface-name
```

For example:

```
[edit forwarding-options]
user@switch# set next-hop-group uuf-nhg group-type layer-2
[edit forwarding-options]
user@switch# set next-hop-group uuf-nhg interface ge-3/1/7.0
```

2. Configure a firewall filter with family address type **ethernet-switching**:

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

For example:

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

3. Configure a term in the firewall filter for the interface that receives unknown unicast packets (the interface specified in Step 1) to discard unknown unicast packets:

```
[edit firewall family ethernet-switching filter filter-name]
user@switch# set term term-name from interface interface-name
user@switch# set term term-name from traffic-type unknown-unicast
user@switch# set term term-name then discard
```

For example:

```
[edit firewall family ethernet-switching filter uuf_filter]
user@switch# set term source-drop from interface ge-3/1/7.0
user@switch# set term source-drop from traffic-type unknown-unicast
user@switch# set term source-drop then discard
```

4. Configure a term in the firewall filter for unknown unicast packets to be flooded to the interface enabled for unknown unicast forwarding by using **next-hop-group** (in step 1):

```
[edit firewall family ethernet-switching filter filter-name]
user@switch# set term term-name from traffic-type unknown-unicast
user@switch# set term term-name then next-hop-group group-name
```

For example:

```
[edit firewall family ethernet-switching filter uuf_filter]
user@switch# set term uuf-flood from traffic-type unknown-unicast
user@switch# set term uuf-flood then next-hop-group uuf-nhg
```

5. Configure a default term for the firewall filter to forward packets other than unknown unicast packets:

```
[edit firewall family ethernet-switching filter filter-name]
user@switch# set term term-name then accept
```

For example:

```
[edit firewall family ethernet-switching filter uuf_filter]
user@switch# set term fwd-default then accept
```

6. Apply the filter as a flood filter on the VLAN that includes the interface which will receive unknown unicast packets:

```
[edit vlans vlan-name]
user@switch# set forwarding-options flood input filter-name
```

For example:

```
[edit vlans v1]
user@switch# set forwarding-options flood input uuf_filter
```

#### Related Documentation

- [Understanding Unknown Unicast Forwarding on page 79](#)
- [Verifying That Unknown Unicast Packets Are Forwarded to a Single Interface](#)



## PART 4

# Configuration Statements and Operational Commands

- [Configuration Statements \(Firewall Filters\) on page 85](#)
- [Configuration Statements \(Policers\) on page 97](#)
- [Operational Commands \(Firewall Filters\) on page 117](#)



## CHAPTER 4

# Configuration Statements (Firewall Filters)

- [family on page 86](#)
- [filter on page 87](#)
- [filter \(Layer 2 and Layer 3 Interfaces\) on page 88](#)
- [filter \(VLANs\) on page 89](#)
- [firewall on page 90](#)
- [from on page 91](#)
- [input \(Forwarding Table\) on page 92](#)
- [interface-specific on page 92](#)
- [output \(Forwarding Table\) on page 93](#)
- [term on page 94](#)
- [then \(Filters\) on page 95](#)

## family

**Supported Platforms** EX4600, MX Series, OCX1100, QFabric System, QFX Series standalone switches

**Syntax**

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

**Hierarchy Level** [edit [firewall](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  
**evpn** options introduced in Junos OS Release 15.1 for the MX Series.

**Description** Configure the fields a firewall filter can match on.

**Options** *family-name*—Type of addressing protocol:

- **ethernet-switching**—Filter Layer 2 Ethernet packets and Layer 3 (IP) packets (allows some Layer 3 filtering). Not supported on OCX Series switches.
- **evpn**—Filter Ethernet VPN (EVPN) packets.
- **inet**—Filter Layer 3 IPv4 packets (provides additional Layer 3 filter options).
- **inet6**—Filter Layer 3 IPv6 packets (provides additional Layer 3 filter options).
- **mpls**—Filter multiprotocol label switched packets. Not supported on OCX Series switches.

The remaining statements are explained separately.

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

**Related Documentation**

- *Firewall Filter Match Conditions and Actions*
- [Configuring Firewall Filters on page 20](#)
- [Overview of Firewall Filters on page 3](#)



## filter

<b>Supported Platforms</b>	EX4600, EX4600, EX4600, EX4600, EX4600, EX4600, EX4600, OCX1100, OCX1100, OCX1100, OCX1100, OCX1100, OCX1100, OCX1100, QFX Series standalone switches
<b>Syntax</b>	<pre>filter <i>filter-name</i> {   interface-specific;   term <i>term-name</i> {     from {       match-conditions;     }     then {       action;       action-modifiers;     }   } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">firewall family</a> <i>family-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure firewall filters.
<b>Options</b>	<p><b><i>filter-name</i></b>—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.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Firewall Filter Match Conditions and Actions</a></li> <li><a href="#">Configuring Firewall Filters on page 20</a></li> <li><a href="#">Overview of Firewall Filters on page 3</a></li> </ul>

## filter (Layer 2 and Layer 3 Interfaces)

---

<b>Supported Platforms</b>	QFabric System
<b>Syntax</b>	filter (input   output) <i>filter-name</i> ;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>family</b> <i>family-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Apply a firewall filter to traffic transiting a port or Layer 3 interface.
<b>Default</b>	All incoming traffic is accepted unmodified on the port or Layer 3 interface, and all outgoing traffic is sent unmodified from the port or Layer 3 interface.
<b>Options</b>	<p><b><i>filter-name</i></b>—Name of a firewall filter defined at the [edit firewall family <i>family-name</i> filter] hierarchy level.</p> <p><b>input</b>—Apply a firewall filter to traffic entering the port or Layer 3 interface.</p> <p><b>output</b>—Apply a firewall filter to traffic exiting the port or Layer 3 interface.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</li><li><a href="#">Configuring Firewall Filters on page 20</a></li><li><a href="#">Overview of Firewall Filters on page 3</a></li></ul>

## filter (VLANs)

---

### Supported Platforms

<b>Syntax</b>	<code>filter (input   output) <i>filter-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit vlans <i>vlan-name</i>],</code> <code>[edit vlans <i>vlan-name</i> forwarding-options]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Apply a firewall filter to traffic entering or exiting a VLAN.
<b>Default</b>	All incoming traffic is accepted unmodified to a VLAN, and all outgoing traffic is sent unmodified from a VLAN.
<b>Options</b>	<p><b><i>filter-name</i></b>—Name of a firewall filter defined at the <code>[edit firewall family <i>family-name</i> filter]</code> hierarchy level.</p> <p><b>input</b>—Apply a firewall filter to VLAN ingress traffic.</p> <p><b>output</b>—Apply a firewall filter to VLAN egress traffic.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Firewall Filters on page 20</a></li> <li>• <a href="#">Overview of Firewall Filters on page 3</a></li> </ul>

## firewall

### Supported Platforms

```
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;
        }
    }
    three-color-policer policer-name {
        action {
            loss-priority high then discard;
        }
        single-rate {
            (color-aware | color-blind);
            committed-information-rate bps;
            committed-burst-size bytes;
            excess-burst-size bytes;
        }
        two-rate {
            (color-aware | color-blind);
            committed-information-rate bps;
            committed-burst-size bytes;
            peak-information-rate bps;
            peak-burst-size bytes;
        }
    }
}
```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

<b>Description</b>	Configure firewall filters and policers.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Firewall Filter Match Conditions and Actions</i></li> <li>• <a href="#">Configuring Firewall Filters on page 20</a></li> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li> <li>• <a href="#">Overview of Firewall Filters on page 3</a></li> </ul>

## from

---

### Supported Platforms

<b>Syntax</b>	from { <i>match-conditions</i> ; }
<b>Hierarchy Level</b>	[edit <b>firewall</b> <i>family</i> <i>family-name</i> <b>filter</b> <i>filter-name</i> <b>term</b> <i>term-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Match packet fields to values specified in a match condition. If the <b>from</b> statement is not included in a firewall filter configuration, all packets are considered to match and the actions and action modifiers in the <b>then</b> statement are implemented.
<b>Options</b>	<b>match-conditions</b> —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 <b>then</b> statement to be implemented.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Firewall Filter Match Conditions and Actions</i></li> <li>• <a href="#">Configuring Firewall Filters on page 20</a></li> <li>• <a href="#">Understanding Firewall Filter Match Conditions on page 9</a></li> </ul>

## input (Forwarding Table)

---

<b>Supported Platforms</b>	M Series, PTX Series, QFX Series standalone switches, T Series
<b>Syntax</b>	input <i>filter-name</i> ;
<b>Hierarchy Level</b>	[edit forwarding-options family (inet   inet6   mpls   vpls) filter], [edit routing-instances <i>routing-instance-name</i> forwarding-options family (inet   inet6   mpls   vpls) filter]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches..
<b>Description</b>	Apply a forwarding table filter to ingress traffic of the forwarding table.
<b>Options</b>	<i>filter-name</i> —Name of the applied filter.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Applying Forwarding Table Filters</a></li></ul>

## interface-specific

---

<b>Supported Platforms</b>	EX4600, OCX1100, QFX Series standalone switches
<b>Syntax</b>	interface-specific;
<b>Hierarchy Level</b>	[edit <a href="#">firewall family</a> <i>family-name</i> <a href="#">filter</a> <i>filter-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure separate counters for each interface to which a filter is applied.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Firewall Filter Match Conditions and Actions</a></li><li>• <a href="#">Configuring Firewall Filters on page 20</a></li><li>• <a href="#">Overview of Firewall Filters on page 3</a></li></ul>

## output (Forwarding Table)

---

<b>Supported Platforms</b>	M Series, PTX Series, QFX Series standalone switches, T Series
<b>Syntax</b>	output <i>filter-name</i> ;
<b>Hierarchy Level</b>	[edit forwarding-options family (inet   inet6   mpls) filter], [edit routing-instances <i>routing-instance-name</i> forwarding-options family (inet   inet6   mpls) filter]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches..
<b>Description</b>	Configure filtering on the egress traffic of the forwarding table.
<b>Options</b>	<i>filter-name</i> —Name of the applied filter.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Applying Forwarding Table Filters</i></li> </ul>

## term

---

### Supported Platforms

**Syntax**    `term term-name {  
          from {  
            match-conditions;  
          }  
          then {  
            action;  
            action-modifiers;  
          }  
          }  
          }`

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

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**    Define a firewall filter term.

**Options**    **term-name**—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.

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](#)
- [Configuring Firewall Filters on page 20](#)
- [Overview of Firewall Filters on page 3](#)



## then (Filters)

---

### Supported Platforms

**Syntax**    `then {  
          action;  
          action-modifiers;  
          }`

**Hierarchy Level**    `[edit firewall family family-name filter filter-name term term-name]`

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**    Configure a firewall filter action.

**Options**    **action**—Actions to accept, discard, or forward packets that match all conditions specified in a filter term.

**action-modifiers**—Additional actions to analyze, classify, count, or police packets that match all conditions specified in a filter term.

**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*
- [Configuring Firewall Filters on page 20](#)
- [Understanding Firewall Filter Match Conditions on page 9](#)



## CHAPTER 5

# Configuration Statements (Policers)

- [action on page 98](#)
- [bandwidth-limit on page 98](#)
- [burst-size-limit on page 99](#)
- [color-aware on page 100](#)
- [color-blind on page 101](#)
- [committed-burst-size on page 102](#)
- [committed-information-rate on page 103](#)
- [excess-burst-size on page 104](#)
- [filter-specific on page 105](#)
- [firewall on page 106](#)
- [if-exceeding on page 107](#)
- [loss-priority high then discard \(Three-Color Policer\) on page 108](#)
- [peak-burst-size on page 109](#)
- [peak-information-rate on page 110](#)
- [policer on page 111](#)
- [single-rate on page 112](#)
- [then \(Policers\) on page 113](#)
- [three-color-policer on page 114](#)
- [two-rate on page 115](#)

## action

---

<b>Supported Platforms</b>	EX4600, OCX1100, QFX Series standalone switches
<b>Syntax</b>	<pre>action {     loss-priority high then discard; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer name</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Discard traffic on a logical interface using tricolor marking policing.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.

## bandwidth-limit

---

<b>Supported Platforms</b>	EX4600, EX4600, EX4600, EX4600, EX4600, EX4600, OCX1100, OCX1100, OCX1100, OCX1100, OCX1100, OCX1100, QFX Series standalone switches
<b>Syntax</b>	<pre>bandwidth-limit <i>bps</i>;</pre>
<b>Hierarchy Level</b>	[edit <a href="#">firewall policer policer-name if-exceeding</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Specify the traffic rate in bits per second.
<b>Options</b>	<b><i>bps</i></b> —Traffic rate in bits per second. Specify <i>bps</i> as a decimal value or as a decimal number followed by one of the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000). <b>Range:</b> 32000 bps (32 Kbps) through 10,000,000,000 bps (10 Gbps)
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li><li>• <a href="#">Overview of Policers on page 37</a></li></ul>

## burst-size-limit

---

### Supported Platforms

**Syntax** `burst-size-limit bytes;`

**Hierarchy Level** `[edit firewall policer policer-name if-exceeding]`

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Specify the maximum allowed burst size to control the amount of traffic bursting.

**Options** *bytes*—Decimal value or a decimal number followed by k (thousand), m (million), or g (giga).

**Range:** 1 through 2,147,450,880 bytes (2147 MB)

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

**Related Documentation**

- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55](#)
- [Overview of Policers on page 37](#)

## color-aware

---

<b>Supported Platforms</b>	EX4600, OCX1100, QFabric System, QFX Series standalone switches
<b>Syntax</b>	color-aware;
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate], [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure the way preclassified packets are metered. In color-aware mode, the switch can assign a higher packet-loss priority, but cannot assign a lower packet loss priority (PLP). For example, suppose an upstream device assigns medium-high PLP to a packet because the packet exceeded its committed information rate (CIR). The switch cannot change the PLP to low even if the packet conforms to the configured CIR of the appropriate interface. On the other hand, if an upstream device assigns low PLP to a packet but the packet exceeds the CIR and committed burst size (CBS) of the switch interface, the switch can increase the PLP to medium-high.
<b>Default</b>	If you omit the <b>color-aware</b> statement, the default behavior is color-aware mode.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Overview of Policers on page 37</a></li><li>• <a href="#">Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 44</a></li><li>• <a href="#">Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 46</a></li><li>• <a href="#">color-blind on page 101</a></li></ul>


## color-blind

---

<b>Supported Platforms</b>	EX4600, OCX1100, QFabric System, QFX Series standalone switches
<b>Syntax</b>	color-blind;
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate], [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure the way preclassified packets are metered. In color-blind mode, the switch ignores any preclassification of packets and can assign a higher or lower packet loss priority (PLP). For example, suppose an upstream device assigns medium-high PLP to a packet because the packet exceeded the CIR on the upstream device. The switch can change the PLP to low if the packet conforms to the CIR of the appropriate interface.
<b>Default</b>	If you omit the <b>color-blind</b> statement, the default behavior is color-aware mode.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Overview of Policers on page 37</a></li> <li>• <a href="#">Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 43</a></li> <li>• <a href="#">Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 46</a></li> <li>• <a href="#">Configuring Color-Blind Egress Policers for Medium-Low PLP on page 54</a></li> <li>• <a href="#">color-aware on page 100</a></li> </ul>


## committed-burst-size

---

Supported Platforms	EX4600, OCX1100, QFabric System, QFX Series standalone switches
Syntax	committed-burst-size <i>bytes</i> ;
Hierarchy Level	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate], [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the maximum number of bytes allowed for incoming traffic to burst above the committed information rate and still be marked with low packet loss priority (green).
<div> <b>NOTE:</b> When you include the <code>committed-burst-size</code> statement in the configuration, you must also include the <code>committed-information-rate</code> statement at the same hierarchy level.</div>	
Options	<b>bytes</b> —Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000). <b>Range:</b> 512 bytes through 268435456 bytes (268 MB)
Required Privilege Level	<b>firewall</b> —To view this statement in the configuration. <b>firewall-control</b> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li><li>• <a href="#">Overview of Policers on page 37</a></li></ul>




## committed-information-rate

<b>Supported Platforms</b>	EX4600, OCX1100, QFabric System, QFX Series standalone switches
<b>Syntax</b>	committed-information-rate <i>bits-per-second</i> ;
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate], [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure the guaranteed bandwidth under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green).
<div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p><b>NOTE:</b> When you include the <code>committed-information-rate</code> statement in the configuration, you must also include the <code>committed-burst-size</code> statement at the same hierarchy level.</p> </div> </div>	
<b>Options</b>	<p><i>bits-per-second</i>—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 32,000 bps through 10,000,000,000 bps (10 gbps)</p>
<b>Required Privilege Level</b>	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li> <li>• <a href="#">Overview of Policers on page 37</a></li> </ul>

## excess-burst-size

---

Supported Platforms	EX4600, OCX1100, QFabric System, QFX Series standalone switches
Syntax	excess-burst-size <i>bytes</i> ;
Hierarchy Level	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the maximum number of bytes allowed for incoming traffic to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red).
<div> <b>NOTE:</b> When you include the <code>excess-burst-size</code> statement in the configuration, you must also include the <code>committed-burst-size</code> and <code>committed-information-rate</code> statements at the same hierarchy level.</div>	
Options	<b>bytes</b> —Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000). <b>Range:</b> 512 bytes through 268435456 bytes (268 MB)
Required Privilege Level	<b>firewall</b> —To view this statement in the configuration. <b>firewall-control</b> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li><li>• <a href="#">Overview of Policers on page 37</a></li></ul>

## filter-specific

---

<b>Supported Platforms</b>	EX4600, OCX1100, QFabric System, QFX Series standalone switches
<b>Syntax</b>	filter-specific;
<b>Hierarchy Level</b>	[edit <a href="#">firewall policer</a> <i>policer-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	<p>Configure a policer to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. If you use a filter-specific policer in multiple terms, both of the following are true:</p> <ul style="list-style-type: none"> <li>• Traffic is policed at the aggregate rate. For example, if you create a policer that has a bandwidth limit of 100 Mbps and use the policer in two terms, the total allowed bandwidth for both terms is 100 Mbps—not 100 Mbps for each term.</li> <li>• The implicit counter counts all the packets that are matched by any of the terms. For example, if you reference the same filter-specific policer in term1 and term2, and term1 matches 1000 packets and term2 matches 500 packets, the implicit counter shows 1500 matches for the policer.</li> </ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li> <li>• <a href="#">Overview of Policers on page 37</a></li> </ul>

## firewall

### Supported Platforms

```
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;
        }
    }
    three-color-policer policer-name {
        action {
            loss-priority high then discard;
        }
        single-rate {
            (color-aware | color-blind);
            committed-information-rate bps;
            committed-burst-size bytes;
            excess-burst-size bytes;
        }
        two-rate {
            (color-aware | color-blind);
            committed-information-rate bps;
            committed-burst-size bytes;
            peak-information-rate bps;
            peak-burst-size bytes;
        }
    }
}
```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

<b>Description</b>	Configure firewall filters and policers.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Firewall Filter Match Conditions and Actions</i></li> <li>• <a href="#">Configuring Firewall Filters on page 20</a></li> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li> <li>• <a href="#">Overview of Firewall Filters on page 3</a></li> </ul>

## if-exceeding

---

### Supported Platforms


<b>Syntax</b>	<pre>if-exceeding {     bandwidth-limit <i>bps</i>;     burst-size-limit <i>bytes</i>; }</pre>
<b>Hierarchy Level</b>	[edit <b>firewall policer</b> <i>policer-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure policer rate limits.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li> <li>• <a href="#">Overview of Policers on page 37</a></li> </ul>

## loss-priority high then discard (Three-Color Policer)

---


<b>Supported Platforms</b>	EX4600, OCX1100, QFX Series standalone switches
<b>Syntax</b>	loss-priority high then discard;
<b>Hierarchy Level</b>	[edit firewall <b>three-color-policer</b> <i>policer-name</i> <b>action</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	<p>For packets with high loss priority, discard the packets. The loss priority setting is not configurable. Include this statement if you do not want the switch to forward packets that have high packet-loss priority.</p> <p>For single-rate three-color policers, Junos OS assigns high loss priority to packets that exceed the committed information rate and the excess burst size.</p> <p>For two-rate three-color policers, Junos OS assigns high loss priority to packets that exceed the peak information rate and the peak burst size.</p>
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li><li>• <a href="#">Overview of Policers on page 37</a></li></ul>

## peak-burst-size

<b>Supported Platforms</b>	EX4600, OCX1100, QFabric System, QFX Series standalone switches
<b>Syntax</b>	peak-burst-size <i>bytes</i> ;
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red).
<div>  <p><b>NOTE:</b> When you include the <code>peak-burst-size</code> statement in the configuration, you must also include the <code>committed-burst-size</code> and <code>peak-information-rate</code> statements at the same hierarchy level.</p> </div>	
<b>Options</b>	<p><b>bytes</b>—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1500 bytes through 100,000,000,000 bytes (100 GB)</p>
<b>Required Privilege Level</b>	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li> <li>• <a href="#">Overview of Policers on page 37</a></li> </ul>

## peak-information-rate

---

Supported Platforms	EX4600, OCX1100, QFabric System, QFX Series standalone switches
Syntax	peak-information-rate <i>bits-per-second</i> ;
Hierarchy Level	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the maximum achievable rate. Packets that exceed the committed information rate (CIR) but are below the peak information rate (PIR) are marked with medium-high packet loss priority (yellow). Packets that exceed the PIR are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR.
<div> <b>NOTE:</b> When you include the <code>peak-information-rate</code> statement in the configuration, you must also include the <code>committed-information-rate</code> and <code>peak-burst-size</code> statements at the same hierarchy level.</div>	
Options	<i>bits-per-second</i> —Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000). <b>Range:</b> 32,000 bps through 10,000,000,000 bps (10 gbps)
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"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li><li>• <a href="#">Overview of Policers on page 37</a></li></ul>



## policer

Supported Platforms	QFabric System
Syntax	<pre> policer <i>policer-name</i> {   filter-specific;   if-exceeding {     bandwidth-limit <i>bps</i>;     burst-size-limit <i>bytes</i>;   }   then {     <i>policer-action</i>;   } } </pre>
Hierarchy Level	[edit <a href="#">firewall</a> ]
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure policer rate limits and actions. To activate a policer, you must include the <b>policer</b> action modifier in the <b>then</b> statement in a firewall filter term.</p> <p>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. If you use the same policer in multiple terms—either within the same filter or across filters—the policer’s implicit counter is used to count packets that are policed in all of these terms. If you want to obtain separate packet counts for each term, use these approaches:</p> <ul style="list-style-type: none"> <li>• Configure a unique policer for each term.</li> <li>• Configure only one policer, but use a unique, explicit counter in each term.</li> </ul>
Options	<p><b><i>policer-name</i></b>—Name that identifies the policer. The name can contain letters, numbers, hyphens (-), and can be up to 64 characters long.</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"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li> <li>• <a href="#">Configuring Firewall Filters on page 20</a></li> <li>• <a href="#">Overview of Policers on page 37</a></li> </ul>

## single-rate

---

Supported Platforms	EX4600, OCX1100, QFX Series standalone switches
Syntax	<pre>single-rate {   (color-aware   color-blind);   committed-information-rate <i>bps</i>;   committed-burst-size <i>bytes</i>;   excess-burst-size <i>bytes</i>; }</pre>
Hierarchy Level	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> ]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Configure a single-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).</p> <p>Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the EBS are assigned medium-high loss priority (yellow). Packets that exceed the EBS are assigned high loss priority (red).</p> <p>Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.</p> <p>The remaining statements are explained separately.</p>
Options	<b><i>policer-name</i></b> —Name of the three-color policer. Use this name when you apply the policer to an interface.
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"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55</a></li><li>• <a href="#">Overview of Policers on page 37</a></li></ul>

## then (Policers)

---

### Supported Platforms

**Syntax**    `then {  
              policer-action;  
          }`

**Hierarchy Level**    `[edit firewall policer policer-name]`

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**    Configure a policer action.

**Options**    *policer-action*—Allowed policer actions are **discard**, **loss-priority high**, and **loss-priority low**. **discard** causes the system to drop traffic that exceeds the rate limits defined by the policer. Use **loss-priority high** to allow the system to forward matching traffic in some cases.



**NOTE:** If you specify a policer in an egress firewall filter, the only supported action is **discard**.

---

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

**Related Documentation**

- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55](#)
- [Configuring Firewall Filters on page 20](#)
- [Overview of Policers on page 37](#)

## three-color-policer

---

**Supported Platforms** [EX4600, OCX1100, QFX Series standalone switches](#)

**Syntax** `three-color-policer policer-name {  
 action {  
 loss-priority high then discard;  
 }  
 single-rate {  
 (color-aware | color-blind);  
 committed-information-rate bps;  
 committed-burst-size bytes;  
 excess-burst-size bytes;  
 }  
 two-rate {  
 (color-aware | color-blind);  
 committed-information-rate bps;  
 committed-burst-size bytes;  
 peak-information-rate bps;  
 peak-burst-size bytes;  
 }  
}`

**Hierarchy Level** `[edit firewall],  
[edit logical-systems logical-system-name firewall]`

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure a three-color policer.

**Options** *policer-name*—Name of the three-color policer. Use this name when you apply the policer to an interface.

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**

- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 55](#)
- [Overview of Policers on page 37](#)

## two-rate

**Supported Platforms** EX4600, OCX1100, QFX Series standalone switches

**Syntax** `two-rate {  
     (color-aware | color-blind);  
     committed-information-rate bps;  
     committed-burst-size bytes;  
     peak-information-rate bps;  
     peak-burst-size bytes;  
 }`

**Hierarchy Level** [edit `firewall three-color-policer policer-name`]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure a two-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), peak information rate (PIR), and peak burst size (PBS).

Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the PIR or the PBS are assigned medium-high loss priority (yellow). Packets that exceed the PIR and the PBS are assigned high loss priority (red).

Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.

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.



## CHAPTER 6

# Operational Commands (Firewall Filters)

- `clear firewall`
- `show firewall`
- `show firewall policer`
- `show interfaces filters`
- `show pfe filter hw summary`

## clear firewall

---

<b>Supported Platforms</b>	EX4600, OCX1100, QFabric System, QFX Series standalone switches
<b>Syntax</b>	clear firewall (all   counter <i>counter-name</i>   filter <i>filter-name</i> )
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	<p>Clear statistics provided by firewall filters.</p> <p>When you clear the counters of a filter, this not only impacts the counters shown by the CLI, but also the ones tracked by SNMP 2.</p>
<b>Options</b>	<p><b>all</b>—Clear the packet and byte counts for all firewall filter counters and clear the packet counts for all policer counters.</p> <p><b>counter <i>counter-name</i></b>—Clear the packet and byte counts for the specified firewall filter counter.</p> <p><b>filter <i>filter-name</i></b>—Clear the packet and byte counts for the specified firewall filter.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Verifying That Firewall Filters Are Operational on page 25</a></li><li>• <a href="#">Verifying That Two-Color Policers Are Operational on page 57</a></li><li>• <a href="#">Overview of Firewall Filters on page 3</a></li><li>• <a href="#">Overview of Policers on page 37</a></li></ul>

## Sample Output

### clear firewall all

```
user@switch> clear firewall all
```

### clear firewall counter

```
user@switch> clear firewall counter port-filter-counter
```

### clear firewall filter

```
user@switch> clear firewall filter ingress-port-filter
```



## show firewall

**Supported Platforms** [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

**Syntax** `show firewall`  
`<counter counter-name>`  
`<filter filter-name>`  
`<log <detail | interface interface-name>>`  
`<terse>`

**Release Information** Command introduced in Junos OS Release 11.1 for the QFX Series.  
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Display statistics about configured firewall filters.

**Options** `counter counter-name`—(Optional) Display statistics about a particular firewall filter counter.

`filter filter-name`—(Optional) Display statistics about a particular firewall filter.

`log`—(Optional) Display log entries for all firewall filter activity.

`terse`—(Optional) Display firewall filter names only.

**Required Privilege Level** view

**Related Documentation**

- [Verifying That Firewall Filters Are Operational on page 25](#)
- [Verifying That Two-Color Policers Are Operational on page 57](#)
- [Overview of Firewall Filters on page 3](#)
- [Overview of Policers on page 37](#)

**List of Sample Output** [show firewall on page 120](#)  
[show firewall filter \*filter-name\* on page 121](#)  
[show firewall counter \*counter-name\* on page 121](#)  
[show firewall log on page 121](#)  
[show firewall log detail on page 121](#)

**Output Fields** [Table 13 on page 119](#) lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

**Table 13: show firewall Output Fields**

Field Name	Field Description	Level of Output
Filter	Name of the filter that is configured at the <code>[edit firewall family <i>family-name</i> filter]</code> hierarchy level.	All levels

Table 13: show firewall Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Counters</b>	Display filter counter information: <ul style="list-style-type: none"> <li>Name—Name of a filter counter that has been configured with the <b>count</b> firewall filter action modifier.</li> <li>Bytes—Number of bytes that match the filter term where the <b>count</b> action modifier was specified.</li> <li>Packets—Number of packets that matched the filter term where the <b>count</b> action modifier was specified.</li> </ul>	All levels
<b>Policers</b>	Display policer information: <ul style="list-style-type: none"> <li>Name—Name of the policer that is configured at the <b>[edit firewall policer]</b> hierarchy level.</li> <li>Packets—Number of packets that matched the filter term where the <b>policer</b> action modifier was specified. This is the number of packets that exceeded the rate limits that the policer specifies.</li> </ul>	All levels
<b>Action</b>	Filter action: <ul style="list-style-type: none"> <li><b>A</b>—Accept</li> <li><b>D</b>—Discard</li> </ul>	All levels
<b>Interface</b>	Interface on which the firewall filter is applied.	All levels
<b>Protocol</b>	Name of the packet protocol.	All levels
<b>Packet Length</b>	Length of the packet.	All levels
<b>Src Addr</b>	Source address of the packet.	All levels
<b>Dest Addr</b>	Destination address of the packet.	All levels

## Sample Output

### show firewall

```

user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                               Bytes          Packets
counter-employee-web              0              0
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                               Bytes          Packets
icmp-counter                      560           10
Policers:
Name                               Packets
icmp-connection-policer          10
tcp-connection-policer           0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest

```

**show firewall filter filter-name**

```

user@switch> show firewall filter ingress-port-limit-tcp-icmp
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                                     Bytes      Packets
icmp-counter                             560         10
Policers:
Name                                     Packets
icmp-connection-policer                  10
tcp-connection-policer                    0

```

**show firewall counter counter-name**

```

user@switch> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                                     Bytes      Packets
icmp-counter                             560         10

```

**show firewall log**

```

user@switch> show firewall log
Log :

Time      Filter  Action Interface  Protocol  Src Addr
Dest Addr
08:00:53  pfe      R      ge-1/0/6.0    ICMP      192.168.3.5
192.168.3.4
08:00:52  pfe      R      ge-1/0/6.0    ICMP      192.168.3.5
192.168.3.4
08:00:51  pfe      R      ge-1/0/6.0    ICMP      192.168.3.5
192.168.3.4
08:00:50  pfe      R      ge-1/0/6.0    ICMP      192.168.3.5
192.168.3.4
08:00:49  pfe      R      ge-1/0/6.0    ICMP      192.168.3.5
192.168.3.4
08:00:48  pfe      R      ge-1/0/6.0    ICMP      192.168.3.5
192.168.3.4
08:00:47  pfe      R      ge-1/0/6.0    ICMP      192.168.3.5
192.168.3.4

```

**show firewall log detail**

```

user@switch> show firewall log detail
Log :

Time of Log: 2010-10-13 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: 2010-10-13 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: 2010-10-13 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: 2010-10-13 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: 2010-10-13 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: 2010-10-13 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 firewall policer

**Supported Platforms** [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

**Syntax** `show firewall policer  
<policer-name>`

**Release Information** Command introduced in Junos OS Release 11.1 for the QFX Series.  
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Display statistics about configured policers.

**Options** **none**—Display the count of policed packets for all configured policers.  
**policer-name**—(Optional) Display the count of policed packets for the specified policer.

**Required Privilege Level** view

**Related Documentation**

- [Verifying That Firewall Filters Are Operational on page 25](#)
- [Verifying That Two-Color Policers Are Operational on page 57](#)
- [Overview of Firewall Filters on page 3](#)
- [Overview of Policers on page 37](#)

**List of Sample Output** [show firewall policer on page 123](#)  
[show firewall policer policer-name on page 124](#)

**Output Fields** [Table 14 on page 123](#) lists the output fields for the **show firewall policer** command. Output fields are listed in the approximate order in which they appear.

**Table 14: show firewall policer Output Fields**

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

## Sample Output

### show firewall policer

```
user@switch> show firewall 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 firewall policer policer-name**

```
user@switch> show firewall policer tcp-connection-policer
Filter: ingress-port-filter
Policers:
Name                                     Packets
tcp-connection-policer                   0
```

## show interfaces filters

<b>Supported Platforms</b>	EX4600, OCX1100, QFabric System, QFX Series standalone switches
<b>Syntax</b>	show interfaces filters <interface-name>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Display firewall filters that are configured on each interface in a switch.
<b>Options</b>	<b>none</b> —Display firewall filter information about all interfaces.  <b>interface-name</b> —(Optional) Display firewall filter information about a particular interface.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show firewall on page 119</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show interfaces filters on page 125</a> <a href="#">show interfaces filters interface-name on page 126</a>
<b>Output Fields</b>	<a href="#">Table 15 on page 125</a> lists the output fields for the <b>show interfaces filters</b> command. Output fields are listed in the approximate order in which they appear.

Table 15: show interfaces filters Output Fields

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

## Sample Output

### show interfaces filters

```
user@switch> show interfaces filters
```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/0/6	up	up			
ge-0/0/6.0	up	up	inet		
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@switch> show interfaces filters ge-0/0/6
```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/0/6	up	up			
ge-0/0/6.0	up	up	inet		



## show pfe filter hw summary

<b>Supported Platforms</b>	<a href="#">QFX Series standalone switches</a>
<b>Syntax</b>	show pfe filter hw summary
<b>Release Information</b>	Command introduced in Junos OS Release 14.1X53-D10 for the QFX Series.
<b>Description</b>	<p>Display a summary of the access control list (ACL; also known as firewall filter) ternary content-addressable memory (TCAM) hardware utilization to show the allocated, used, and free TCAM entry space.</p> <p>Command supported on standalone QFX Series switches, QFX5100-only (pure QFX5100) Virtual Chassis Fabric (VCF), QFX5100-only (pure QFX5100) Virtual Chassis (VC), and QFX3500-only (pure QFX3500) VC.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Planning the Number of Firewall Filters to Create on page 14</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show pfe filter hw summary on page 128</a>
<b>Output Fields</b>	<a href="#">Table 16 on page 127</a> lists the output fields for the <b>show pfe filter hw summary</b> command. Output fields are listed in the approximate order in which they appear.

**Table 16: show pfe filter hw summary Output Fields**

Field Name	Field Description
<b>Group</b>	<p>ACL ingress and egress filter groups:</p> <ul style="list-style-type: none"> <li>• iRACL group—ingress routing ACL filter group</li> <li>• iVACL group—ingress VLAN ACL filter group</li> <li>• iPACL group—ingress port ACL filter group</li> <li>• ePACL group—egress port ACL filter group</li> <li>• eVACL group—egress VLAN ACL filter group</li> <li>• eRACL group—egress routing ACL filter group</li> <li>• eRACL IPv6 group—egress IPv6 routing ACL filter group</li> </ul>
<b>Group-ID</b>	Internal identification number of the filter group.
<b>Allocated</b>	Number of TCAM filter entries allocated to the filter group.
<b>Used</b>	Number of TCAM filter entries used by the filter group.
<b>Free</b>	Number of TCAM filter entries available for use by the filter group.

## Sample Output

### show pfe filter hw summary

```
user@switch> show pfe filter hw summary
```

Group	Group-ID	Allocated	Used	Free
-----				
> Ingress filter groups:				
iRACL group	14	512	4	508
iVACL group	13	512	2	510
iPACL group	12	256	2	254
> Egress filter groups:				
ePACL group	20	256	3	253
eVACL group	21	256	4	252
eRACL group	22	256	245	11
eRACL IPV6 group	24	256	3	253