



Junos[®] OS for EX Series Ethernet Switches

Security Feature Guide for EX4600 Switches

Release

15.1



Modified: 2016-11-30

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
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.

Junos[®] OS for EX Series Ethernet Switches Security Feature Guide for EX4600 Switches
Release 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	xv
	Documentation and Release Notes	xv
	Supported Platforms	xv
	Using the Examples in This Manual	xv
	Merging a Full Example	xvi
	Merging a Snippet	xvi
	Documentation Conventions	xvii
	Documentation Feedback	xix
	Requesting Technical Support	xix
	Self-Help Online Tools and Resources	xix
	Opening a Case with JTAC	xx
Part 1	Firewall Filters	
Chapter 1	Configuring Firewall Filters	3
	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	8
	Understanding Firewall Filter Match Conditions	9
	Filter Match Conditions	9
	Numeric Filter Match Conditions	10
	Interface Filter Match Conditions	10
	IP Address Filter Match Conditions	11
	MAC Address Filter Match Conditions	11
	Bit-Field Filter Match Conditions	12
	Firewall Filter Match Conditions and Actions	13
	Understanding How a Firewall Filter Tests a Protocol	28
	Understanding Firewall Filter Planning	28
	Planning the Number of Firewall Filters to Create	30
	Understanding How Many Firewall Filters Are Supported	30
	Egress Filters	31
	Avoid Configuring too Many Filters	31
	Configuring TCAM Error Messages	32
	Policers can Limit Egress Filters	32
	Planning for Filter-Specific Policers	33
	Planning for Filter-Based Forwarding	33

Understanding Firewall Filter Processing Points for Bridged and Routed Packets	34
Configuring Firewall Filters	35
Configuring a Firewall Filter	35
Applying a Firewall Filter to a Port	37
Applying a Firewall Filter to a VLAN	37
Applying a Firewall Filter to a Layer 3 (Routed) Interface	38
Verifying That Firewall Filters Are Operational	38
Applying Firewall Filters to Interfaces	39
Understanding Filter-Based Forwarding	40
Applying Firewall Filters to Interfaces	40
Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device	41
Configuring MPLS Firewall Filters and Policers	44
Configuring MPLS Firewall Filters	45
Examples: Configuring MPLS Firewall Filters	45
Configuring Policers for LSPs	46
LSP Policer Limitations	46
Configuring a Firewall Filter to De-Encapsulate GRE Traffic on a QFX5100 or OCX Switch	47
Configuring a Filter to De-Encapsulate GRE Traffic	47
Applying the Filter to an Interface	48
Monitoring Firewall Filter Traffic	49
Monitoring Traffic for All Firewall Filters and Policers That Are Configured	49
Monitoring Traffic for a Specific Firewall Filter	49
Monitoring Traffic for a Specific Policer	50
Troubleshooting Firewall Filter Configuration	50
Firewall Filter Configuration Returns a No Space Available in TCAM Message	50
Filter Counts Previously Dropped Packet	52
Matching Packets Not Counted	53
Counter Reset When Editing Filter	53
Cannot Include loss-priority and policer Actions in Same Term	53
Cannot Egress Filter Certain Traffic Originating on QFX Switch	54
Firewall Filter Match Condition Not Working with Q-in-Q Tunneling	54
Egress Firewall Filters with Private VLANs	54
Egress Filtering of L2PT Traffic Not Supported	55
Cannot Drop BGP Packets in Certain Circumstances	55
Invalid Statistics for Policer	55
Policers can Limit Egress Filters	55
Part 2	
Chapter 2	
Policers	
Configuring Policers	59
Overview of Policers	59
Policer Overview	60
Policer Types	60
Policer Actions	61

Policer Colors	62
Filter-Specific Policers	62
Suggested Naming Convention for Policers	63
Policer Counters	63
Policer Algorithms	63
How Many Policers Are Supported?	64
Policers Can Limit Egress Firewall Filters	64
Understanding Policers with Link Aggregation Groups	65
Understanding Color-Blind Mode for Single-Rate Tricolor Marking	65
Understanding Color-Aware Mode for Single-Rate Tricolor Marking	66
Summary of PLP Changes	66
Effect on Green Packets (Low PLP)	67
Effect on Yellow Packets (Medium PLP)	67
Effect on Red Packets (High PLP)	67
Understanding Color-Blind Mode for Two-Rate Tricolor Marking	68
Understanding Color-Aware Mode for Two-Rate Tricolor Marking	68
Summary of PLP Changes	68
Effect on Green Packets (Low PLP)	69
Effect on Yellow Packets (Medium PLP)	69
Effect on Red Packets (High PLP)	70
Example: Using Two-Color Policers and Prefix Lists	70
Example: Using Policers to Manage Oversubscription	73
Assigning Forwarding Classes and Loss Priority	75
Configuring Color-Blind Egress Policers for Medium-Low PLP	76
Configuring Two-Color and Three-Color Policers to Control Traffic Rates	77
Configuring Two-Color Policers	77
Configuring Three-Color Policers	78
Specifying Policers in a Firewall Filter Configuration	78
Applying a Firewall Filter That Includes a Policer	79
Verifying That Two-Color Policers Are Operational	79
Verifying That Three-Color Policers Are Operational	80
Troubleshooting Policer Configuration	80
Incomplete Count of Packet Drops	80
Counter Reset When Editing Filter	81
Invalid Statistics for Policer	81
Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured	81
Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured	82
Policers Can Limit Egress Filters	83

Part 3

Chapter 3

Port Security

Configuring Port Security	87
Overview of Access Port Protection	87
Mitigation of Ethernet Switching Table Overflow Attacks	88
Mitigation of Rogue DHCP Server Attacks	88
Protection Against ARP Spoofing Attacks	89
Protection Against DHCP Snooping Database Alteration Attacks	89

Protection Against DHCP Starvation Attacks	89
Understanding Port Security Features to Protect the Access Ports on Your Device	
Against the Loss of Information and Productivity	90
Understanding DHCP Snooping for Monitoring DHCP Messages Received from	
Untrusted Devices	92
DHCP Snooping Basics	92
DHCP Snooping Process	93
DHCPv6 Snooping	94
Rapid Commit for DHCPv6	95
DHCP Server Access	95
Switching Device, DHCP Clients, and DHCP Server Are All on the Same	
VLAN	95
Switching Device Acts as DHCP Server	96
Switching Device Acts as Relay Agent	97
Static IP Address Additions to the DHCP Snooping Database	98
Snooping DHCP Packets That Have Invalid IP Addresses	98
Prioritizing Snooped Packets	99
Verifying That DHCP Snooping Is Working Correctly	99
Understanding MAC Limiting and MAC Move Limiting for Port Security	100
MAC Limiting	100
MAC Move Limiting	101
Actions for MAC Limiting	101
MAC Addresses That Exceed the MAC Limit or MAC Move Limit	102
Verifying That MAC Limiting Is Working Correctly	102
Verifying That MAC Limiting for Dynamic MAC Addresses Is Working	
Correctly	103
Verifying That Allowed MAC Addresses Are Working Correctly	103
Verifying That Interfaces Are Shut Down	104
Customizing the Ethernet Switching Table Display to View Information for	
a Specific Interface	104
Verifying That MAC Move Limiting Is Working Correctly	105
Verifying That the Port Error Disable Setting Is Working Correctly	106
Understanding Dynamic ARP Inspection for Protecting Switching Devices Against	
ARP Spoofing	107
Address Resolution Protocol	107
ARP Spoofing	107
Dynamic ARP Inspection	108
Prioritizing Inspected Packets	109
Verifying That DAI Is Working Correctly	109
Understanding Trusted and Untrusted Ports	110
Understanding Trusted DHCP Servers for Port Security	110
Verifying That a Trusted DHCP Server Is Working Correctly	111
Understanding DHCP Option 82 for Port Security	112
DHCP Option 82 Processing	112
Suboption Components of Option 82	113
Configurations That Support Option 82	113
Understanding Static ARP Entries	114
Monitoring Port Security	115

Part 4	Device Security	
Chapter 4	Configuring Device Security	119
	Understanding Storm Control	119
	Example: Configuring Storm Control to Prevent Network Outages	120
	Verifying That the Port Error Disable Setting Is Working Correctly	123
	Understanding Unicast RPF	124
	Unicast RPF for Switches Overview	124
	Unicast RPF Implementation	125
	Unicast RPF Packet Filtering	125
	Bootstrap Protocol (BOOTP) and DHCP Requests	125
	Default Route Handling	125
	When to Enable Unicast RPF	125
	When Not to Enable Unicast RPF	126
	Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches	127
	Configuring Unicast RPF (CLI Procedure)	128
	Verifying Unicast RPF Status	129
	Disabling Unicast RPF (CLI Procedure)	132
	Understanding Unknown Unicast Forwarding	133
	Configuring Unknown Unicast Forwarding (CLI Procedure)	133
	Configuring Unknown Unicast Forwarding on EX4300 Switches	134
	Configuring Unknown Unicast Forwarding on EX9200 Switches	134
Part 5	Media Access Control Security (MACsec)	
Chapter 5	Configuring Media Access Control Security (MACsec)	139
	Understanding Media Access Control Security (MACsec)	139
	How MACsec Works	140
	Understanding Connectivity Associations and Secure Channels	140
	Understanding MACsec Security Modes	141
	Understanding Static Connectivity Association Key Security Mode (Recommended Security Mode for Switch-to-Switch Links)	141
	Understanding Dynamic Secure Association Key Security Mode (Switch-to-Host Links)	142
	Understanding Static Secure Association Key Security Mode (Supported for Switch-to-Switch Links)	142
	Understanding the Requirements to Enable MACsec on a Switch-to-Host Link	143
	MACsec Hardware and Software Support Summary	143
	Understanding MACsec Hardware Requirements for EX Series and QFX Series Switches	144
	Understanding MACsec Software Requirements for EX Series and QFX Series Switches	145
	Understanding the MACsec Feature License Requirement	146
	MACsec Limitations	146
	Configuring Media Access Control Security (MACsec)	147
	Acquiring and Downloading the Junos OS Software	147
	Acquiring and Downloading the MACsec Feature License	149

	Configuring the PIC Mode of the MACsec-capable Interfaces (EX4200 switches only)	149
	Configuring MACsec Using Static Connectivity Association Key Security Mode (Recommended for Enabling MACsec on Switch-to-Switch Links)	151
	Configuring MACsec on the Switch Using Dynamic Secure Association Key Security Mode to Secure a Switch-to-Host Link	155
	Configuring MACsec Using Static Secure Association Key Security Mode to Secure a Switch-to-Switch Link	159
Part 6	Configuration Statements and Operational Commands	
Chapter 6	Configuration Statements for Firewall Filters	167
	family	168
	filter	169
	filter (Layer 2 and Layer 3 Interfaces)	170
	filter (VLANs)	171
	firewall	172
	from	173
	interface-specific	174
	term	175
	then (Filters)	176
Chapter 7	Configuration Statements for Policers	177
	action	178
	bandwidth-limit	178
	burst-size-limit	179
	color-aware	180
	color-blind	181
	committed-burst-size	182
	committed-information-rate	183
	excess-burst-size	184
	filter-specific	185
	firewall	186
	if-exceeding	187
	loss-priority high then discard (Three-Color Policer)	188
	peak-burst-size	189
	peak-information-rate	190
	policer	191
	single-rate	192
	then (Policers)	193
	three-color-policer	194
	two-rate	195
Chapter 8	Configuration Statements for Port Security	197
	circuit-id	198
	dhcp-snooping-file	199
	fc-map	200
	fcoe-trusted	202
	mac-move-limit	203

	no-allowed-mac-log	204
	no-gratuitous-arp-request	205
	persistent-learning	205
	port-error-disable	206
	vendor-id	208
	write-interval	209
Chapter 9	Configuration Statements for Port Security (ELS CLI Only)	211
	accept-source-mac	212
	arp-inspection	214
	dhcp-security	216
	dhcp-service	219
	group (DHCP Security)	220
	interface (DHCP Security)	221
	interface-mac-limit	222
	no-dhcp-snooping	224
	no-option82	225
	option-82	226
	overrides (DHCP Security)	227
	recovery-timeout	228
	static-ip	230
	switch-options	231
	trusted	232
	untrusted	232
Chapter 10	Configuration Statements for Device Security	233
	action-shutdown	234
	interface (Unknown Unicast Forwarding)	235
	no-broadcast	236
	no-multicast	237
	no-unknown-unicast	238
	rpf-check	239
	unknown-unicast-forwarding	240
Chapter 11	Configuration Statements for Device Security (ELS CLI Only)	241
	bandwidth-level	242
	bandwidth-percentage	243
	no-registered-multicast	244
	no-unregistered-multicast	245
	storm-control	246
	storm-control-profiles	247
Chapter 12	Configuration Statements for Media Access Control Security (MACsec)	249
	cak	250
	ckn	251
	connectivity-association	252
	connectivity-association (MACsec Interfaces)	253
	direction	254
	encryption (MACsec)	255
	exclude-protocol	256

	id	257
	include-sci	258
	interfaces (MACsec)	259
	key (MACsec)	260
	key-server-priority (MACsec)	261
	mac-address (MACsec)	262
	macsec	263
	mka	264
	must-secure	265
	no-encryption (MACsec)	266
	offset	267
	port-id	268
	pre-shared-key	269
	replay-protect	270
	replay-window-size	271
	secure-channel	272
	security-association	273
	security-mode	274
	transmit-interval (MACsec)	275
Chapter 13	Operational Commands for Firewall Filters	277
	clear firewall	278
	show firewall	279
	show firewall policer	283
	show interfaces filters	285
Chapter 14	Operational Commands for Media Access Control Security (MACsec) ..	287
	clear security mka statistics	288
	show security macsec connections	289
	show security macsec statistics	291
	show security mka sessions	295
	show security mka statistics	297
Chapter 15	Operational Commands for Port Security	299
	clear arp inspection statistics	300
	clear dhcp snooping binding	301
	clear ethernet-switching port-error	302
	show arp inspection statistics	303
	show dhcp snooping binding	304

List of Figures

Part 1	Firewall Filters	
Chapter 1	Configuring Firewall Filters	3
	Figure 1: Evaluation of Terms Within a Firewall Filter	7
	Figure 2: Application of Firewall Filters to Control Packet Flow	9
Part 2	Policers	
Chapter 2	Configuring Policers	59
	Figure 3: Flow of Tricolor Marking Policer Operation	60
Part 3	Port Security	
Chapter 3	Configuring Port Security	87
	Figure 4: DHCP Server Connected Directly to a Switching Device	96
	Figure 5: DHCP Server Connected Directly to Switching Device 2, with Switching Device 2 Connected to Switching Device 1 Through a Trusted Trunk Port . . .	96
	Figure 6: Switching Device Is the DHCP Server	97
	Figure 7: Switching Device Acting as Relay Agent Through Router to DHCP Server	98
	Figure 8: Switch Relays DHCP Requests to Server	114
Part 4	Device Security	
Chapter 4	Configuring Device Security	119
	Figure 9: Symmetrically Routed Interfaces	126
	Figure 10: Asymmetrically Routed Interfaces	127

List of Tables

	About the Documentation	xv
	Table 1: Notice Icons	xvii
	Table 2: Text and Syntax Conventions	xvii
Part 1	Firewall Filters	
Chapter 1	Configuring Firewall Filters	3
	Table 3: Supported Firewall Filter Numbers for Specific Switches	6
	Table 4: Actions for Firewall Filters	12
	Table 5: Supported Match Conditions for Firewall Filters	13
	Table 6: Actions for Firewall Filters	24
	Table 7: Action Modifiers for Firewall Filters	25
	Table 8: Supported Firewall Filter Numbers	30
Part 2	Policers	
Chapter 2	Configuring Policers	59
	Table 9: Policer Actions	61
	Table 10: Color-Blind Mode TCM Color-to-PLP Mapping	66
	Table 11: Color-Aware Mode Single-Rate PLP Mapping	66
	Table 12: Color-Blind Mode TCM Color-to-PLP Mapping	68
	Table 13: Color-Aware Mode Two-Rate PLP Mapping	68
	Table 14: Servers Connected to Switch	73
	Table 15: Unicast Forwarding Classes	75
Part 3	Port Security	
Chapter 3	Configuring Port Security	87
	Table 16: DHCPv6 Messages and Equivalent DHCPv4 Messages	94
Part 5	Media Access Control Security (MACsec)	
Chapter 5	Configuring Media Access Control Security (MACsec)	139
	Table 17: MACsec Hardware and Software Support Summary for EX Series and QFX Series Switches	144
Part 6	Configuration Statements and Operational Commands	
Chapter 13	Operational Commands for Firewall Filters	277
	Table 18: show firewall Output Fields	279
	Table 19: show firewall policer Output Fields	283
	Table 20: show interfaces filters Output Fields	285

Chapter 14	Operational Commands for Media Access Control Security (MACsec) . .	287
	Table 21: show security macsec connections Output Fields	289
	Table 22: show security macsec statistics Output Fields	291
	Table 23: show security mka sessions Output Fields	295
	Table 24: show security mka statistics Output Fields	297
Chapter 15	Operational Commands for Port Security	299
	Table 25: show arp inspection statistics Output Fields	303
	Table 26: show dhcp snooping binding Output Fields	304

About the Documentation

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

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>.

Supported Platforms

For the features described in this document, the following platforms are supported:

- EX Series

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 [CLI Explorer](#).

Documentation Conventions

Table 1 on page xvii 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 xvii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure

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

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none">Introduces or emphasizes important new terms.Identifies guide names.Identifies RFC and Internet draft titles.	<ul style="list-style-type: none">A policy <i>term</i> is a named structure that defines match conditions and actions.<i>Junos OS CLI User Guide</i>RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <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.	broadcast multicast (<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.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
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.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.

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

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. 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. 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 8](#)
- [Understanding Firewall Filter Match Conditions on page 9](#)
- [Firewall Filter Match Conditions and Actions on page 13](#)
- [Understanding How a Firewall Filter Tests a Protocol on page 28](#)
- [Understanding Firewall Filter Planning on page 28](#)
- [Planning the Number of Firewall Filters to Create on page 30](#)
- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 34](#)
- [Configuring Firewall Filters on page 35](#)
- [Verifying That Firewall Filters Are Operational on page 38](#)
- [Applying Firewall Filters to Interfaces on page 39](#)
- [Understanding Filter-Based Forwarding on page 40](#)
- [Applying Firewall Filters to Interfaces on page 40](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on page 41](#)
- [Configuring MPLS Firewall Filters and Policers on page 44](#)
- [Configuring a Firewall Filter to De-Encapsulate GRE Traffic on a QFX5100 or OCX Switch on page 47](#)
- [Monitoring Firewall Filter Traffic on page 49](#)
- [Troubleshooting Firewall Filter Configuration on page 50](#)

Overview of Firewall Filters

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, QFX5110, QFX5200, 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 6](#).

Table 3: Supported Firewall Filter Numbers for Specific Switches

Filter Type	QFX3500, QFX3600	QFX5100, EX4600	QFX5110	QFX5200
Ingress	768	1536	provide number	768
Egress	1024	1024	provide number	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 30](#) for detailed information about this topic.

Related Documentation

- [Understanding Firewall Filter Planning on page 28](#)
- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 34](#)
- [Understanding How Firewall Filters Are Evaluated on page 6](#)
- [Understanding Firewall Filter Match Conditions on page 9](#)
- [Overview of Policers on page 59](#)
- [Configuring Firewall Filters on page 35](#)

Understanding How Firewall Filters Are Evaluated

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 59](#)
- [Configuring Firewall Filters on page 35](#)

Understanding How Firewall Filters Control Packet Flows

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 9](#) 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 34](#)
- [Understanding How Firewall Filters Are Evaluated on page 6](#)
- [Configuring Firewall Filters on page 35](#)

Understanding Firewall Filter Match Conditions

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 10](#)
- [Interface Filter Match Conditions on page 10](#)
- [IP Address Filter Match Conditions on page 11](#)
- [MAC Address Filter Match Conditions on page 11](#)
- [Bit-Field Filter Match Conditions on page 12](#)

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 28](#)
- [Firewall Filter Match Conditions and Actions on page 13](#)
- [Configuring Firewall Filters on page 35](#)

Firewall Filter Match Conditions and Actions

Each term in a firewall filter consists of *match conditions* and an *action*. Match conditions are the fields and values that a packet must contain to be considered a match. You can define single or multiple match conditions in *match statements*. You can also include no match statement, in which case the term matches all packets.

When a packet matches a filter, a switch takes the action specified in the term. In addition, you can specify action modifiers to count, mirror, rate-limit, and classify packets. If no match conditions are specified for the term, the switch accepts the packet by default.



NOTE: This topic does not apply to QFX10000 switches. For information about match conditions and actions on those switches, see *Firewall Filter Match Conditions and Actions for QFX10000 Switches*.

This topic describes the various match conditions, actions, and action modifiers that you can define in a firewall filter.

- [Table 5 on page 13](#) describes the match conditions you can specify when configuring a firewall filter. Some of the numeric range and bit-field match conditions allow you to specify a text synonym. To see a list of all the synonyms for a match condition, type ? at the appropriate place in a statement.
- [Table 6 on page 24](#) shows the actions that you can specify in a term.
- [Table 7 on page 25](#) shows the action modifiers you can use to count, mirror, rate-limit, and classify packets.



NOTE: On switches that do not support Layer 2 features (such as the OCX1100), you can use only those match conditions that are valid for IPv4 and IPv6 interfaces.

Table 5: Supported Match Conditions for Firewall Filters

Match Condition	Description	Direction and Interface
arp-type	ARP request packet or ARP reply packet.	Egress and ingress ports.
destination-address ip-address	IP destination address field, which is the address of the final destination node.	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces. Egress IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.
destination-mac-address mac-address	Destination media access control (MAC) address of the packet.	Ingress ports, VLANs and IPv4 (inet) interfaces. Egress ports and VLANs.

Table 5: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
destination-port value	<p>TCP or UDP destination port field. Typically, you specify this match in conjunction with the protocol match statement. For the following well-known ports you can specify text synonyms (the port numbers are also listed):</p> <p>afs (1483), bgp (179), biff (512), bootpc (68), bootps (67),</p> <p>cmd (514), cvspserver (2401),</p> <p>dhcp (67), domain (53),</p> <p>eklogin (2105), ekshell (2106), exec (512),</p> <p>finger (79), ftp (21), ftp-data (20),</p> <p>http (80), https (443),</p> <p>ident (113), imap (143),</p> <p>kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544),</p> <p>ldap (389), login (513),</p> <p>mobileip-agent (434), mobileip-mn (435), msdp (639),</p> <p>netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123),</p> <p>pop3 (110), pptp (1723), printer (515),</p> <p>radacct (1813), radius (1812), rip (520), rkinit (2108),</p> <p>smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514),</p> <p>tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525),</p> <p>who (513),</p> <p>xdmcp (177),</p> <p>zephyr-clt (2103), zephyr-hm (2104)</p>	<p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>

Table 5: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
destination-port range-optimize <i>range</i>	Match a range of TCP or UDP port ranges while using the available memory more efficiently. Using this condition allows you to configure more firewall filters than if you configure individual destination ports. (Not supported with filter-based forwarding.)	Egress and ingress IPv4 (inet) interfaces.
destination-prefix-list <i>prefix-list</i>	IP destination prefix list field. You can define a list of IP address prefixes under a prefix-list alias for frequent use. Define this list at the [edit policy-options] hierarchy level.	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces. Egress IPv4 (inet) interfaces and IPv6 (inet6) interfaces.
dot1q-tag <i>number</i>	802.1Q VLAN ID field in the Ethernet frame. The tag values can be 1–4094.	Ingress ports and VLANs. Egress ports and VLANs (<i>Number</i> must be the VLAN ID of the VLAN you want to match).
dot1q-user-priority <i>number</i>	<p>802.1Q priority field in the Ethernet frame (used for class-of-service priorities). Values can be 0–7.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> • best-effort (0)—Best effort • background (1)—Background • standard (2)—Standard or spare • excellent-load (3)—Excellent load • controlled-load (4)—Controlled load • video (5)—Video • voice (6)—Voice • network-control (7)—Network control reserved traffic 	Ingress ports and VLANs. Egress ports and VLANs.

Table 5: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
dscp value	<p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most-significant 6 bits of this byte form the DSCP.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> be—best effort (default) ef (46)—as defined in RFC 3246, <i>Assured Forwarding PHB</i>. af11 (10), af12 (12), af13 (14); af21 (18), af22 (20), af23 (22); af31 (26), af32 (28), af33 (30); af41 (34), af42 (36), af43 (38) These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in RFC 2597, <i>Assured Forwarding PHB</i>. cs0, cs1, cs2, cs3, cs4, cs5, cs6, cs7, cs5 	<p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>

Table 5: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
ether-type value	<p>Ethernet type field of a packet. The EtherType value specifies what protocol is being transported in the Ethernet frame. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> • aarp (0x80F3)—EtherType value AARP • appletalk (0x809B)—EtherType value AppleTalk • arp (0x0806)—EtherType value ARP • fcoe (0x8906)—EtherType value FCoE • fip (0x8914)—EtherType value FIP • ipv4 (0x0800)—EtherType value IPv4 • ipv6 (0x08DD)—EtherType value IPv6 • mpls-multicast (0x8848)—EtherType value MPLS multicast • mpls-unicast (0x8847)—EtherType value MPLS unicast • oam (0x88A8)—EtherType value OAM • ppp (0x880B)—EtherType value PPP • pppoe-discovery (0x8863)—EtherType value PPPoE Discovery Stage • pppoe-session (0x8864)—EtherType value PPPoE Session Stage • sna (0x80D5)—EtherType value SNA 	<p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>
exp	Match on MPLS EXP bits.	<p>Ingress MPLS interfaces.</p> <p>Egress MPLS interfaces.</p>
fragment-flags value	<p>IP fragmentation flags. In place of the numeric value, you can specify one of the following text synonyms (the hexadecimal values are also listed):</p> <ul style="list-style-type: none"> • is-fragment • dont-fragment (0x4000) • more-fragments (0x2000) • reserved (0x8000) 	Ingress ports and VLANs.

Table 5: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
icmp-code value	<p>ICMP code field. Because the meaning of the value depends upon the associated icmp-type, you must specify a value for icmp-type along with a value for icmp-code. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> <i>IPv4</i>: parameter-problem—ip-header-bad (0), required-option-missing (1) <i>IPv6</i>: parameter-problem—ip6-header-bad (0), unrecognized-next-header (1), unrecognized-option (2) redirect—redirect-for-network (0), redirect-for-host (1), redirect-for-tos-and-net (2), redirect-for-tos-and-host (3) time-exceeded—ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) <i>IPv4</i>: unreachable—network-unreachable (0), host-unreachable (1), protocol-unreachable (2), port-unreachable (3), fragmentation-needed (4), source-route-failed (5), destination-network-unknown (6), destination-host-unknown (7), source-host-isolated (8), destination-network-prohibited (9), destination-host-prohibited (10), network-unreachable-for-TOS (11), host-unreachable-for-TOS (12), communication-prohibited-by-filtering (13), host-precedence-violation (14), precedence-cutoff-in-effect (15) <i>IPv6</i>: unreachable—address-unreachable (3), administratively-prohibited (1), no-route-to-destination (0), port-unreachable (4) 	<p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
hop-limitvalue	Match the the specified hop limit or set of hop limits. Specify a single value or a range of values from 0 through 255.	Ingress and egress IPv6 (inet6) interfaces.

Table 5: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
icmp-type <i>value</i>	<p>ICMP message type field. Typically, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <p><i>IPv4:</i> echo-reply (0), destination unreachable (3), source-quench (4), redirect (5), echo-request (8), IPv4 (inet)-advertisement (9), IPv4 (inet)-solicit (10), time-exceeded (11), parameter-problem (12), timestamp (13), timestamp-reply (14), info-request (15), info-reply (16), mask-request (17), mask-reply (18)</p> <p><i>IPv6:</i> destination-unreachable (1), packet-too-big (2), time-exceeded (3), parameter-problem (4), echo-request (128), echo-reply (129), membership-query (130), membership-report (131), membership-termination (132), router-solicit (133), router-advertisement (134), neighbor-solicit (135), neighbor-advertisement (136), redirect (137), router-renumbering (138), node-information-request (139), node-information-reply (140)</p> <p>See also icmp-code <i>variable</i>.</p>	<p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
interface <i>interface-name</i>	<p>Interface on which the packet is received, including the logical unit. You can include the wildcard character (*) as part of an interface name or logical unit.</p> <p>NOTE: An interface from which a packet is sent cannot be used as a match condition.</p>	<p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces and IPv6 (inet6) interfaces.</p>
ip-destination-address <i>address</i>	IPv4 address that is the final destination node address for the packet.	Ingress ports and VLANs.
ip6-destination-address <i>address</i>	IPv6 address that is the final destination node address for the packet.	Ingress ports and VLANs. (You cannot simultaneously apply a filter with this match criterion to a Layer 2 port and VLAN that includes that port.)
ip-options	Specify any to create a match if anything is specified in the options field in the IP header.	<p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>

Table 5: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
ip-precedence <i>ip-precedence-field</i>	IP precedence field. In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): critical-ecp (0xa0), flash (0x60), flash-override (0x80), immediate (0x40), internet-control (0xc0), net-control (0xe0), priority (0x20), or routine (0x00).	Ingress ports, VLANs, and IPv4 (inet) interfaces. Egress IPv4 (inet) interfaces.
ip-protocol <i>number</i>	IP protocol field.	Ingress ports, VLANs, and IPv4 (inet) interfaces. Egress IPv4 (inet) interfaces.
ip-source-address <i>address</i>	IPv4 address of the source node sending the packet.	Ingress ports and VLANs.
ip6-source-address <i>address</i>	IPv6 address of the source node sending the packet.	Ingress ports and VLANs. (You cannot simultaneously apply a filter with this match criterion to a Layer 2 port and VLAN that includes that port.)
ip-version <i>address</i>	IP version of the packet. Use this condition to match IPv4 or IPv6 header fields in traffic that arrives on a Layer 2 port or VLAN interface.	Ingress ports and VLANs.
is-fragment	Using this condition causes a match if the More Fragments flag is enabled in the IP header or if the fragment offset is not zero.	Ingress ports, VLANs, and IPv4 (inet) interfaces. Egress IPv4 (inet) interfaces.
l2-encap-type <i>llc-non-snap</i>	Match on logical link control (LLC) layer packets for non-Subnet Access Protocol (SNAP) Ethernet Encapsulation type.	Ingress ports and VLANs. Egress ports and VLANs.
label	Match on MPLS label bits.	Ingress MPLS interfaces. Egress MPLS interfaces.
learn-vlan-id <i>number</i>	Matches the ID of a normal VLAN or the ID of the outer (service) VLAN (for Q-in-Q VLANs). The acceptable values are 1-4095.	Ingress ports and VLANs. Egress ports and VLANs.

Table 5: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
next-header	<p>IPv4 or IPv6 protocol value. In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <p>hop-by-hop (0), icmp (1), icmp6 (58), igmp (2), ipip (4), tcp (6), egp (8), udp (17), ipv6 (41), routing (43), fragment (44), rsvp (46), gre (47), esp (50), ah (51), icmp6 (58), no-next-header (59), dstopts (60), ospf (89), pim (103), vrrp (112), sctp (132)</p>	<p>Ingress ports, VLANs, and IPv6 (inet6) interfaces.</p> <p>Egress IPv6 (inet6) interfaces.</p>
packet-length	<p>Packet length in bytes. You must enter a value between 0 and 65535.</p>	<p>Ingress ports, VLANs, IPv4 (inet), and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
payload-protocol	<p>IPv4 or IPv6 protocol value. In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <p>hop-by-hop (0), icmp (1), icmp6 (58), igmp (2), ipip (4), tcp (6), egp (8), udp (17), ipv6 (41), routing (43), fragment (44), rsvp (46), gre (47), esp (50), ah (51), icmp6 (58), no-next-header (59), dstopts (60), ospf (89), pim (103), vrrp (112), sctp (132)</p>	<p>Ingress ports, VLANs, and IPv6 (inet6) interfaces.</p> <p>Egress IPv6 (inet6) interfaces.</p>
precedence value	<p>IP precedence bits in the type-of-service (ToS) byte in the IP header. (This byte can also be used for the DiffServ DSCP.) In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <ul style="list-style-type: none"> • routine (0) • priority (1) • immediate (2) • flash (3) • flash-override (4) • critical-ecp (5) • internet-control (6) • net-control (7) 	<p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>

Table 5: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
protocol type	<p>IPv4 or IPv6 protocol value. In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <p>hop-by-hop (0), icmp (1), icmp6, igmp (2), ipip (4), tcp (6), egp (8), udp (17), ipv6 (41), routing (43), fragment (44), rsvp (46), gre (47), esp (50), ah (51), icmp6 (58), no-next-header (59), dstopts (60), ospf (89), pim (103), vrrp (112), sctp (132)</p>	<p>Ingress ports, VLANs and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
rat-type tech-type-value	<p>Match the radio-access technology (RAT) type specified in the 8-bit Tech-Type field of Proxy Mobile IPv4 (PMIPv4) access technology type extension. The technology type specifies the access technology through which the mobile device is connected to the access network. Specify a single value, a range of values, or a set of values. You can specify a technology type as a numeric value from 0 through 255 or as a system keyword.</p> <ul style="list-style-type: none"> Numeric value 1 matches IEEE 802.3. Numeric value 2 matches IEEE 802.11a/b/g. Numeric value 3 matches IEEE 802.16e Numeric value 4 matches IEEE 802.16m. Text string eutran matches 4G. Text string geran matches 2G. Text string utran matches 3G. 	Egress and ingress IPv4 (inet) interfaces.
sample	Sample the packet traffic. Apply this option only if you have enabled traffic sampling.	Egress and ingress IPv4 (inet) interfaces.
source-address ip-address	IP source address field, which is the address of the node that sent the packet.	<p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
source-mac-address <i>mac-address</i>	Source media access control (MAC) address of the packet.	<p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>

Table 5: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
source-port <i>value</i>	TCP or UDP source port. Typically, you specify this match in conjunction with the protocol match statement. In place of the numeric field, you can specify one of the text synonyms listed under destination-port .	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces. Egress IPv4 (inet) interfaces.
source-port range-optimize <i>range</i>	Match a range of TCP or UDP port ranges while using the available memory more efficiently. Using this condition allows you to configure more firewall filters than if you configure individual source ports. (Not supported with filter-based forwarding.)	Egress and ingress IPv4 (inet) interfaces.
source-prefix-list <i>prefix-list</i>	IP source prefix list. You can define a list of IP address prefixes under a prefix-list alias for frequent use. Define this list at the [edit policy-options] hierarchy level.	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces. Egress IPv4 (inet) interfaces.
tcp-established	Match packets of an established TCP connection. This condition matches packets other than those used to set up a TCP connection—that is, three-way handshake packets are not matched. When you specify tcp-established , a switch does not implicitly verify that the protocol is TCP. You must also specify the protocol tcp match condition.	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces. Egress IPv4 (inet) interfaces.
tcp-flags <i>value</i>	One or more TCP flags: <ul style="list-style-type: none"> • ack (0x10) • fin (0x01) • push (0x08) • rst (0x04) • syn (0x02) • urgent (0x20) 	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces. Egress IPv4 (inet) interfaces.
tcp-initial	Match the first TCP packet of a connection. A match occurs when the TCP flag SYN is set and the TCP flag ACK is not set. When you specify tcp-initial , a switch does not implicitly verify that the protocol is TCP. You must also specify the protocol tcp match condition.	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces. Egress IPv4 (inet) interfaces.

Table 5: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
traffic-class	<p>8-bit field that specifies the class-of-service (CoS) priority of the packet. The traffic-class field is used to specify a DiffServ code point (DSCP) value. This field was previously used as the type-of-service (ToS) field in IPv4, and, the semantics of this field (for example, DSCP) are identical to those of IPv4.</p> <p>You can specify one of the following text synonyms (the field values are also listed):</p> <p>af11 (10), af12 (12), af13 (14), af21 (18), af22 (20), af23 (22), af31 (26), af32 (28), af33 (30), af41 (34), af42 (36), af43 (38), cs0 (0), cs1 (8), cs2 (16), cs3 (24), cs4 (32), cs5 (40), cs6 (48), cs7 (56), ef (46)</p>	<p>Ingress ports, VLANs, and IPv6 (inet6) interfaces.</p> <p>Egress IPv6 (inet6) interfaces.</p>
ttl value	IP Time-to-live (TTL) field in decimal. The value can be 1-255.	<p>Ingress IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
user-vlan-1p-priority value	Matches the specified 802.1p VLAN priority in the range 0-7.	<p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
user-vlan-id number	Matches the ID of the inner (customer) VLAN for a Q-in-Q VLAN. The acceptable values are 1-4095.	<p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>

Use **then** statements to define actions that should occur if a packet matches all conditions in a **from** statement. [Table 6 on page 24](#) shows the actions that you can specify in a term. (If you do not include a **then** statement, the system accepts packets that match the filter.)

Table 6: Actions for Firewall Filters

Action	Description
accept	Accept a packet. This is the default action for packets that match a term.
discard	Discard a packet silently without sending an Internet Control Message Protocol (ICMP) message.

Table 6: Actions for Firewall Filters (*continued*)

Action	Description
reject <i>message-type</i>	<p>Discard a packet and send a “destination unreachable” ICMPv4 message (type 3). To log rejected packets, configure the syslog action modifier.</p> <p>You can specify one of the following message types: administratively-prohibited (default), bad-host-tos, bad-network-tos, host-prohibited, host-unknown, host-unreachable, network-prohibited, network-unknown, network-unreachable, port-unreachable, precedence-cutoff, precedence-violation, protocol-unreachable, source-host-isolated, source-route-failed, or tcp-reset.</p> <p>If you specify tcp-reset, the system sends a TCP reset if the packet is a TCP packet; otherwise nothing is sent.</p> <p>If you do not specify a message type, the ICMP notification “destination unreachable” is sent with the default message “communication administratively filtered.”</p> <p>NOTE: The reject action is supported on ingress interfaces only.</p>
routing-instance <i>instance-name</i>	Forward matched packets to a virtual routing instance.
vlan <i>VLAN-name</i>	<p>Forward matched packets to a specific VLAN.</p> <p>NOTE: The vlan action is supported on ingress interfaces only.</p> <p>NOTE: This action is not supported on OCX series switches.</p>

You can also specify the action modifiers listed in [Table 7 on page 25](#) to count, mirror, rate-limit, and classify packets.

Table 7: Action Modifiers for Firewall Filters

Action Modifier	Description
analyzer <i>analyzer-name</i>	<p>(Non-ELS platforms) Mirror traffic (copy packets) to an analyzer configured at the [edit ethernet-switching-options analyzer] hierarchy level.</p> <p>You can specify port mirroring for ingress port, VLAN, and IPv4 (inet) firewall filters only.</p>
count <i>counter-name</i>	Count the number of packets that match the term.
decapsulate [<i>gre</i> <i>routing-instance</i>]	De-encapsulate GRE packets or forward de-encapsulated GRE packets to the specified routing instance

Table 7: Action Modifiers for Firewall Filters (*continued*)

Action Modifier	Description
dscp value	<p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most-significant 6 bits of this byte form the DSCP.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> be—best effort (default) ef (46)—as defined in RFC 3246, <i>An Expedited Forwarding PHB</i>. af11 (10), af12 (12), af13 (14); af21 (18), af22 (20), af23 (22); af31 (26), af32 (28), af33 (30); af41 (34), af42 (36), af43 (38) <p>These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in RFC 2597, <i>Assured Forwarding PHB</i>.</p> <ul style="list-style-type: none"> cs0, cs1, cs2, cs3, cs4, cs5, cs6, cs7, cs5
forwarding-class class	<p>Classify the packet in one of the following default forwarding classes, or in a user-defined forwarding class:</p> <ul style="list-style-type: none"> best-effort fcoe mcast network-control no-loss <p>NOTE: To configure a forwarding class, you must also configure loss priority.</p>
interface	<p>Switch the traffic to the specified interface without performing a lookup on it. This action is valid only when the filter is applied on ingress.</p>
log	<p>Log the packet's header information in the Routing Engine. To view this information, enter the show firewall log operational mode command.</p> <p>NOTE: The log action modifier is supported on ingress interfaces only.</p>
loss-priority (low medium-low medium-high high)	<p>Set the packet loss priority (PLP).</p> <p>NOTE: The loss-priority action modifier is supported on ingress interfaces only.</p> <p>NOTE: The loss-priority action modifier is not supported in combination with the policer action.</p>

Table 7: Action Modifiers for Firewall Filters (*continued*)

Action Modifier	Description
policer <i>policer-name</i>	<p>Send packets to a policer (for the purpose of applying rate limiting).</p> <p>You can specify a policer for ingress port, VLAN, IPv4 (inet), IPv6 (inet6), and MPLS filters.</p> <p>NOTE: The policer action modifier is not supported in combination with the loss-priority action.</p>
port-mirror	<p>(ELS platforms) Mirror traffic (copy packets) to an output interface configured in a port-mirroring instance at the [edit forwarding-options port-mirroring] hierarchy level.</p> <p>You can specify port mirroring for ingress port, VLAN, and IPv4 (inet) firewall filters only.</p>
port-mirror-instance <i>port-mirror-instance-name</i>	<p>(ELS platforms) Mirror traffic to a port-mirroring instance configured at the [edit forwarding-options port-mirroring] hierarchy level.</p> <p>You can specify port mirroring for ingress port, VLAN, and IPv4 (inet) firewall filters only.</p> <p>NOTE: This action modifier is not supported on OCX series switches.</p>
syslog	<p>Log an alert for this packet.</p> <p>NOTE: The syslog action modifier is supported on ingress interfaces only.</p>
three-color-policer <i>three-color-policer-name</i>	<p>Send packets to a three-color policer (for the purpose of applying rate limiting).</p> <p>You can specify a three-color policer for ingress and egress port, VLAN, IPv4 (inet), IPv6 (inet6), and MPLS filters.</p> <p>NOTE: The policer action modifier is not supported in combination with the loss-priority action.</p>

**Related
Documentation**

- [Understanding How Firewall Filters Are Evaluated on page 6](#)
- [Understanding How a Firewall Filter Tests a Protocol on page 28](#)
- [Overview of Policers on page 59](#)
- [Understanding Port Mirroring](#)
- [Configuring Firewall Filters on page 35](#)

Understanding How a Firewall Filter Tests a Protocol

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 35](#)

Understanding Firewall Filter Planning

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 30](#) for information about how many firewall filters you can apply.

- Related Documentation**
- [Overview of Policers on page 59](#)
 - [Understanding How Firewall Filters Are Evaluated on page 6](#)
 - [Configuring Firewall Filters on page 35](#)

Planning the Number of Firewall Filters to Create

- [Understanding How Many Firewall Filters Are Supported on page 30](#)
- [Egress Filters on page 31](#)
- [Avoid Configuring too Many Filters on page 31](#)
- [Configuring TCAM Error Messages on page 32](#)
- [Policies can Limit Egress Filters on page 32](#)
- [Planning for Filter-Specific Policies on page 33](#)
- [Planning for Filter-Based Forwarding on page 33](#)

Understanding How Many Firewall Filters Are Supported

QFX3500, QFX3600, QFX5100, QFX5110, 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 8 on page 30](#).

Table 8: Supported Firewall Filter Numbers

Filter Type	QFX3500, QFX3600	QFX5100, EX4600	QFX5110	QFX5200	QFX10000
Ingress	768	1536	provide number	768	8192
Egress	1024	1024	provide number	1024	8192

On QFX5100, QFX5110, and QFX5200 switches, you can see how many filters have been programmed of each type by entering **show pfe filter hw summary**.

The 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 28](#)
- [Configuring Firewall Filters on page 35](#)
- [Understanding Filter-Based Forwarding on page 40](#)

Understanding Firewall Filter Processing Points for Bridged and Routed Packets

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 8](#)
- [Configuring Firewall Filters on page 35](#)

Configuring Firewall Filters

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 35](#)
- [Applying a Firewall Filter to a Port on page 37](#)
- [Applying a Firewall Filter to a VLAN on page 37](#)
- [Applying a Firewall Filter to a Layer 3 \(Routed\) Interface on page 38](#)

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 on page 13](#)
- [Verifying That Firewall Filters Are Operational on page 38](#)
- [Monitoring Firewall Filter Traffic on page 49](#)
- [Configuring Port Mirroring](#)

Verifying That Firewall Filters Are Operational

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 35](#)
 - [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77](#)
 - [Monitoring Firewall Filter Traffic on page 49](#)

Applying Firewall Filters to Interfaces

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 35](#)

Understanding 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.

You might want to use filter-based forwarding to route specific types of traffic through a firewall or other security device before the traffic continues on its path. You can also use filter-based forwarding to give certain types of traffic preferential treatment. For example, you might want to ensure that the highest-priority traffic is forwarded over a 40-Gigabit Ethernet link. You might also use filter-based forwarding to obtain more control over load balancing than dynamic routing protocols provide.



NOTE: You can create as many as 128 filters or terms that direct packets to a given virtual routing instance.

Related Documentation

- [Understanding Virtual Router Routing Instances](#)
- [Overview of Firewall Filters on page 3](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on page 41](#)

Applying Firewall Filters to Interfaces

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 35](#)

Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device

You can configure filter-based forwarding by using a firewall filter to forward matched traffic to a specific virtual routing instance.

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

- [Requirements on page 41](#)
- [Overview and Topology on page 41](#)
- [Configuration on page 41](#)
- [Verification on page 43](#)

Requirements

This example requires Junos OS Release 15.1X53-D10 or later on a QFX10000 switch..

Overview and Topology

In this example, traffic from one application server that is destined for a different application server is matched by a firewall filter based on the IP address of the source application server. Any matching packets are routed to a virtual routing instance that sends the traffic to a security device. In this case, the security device must be able to forward the traffic to the destination application server. For this example, assume that the address of the destination application server is 192.168.0.1.



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

Configuration

To configure filter-based forwarding:

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste them into the CLI at the **[edit]** hierarchy level.

```
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.0.1/24
set interfaces xe-0/0/3 unit 0 family inet address 10.1.3.1/24
set firewall family inet filter f1 term t1 from source-address 10.1.0.50/32
set firewall family inet filter f1 term t1 from protocol tcp
set interfaces xe-0/0/0 unit 0 family inet filter input f1
set routing-instances vrf01 instance-type virtual-router
set routing-instances vrf01 interface xe-0/0/3.0
set routing-instances vrf01 routing-options static route 192.168.0.1/24 next-hop 10.1.3.254
set firewall family inet filter f1 term t1 then routing-instance vrf01
```

**Step-by-Step
Procedure**

To configure filter-based forwarding:

1. Configure an interface to connect to the application server:

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family inet address 10.1.0.1/24
```
2. Configure an interface to connect to the security device:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 10.1.3.1/24
```
3. Create a firewall filter that matches packets based on the address of the application server that the traffic will be sent from. Also configure the filter so that it matches only TCP packets:

```
[edit firewall]
user@switch# set family inet filter f1 term t1 from source-address 10.1.0.50/32
user@switch# set firewall family inet filter f1 term t1 from protocol tcp
```
4. Apply the filter to the interface that connects to the source application server and configure it to match incoming packets:

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family inet filter input f1
```
5. Create a virtual router:

```
[edit]
user@switch# set routing-instances vrf01 instance-type virtual-router
```
6. Associate the virtual router with the interface that connects to the security device:

```
[edit routing-instances]
user@switch# set vrf01 interface xe-0/0/3.0
```
7. Configure the routing information for the virtual routing instance:

```
[edit routing-instances]
user@switch# set vrf01 routing-options static route 192.168.0.1/24 next-hop 10.1.3.254
```
8. Set the filter to forward packets to the virtual router:

```
[edit firewall]
user@switch# set family inet filter f1 term t1 then routing-instance vrf01
```

Results

Check the results of the configuration:

```
user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        filter {
          input f1;
        }
        address 10.1.0.1/24;
      }
    }
  }
  xe-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.3.1/24;
      }
    }
  }
}
```

```

    }
  }
}
firewall {
  family inet {
    filter f1 {
      term t1 {
        from {
          source-address {
            10.1.0.50/32;
          }
          protocol tcp;
        }
        then {
          routing-instance vrf01;
        }
      }
    }
  }
}
routing-instances {
  vrf01 {
    instance-type virtual-router;
    interface xe-0/0/1.0;
    routing-options {
      static {
        route 192.168.0.1/24 next-hop 10.1.3.254;
      }
    }
  }
}
}

```

Verification

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

- [Verifying That Filter-Based Forwarding Was Configured on page 43](#)

Verifying That Filter-Based Forwarding Was Configured

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

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

```

user@switch> show interfaces filters xe-0/0/0.0
Interface      Admin Link Proto Input Filter      Output Filter
xe-0/0/0.0     up    down inet f1

```

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

```

user@switch> show route forwarding-table

Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          user   1 0:12:f2:21:cf:0 ucst  331   4 me0.0

```

default	perm	0		rjct	36	3
0.0.0.0/32	perm	0		dscd	34	1
10.1.0.0/24	ifdn	0		rslv	613	1 xe-0/0/0.0
10.1.0.0/32	iddn	0	10.1.0.0	recv	611	1 xe-0/0/0.0
10.1.0.1/32	user	0		rjct	36	3
10.1.0.1/32	intf	0	10.1.0.1	loc1	612	2
10.1.0.1/32	iddn	0	10.1.0.1	loc1	612	2
10.1.0.255/32	iddn	0	10.1.0.255	bcst	610	1 xe-0/0/0.0
10.1.1.0/26	ifdn	0		rslv	583	1 vlan.0
10.1.1.0/32	iddn	0	10.1.1.0	recv	581	1 vlan.0
10.1.1.1/32	user	0		rjct	36	3
10.1.1.1/32	intf	0	10.1.1.1	loc1	582	2
10.1.1.1/32	iddn	0	10.1.1.1	loc1	582	2
10.1.1.63/32	iddn	0	10.1.1.63	bcst	580	1 vlan.0
255.255.255.255/32	perm	0		bcst	32	1

Routing table: vrf01.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	559	2	
0.0.0.0/32	perm	0		dscd	545	1	
10.1.3.0/24	ifdn	0		rslv	617	1 xe-0/0/3.0	
10.1.3.0/32	iddn	0	10.1.3.0	recv	615	1 xe-0/0/3.0	
10.1.3.1/32	user	0		rjct	559	2	
192.168.0.1/24	user	0	10.1.3.254	ucst	616	2 xe-0/0/3.0	
192.168.0.1/24	user	0	10.1.3.254	ucst	616	2 xe-0/0/3.0	
10.1.3.255/32	iddn	0	10.1.3.255	bcst	614	1 xe-0/0/3.0	
224.0.0.0/4	perm	0		mdsc	546	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	529	1	
255.255.255.255/32	perm	0		bcst	543	1	

Routing table: default.iso

ISO:

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

Routing table: vrf01.iso

ISO:

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

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

- Related Documentation**
- [Configuring Firewall Filters on page 35](#)
 - [Understanding Filter-Based Forwarding on page 40](#)
 - [Understanding Virtual Router Routing Instances](#)

Configuring MPLS Firewall Filters and Policers

You can configure an MPLS firewall filter to count packets based on the EXP bits for the top-level MPLS label in a packet. You can also configure policers for MPLS LSPs.

The following sections discuss MPLS firewall filters and policers:

- [Configuring MPLS Firewall Filters on page 45](#)
- [Examples: Configuring MPLS Firewall Filters on page 45](#)
- [Configuring Policers for LSPs on page 46](#)

Configuring MPLS Firewall Filters

You can configure an MPLS firewall filter to count packets based on the EXP bits for the top-level MPLS label in a packet. You can then apply this filter to a specific interface on input or output. You can also configure a policer for the MPLS filter to police (that is, rate-limit) the traffic on the interface to which the filter is attached. You cannot apply MPLS firewall filters to loopback interfaces.

You can configure the following match conditions for MPLS filters at the **[edit firewall family mpls filter *filter-name* term *term-name* from]** hierarchy level:

- **exp**
- **label**

These **exp** match condition can accept EXP bits in the range 0 through 7. You can configure the following choices:

- A single EXP bit—for example, **exp 3**;
- Several EXP bits—for example, **exp 0, 4**;
- A range of EXP bits—for example, **exp [0-5]**;

The **label** match condition can accept a range of values from 0 to 1048575.

If you do not specify a match criterion (that is, you do not configure the **from** statement and use only the **then** statement with the **count** action keyword), all the MPLS packets passing through the interface on which the filter is applied will be counted.

You also can configure any of the following action keywords at the **[edit firewall family mpls filter *filter-name* term *term-name* then]** hierarchy level:

- **accept**
- **count**
- **discard**
- **policer**
- **three-color-policer**

Examples: Configuring MPLS Firewall Filters

The following examples illustrate how you might configure an MPLS firewall filter and then apply the filter to an interface. This filter is configured to count MPLS packets with EXP bits set to either 0 or 4.

The following shows a configuration for an MPLS firewall filter:

```
[edit firewall]
family mpls {
  filter expf {
    term expt0 {
      from {
        exp 0,4;
      }
      then {
        count counter0;
        accept;
      }
    }
  }
}
```

Configuring Policers for LSPs

MPLS LSP policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. LSP policing is supported on regular LSPs, LSPs configured with DiffServ-aware traffic engineering, and multiclass LSPs. You can configure multiple policers for each multiclass LSP. For regular LSPs, each LSP policer is applied to all of the traffic traversing the LSP. The policer's bandwidth limitations become effective as soon as the total sum of traffic traversing the LSP exceeds the configured limit.

You configure the multiclass LSP and DiffServ-aware traffic engineering LSP policers in a filter. The filter can be configured to distinguish between the different class types and apply the relevant policer to each class type. The policers distinguish between class types based on the EXP bits.

You configure LSP policers under the **family any** filter. The **family any** filter is used because the policer is applied to traffic entering the LSP. This traffic might be from different families: IPv6, MPLS, and so on. You do not need to know what sort of traffic is entering the LSP, as long as the match conditions apply to all types of traffic.

LSP Policer Limitations

When configuring MPLS LSP policers, be aware of the following limitations:

- LSP policers are supported for packet LSPs only.
- LSP policers are supported for unicast next hops only. Multicast next hops are not supported.
- LSP policers are not supported on aggregated interfaces.
- The LSP policer runs before any output filters.
- Traffic sourced from the Routing Engine (for example, ping traffic) does not take the same forwarding path as transit traffic. This type of traffic cannot be policed.

Related Documentation • [Overview of Policers on page 59](#)

Configuring a Firewall Filter to De-Encapsulate GRE Traffic on a QFX5100 or OCX Switch

Generic routing encapsulation (GRE) provides a private, secure path for transporting packets through a network by encapsulating (or tunneling) the packets. GRE tunneling is performed by tunnel endpoints that encapsulate or de-encapsulate traffic.

You can use a firewall filter to de-encapsulate GRE traffic on a QFX5100 or OCX switch. This feature provides significant benefits in terms of scalability, performance, and flexibility because you don't need to create a tunnel interface to perform the de-encapsulation. For example, you can terminate many tunnels from multiple source IP addresses with one firewall term.



NOTE: QFX5100 and OCX switches support as many as 512 GRE tunnels, including tunnels created with a firewall filter. That is, you can create a total of 512 GRE tunnels, regardless of which method you use.

This topic describes:

1. [Configuring a Filter to De-Encapsulate GRE Traffic on page 47](#)
2. [Applying the Filter to an Interface on page 48](#)

Configuring a Filter to De-Encapsulate GRE Traffic

To configure a firewall filter to de-encapsulate GRE traffic:

1. Create an IPv4 firewall filter and (optionally) specify a source address for the tunnel:

[edit]

```
user@switch# set firewall family inet filter filter-name term term-name from
source-address address
```

You must create an IPv4 filter by using **family inet** because the outer header of a GRE packet must be IPv4. If you specify a source address, it should be an address on a device that will encapsulate traffic into GRE packets.



NOTE: To terminate many tunnels from multiple source IP addresses with one firewall term, do not configure a source address. In this case, the filter will de-encapsulate any GRE packets received by the interface that you apply the filter to.

2. Specify a destination address for the tunnel:

[edit]

```
user@switch# set firewall family inet filter filter-name term term-name from
destination-address address
```

This should be an address on an interface of the switch on which you want the tunnel or tunnels to terminate and the GRE packets to be de-encapsulated. You should also configure this address as a tunnel endpoint on all the tunnel source routers that you want to form tunnels with the switch.

3. Specify that the filter should match and accept GRE traffic:

```
[edit ]
user@switch# set firewall family inet filter filter-name term term-name from protocol
gre
```

4. Specify that the filter should de-encapsulate GRE traffic:

```
[edit ]
user@switch# set firewall family inet filter filter-name term term-name then decapsulate
gre
```

Based on the configuration you have performed so far, the switch forwards the de-encapsulated packets by comparing the inner header to the default routing table (**inet0**). If you want the switch to use a virtual routing instance to forward the de-encapsulated packets, perform the following steps:

5. Specify the name of the virtual routing instance:

```
[edit ]
user@switch# set firewall family inet filter filter-name term term-name then decapsulate
routing-instance instance-name
```

6. Specify that the virtual routing instance is a virtual router:

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

7. Specify the interfaces that belong to the virtual router:

```
[edit ]
user@switch# set routing-instances instance-name interface interface-name
```

Applying the Filter to an Interface

After you create the firewall filter, you must also apply it to an interface that will receive GRE traffic. Be sure to apply it in the input direction. For example, enter

```
[edit ]
user@switch# set interfaces interface-name unit logical-unit-number family inet filter
input filter-name
```

Because the outer header of a GRE packet must be IPv4, you must apply the filter to an IPv4 interface and specify **family inet**.

Related Documentation

- [Understanding Generic Routing Encapsulation](#)
- [Configuring Generic Routing Encapsulation Tunneling](#)
- [Configuring Firewall Filters on page 35](#)

Monitoring Firewall Filter Traffic

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

- [Monitoring Traffic for All Firewall Filters and Policers That Are Configured on page 49](#)
- [Monitoring Traffic for a Specific Firewall Filter on page 49](#)
- [Monitoring Traffic for a Specific Policer on page 50](#)

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 <i>policer-name</i> 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 <i>policer-name</i> command displays the number of packets that exceeded the rate limits for the specified policer.
Related Documentation	<ul style="list-style-type: none">• Configuring Firewall Filters on page 35• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77• Verifying That Firewall Filters Are Operational on page 38

Troubleshooting Firewall Filter Configuration

Use the following information to troubleshoot your firewall filter configuration.

- [Firewall Filter Configuration Returns a No Space Available in TCAM Message on page 50](#)
- [Filter Counts Previously Dropped Packet on page 52](#)
- [Matching Packets Not Counted on page 53](#)
- [Counter Reset When Editing Filter on page 53](#)
- [Cannot Include loss-priority and policer Actions in Same Term on page 53](#)
- [Cannot Egress Filter Certain Traffic Originating on QFX Switch on page 54](#)
- [Firewall Filter Match Condition Not Working with Q-in-Q Tunneling on page 54](#)
- [Egress Firewall Filters with Private VLANs on page 54](#)
- [Egress Filtering of L2PT Traffic Not Supported on page 55](#)
- [Cannot Drop BGP Packets in Certain Circumstances on page 55](#)
- [Invalid Statistics for Policer on page 55](#)
- [Policers can Limit Egress Filters on page 55](#)

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 **adminVLAN**, 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

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

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

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

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.

**Related
Documentation**

- [Understanding FIP Snooping, FBF, and MVR Filter Scalability](#)
- [Configuring Firewall Filters on page 35](#)
- [Verifying That Firewall Filters Are Operational on page 38](#)

PART 2

Policers

- [Configuring Policers on page 59](#)

CHAPTER 2

Configuring Policers

- [Overview of Policers on page 59](#)
- [Understanding Policers with Link Aggregation Groups on page 65](#)
- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 65](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 66](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 68](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 68](#)
- [Example: Using Two-Color Policers and Prefix Lists on page 70](#)
- [Example: Using Policers to Manage Oversubscription on page 73](#)
- [Assigning Forwarding Classes and Loss Priority on page 75](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 76](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77](#)
- [Verifying That Two-Color Policers Are Operational on page 79](#)
- [Verifying That Three-Color Policers Are Operational on page 80](#)
- [Troubleshooting Policer Configuration on page 80](#)

Overview of Policers

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 60](#)
- [Policer Types on page 60](#)
- [Policer Actions on page 61](#)
- [Policer Colors on page 62](#)
- [Filter-Specific Policers on page 62](#)
- [Suggested Naming Convention for Policers on page 63](#)
- [Policer Counters on page 63](#)
- [Policer Algorithms on page 63](#)


- [How Many Policers Are Supported? on page 64](#)
- [Policies Can Limit Egress Firewall Filters on page 64](#)

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 60](#) 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 9 on page 61](#) 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 9 on page 61](#) describes the policer actions.

Table 9: 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 9: 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 30](#) 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 ***policer#-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). QFX5110 switches support <<provide value>> ingress policers and <<provide value>> 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 65](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 68](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 66](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 68](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77](#)

Understanding Policers with Link Aggregation Groups

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 59](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77](#)

Understanding Color-Blind Mode for Single-Rate Tricolor Marking

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 10 on page 66](#).

Table 10: 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 59](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 76](#)

Understanding Color-Aware Mode for Single-Rate Tricolor Marking

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 11 on page 66 shows how a packet's incoming priority can be modified with single-rate marking.

Table 11: 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 59](#)
 - [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 76](#)

Understanding Color-Blind Mode for Two-Rate Tricolor Marking

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 12: Color-Blind Mode TCM Color-to-PLP Mapping

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

- Related Documentation**
- [Overview of Policers on page 59](#)
 - [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 76](#)

Understanding Color-Aware Mode for Two-Rate Tricolor Marking

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 13 on page 68](#) shows how a packet's incoming priority can be modified with two-rate marking.

Table 13: Color-Aware Mode Two-Rate PLP Mapping

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

Table 13: Color-Aware Mode Two-Rate PLP Mapping (*continued*)

Incoming PLP	Packet Metered Against	Possible Cases	Outgoing PLP
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 59](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 76](#)

Example: Using Two-Color Policers and Prefix Lists

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 59](#)
 - [prefix-list](#)

Example: Using Policers to Manage Oversubscription

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 14 on page 73](#).

Table 14: 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 59](#)

Assigning Forwarding Classes and Loss Priority

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 15 on page 75](#)

Table 15: Unicast Forwarding Classes

Unicast Forwarding Class	For CoS Traffic Type
be	Best-effort traffic
no-loss	Guaranteed delivery for TCP traffic
fcoe	Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic
nc	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 35](#). (Assigning forwarding classes and PLP is supported only on ingress filters.)

Related Documentation

- [Configuring Firewall Filters on page 35](#)
- [Verifying That Firewall Filters Are Operational on page 38](#)
- [Monitoring Firewall Filter Traffic on page 49](#)
- [Overview of Policers on page 59](#)
- [Understanding CoS Classifiers](#)
- [Understanding CoS Forwarding Classes](#)

Configuring Color-Blind Egress Policers for Medium-Low PLP

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 59](#)
- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 65](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 68](#)
- [Configuring Firewall Filters on page 35](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77](#)

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

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 77](#)
2. [Configuring Three-Color Policers on page 78](#)
3. [Specifying Policers in a Firewall Filter Configuration on page 78](#)
4. [Applying a Firewall Filter That Includes a Policer on page 79](#)

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 Policar

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 35.



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 35](#)
- [Overview of Policers on page 59](#)
- [Verifying That Two-Color Policers Are Operational on page 79](#)
- [Verifying That Three-Color Policers Are Operational on page 80](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 76](#)

Verifying That Two-Color Policers Are Operational

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
Policies:
Name
icmp-connection-policer
tcp-connection-policer
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

```
Packets
10
539
```

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	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77• Configuring Firewall Filters on page 35• Monitoring Firewall Filter Traffic on page 49

Verifying That Three-Color Policers Are Operational

Purpose	Verify that three-color policers in firewall filter configurations are working properly.
Action	<p>Use the following operational mode commands to verify that a three-color policer is working properly:</p> <ul style="list-style-type: none">• show class-of-service forwarding-table classifiers• show interfaces <i>interface-name</i> extensive• show interfaces queue <i>interface-name</i>
Related Documentation	<ul style="list-style-type: none">• Overview of Policers on page 59• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77

Troubleshooting Policer Configuration

- [Incomplete Count of Packet Drops on page 80](#)
- [Counter Reset When Editing Filter on page 81](#)
- [Invalid Statistics for Policer on page 81](#)
- [Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 81](#)
- [Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 82](#)
- [Policers Can Limit Egress Filters on page 83](#)

Incomplete Count of Packet Drops

Problem	<p>Description: Under certain circumstances, Junos OS might display a misleading number of packets dropped by an ingress policer.</p> <p>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</p>
----------------	---

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

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

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 30](#) 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 plicer 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 plicer 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 plicer 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 plicer 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 plicer. The plicer is committed without the counters.

PART 3

Port Security

- [Configuring Port Security on page 87](#)

CHAPTER 3

Configuring Port Security

- Overview of Access Port Protection on page 87
- Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 90
- Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92
- Verifying That DHCP Snooping Is Working Correctly on page 99
- Understanding MAC Limiting and MAC Move Limiting for Port Security on page 100
- Verifying That MAC Limiting Is Working Correctly on page 102
- Verifying That MAC Move Limiting Is Working Correctly on page 105
- Verifying That the Port Error Disable Setting Is Working Correctly on page 106
- Understanding Dynamic ARP Inspection for Protecting Switching Devices Against ARP Spoofing on page 107
- Verifying That DAI Is Working Correctly on page 109
- Understanding Trusted and Untrusted Ports on page 110
- Understanding Trusted DHCP Servers for Port Security on page 110
- Verifying That a Trusted DHCP Server Is Working Correctly on page 111
- Understanding DHCP Option 82 for Port Security on page 112
- Understanding Static ARP Entries on page 114
- Monitoring Port Security on page 115

Overview of Access Port Protection

Port security features can protect a switch against various types of attacks. Protection methods against some common attacks are:

- Mitigation of Ethernet Switching Table Overflow Attacks on page 88
- Mitigation of Rogue DHCP Server Attacks on page 88
- Protection Against ARP Spoofing Attacks on page 89
- Protection Against DHCP Snooping Database Alteration Attacks on page 89
- Protection Against DHCP Starvation Attacks on page 89

Mitigation of Ethernet Switching Table Overflow Attacks

In an overflow attack on an Ethernet switching table, an intruder sends so many requests from new MAC addresses that the table cannot learn all the addresses. The attack forces the switch to send broadcast messages when it needs to send traffic to addresses for which it lacks MAC addresses. In addition to generating unnecessary traffic, the attacker might be able to sniff the broadcast packets.

To mitigate such attacks, you can configure a limit for learned MAC addresses or allow only specific MAC addresses. Use the MAC limit feature to control the total number of MAC addresses that can be added to the Ethernet switching table for the specified interface or interfaces. By setting the MAC addresses that are explicitly allowed, you ensure that the addresses of network devices whose network access is critical are guaranteed to be included in the Ethernet switching table.

Mitigation of Rogue DHCP Server Attacks

By default, all access ports are untrusted, and all trunk ports are trusted with regard to DHCP. Trusted ports allow DHCP servers to provide IP addresses and other information to requesting devices. If someone connects an unauthorized DHCP server to a trusted port, the unauthorized server can start issuing IP addresses and configuration information to the network's DHCP clients. The information provided to the clients by this server can disrupt their network access. The unauthorized server might also assign itself as the default gateway device for the network. An attacker can then sniff the network traffic and perpetrate a man-in-the-middle attack—that is, it misdirects traffic intended for a legitimate network device to a device of its choice.

To mitigate this problem, set the interface to which the unauthorized server is connected as untrusted. That action blocks all ingress DHCP server messages from that interface.



NOTE: The switch logs all DHCP server packets that are received on untrusted ports. For example:

```
5 untrusted DHCPOFFER received, interface xe-0/0/2.0[65], vlan v1[10] server
ip/mac 12.12.12.1/00:00:00:00:01:12 offer ip/client mac
12.12.12.253/00:AA:BB:CC:DD:01
```

You can use these messages to detect unauthorized DHCP servers on the network.



NOTE: If you attach a DHCP server to an access port, you must configure the port as trusted.

Protection Against ARP Spoofing Attacks

In ARP spoofing, an attacker sends faked ARP messages on the network. The attacker associates its own MAC address with the IP address of a network device connected to the switch. Any traffic sent to that IP address is instead sent to the attacker. Now the attacker can create various types of problems, including sniffing the packets that were meant for another host and perpetrating man-in-the-middle attacks. (In a man-in-the-middle attack, the attacker intercepts messages between two hosts, reads them, and perhaps alters them, all without the original hosts knowing that their communications have been compromised.)

To protect against ARP spoofing on your switch, enable both DHCP snooping and dynamic ARP inspection (DAI). DHCP snooping builds and maintains the DHCP snooping table. That table contains the MAC addresses, IP addresses, lease times, binding types, VLAN information, and interface information for the untrusted interfaces on the switch. DAI uses the information in the DHCP snooping table to validate ARP packets. Invalid ARP packets are blocked, and when they are blocked, a system log message is recorded that includes the type of ARP packet and the sender's IP address and MAC address.

See Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks.

Protection Against DHCP Snooping Database Alteration Attacks

In an attack designed to alter the DHCP snooping database, an intruder introduces a DHCP client on one of the switch's untrusted access interfaces that has a MAC address identical to that of a client on another untrusted port. The intruder acquires the DHCP lease, which results in changes to the entries in the DHCP snooping table. Subsequently, what would have been valid ARP requests from the legitimate client are blocked.

To protect against this type of alteration of the DHCP snooping database, configure MAC addresses that are explicitly allowed on the interface. *See Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks.*

Protection Against DHCP Starvation Attacks

In a DHCP starvation attack, an attacker floods an Ethernet LAN with DHCP requests from spoofed (counterfeit) MAC addresses so that trusted DHCP servers cannot keep up with requests from legitimate DHCP clients. The address space of those servers is completely used up, so they can no longer assign IP addresses and lease times to clients. DHCP requests from those clients are either dropped—that is, the result is a denial of service (DoS)—or directed to a rogue DHCP server set up by the attacker to imitate a legitimate DHCP server.

To protect the switch from DHCP starvation attacks, use the MAC limiting feature. Specify the maximum number of MAC addresses that the switch can learn on the access interfaces to which DHCP clients connect. The DHCP server or servers can then supply only the specified number of IP addresses over each of those interfaces. If a DHCP starvation attack occurs after the maximum number of IP addresses has been assigned, the attack fails.

Related Documentation

- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 100](#)
- [Configuring MAC Limiting](#)
- [Verifying That MAC Limiting Is Working Correctly on page 102](#)
- [Understanding DHCP Option 82 for Port Security on page 112](#)
- [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks](#)
- [Understanding Dynamic ARP Inspection for Protecting Switching Devices Against ARP Spoofing on page 107](#)

Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity

Ethernet LANs are vulnerable to attacks such as address spoofing (forging) and Layer 2 denial of service (DoS) on network devices. Port security features help protect the access ports on your device against the loss of information and productivity that such attacks can cause.

The Juniper Networks Junos operating system (Junos OS) provides features to help secure ports on a device. Ports can be categorized as either trusted or untrusted. You apply policies appropriate to each category to protect ports against various types of attacks.

Basic port security features are enabled in the device's default configuration. You can configure additional features with minimal configuration steps.

Depending on the particular feature, you can configure the feature either on VLANs or bridge domain interfaces.

Port security features supported on switching devices are:

- DHCP snooping—Filters and blocks ingress Dynamic Host Configuration Protocol (DHCP) server messages on untrusted ports, and builds and maintains a database of DHCP lease information, which is called the DHCP snooping database.



NOTE: DHCP snooping is not enabled in the default configuration of the switching device. DHCP snooping is enabled on a VLAN or bridge domain. The details of enabling DHCP snooping depend on the particular device.

- DHCPv6 snooping—DHCP snooping for IPv6.
- DHCP option 82—Also known as the DHCP Relay Agent Information option. This DHCPv4 feature helps protect the switching device against attacks such as spoofing of IP addresses and MAC addresses and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.
- DHCPv6 option 37—Option 37 is the remote ID option for DHCPv6 and is used to insert information about the network location of the remote host into DHCPv6 packets. You enable option 37 on a VLAN.



NOTE: DHCPv6 snooping with option 37 is not supported on the MX Series.

- DHCPv6 option 18—Option 18 is the circuit ID option for DHCPv6 and is used to insert information about the client port into DHCPv6 packets. This option includes other details that can be optionally configured, such as the prefix and the interface description.
- DHCPv6 option 16—Option 16 is the vendor ID option for DHCPv6 and is used to insert information about the vendor of the client hardware into DHCPv6 packets.
- Dynamic ARP inspection (DAI)—Prevents Address Resolution Protocol (ARP) spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made on the basis of the results of those comparisons. You enable DAI on a VLAN.
- IPv6 neighbor discovery inspection—Prevents IPv6 address spoofing attacks. Neighbor discovery requests and replies are compared against entries in the DHCPv6 snooping database, and filtering decisions are made on the basis of the results of those comparisons. You enable neighbor discovery inspection on a VLAN.
- IP source guard—Mitigates the effects of IP address spoofing attacks on the Ethernet LAN. With IP source guard enabled, the source IP address in the packet sent from an untrusted access interface is validated against the DHCP snooping database. If the packet cannot be validated, it is discarded. You enable IP source guard on a VLAN or bridge domain.



NOTE: IP source guard is not supported on the QFX Series.

- IPv6 source guard—IP source guard for IPv6.



NOTE: IPv6 source guard is not supported on the QFX Series.

- MAC limiting—Protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You can enable MAC limiting on an interface.
- MAC move limiting—Tracks MAC movement and detects MAC spoofing on access ports. You enable this feature on a VLAN or bridge domain.
- Persistent MAC learning—Also known as sticky MAC. Persistent MAC learning enables interfaces to retain dynamically learned MAC addresses across switch reboots. You enable this feature on an interface.
- Trusted DHCP server—Configuring the DHCP server on a trusted port protects against rogue DHCP servers sending leases. You enable this feature on an interface (port). By default, access ports are untrusted, and trunk ports are trusted. (Access ports are the switch ports that connect to Ethernet endpoints such as user PCs and laptops, servers, and printers. Trunk ports are the switch ports that connect an Ethernet switch to other switches or to routers.)

- Related Documentation**
- [Security Features for EX Series Switches Overview](#)
 - [Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92](#)
 - [Understanding DHCP Snooping for Port Security](#)
 - [Understanding IPv6 Neighbor Discovery Inspection](#)
 - [Understanding Dynamic ARP Inspection for Protecting Switching Devices Against ARP Spoofing on page 107](#)
 - [Understanding IP Source Guard for Port Security on EX Series Switches](#)
 - [Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches](#)
 - [Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks](#)

Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices

DHCP snooping enables the switching device, which can be either a switch or a router, to monitor DHCP messages received from untrusted devices connected to the switching device. When DHCP snooping is enabled on a VLAN, the system examines DHCP messages sent from untrusted hosts associated with the VLAN and extracts their IP addresses and lease information. This information is used to build and maintain the DHCP snooping database. Only hosts that can be verified using this database are allowed access to the network.

- [DHCP Snooping Basics on page 92](#)
- [DHCP Snooping Process on page 93](#)
- [DHCPv6 Snooping on page 94](#)
- [Rapid Commit for DHCPv6 on page 95](#)
- [DHCP Server Access on page 95](#)
- [Static IP Address Additions to the DHCP Snooping Database on page 98](#)
- [Snooping DHCP Packets That Have Invalid IP Addresses on page 98](#)
- [Prioritizing Snooped Packets on page 99](#)

DHCP Snooping Basics

The Dynamic Host Configuration Protocol (DHCP) allocates IP addresses dynamically, *leasing* addresses to devices so that the addresses can be reused when no longer needed. Hosts and end devices that require IP addresses obtained through DHCP must communicate with a DHCP server across the LAN.

DHCP snooping acts as a guardian of network security by keeping track of valid IP addresses assigned to downstream network devices by a trusted DHCP server (the server is connected to a trusted network port).

By default, all trunk ports on the switch are trusted and all access ports are untrusted for DHCP snooping.

When DHCP snooping is enabled, the lease information from the switching device is used to create the DHCP snooping table, also known as the binding table. The table shows the IP-MAC binding, as well as the lease time for the IP address, type of binding, VLAN name, and interface for each host.



NOTE: DHCP snooping is disabled in the default configuration of the switching device. You must explicitly enable DHCP snooping by setting `examine-dhcp` at the `[edit ethernet-switching-options secure-access-port]` hierarchy level.

Entries in the DHCP snooping database are updated in these events:

- When a DHCP client releases an IP address (sends a DHCPRELEASE message). In this event, the associated mapping entry is deleted from the database.
- If you move a network device from one VLAN to another. In this event, typically the device needs to acquire a new IP address. Therefore, its entry in the database, including its VLAN ID, is updated.
- When the lease time (timeout value) assigned by the DHCP server expires. In this event, the associated entry is deleted from the database.



TIP: By default, the IP-MAC bindings are lost when the switching device is rebooted and DHCP clients (the network devices, or hosts) must reacquire bindings. However, you can configure the bindings to persist by setting the `dhcp-snooping-file` statement to store the database file either locally or remotely.

You can configure the switching device to snoop DHCP server responses from particular VLANs only. This prevents spoofing of DHCP server messages.

You configure DHCP snooping per VLAN, not per interface (port). DHCP snooping is disabled by default on switching devices.

DHCP Snooping Process

The basic process of DHCP snooping consists of the following steps:



NOTE: When DHCP snooping is enabled for a VLAN, all DHCP packets sent from the network devices in that VLAN are subjected to DHCP snooping. The final IP-MAC binding occurs when the DHCP server sends DHCPACK to the DHCP client.

1. The network device sends a DHCPDISCOVER packet to request an IP address.
2. The switching device forwards the packet to the DHCP server.

3. The server sends a DHCPOFFER packet to offer an address. If the DHCPOFFER packet is from a trusted interface, the switching device forwards the packet to the DHCP client.
4. The network device sends a DHCPREQUEST packet to accept the IP address. The switching device adds an IP-MAC placeholder binding to the database. The entry is considered a placeholder until a DHCPACK packet is received from the server. Until then, the IP address could still be assigned to some other host.
5. The server sends a DHCPACK packet to assign the IP address or a DHCPNAK packet to deny the address request.
6. The switching device updates the DHCP snooping database according to the type of packet received:
 - If the switching device receives a DHCPACK packet, it updates lease information for the IP-MAC bindings in its database.
 - If the switching device receives a DHCPNACK packet, it deletes the placeholder.



NOTE: The DHCP snooping database is updated only after the DHCPREQUEST packet has been sent.

For general information about the messages that the DHCP client and DHCP server exchange during the assignment of an IP address for the client, see the *Junos OS Administration Library for Routing Devices*.

DHCPv6 Snooping

DHCPv6 snooping is the equivalent of DHCP snooping for IPv6. The process for DHCPv6 snooping is similar to that for DHCP snooping, but uses different names for the messages exchanged between the client and server to assign IPv6 addresses. [Table 16 on page 94](#) shows DHCPv6 messages and their DHCP equivalents.

Table 16: DHCPv6 Messages and Equivalent DHCPv4 Messages

Sent by	DHCPv6 Messages	Equivalent DHCP Messages
Client	SOLICIT	DHCPDISCOVER
Server	ADVERTISE	DHCPOFFER
Client	REQUEST, RENEW, REBIND	DHCPREQUEST
Server	REPLY	DHCPACK/DHCPNAK
Client	RELEASE	DHCPRELEASE
Client	INFORMATION-REQUEST	DHCPINFORM
Client	DECLINE	DHCPDECLINE

Table 16: DHCPv6 Messages and Equivalent DHCPv4 Messages (*continued*)

Sent by	DHCPv6 Messages	Equivalent DHCP Messages
Client	CONFIRM	none
Server	RECONFIGURE	DHCPFORCERENEW
Client	RELAY-FORW, RELAY-REPLY	none

Rapid Commit for DHCPv6

DHCPv6 provides for a Rapid Commit option (DHCPv6 option 14), which, when supported by the server and set by the client, shortens the exchange from a four-way relay to a two-message handshake. For more information about enabling the Rapid Commit option, see *Enabling DHCPv6 Rapid Commit Support*.

In the rapid commit process:

1. The DHCPv6 client sends out a SOLICIT message that contains a request that rapid assignment of address, prefix, and other configuration parameters be preferred.
2. If the DHCPv6 server supports rapid assignment, it responds with a REPLY message, which contains the assigned IPv6 address and prefix and other configuration parameters.

DHCP Server Access

You can configure a switching device's access to the DHCP server in three ways:

- [Switching Device, DHCP Clients, and DHCP Server Are All on the Same VLAN on page 95](#)
- [Switching Device Acts as DHCP Server on page 96](#)
- [Switching Device Acts as Relay Agent on page 97](#)

Switching Device, DHCP Clients, and DHCP Server Are All on the Same VLAN

When the switching device, DHCP clients, and DHCP server are *all members of the same VLAN*, the DHCP server can be connected to the switching device in one of two ways:

- The server is directly connected to the same switching device as the one connected to the DHCP clients (the hosts, or network devices, that are requesting IP addresses from the server). The VLAN is enabled for DHCP snooping to protect the untrusted access ports. The trunk port is configured by default as a trusted port. See [Figure 4 on page 96](#).
- The server is connected to an intermediary switching device (Switching Device 2). The DHCP clients are connected to Switching Device 1, which is connected through a trunk port to Switching Device 2. Switching Device 2 is being used as a transit device. The VLAN is enabled for DHCP snooping to protect the untrusted access ports. The trunk port is configured by default as a trusted port. As shown in [Figure 5 on page 96](#), ge-0/0/11 is a trusted trunk port.

Figure 4: DHCP Server Connected Directly to a Switching Device

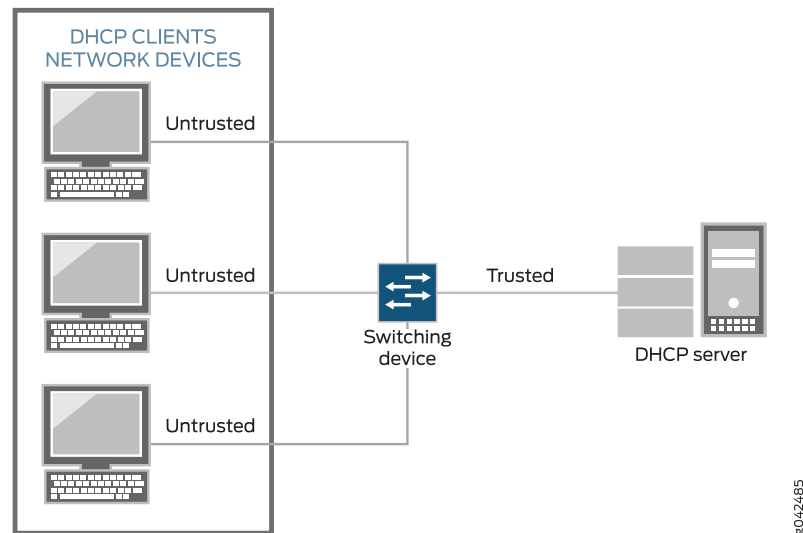



Figure 5: DHCP Server Connected Directly to Switching Device 2, with Switching Device 2 Connected to Switching Device 1 Through a Trusted Trunk Port




Switching Device Acts as DHCP Server



NOTE: The switching device acting as a DHCP server is not supported on the QFX Series.

The switching device itself is configured as a DHCP server; this is known as a *local configuration*. See [Figure 6 on page 97](#).

Figure 6: Switching Device Is the DHCP Server




Switching Device Acts as Relay Agent

The switching device functions as a relay agent when the DHCP clients or the DHCP server is connected to the device through a Layer 3 interface. The Layer 3 interfaces on the switching device are configured as routed VLAN interfaces (RVIs), which are also known as integrated routing and bridging (IRB) interfaces. The trunk interfaces are trusted by default.

These two scenarios illustrate the switching device acting as a relay agent:

- The DHCP server and clients are in different VLANs.
- The switching device is connected to a router that is in turn connected to the DHCP server. See [Figure 7 on page 98](#).

Figure 7: Switching Device Acting as Relay Agent Through Router to DHCP Server



Static IP Address Additions to the DHCP Snooping Database

You can add specific static IP addresses to the database as well as have the addresses dynamically assigned through DHCP snooping. To add static IP addresses, you supply the IP address, the MAC address of the device, the interface on which the device is connected, and the VLAN with which the interface is associated. No lease time is assigned to the entry. The statically configured entry never expires.

Snooping DHCP Packets That Have Invalid IP Addresses

If you enable DHCP snooping on a VLAN and then devices on that VLAN send DHCP packets that request invalid IP addresses, these invalid IP addresses are stored in the DHCP snooping database until they are deleted when their default timeout is reached. To eliminate this unnecessary consumption of space in the DHCP snooping database, the switching device drops the DHCP packets that request invalid IP addresses, preventing the snooping of these packets. The invalid IP addresses are:

- 0.0.0.0
- 128.0.x.x
- 191.255.x.x
- 192.0.0.x
- 223.255.255.x
- 224.x.x.x
- 240.x.x.x to 255.255.255.255

Prioritizing Snooped Packets



NOTE: Prioritizing snooped packets is not supported on the QFX Series and the EX4600 switch.

You can use class-of-service (CoS) forwarding classes and queues to prioritize DHCP snooped packets for a specified VLAN. This type of configuration places the DHCP snooped packets for that VLAN in a specified egress queue, so that the security procedure does not interfere with the transmission of high-priority traffic. For additional information, see *Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic*.

Related Documentation

- [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 90](#)
- [Understanding Trusted DHCP Servers for Port Security on page 110](#)
- [Configuring a DHCP Server on a Trusted Interface to Protect it Against Rogue DHCP Servers Sending Leases \(CLI Procedure\)](#)
- [Configuring Static IP Addresses for DHCP Bindings on Access Ports for MX Series Routers \(CLI Procedure\)](#)

Verifying That DHCP Snooping Is Working Correctly

Purpose Verify that DHCP snooping is working on the switch and that the DHCP snooping database is correctly populated with both dynamic and static bindings.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch> show dhcp snooping binding
```

DHCP Snooping Information:

MAC address	IP address	Lease (seconds)	Type	VLAN	Interface
00:05:85:3A:82:77	192.0.2.17	600	dynamic	employee	ge-0/0/1.0
00:05:85:3A:82:79	192.0.2.18	653	dynamic	employee	ge-0/0/1.0
00:05:85:3A:82:80	192.0.2.19	720	dynamic	employee	ge-0/0/2.0
00:05:85:3A:82:81	192.0.2.20	932	dynamic	employee	ge-0/0/2.0
00:05:85:3A:82:83	192.0.2.21	1230	dynamic	employee	ge-0/0/2.0
00:05:85:27:32:88	192.0.2.22	—	static	data	ge-0/0/4.0

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease

expires. Static IP addresses have no assigned lease time. The statically configured entry never expires.

If the DHCP server had been configured as untrusted, no entries would be added to the DHCP snooping database and nothing would be shown in the output of the **show dhcp snooping binding** command.

**Related
Documentation**

- *Enabling DHCP Snooping (CLI Procedure)*
- *Enabling DHCP Snooping (J-Web Procedure)*
- *Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure)*
- *Example: Configuring Basic Port Security Features*
- *Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch*
- *Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks*
- [Monitoring Port Security on page 115](#)
- *Troubleshooting Port Security*

Understanding MAC Limiting and MAC Move Limiting for Port Security

MAC limiting protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You enable this feature on Layer 2 interfaces (ports). MAC move limiting detects MAC movement and MAC spoofing on access interfaces. You enable this feature on VLANs.

- [MAC Limiting on page 100](#)
- [MAC Move Limiting on page 101](#)
- [Actions for MAC Limiting on page 101](#)
- [MAC Addresses That Exceed the MAC Limit or MAC Move Limit on page 102](#)

MAC Limiting

MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface or on all the Layer 2 access interfaces on the switch. Junos OS provides two MAC limiting methods:

- **Maximum number of MAC addresses**—You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses can be ignored, dropped, or logged. You can also specify that the interface be shut down or temporarily disabled.
- **Allowed MAC addresses**—You configure specific “allowed” MAC addresses for the access interface. Any MAC address that is not in the list of configured addresses is not learned, and the switch logs an appropriate message. Allowed MAC binds MAC

addresses to a VLAN so that the address does not get registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.



NOTE: If you do not want the system to log messages about invalid MAC addresses received by an interface that has been configured for allowed MAC addresses, disable the logging by configuring the `no-allowed-mac-log` statement.

You configure MAC limiting per interface, not per VLAN. You can specify the maximum number of dynamic MAC addresses that can be learned on a single Layer 2 access interface (including tagged-access interfaces) or on all Layer 2 access interfaces.

MAC Move Limiting

MAC move limiting causes the switch to track the number of times a MAC address can move to a new interface (port). It can help to prevent MAC spoofing, and it can also detect and prevent loops.

If a MAC address moves more than the configured number of times within 1 second, the switch performs the configured action. You can configure MAC move limiting to apply to all VLANs or to a specific VLAN.



CAUTION: Mac move limiting does not work properly on a QFX5100 switch used as a Node device in a QFabric system. Do not use this feature on a QFX5100 switch in a QFabric system.

Actions for MAC Limiting

You can choose to have one of the following actions performed when the limit of MAC addresses or the limit of MAC moves is exceeded:

- **drop**—Drop the packet and generate a system log entry. This is the default.
- **log**—Do not drop the packet but generate a system log entry.
- **none**—Take no action.
- **shutdown**—Disable the interface and generate an alarm. If you configure the switch with the `port-error-disable` statement, the disabled interface recovers automatically upon expiration of the specified timeout. If this is not configured, you can bring up the disabled interfaces by running the `clear ethernet-switching port-error` command.

See descriptions of results of these various action settings in “[Verifying That MAC Limiting Is Working Correctly](#)” on page 102.

If you set a MAC limit to apply to all interfaces on the switch, you can override that setting for a particular interface by specifying action **none**. See *mac-limit* for more information.

MAC Addresses That Exceed the MAC Limit or MAC Move Limit

If you have configured the **port-error-disable** statement, you can view which interfaces are temporarily disabled because the MAC limit or MAC move limit was exceeded. Use the **show ethernet-switching interfaces** command.

The log messages that indicate the MAC limit or MAC move limit has been exceeded include the offending MAC addresses.

Related Documentation

- [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 90](#)
- [Configuring MAC Limiting](#)
- [Configuring MAC Move Limiting \(CLI Procedure\)](#)
- [Verifying That MAC Limiting Is Working Correctly on page 102](#)
- [Verifying That MAC Move Limiting Is Working Correctly on page 105](#)
- [Example: Configuring Basic Port Security Features](#)
- [no-allowed-mac-log on page 204](#)

Verifying That MAC Limiting Is Working Correctly

MAC limiting protects against flooding of the Ethernet switching table by setting a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface (port).

Junos OS provides two MAC limiting methods:

- Maximum number of MAC addresses—You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses can be ignored, dropped, or logged. You can also specify that the interface be shut down or temporarily disabled.
- Allowed MAC addresses—You configure specific “allowed” MAC addresses for the access interface. Any MAC address that is not in the list of configured addresses is not learned, and the switch logs an appropriate message. The allowed MAC method binds MAC addresses to a VLAN so that the address is not registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.

This topic includes the following tasks:

1. [Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly on page 103](#)
2. [Verifying That Allowed MAC Addresses Are Working Correctly on page 103](#)
3. [Verifying That Interfaces Are Shut Down on page 104](#)
4. [Customizing the Ethernet Switching Table Display to View Information for a Specific Interface on page 104](#)

Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly

Purpose Verify that MAC limiting for dynamic MAC addresses is working.

Action Display the MAC addresses that have been learned. The following sample output shows the results of sending two packets from hosts connected to **xe-1:0/0/1** and five packets from hosts connected to **xe-1:0/0/2**, with both interfaces configured with a MAC limit of 4 and the action **drop**:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 7 entries, 6 learned
```

VLAN	MAC address	Type	Age	Interfaces
employee-vlan	*	Flood	-	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:77	Learn	0	xe-1:0/0/1.0
employee-vlan	00:05:85:3A:82:79	Learn	0	xe-1:0/0/1.0
employee-vlan	00:05:85:3A:82:80	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:81	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:83	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:85	Learn	0	xe-1:0/0/2.0

Meaning The output shows that the fifth packet received on the **xe-1:0/0/2** interface was dropped because it exceeded the MAC limit for that interface. The address was not learned, and thus an asterisk (*) rather than an address appears in the MAC address column in the first line of the sample output.

Verifying That Allowed MAC Addresses Are Working Correctly

Purpose Verify that allowed MAC addresses are working.

Action Display the MAC cache information after allowed MAC addresses have been configured on an interface. The following sample shows the MAC cache after four allowed MAC addresses had been configured on interface **xe-1:0/0/2** and a fifth MAC address appeared on the interface.

```
user@switch> show ethernet-switching table
Ethernet-switching table: 5 entries, 4 learned
```

VLAN	MAC address	Type	Age	Interfaces
employee-vlan	00:05:85:3A:82:80	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:81	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:83	Learn	0	xe-1:0/0/2.0
employee-vlan	00:05:85:3A:82:85	Learn	0	xe-1:0/0/2.0
employee-vlan	*	Flood	-	xe-1:0/0/2.0

Meaning Because the fifth address was not allowed it was not learned, and an asterisk (*) rather than an address appears in the MAC address column in the last line of the sample output.

Verifying That Interfaces Are Shut Down

Purpose Verify that an interface is shut down when the MAC limit is exceeded.

Action For more information about interfaces that have been shut down because the MAC limit was exceeded, use the **show ethernet-switching interfaces** command.

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

Interface	State	VLAN members	Tag	Tagging	Blocking
-----------	-------	--------------	-----	---------	----------

bme0.32770	down	mgmt	untagged	unblocked	
xe-0/0/0.0	down	v1	untagged	MAC limit exceeded	
xe-0/0/1.0	up	v1	untagged	unblocked	
xe-0/0/2.0	up	v1	untagged	unblocked	
me0.0	up	mgmt	untagged	unblocked	



NOTE: You can configure interfaces to recover automatically when the MAC limit has been exceeded by specifying the **port-error-disable** statement with a **disable timeout** value. The switch automatically restores the disabled interface to service when the disable timeout expires. The **port-error-disable** configuration does not apply to preexisting error conditions—it affects only error conditions that are detected after the **port-error-disable** statement has been enabled and the configuration has been committed. To clear a preexisting error condition and restore the interface to service, use the **clear ethernet-switching port-error** command.

Customizing the Ethernet Switching Table Display to View Information for a Specific Interface

Purpose You can use the **show ethernet-switching table** command to view information for a specific interface.

Action For example, to display the MAC addresses that have been learned on the **xe-0/0/2.0** interface, enter:

```
user@switch> show ethernet-switching table interface xe-0/0/2.0
```

Ethernet-switching table: 1 unicast entries

VLAN	MAC address	Type	Age	Interfaces
v1	*	Flood	-	All-members
v1	00:00:06:00:00:00	Learn	0	xe-0/0/2.0

Meaning The MAC limit value for the **xe-0/0/2** interface had been set to **1**, and the output shows that only one MAC address was learned and added to the MAC cache.

- Related Documentation**
- [Configuring MAC Limiting](#)
 - [Monitoring Port Security on page 115](#)
 - [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\)](#)
 - [Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks](#)
 - [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks](#)

Verifying That MAC Move Limiting Is Working Correctly

Purpose Verify that MAC move limiting is working on the switch.

Action Display the MAC addresses in the Ethernet switching table when MAC move limiting has been configured for a VLAN. The following sample shows the results after two of the hosts on **ge-0/0/2** sent packets after the MAC addresses for those hosts had moved to other interfaces more than five times in 1 second. The VLAN, **employee-vlan**, was set to a MAC move limit of **5** with the action **drop**:

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

```
Ethernet-switching table: 7 entries, 4 learned
```

VLAN	MAC address	Type	Age	Interfaces
employee-vlan	00:05:85:3A:82:77	Learn	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:79	Learn	0	ge-0/0/1.0
employee-vlan	00:05:85:3A:82:80	Learn	0	ge-0/0/2.0
employee-vlan	00:05:85:3A:82:81	Learn	0	ge-0/0/2.0
employee-vlan	*	Flood	-	ge-0/0/2.0
employee-vlan	*	Flood	-	ge-0/0/2.0

Meaning The last two lines of the sample output show that MAC addresses for two hosts on **ge-0/0/2** were not learned, because the hosts had been moved back and forth from the original interfaces more than five times in 1 second.

- Related Documentation**
- [Configuring MAC Move Limiting \(CLI Procedure\)](#)
 - [Configuring MAC Move Limiting \(J-Web Procedure\)](#)
 - [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\)](#)
 - [Example: Configuring Basic Port Security Features](#)
 - [Monitoring Port Security on page 115](#)

Verifying That the Port Error Disable Setting Is Working Correctly

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 on page 100](#)
- [port-error-disable on page 206](#)

Understanding Dynamic ARP Inspection for Protecting Switching Devices Against ARP Spoofing

Dynamic ARP inspection (DAI) protects switching devices against ARP spoofing.

DAI inspects Address Resolution Protocol (ARP) packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP spoofing (also known as ARP poisoning or ARP cache poisoning). ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons. When an attacker tries to use a forged ARP packet to spoof an address, the switch compares the address with entries in the database. If the media access control (MAC) address or IP address in the ARP packet does not match a valid entry in the DHCP snooping database, the packet is dropped.

ARP packets are sent to the Routing Engine and are rate-limited to protect the switching device from CPU overload.

- [Address Resolution Protocol on page 107](#)
- [ARP Spoofing on page 107](#)
- [Dynamic ARP Inspection on page 108](#)
- [Prioritizing Inspected Packets on page 109](#)

Address Resolution Protocol

Sending IP packets on a multi-access network requires mapping an IP address to an Ethernet MAC address.

Ethernet LANs use ARP to map MAC addresses to IP addresses.

The switching device maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

ARP Spoofing

ARP spoofing is one way to initiate man-in-the-middle attacks. The attacker sends an ARP packet that spoofs the MAC address of another device on the LAN. Instead of the switching device sending traffic to the proper network device, it sends the traffic to the device with the spoofed address that is impersonating the proper device. If the impersonating device is the attacker's machine, the attacker receives all the traffic from the switch that must have gone to another device. The result is that traffic from the switching device is misdirected and cannot reach its proper destination.

One type of ARP spoofing is gratuitous ARP, which is when a network device sends an ARP request to resolve its own IP address. In normal LAN operation, gratuitous ARP messages indicate that two devices have the same MAC address. They are also broadcast when a network interface card (NIC) in a device is changed and the device is rebooted, so that other devices on the LAN update their ARP caches. In malicious situations, an

attacker can poison the ARP cache of a network device by sending an ARP response to the device that directs all packets destined for a certain IP address to go to a different MAC address instead.

To prevent MAC spoofing through gratuitous ARP and through other types of spoofing, the switches examine ARP responses through DAI.

Dynamic ARP Inspection

DAI examines ARP requests and responses on the LAN and validates ARP packets. The switch intercepts ARP packets from an access port and validates them against the DHCP snooping database. If no IP-MAC entry in the database corresponds to the information in the ARP packet, DAI drops the ARP packet and the local ARP cache is not updated with the information in that packet. DAI also drops ARP packets when the IP address in the packet is invalid. ARP probe packets are not subjected to dynamic ARP inspection. The switch always forwards such packets.

Junos OS for EX Series switches and the QFX Series uses DAI for ARP packets received on access ports because these ports are untrusted by default. Trunk ports are trusted by default, and therefore ARP packets bypass DAI on them.

You configure DAI for each VLAN, not for each interface (port). By default, DAI is disabled for all VLANs.

If you set an interface to be a DHCP trusted port, it is also trusted for ARP packets.



NOTE:

- If your switching device is an EX Series switch and uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style, see *Enabling a Trusted DHCP Server (CLI Procedure)* for information about configuring an access interface to be a DHCP trusted port.
- If your switching device is an EX Series switch and is *not* using Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style, see *Enabling a Trusted DHCP Server (CLI Procedure)* for information about configuring an access interface to be a DHCP trusted port.

For packets directed to the switching device to which a network device is connected, ARP queries are broadcast on the VLAN. The ARP responses to those queries are subjected to the DAI check.

For DAI, all ARP packets are trapped to the Packet Forwarding Engine. To prevent CPU overloading, ARP packets destined for the Routing Engine are rate-limited.

If the DHCP server goes down and the lease time for an IP-MAC entry for a previously valid ARP packet runs out, that packet is blocked.

Prioritizing Inspected Packets



NOTE: Prioritizing inspected packets is not supported on the QFX Series and the EX4600 switch.

You can use class-of-service (CoS) forwarding classes and queues to prioritize DAI packets for a specified VLAN. This type of configuration places inspected packets for that VLAN in the egress queue, that you specify, ensuring that the security procedure does not interfere with the transmission of high-priority traffic.

Related Documentation

- [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 90](#)
- [Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92](#)
- [Example: Configuring Basic Port Security Features](#)
- [Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch](#)
- [Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks](#)
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing](#)
- [Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic](#)
- [Enabling Dynamic ARP Inspection \(CLI Procedure\)](#)
- [Enabling Dynamic ARP Inspection \(CLI Procedure\)](#)
- [Enabling Dynamic ARP Inspection \(J-Web Procedure\)](#)

Verifying That DAI Is Working Correctly

Purpose Verify that dynamic ARP inspection (DAI) is working on the switch.

Action Send some ARP requests from network devices connected to the switch.
Display the DAI information:

```
user@switch> show arp inspection statistics
ARP inspection statistics:
Interface          Packets received  ARP inspection pass  ARP inspection failed
-----
ge-0/0/1.0          7                 5                   2
ge-0/0/2.0          10                10                  0
ge-0/0/3.0          12                12                  0
```

Meaning The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection

on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

Related Documentation

- *Enabling Dynamic ARP Inspection (CLI Procedure)*
- *Enabling Dynamic ARP Inspection (J-Web Procedure)*
- *Example: Configuring Basic Port Security Features*
- *Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch*
- *Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks*
- [Monitoring Port Security on page 115](#)

Understanding Trusted and Untrusted Ports

By default, all access ports are untrusted and all trunk ports are trusted in regard to DHCP. Trusted ports allow DHCP servers to provide IP addresses and other information to requesting devices. Untrusted ports drop traffic from DHCP servers to prevent unauthorized servers from providing any configuration information to clients.

If you attach a DHCP server to an access port, you must configure the port as trusted. Before you do so, ensure that the server is physically secure—that is, that access to the server is monitored and controlled.

Related Documentation

- [Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92](#)
- *Example: Configuring Basic Port Security Features*

Understanding Trusted DHCP Servers for Port Security

Any interface on the switching device that connects to a DHCP server can be configured as a trusted port. Configuring a DHCP server on a trusted port protects against rogue DHCP servers sending leases.

Ensure that the DHCP server interface is physically secure—that is, that access to the server is monitored and controlled at the site—before you configure the port as trusted.

Related Documentation

- *Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks*
- *Enabling a Trusted DHCP Server (CLI Procedure)*

Verifying That a Trusted DHCP Server Is Working Correctly

Purpose Verify that a DHCP trusted server is working on the switch. See what happens when the DHCP server is trusted and then untrusted.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch> show dhcp snooping binding
```

DHCP Snooping Information:

MAC Address	IP Address	Lease	Type	VLAN	Interface
00:05:85:3A:82:77	192.0.2.17	600	dynamic	employee-vlan	ge-0/0/1.0
00:05:85:3A:82:79	192.0.2.18	653	dynamic	employee-vlan	ge-0/0/1.0
00:05:85:3A:82:80	192.0.2.19	720	dynamic	employee-vlan	ge-0/0/2.0
00:05:85:3A:82:81	192.0.2.20	932	dynamic	employee-vlan	ge-0/0/2.0
00:05:85:3A:82:83	192.0.2.21	1230	dynamic	employee-vlan	ge-0/0/2.0
00:05:85:27:32:88	192.0.2.22	3200	dynamic	employee-vlan	ge-0/0/2.0

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires.

If the DHCP server had been configured as untrusted, no entries would be added to the DHCP snooping database and nothing would be shown in the output of the **show dhcp snooping binding** command.

- Related Documentation**
- *Enabling a Trusted DHCP Server (CLI Procedure)*
 - *Enabling a Trusted Port for DHCP*
 - *Example: Configuring Basic Port Security Features*
 - *Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks*
 - [Monitoring Port Security on page 115](#)
 - *Troubleshooting Port Security*

Understanding DHCP Option 82 for Port Security

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Hosts on untrusted access interfaces on Ethernet LAN switches send requests for IP addresses in order to access the Internet. The switch forwards or relays these requests to DHCP servers, and the servers send offers for IP address leases in response. Attackers can use these messages to perpetrate address spoofing and starvation.

Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client. The Juniper Networks Junos operating system (Junos OS) implementation of DHCP option 82 supports RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.

This topic covers:

- [DHCP Option 82 Processing on page 112](#)
- [Suboption Components of Option 82 on page 113](#)
- [Configurations That Support Option 82 on page 113](#)

DHCP Option 82 Processing

If DHCP option 82 is enabled on the switch, then when a DHCP client that is connected to the switch on an untrusted interface sends a DHCP request, the switch inserts information about the client's network location into the packet header of that request. The switch then sends the request to the DHCP server. The DHCP server reads the option 82 information in the packet header and uses it to implement the IP address or another parameter for the client. See “[Suboption Components of Option 82](#)” on [page 113](#) for details about option 82 information.

You can enable DHCP option 82 on a single VLAN or on all VLANs on the switch. You can also configure it on Layer 3 interfaces (in routed VLAN interfaces, or RVIs) when the switch is functioning as a relay agent.

When option 82 is enabled on the switch, then this sequence of events occurs when a DHCP client sends a DHCP request:

1. The switch receives the request and inserts the option 82 information in the packet header.
2. The switch forwards or relays the request to the DHCP server.
3. The server uses the DHCP option 82 information to formulate its reply and sends a response back to the switch. It does not alter the option 82 information.
4. The switch strips the option 82 information from the response packet.
5. The switch forwards the response packet to the client.



NOTE: To use the DHCP option 82 feature, you must ensure that the DHCP server is configured to accept option 82. If it is not configured to accept option 82, then when it receives requests containing option 82 information, it does not use the information in setting parameters and it does not echo the information in its response message.

Suboption Components of Option 82

When configuring DHCP option 82, you can use the following suboptions:

- **circuit ID**—Identifies the circuit (interface and/or VLAN) on the switch on which the request was received. The circuit ID contains the interface name and/or VLAN name, with the two elements separated by a colon—for example, `xe-0/0/10:vlan1`. If the request packet is received on a Layer 3 interface, the circuit ID is just the interface name—for example, `xe-0/0/10`.

Use the **prefix** option to add an optional prefix to the circuit ID. If you enable the **prefix** option, the hostname for the switch is used as the prefix; for example, `switch1:xe-0/0/10:vlan1`.

You can also specify that the interface description be used rather than the interface name and that the VLAN ID be used rather than the VLAN name.


- **remote ID**—Identifies the host. By default, the remote ID is the MAC address of the switch. You can specify that the remote ID be the hostname of the switch, the interface description, or a character string of your choice. You can also add an optional prefix to the remote ID.
- **vendor ID**—Identifies the vendor of the host. If you specify the **vendor-id** option but do not enter a value, the default value **Juniper** is used. To specify a value, you type a character string.

Configurations That Support Option 82

You can use option 82 with the following configurations:

- The DHCP client and the DHCP server are on the same VLAN. In this case the switch forwards the requests from the clients on untrusted access interfaces to the server on a trusted interface. For this configuration, you set DHCP option 82 at the **[edit ethernet-switching-options secure-access-port vlan]** hierarchy level.
- The DHCP client or the DHCP server is connected to the switch through a Layer 3 interface and the switch is configured to relay DHCP requests. [Figure 8 on page 114](#) illustrates a scenario for the switch-as-relay-agent; in this instance, the switch relays requests through a router to the server.

Figure 8: Switch Relays DHCP Requests to Server



For the configuration shown in [Figure 8 on page 114](#), you set DHCP option 82 at the **[edit forwarding-options helpers bootp]** hierarchy level.

Related Documentation

- [Overview of Access Port Protection on page 87](#)
- [DHCP and BOOTP Relay Overview](#)
- [dhcp-option82](#)
- [Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server](#)
- [Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\)](#)
- [Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server \(CLI Procedure\)](#)

Understanding Static ARP Entries

You can create explicit mappings between IP addresses and MAC addresses, which are called static ARP table entries. Unlike dynamically learned ARP entries, static entries do not age out. You might want to create static ARP entries in a troubleshooting situation or if your device is unable to learn a MAC address dynamically for any reason.

Related Documentation

- [Configuring Static ARP Entries](#)

- `arp`

Monitoring Port Security

Purpose



NOTE: This topic applies only to the J-Web Application package.

Use the monitoring functionality to view these port security details:

- DHCP snooping database for a VLAN or all VLANs
- ARP inspection details for all interfaces

Action

To monitor port security in the J-Web interface, select **Monitor > Security > Port Security**.

To monitor and manipulate the DHCP snooping database and ARP inspection statistics in the CLI, enter the following commands:

- **show dhcp snooping binding**
- **clear dhcp snooping binding**—In addition to clearing the whole database, you can clear database entries for specified VLANs or MAC addresses.
- **show arp inspection statistics**
- **clear arp inspection statistics**



NOTE: On EX4300 switches, to monitor and manipulate the DHCP snooping database and ARP inspection statistics in the CLI, enter the following commands:

- **show dhcp-security binding**
- **clear dhcp-security binding**—In addition to clearing the whole database, you can clear database entries for specified VLANs or IP Address.
- **show dhcp-security arp inspection statistics**
- **clear arp inspection statistics**

Meaning

The J-Web Port Security Monitoring page comprises two sections:

- **DHCP Snooping Details**—Displays the DHCP snooping database for all the VLANs for which DHCP snooping is enabled. To view the DHCP snooping database for a specific VLAN, select the specific VLAN from the list.
- **ARP Inspection Details**—Displays the ARP inspection details for all interfaces. The information includes details of the number of packets that passed ARP inspection and the number of packets that failed the inspection. The pie chart graphically represents

these statistics when you select an interface. To view ARP inspection statistics for a specific interface, select the interface from the list.

You can use the following options on the page to clear DHCP snooping and ARP inspection details:

- **Clear All**—Clears the DHCP snooping database, either for all VLANs if the option **ALL** has been selected in the Select VLANs list or for the specific VLAN that has been selected in that list.
- **Clear**—Deletes a specific IP address from the DHCP snooping database.

To clear ARP inspection details on the page, click **Clear All** in the ARP inspection details section.



NOTE: Clear All button in the ARP inspection details section is not supported on EX4300 switches.

Use the CLI commands to show and clear DHCP snooping database and ARP inspection statistics details.

**Related
Documentation**

- *Configuring Port Security (CLI Procedure)*
- *Configuring Port Security (J-Web Procedure)*
- *Example: Configuring Basic Port Security Features*

PART 4

Device Security

- [Configuring Device Security on page 119](#)

CHAPTER 4

Configuring Device Security

- [Understanding Storm Control on page 119](#)
- [Example: Configuring Storm Control to Prevent Network Outages on page 120](#)
- [Verifying That the Port Error Disable Setting Is Working Correctly on page 123](#)
- [Understanding Unicast RPF on page 124](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 128](#)
- [Verifying Unicast RPF Status on page 129](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 132](#)
- [Understanding Unknown Unicast Forwarding on page 133](#)
- [Configuring Unknown Unicast Forwarding \(CLI Procedure\) on page 133](#)

Understanding Storm Control

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.



NOTE: On a QFX10002 switch, if storm control is configured on a VLAN port associated with an IRB interface, unregistered multicast traffic is classified as registered multicast traffic if IGMP snooping is enabled. If IGMP snooping is disabled, the traffic is classified as unknown unicast traffic.

**Related
Documentation**

- [action-shutdown on page 234](#)
- [port-error-disable on page 206](#)
- [storm-control on page 246](#)

Example: Configuring Storm Control to Prevent Network Outages

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 121](#)
- [Overview and Topology on page 121](#)
- [Configuration on page 122](#)

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 119](#)
 - [Configuring Autorecovery for MAC Limited or Storm Control Interfaces \(CLI Procedure\)](#)

Verifying That the Port Error Disable Setting Is Working Correctly

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 on page 100](#)
 - [port-error-disable on page 206](#)

Understanding Unicast RPF

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 127.](#)

This topic covers:

- [Unicast RPF for Switches Overview on page 124](#)
- [Unicast RPF Implementation on page 125](#)
- [When to Enable Unicast RPF on page 125](#)
- [When Not to Enable Unicast RPF on page 126](#)
- [Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches on page 127](#)

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 125.](#))

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 125](#)
- [Bootstrap Protocol \(BOOTP\) and DHCP Requests on page 125](#)
- [Default Route Handling on page 125](#)

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.

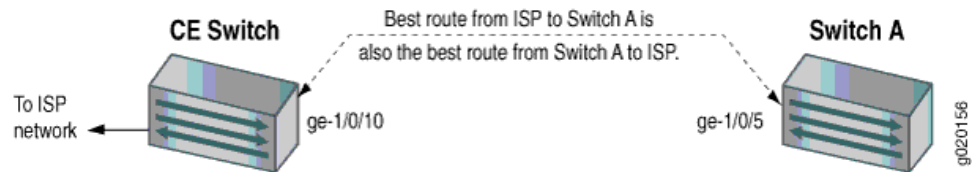
When to Enable Unicast RPF

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 9 on page 126](#). 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 9: 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 10 on page 127](#). 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 10: 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 128](#)
 - [Disabling Unicast RPF \(CLI Procedure\) on page 132](#)

Configuring Unicast RPF (CLI Procedure)

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 129](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 132](#)
- [Troubleshooting Unicast RPF](#)
- [Understanding Unicast RPF on page 124](#)

Verifying Unicast RPF Status

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:
    Receive          Transmit
    Total octets     0                0
    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 128](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 132](#)
- *Troubleshooting Unicast RPF*

Disabling Unicast RPF (CLI Procedure)

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 129](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 128](#)

- [Understanding Unicast RPF on page 124](#)

Understanding Unknown Unicast Forwarding

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 133](#)
- [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)



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 134](#)
- [Configuring Unknown Unicast Forwarding on EX9200 Switches on page 134](#)

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 133](#)
- [Verifying That Unknown Unicast Packets Are Forwarded to a Single Interface](#)

PART 5

Media Access Control Security (MACsec)

- [Configuring Media Access Control Security \(MACsec\) on page 139](#)

CHAPTER 5

Configuring Media Access Control Security (MACsec)

- [Understanding Media Access Control Security \(MACsec\) on page 139](#)
- [Configuring Media Access Control Security \(MACsec\) on page 147](#)

Understanding Media Access Control Security (MACsec)

Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication for all traffic on Ethernet links. MACsec provides point-to-point security on Ethernet links between directly connected nodes and is capable of identifying and preventing most security threats, including denial of service, intrusion, man-in-the-middle, masquerading, passive wiretapping, and playback attacks.

MACsec allows you to secure an Ethernet link for almost all traffic, including frames from the Link Layer Discovery Protocol (LLDP), Link Aggregation Control Protocol (LACP), Dynamic Host Configuration Protocol (DHCP), Address Resolution Protocol (ARP), and other protocols that are not typically secured on an Ethernet link because of limitations with other security solutions. MACsec can be used in combination with other security protocols such as IP Security (IPsec) and Secure Sockets Layer (SSL) to provide end-to-end network security.

MACsec is standardized in IEEE 802.1AE. The IEEE 802.1AE standard can be seen on the IEEE organization website at [IEEE 802.1: BRIDGING & MANAGEMENT](#).

This topic contains the following sections:

- [How MACsec Works on page 140](#)
- [Understanding Connectivity Associations and Secure Channels on page 140](#)
- [Understanding MACsec Security Modes on page 141](#)
- [Understanding the Requirements to Enable MACsec on a Switch-to-Host Link on page 143](#)
- [MACsec Hardware and Software Support Summary on page 143](#)
- [Understanding MACsec Hardware Requirements for EX Series and QFX Series Switches on page 144](#)
- [Understanding MACsec Software Requirements for EX Series and QFX Series Switches on page 145](#)

- [Understanding the MACsec Feature License Requirement on page 146](#)
- [MACsec Limitations on page 146](#)

How MACsec Works

MACsec provides industry-standard security through the use of secured point-to-point Ethernet links. The point-to-point links are secured after matching security keys—a user-configured pre-shared key when you enable MACsec using static connectivity association key (CAK) security mode, a user-configured static secure association key when you enable MACsec using static secure association key (SAK) security mode, or a dynamic key included as part of the AAA handshake with the RADIUS server when you enable MACsec using dynamic security mode—are exchanged and verified between the interfaces at each end of the point-to-point Ethernet link. Other user-configurable parameters, such as MAC address or port, must also match on the interfaces on each side of the link to enable MACsec. See [“Configuring Media Access Control Security \(MACsec\)” on page 147](#).

Once MACsec is enabled on a point-to-point Ethernet link, all traffic traversing the link is MACsec-secured through the use of data integrity checks and, if configured, encryption.

The data integrity checks verify the integrity of the data. MACsec appends an 8-byte header and a 16-byte tail to all Ethernet frames traversing the MACsec-secured point-to-point Ethernet link, and the header and tail are checked by the receiving interface to ensure that the data was not compromised while traversing the link. If the data integrity check detects anything irregular about the traffic, the traffic is dropped.

MACsec can also be used to encrypt all traffic on the Ethernet link. The encryption used by MACsec ensures that the data in the Ethernet frame cannot be viewed by anybody monitoring traffic on the link. MACsec encryption is optional and user-configurable; you can enable MACsec to ensure the data integrity checks are performed while still sending unencrypted data “in the clear” over the MACsec-secured link, if desired.

MACsec is configured on point-to-point Ethernet links between MACsec-capable interfaces. If you want to enable MACsec on multiple Ethernet links, you must configure MACsec individually on each point-to-point Ethernet link.

Understanding Connectivity Associations and Secure Channels

MACsec is configured in connectivity associations. MACsec is enabled when a connectivity association is assigned to an interface.

When you are configuring MACsec using static secure association key (SAK) security mode, you must configure secure channels within a connectivity association. The secure channels are responsible for transmitting and receiving data on the MACsec-enabled link, and also responsible for transmitting SAKs across the link to enable and maintain MACsec. A single secure channel is uni-directional—it can only be used to apply MACsec to inbound or outbound traffic. A typical connectivity association when MACsec is enabled using SAK security mode contains two secure channels—one secure channel for inbound traffic and another secure channel for outbound traffic.

When you enable MACsec using static CAK or dynamic security mode, you have to create and configure a connectivity association. Two secure channels—one secure channel for inbound traffic and another secure channel for outbound traffic—are automatically created. The automatically-created secure channels do not have any user-configurable parameters; all configuration is done in the connectivity association outside of the secure channels.

Understanding MACsec Security Modes

Understanding Static Connectivity Association Key Security Mode (Recommended Security Mode for Switch-to-Switch Links)

When you enable MACsec using static connectivity association key (CAK) security mode, two security keys—a connectivity association key (CAK) that secures control plane traffic and a randomly-generated secure association key (SAK) that secures data plane traffic—are used to secure the point-to-point Ethernet link. Both keys are regularly exchanged between both devices on each end of the point-to-point Ethernet link to ensure link security.

You initially establish a MACsec-secured link using a pre-shared key when you are using static CAK security mode to enable MACsec. A pre-shared key includes a connectivity association name (CKN) and its own connectivity association key (CAK). The CKN and CAK are configured by the user in the connectivity association and must match on both ends of the link to initially enable MACsec.

Once matching pre-shared keys are successfully exchanged, the MACsec Key Agreement (MKA) protocol is enabled. The MKA protocol is responsible for maintaining MACsec on the link, and decides which switch on the point-to-point link becomes the key server. The key server then creates an SAK that is shared with the switch at the other end of the point-to-point link only, and that SAK is used to secure all data traffic traversing the link. The key server will continue to periodically create and share a randomly-created SAK over the point-to-point link for as long as MACsec is enabled.

You enable MACsec using static CAK security mode by configuring a connectivity association on both ends of the link. All configuration is done within the connectivity association but outside of the secure channel. Two secure channels—one for inbound traffic and one for outbound traffic—are automatically created when using static CAK security mode. The automatically-created secure channels do not have any user-configurable parameters that cannot already be configured in the connectivity association.

We recommend enabling MACsec on switch-to-switch links using static CAK security mode. Static CAK security mode ensures security by frequently refreshing to a new random security key and by only sharing the security key between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available when you enable MACsec using static CAK security mode.

See [“Configuring Media Access Control Security \(MACsec\)” on page 147](#) for step-by-step instructions on enabling MACsec using static CAK security mode.

Understanding Dynamic Secure Association Key Security Mode (Switch-to-Host Links)

Dynamic secure association key security mode is used to enable MACsec on a switch-to-host link.

To enable MACsec on a link connecting an endpoint device—such as a server, phone, or personal computer—to a switch, the endpoint device must support MACsec and must be running software that allows it to enable a MACsec-secured connection. When configuring MACsec on a switch-to-host link, the MACsec Key Agreement (MKA) keys, which are included as part of 802.1X authentication, are retrieved from a RADIUS server as part of the AAA handshake. A master key is passed from the RADIUS server to the switch and from the RADIUS server to the host in independent authentication transactions. The master key is then passed between the switch and the host to create a MACsec-secured connection.

A secure association using dynamic secure association security mode must be configured on the switch's Ethernet interface that connects to the host in order for the switch to create a MACsec-secured connection after receiving the MKA keys from the RADIUS server.

The RADIUS server must be using Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) in order to support MACsec. The RADIUS servers that support other widely-used authentication frameworks, such as password-only or md5, cannot be used to support MACsec. In order to enable MACsec on a switch to secure a connection to a host, you must be using 802.1X authentication on the RADIUS server. MACsec must be configured into dynamic mode. MACsec is still enabled using connectivity associations when enabled on a switch-to-host link, as it is on a switch-to-switch link.

Understanding Static Secure Association Key Security Mode (Supported for Switch-to-Switch Links)

When you enable MACsec using static secure association key (SAK) security mode, one of up to two manually configured SAKs is used to secure data traffic on the point-to-point Ethernet link. All SAK names and values are configured by the user; there is no key server or other tool that creates SAKs. Security is maintained on the point-to-point Ethernet link by periodically rotating between the two security keys. Each security key name and value must have a corresponding matching value on the interface at the other end of the point-to-point Ethernet link to maintain MACsec on the link.

You configure SAKs within secure channels when you enable MACsec using static SAK security mode. You configure secure channels within connectivity associations. A typical connectivity association for MACsec using static SAK security mode contains two secure channels—one for inbound traffic and one for outbound traffic—that have each been configured with two manually-configured SAKs. You must attach the connectivity association with the secure channel configurations to an interface to enable MACsec using static SAK security mode.

We recommend enabling MACsec using static CAK security mode. You should only use static SAK security mode if you have a compelling reason to use it instead of static CAK security mode.

See “Configuring Media Access Control Security (MACsec)” on page 147 for step-by-step instructions on enabling MACsec using SAKs.

Understanding the Requirements to Enable MACsec on a Switch-to-Host Link

When configuring MACsec on a switch-to-host link, the MACsec Key Agreement (MKA) keys, which are included as part of 802.1X authentication, are retrieved from a RADIUS server as part of the AAA handshake. A master key is passed from the RADIUS server to the switch and from the RADIUS server to the host in independent authentication transactions. The master key is then passed between the switch and the host to create a MACsec-secured connection.

The following requirements must be met in order to enable MACsec on a link connecting a host device to a switch.

The host device:

- must support MACsec and must be running software that allows it to enable a MACsec-secured connection with the switch.

The switch:

- must be an EX4200, EX4300, or EX4550 switch running Junos OS Release 14.1X53-D10 or later, or an EX9200 switch running Junos OS Release 15.1R1 or later.
- must be configured into dynamic secure association key security mode.
- must be using 802.1X authentication to communicate with the RADIUS server.

The RADIUS server:

- must be using the Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication framework.



NOTE: RADIUS servers that support other widely-used authentication frameworks, such as password-only or md5, cannot be used to support MACsec.

- must be using 802.1X authentication.
- can be multiple hops from the switch and the host device.

MACsec Hardware and Software Support Summary

Table 17 on page 144 summarizes MACsec hardware and software support for EX Series and QFX Series switches.

MACsec hardware and software support is discussed in greater detail in the remaining sections.

Table 17: MACsec Hardware and Software Support Summary for EX Series and QFX Series Switches

Switch	MACsec-capable Interfaces	Switch-to-Switch Support Introduction	Switch-to-Host Support Introduction	Required Software Package
EX4200	All uplink port connections on the SFP+ MACsec uplink module.	13.2X50-D15	14.1X53-D10	controlled
EX4300	All access and uplink ports.	13.2X50-D15	14.1X53-D10	controlled
EX4550	All EX4550 optical interfaces that use the LC connection type.	13.2X50-D15	14.1X53-D10	controlled
EX4600	All twenty-four fixed 1GbE SFP/10GbE SFP+ interfaces and all interfaces that support the copper Gigabit Interface Converter (GBIC). All eight SFP+ interfaces on the EX4600-EM-8F expansion module.	14.1X53-D15	Not supported	controlled
EX9200	All forty SFP interfaces on the EX9200-40F-M.	15.1R1	15.1R1	Junos image <i>NOTE:</i> MACsec is available on the Junos OS image in EX9200 switches only. MACsec is not available on the limited Junos OS image package.
QFX5100	All eight SFP+ interfaces on the EX4600-EM-8F expansion module.	14.1X53-D15	Not supported	controlled

Understanding MACsec Hardware Requirements for EX Series and QFX Series Switches

MACsec is currently supported on the following EX Series and QFX Series switch interfaces:

- The uplink port connections on the SFP+ MACsec uplink module that can be installed on EX4200 series switches.
- All access and uplink ports on EX4300 switches.
- All EX4550 optical interfaces that use the LC connection type. See *Pluggable Transceivers Supported on EX4550 Switches*.

- All twenty-four fixed 1GbE SFP/10GbE SFP+ interfaces on an EX4600 switch and all interfaces that support the copper Gigabit Interface Converter (GBIC).
- All eight SFP+ interfaces on the EX4600-EM-8F expansion module, when installed in an EX4600 or QFX5100-24Q switch.



NOTE: MACsec is not supported on EX4600 or QFX5100-24Q switches in Junos OS Release 15.1.

See [Feature Explorer](#) for a full listing of Junos OS releases that support MACsec.

- All forty SFP interfaces on the EX9200-40F-M line card, when the line card is installed in an EX9200 series switch.

MACsec can be configured on supported switch interfaces when those switches are configured in a Virtual Chassis or Virtual Chassis Fabric (VCF), including when MACsec-supported interfaces are on member switches in a mixed Virtual Chassis or VCF that includes switch interfaces that do not support MACsec. MACsec, however, cannot be enabled on Virtual Chassis ports (VCPs) to secure traffic travelling between member switches in a Virtual Chassis or VCF.

Understanding MACsec Software Requirements for EX Series and QFX Series Switches

See [Feature Explorer](#) for a full listing of Junos OS releases and platforms that support MACsec.

MACsec was initially released on EX4200, EX4300, and EX4550 switches in Junos OS Release 13.2X50-D15.

MACsec support for dynamic security mode, which allows MACsec to be configured on switch-to-host links, for EX4200, EX4300, and EX4550 switches was introduced in Junos OS Release 14.1X53-D10.

The switches on each end of a MACsec-secured switch-to-switch link must either both be using Junos OS Release 14.1X53-D10 or later, or must both be using an earlier version of Junos, in order to establish a MACsec-secured connection when using static CAK security mode.

MACsec support for EX4600 switches and QFX5100-24Q switches was introduced in Junos OS Release 14.1X53-D15. The EX4600 and QFX5100-24Q switches supports MACsec on switch-to-switch links only.



NOTE: MACsec is not supported on EX4600 or QFX5100-24Q switches in Junos OS Release 15.1.

See [Feature Explorer](#) for a full listing of Junos OS releases and platforms that support MACsec.

MACsec support for EX9200 switches for both switch-to-switch links and for switch-to-host links was introduced in Junos OS Release 15.1R1.

You must download the controlled version of your Junos OS software to enable MACsec on EX4200, EX4300, EX4550, EX4600, or QFX5100-24Q switches. MACsec software support is not available in the domestic version of your Junos OS software on these platforms.

You must download the standard Junos image to enable MACsec on EX9200 switches. MACsec is not supported on the limited image.

The controlled version of Junos OS software for EX4200, EX4300, EX4550, EX4600, or QFX5100-24Q switches includes all features and functionality available in the domestic version of Junos OS, while also supporting MACsec. The domestic version of Junos OS software is shipped on all switches that support MACsec, so you must download and install a controlled version of Junos OS software for your switch before you can enable MACsec.

The controlled version of Junos OS software for EX4200, EX4300, EX4550, EX4600, or QFX5100-24Q switches contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of the controlled version of Junos OS software is strictly controlled under United States export laws. The export, import, and use of the controlled version of Junos OS software is also subject to controls imposed under the laws of other countries. If you have questions about acquiring the controlled version of your Junos OS software, contact Juniper Networks Trade Compliance group at compliance_helpdesk@juniper.net.

The standard version of Junos OS software for EX9200 switches contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of this Junos OS software is strictly controlled under United States export laws. The export, import, and use of this Junos OS software is also subject to controls imposed under the laws of other countries. If you have questions about acquiring this version of your Junos OS software, contact Juniper Networks Trade Compliance group at compliance_helpdesk@juniper.net.

Understanding the MACsec Feature License Requirement

A feature license is required to configure MACsec on a switch.

To purchase a feature license for MACsec, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with a feature license file and a license key. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show chassis hardware** command.

The MACsec feature license is an independent feature license; the feature licenses that must be purchased to enable other groups of features on your switches cannot be purchased to enable MACsec.

MACsec Limitations

All types of Spanning Tree Protocol frames cannot currently be encrypted using MACsec.

- Related Documentation**
- [Configuring Media Access Control Security \(MACsec\) on page 147](#)

Configuring Media Access Control Security (MACsec)

Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication for almost all types of traffic on Ethernet links. MACsec provides point-to-point security on Ethernet links between directly-connected nodes and is capable of identifying and preventing most security threats, including denial of service, intrusion, man-in-the-middle, masquerading, passive wiretapping, and playback attacks. MACsec is standardized in IEEE 802.1AE.

You can configure MACsec to secure point-to-point Ethernet links connecting EX Series or QFX Series switches, or on Ethernet links connecting a switch to a host device such as a PC, phone, or server. Each point-to-point Ethernet link that you want to secure using MACsec must be configured independently. You can enable MACsec on switch-to-switch links using static secure association key (SAK) security mode or static connectivity association key (CAK) security mode. Both processes are provided in this document.



BEST PRACTICE: We recommend enabling MACsec using static CAK security mode on switch-to-switch links. Static CAK security mode ensures security by frequently refreshing to a new random secure association key (SAK) and by only sharing the SAK between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available for MACsec-secured switch-to-switch connections that are enabled using static CAK security mode.

The configuration steps for both processes are provided in this document.

- [Acquiring and Downloading the Junos OS Software on page 147](#)
- [Acquiring and Downloading the MACsec Feature License on page 149](#)
- [Configuring the PIC Mode of the MACsec-capable Interfaces \(EX4200 switches only\) on page 149](#)
- [Configuring MACsec Using Static Connectivity Association Key Security Mode \(Recommended for Enabling MACsec on Switch-to-Switch Links\) on page 151](#)
- [Configuring MACsec on the Switch Using Dynamic Secure Association Key Security Mode to Secure a Switch-to-Host Link on page 155](#)
- [Configuring MACsec Using Static Secure Association Key Security Mode to Secure a Switch-to-Switch Link on page 159](#)

Acquiring and Downloading the Junos OS Software

MACsec was initially released on EX4200, EX4300, and EX4550 switches in Junos OS Release 13.2X50-D15. MACsec was released on EX4600 and QFX5100-24Q switches in Junos OS Release 14.1X53-D15, and on EX9200 series switches in Junos OS Release 15.1R1. The switches on each end of a MACsec-secured switch-to-switch link must either both

be using Junos OS Release 14.1X51-D10 or later, or must both be using an earlier version of Junos, in order to establish a MACsec-secured connection when using static CAK security mode.

You must download the controlled version of your Junos OS software to enable MACsec on EX4200, EX4300, EX4550, EX4600, and QFX5100-24Q switches.

You must download the standard version of your Junos OS software to enable MACsec on EX9200 switches. MACsec is not supported in the limited version of Junos OS on EX9200 switches.

See [“Understanding Media Access Control Security \(MACsec\)” on page 139](#) for additional information on the versions of Junos OS software that are required for MACsec.

You can identify whether a software package is the controlled or standard version of Junos OS by viewing the package name. A software package for a controlled version of Junos OS is named using the following format:

package-name-m.nZx.y-controlled-signed.tgz

A software package for a standard version of Junos OS on an EX9200 switch is named using the following format:

package-name-m.nZx.y-.tgz

If you are unsure which version of Junos OS is running on your switch, enter the **show version** command. If the “JUNOS Crypto Software Suite” description appears in the output, you are running the controlled version of Junos OS.

The controlled version of Junos OS software for EX4200, EX4300, EX4550, EX4600, or QFX5100-24Q switches contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of the controlled version of Junos OS software is strictly controlled under United States export laws. The export, import, and use of the controlled version of Junos OS software is also subject to controls imposed under the laws of other countries. If you have questions about acquiring the controlled version of your Junos OS software, contact Juniper Networks Trade Compliance group at compliance_helpdesk@juniper.net.

The standard version of Junos OS software for EX9200 switches contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of this Junos OS software is strictly controlled under United States export laws. The export, import, and use of this Junos OS software is also subject to controls imposed under the laws of other countries. If you have questions about acquiring this version of your Junos OS software, contact Juniper Networks Trade Compliance group at compliance_helpdesk@juniper.net.

The process for installing the controlled or standard version of Junos OS software onto your switch is identical to installing any other version of Junos OS software. You must enter the **request system software add** statement to download the Junos OS image, and the **request system reboot** statement to reboot the switch to complete the upgrade procedure. See *Downloading Software Packages from Juniper Networks, Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*, and *Installing Software*

on an EX Series Switch with Redundant Routing Engines (CLI Procedure) for detailed information about acquiring and installing Junos OS software images for your switches.

Acquiring and Downloading the MACsec Feature License

A feature license is required to configure MACsec on an EX Series or a QFX Series switch.

The MACsec feature license is an independent feature license; the enhanced feature licenses (EFLs) or advanced feature licenses (AFLs) that must be purchased to enable some features on EX Series or QFX Series switches cannot be purchased to enable MACsec.

To purchase a software license for MACsec, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with a feature license file and a license key. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show chassis hardware** command.

For a Virtual Chassis deployment, two MACsec license keys are recommended for redundancy—one for the device in the master role and the other for the device in the backup role.

To add one or more new MACsec license keys on the switch, follow this procedure:

1. Add the license key or keys:
 - To add one or more license keys from a file or URL, specify the filename of the file or the URL where the key is located:


```
user@switch> request system license add filename |url
```
 - To add a license key from the terminal:


```
user@switch> request system license add terminal
```
2. When prompted, enter the license key, separating multiple license keys with a blank line.

If the license key you enter is invalid, an error appears in the CLI output when you press Ctrl+d to exit the license entry mode.

A MACsec feature license is installed and maintained like any other switch license. See *Managing Licenses for the EX Series Switch (CLI Procedure)* or *Adding New Licenses (CLI Procedure)* for more detailed information on configuring and managing your MACsec software license.

Configuring the PIC Mode of the MACsec-capable Interfaces (EX4200 switches only)

To configure MACsec on an EX4200 switch, you must install the SFP+ MACsec uplink module. The interfaces on the SFP+ MACsec uplink module are the only MACsec-capable interfaces available for EX4200 switches. All four ports on the uplink module are MACsec-capable.

The SFP+ MACsec uplink module provides two ports for 10-gigabit small form-factor pluggable (SFP+) transceivers when configured to operate in 10-gigabit mode or four

ports for 1-gigabit small form-factor pluggable (SFP) transceivers when configured to operate in 1-gigabit mode.

The PIC mode is set to **10g**, by default. You only need to perform this procedure if you want to operate your uplink in 1-gigabit mode, or if you previously set the uplink module to 1-gigabit mode and would like to return it to 10-gigabit mode.

To configure the PIC mode:

```
[edit chassis]
```

```
user@switch# set fpc fpc-slot-number pic 1 sfpplus pic-mode (1g | 10g)
```

where *fpc-slot-number* is the FPC slot number, *pic-slot-number* is the PIC slot number, and the **[1g | 10g]** option configures the MACsec capability of the four SFP+ ports on the MACsec uplink module.

The *fpc-slot-number* is always 0 on standalone EX4200 switches, and is the member ID of the member switch in an EX4200 Virtual Chassis.

The PIC slot number is always 1 for the uplink module port slot on an EX4200 switch, so **pic 1** is always the specified PIC slot number.

The PIC mode is set to **10g** by default. When the PIC mode is set to **10g**, uplink ports 0 and 2 on the MACsec uplink module support MACsec at 10-Gbps speeds. Ports 1 and 3 cannot be used to send any traffic.

When the PIC mode is set to **1g**, all four SFP+ ports on the MACsec uplink module support MACsec at 1-Gbps speeds.

Configuring MACsec Using Static Connectivity Association Key Security Mode (Recommended for Enabling MACsec on Switch-to-Switch Links)

You can enable MACsec using static connectivity association key (CAK) security mode or static secure association keys (SAK) security mode on a point-to-point Ethernet link connecting switches. This procedure shows you how to configure MACsec using static CAK security mode.



BEST PRACTICE: We recommend enabling MACsec using static CAK security mode on switch-to-switch links. Static CAK security mode ensures security by frequently refreshing to a new random secure association key (SAK) and by only sharing the SAK between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available for MACsec-secured switch-to-switch connections that are enabled using static CAK security mode.

When you enable MACsec using static CAK security mode, a pre-shared key is exchanged between the switches on each end of the point-to-point Ethernet link. The pre-shared key includes a connectivity association name (CKN) and a connectivity association key (CAK). The CKN and CAK are configured by the user in the connectivity association and must match on both ends of the link to initially enable MACsec.

After the pre-shared keys are exchanged and verified, the MACsec Key Agreement (MKA) protocol, which enables and maintains MACsec on the link, is enabled. The MKA is responsible for selecting one of the two switches on the point-to-point link as the key server. The key server then creates a randomized security key that is shared only with the other device over the MACsec-secured link. The randomized security key enables and maintains MACsec on the point-to-point link. The key server will continue to periodically create and share a randomly-created security key over the point-to-point link for as long as MACsec is enabled.

You enable MACsec using static CAK security mode by configuring a connectivity association on both ends of the link. All configuration is done within the connectivity association but outside of the secure channel. Two secure channels—one for inbound traffic and one for outbound traffic—are automatically created when using static CAK security mode. The automatically-created secure channels do not have any user-configurable parameters that cannot already be configured in the connectivity association.

To configure MACsec using static CAK security mode to secure a switch-to-switch Ethernet link:

1. Create a connectivity association. You can skip this step if you are configuring an existing connectivity association.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name
```

For instance, to create a connectivity association named **ca1**, enter:

```
[edit security macsec]
user@switch# set connectivity-association ca1
```

2. Configure the MACsec security mode as **static-cak** for the connectivity association:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name security-mode
static-cak
```

For instance, to configure the MACsec security mode to **static-cak** on connectivity association **ca1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 security-mode static-cak
```

3. Create the pre-shared key by configuring the connectivity association key name (CKN) and connectivity association key (CAK):

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name pre-shared-key
ckn hexadecimal-number
user@switch# set connectivity-association connectivity-association-name pre-shared-key
cak hexadecimal-number
```

A pre-shared key is exchanged between directly-connected links to establish a MACsec-secure link. The pre-shared-key includes the CKN and the CAK. The CKN is a 64-digit hexadecimal number and the CAK is a 32-digit hexadecimal number. The CKN and the CAK must match on both ends of a link to create a MACsec-secured link.



NOTE: To maximize security, we recommend configuring all 64 digits of a CKN and all 32 digits of a CAK.

If you do not configure all 64 digits of a CKN or all 32 digits of a CAK, however, all remaining digits will be auto-configured to 0.

After the pre-shared keys are successfully exchanged and verified by both ends of the link, the MACsec Key Agreement (MKA) protocol is enabled and manages the secure link. The MKA protocol then elects one of the two directly-connected switches as the key server. The key server then shares a random security with the other device over the MACsec-secure point-to-point link. The key server will continue to periodically create and share a random security key with the other device over the MACsec-secured point-to-point link as long as MACsec is enabled.

To configure a CKN of **37c9c2c45ddd012aa5bc8ef284aa23ff6729ee2e4acb66e91fe34ba2cd9fe311** and CAK of **228ef255aa23ff6729ee664acb66e91f** on connectivity association **ca1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 pre-shared-key ckn
37c9c2c45ddd012aa5bc8ef284aa23ff6729ee2e4acb66e91fe34ba2cd9fe311
user@switch# set connectivity-association ca1 pre-shared-key cak
228ef255aa23ff6729ee664acb66e91f
```



NOTE: MACsec is not enabled until a connectivity association is attached to an interface. See the final step of this procedure to attach a connectivity association to an interface.

4. (Required on switches when connecting to EX4300 switches only) Enable SCI tagging:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set include-sci
```

You must enable SCI tagging on a switch that is enabling MACsec on an Ethernet link connecting to an EX4300 switch.

SCI tags are automatically appended to packets leaving a MACsec-enabled interface on an EX4300 switch. This option is, therefore, not available on EX4300 switches.

You should only use this option when enabling MACsec on a link to an EX4300 switch. SCI tags are eight octets long, so appending an SCI tag to all traffic on the link adds a significant amount of unneeded overhead.

5. (Optional) Set the MKA key server priority:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set mka key-server-priority priority-number
```

Specifies the key server priority used by the MKA protocol to select the key server. The switch with the lower *priority-number* is selected as the key server.

The default *priority-number* is 16.

If the **key-server-priority** is identical on both sides of the point-to-point link, the MKA protocol selects the interface with the lower MAC address as the key server. Therefore, if this statement is not configured in the connectivity associations at each end of a MACsec-secured point-to-point link, the interface with the lower MAC address becomes the key server.

To change the key server priority to 0 to increase the likelihood that the current device is selected as the key server when MACsec is enabled on the interface using connectivity association **ca1**:

```
[edit security macsec connectivity-association ca1]
user@switch# set mka key-server-priority 0
```

To change the key server priority to 255 to decrease the likelihood that the current device is selected as the key server in connectivity association **ca1**:

```
[edit security macsec connectivity-association ca1]
user@switch# set mka key-server-priority 255
```

6. (Optional) Set the MKA transmit interval:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set mka transmit-interval interval
```

The MKA transmit interval setting sets the frequency for how often the MKA protocol data unit (PDU) is sent to the directly connected device to maintain MACsec connectivity on the link. A lower *interval* increases bandwidth overhead on the link; a higher *interval* optimizes MKA protocol communication.

The default *interval* is 2000ms. We recommend increasing the interval to 6000 ms in high-traffic load environments. The transmit interval settings must be identical on both ends of the link when MACsec using static CAK security mode is enabled.

For instance, if you wanted to increase the MKA transmit interval to 6000 milliseconds when connectivity association *ca1* is attached to an interface:

```
[edit security macsec connectivity-association ca1]
user@switch# set mka transmit-interval 6000
```

7. (Optional) Disable MACsec encryption:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set no-encryption
```

Encryption is enabled for all traffic entering or leaving the interface when MACsec is enabled using static CAK security mode, by default.

When encryption is disabled, traffic is forwarded across the Ethernet link in clear text. You are able to view unencrypted data in the Ethernet frame traversing the link when you are monitoring it. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic sent or received on the link has not been tampered with and does not represent a security threat.

8. (Optional) Set an offset for all packets traversing the link:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set offset (0 | 30 | 50)
```

For instance, if you wanted to set the offset to 30 in the connectivity association named *ca1*:

```
[edit security macsec connectivity-association ca1]
user@switch# set offset 30
```

The default offset is 0. All traffic in the connectivity association is encrypted when encryption is enabled and an **offset** is not set.

When the offset is set to 30, the IPv4 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic. When the offset is set to 50, the IPv6 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

9. (Optional) Enable replay protection.

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set replay-protect replay-window-size number-of-packets
```

When MACsec is enabled on a link, an ID number is assigned to each packet on the MACsec-secured link.

When replay protection is enabled, the receiving interface checks the ID number of all packets that have traversed the MACsec-secured link. If a packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance,

if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

You can require that all packets arrive in order by setting the replay window size to 0.

To enable replay protection with a window size of five on connectivity association **ca1**:

```
[edit security macsec connectivity-association ca1]
user@switch# set replay-protect replay-window-size 5
```

10. (Optional) Exclude a protocol from MACsec:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set exclude-protocol protocol-name
```

For instance, if you did not want Link Level Discovery Protocol (LLDP) to be secured using MACsec:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set exclude-protocol lldp
```

When this option is enabled, MACsec is disabled for all packets of the specified protocol—in this case, LLDP—that are sent or received on the link.

11. Assign the connectivity association to an interface:

```
[edit security macsec]
user@switch# set interfaces interface-names connectivity-association
connectivity-association-name
```

Assigning the connectivity association to an interface is the final configuration step to enabling MACsec on an interface.

For instance, to assign connectivity association **ca1** to interface **xe-0/0/1**:

```
[edit security macsec]
user@switch# set interfaces xe-0/1/0 connectivity-association ca1
```

MACsec using static CAK security mode is not enabled until a connectivity association on the opposite end of the link is also configured, and contains pre-shared keys that match on both ends of the link.

Configuring MACsec on the Switch Using Dynamic Secure Association Key Security Mode to Secure a Switch-to-Host Link

Before you begin to enable MACsec on a switch-to-host link:

- Confirm that MACsec on switch-to-host links is supported on your switch. See [“Understanding Media Access Control Security \(MACsec\)” on page 139](#).
- Configure a RADIUS server. The RADIUS server:
 - must be configured as the user database for 802.1X authentication.

- must be using the Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication framework.
- must have connectivity to the switch and to the host. The RADIUS server can be multiple hops from the switch or the host.

See Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch.

- Enable MACsec on the host device.

The procedures for enabling MACsec on the host device varies by host device, and is beyond the scope of this document.

To configure MACsec using dynamic security mode to secure a switch-to-host Ethernet link:

1. Create a connectivity association. You can skip this step if you are configuring an existing connectivity association.

```
[edit security macsec]
```

```
user@switch# set connectivity-association connectivity-association-name
```

For instance, to create a connectivity association named `ca-dynamic1`, enter:

```
[edit security macsec]
```

```
user@switch# set connectivity-association ca-dynamic1
```

2. Configure the MACsec security mode as dynamic for the connectivity association:

```
[edit security macsec]
```

```
user@switch# set connectivity-association connectivity-association-name security-mode dynamic
```

For instance, to configure the MACsec security mode to dynamic on connectivity association `ca-dynamic1`:

```
[edit security macsec]
```

```
user@switch# set connectivity-association ca-dynamic1 security-mode dynamic
```

3. (Optional) Configure the **must-secure** option:

```
[edit security macsec]
```

```
user@switch# set connectivity-association connectivity-association-name mka must-secure
```

When the **must-secure** option is enabled, all traffic that is not MACsec-secured that is received on the interface is dropped.

When the **must-secure** option is disabled, all traffic from devices that support MACsec is MACsec-secured while traffic received from devices that do not support MACsec is forwarded through the network.

The **must-secure** option is particularly useful in scenarios where multiple devices, such as a phone and a PC, are accessing the network through the same Ethernet interface. If one of the devices supports MACsec while the other device does not support MACsec, the device that doesn't support MACsec can continue to send and receive traffic over the network—provided the **must-secure** option is disabled—while traffic to and from the device that supports MACsec is MACsec-secured. In this scenario, traffic to the device that is not MACsec-secured must be VLAN-tagged.

The **must-secure** option is disabled, by default.

4. (Required only if the host device requires SCI tagging) Enable SCI tagging:


```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set include-sci
```

You should only use this option when connecting a switch to a host that requires SCI tags. SCI tags are eight octets long, so appending an SCI tag to all traffic on the link adds a significant amount of unneeded overhead.

5. (Optional) Set the MKA key server priority:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set mka key-server-priority priority-number
```

Specifies the key server priority used by the MKA protocol to select the key server. The switch with the lower *priority-number* is selected as the key server.

The default *priority-number* is 16. If the **key-server-priority** is identical on both sides of the point-to-point link, the MKA protocol selects the interface with the lower MAC address as the key server. Therefore, if this statement is not configured in the connectivity associations at each end of a MACsec-secured point-to-point link, the interface with the lower MAC address becomes the key server.

To change the key server priority to 0 to increase the likelihood that the current device is selected as the key server when MACsec is enabled on the interface using connectivity association *ca1*:

```
[edit security macsec connectivity-association ca-dynamic1]
user@switch# set mka key-server-priority 0
```

To change the key server priority to 255 to decrease the likelihood that the current device is selected as the key server in connectivity association *ca-dynamic1*:

```
[edit security macsec connectivity-association ca-dynamic1]
user@switch# set mka key-server-priority 255
```

6. (Optional) Set the MKA transmit interval:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set mka transmit-interval interval
```

The MKA transmit interval setting sets the frequency for how often the MKA protocol data unit (PDU) is sent to the directly connected device to maintain MACsec connectivity on the link. A lower interval increases bandwidth overhead on the link; a higher interval optimizes MKA protocol communication.

The default interval is 2000ms. We recommend increasing the interval to 6000 ms in high-traffic load environments. The transmit interval settings must be identical on both ends of the link.

For instance, if you wanted to increase the MKA transmit interval to 6000 milliseconds when connectivity association *ca-dynamic1* is attached to an interface:

```
[edit security macsec connectivity-association ca-dynamic1]
user@switch# set mka transmit-interval 6000
```

7. (Optional) Disable MACsec encryption:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set no-encryption
```

Encryption is enabled for all traffic entering or leaving the interface when MACsec is enabled using dynamic security mode, by default. When encryption is disabled, traffic is forwarded across the Ethernet link in clear text. You are able to view unencrypted data in the Ethernet frame traversing the link when you are monitoring it. The MACsec

header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic sent or received on the link has not been tampered with and does not represent a security threat.

8. (Optional) Set an offset for all packets traversing the link:

```
[edit security macsec connectivity-association connectivity-association-name]  
user@switch# set offset (0 | 30 | 50)
```

For instance, if you wanted to set the offset to 30 in the connectivity association named ca-dynamic1:

```
[edit security macsec connectivity-association ca-dynamic1]  
user@switch# set offset 30
```

The default offset is 0. All traffic in the connectivity association is encrypted when encryption is enabled and an offset is not set.

When the offset is set to 30, the IPv4 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic. When the offset is set to 50, the IPv6 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

9. (Optional) Enable replay protection.

```
[edit security macsec connectivity-association connectivity-association-name]  
user@switch# set replay-protect replay-window-size number-of-packets
```

When MACsec is enabled on a link, an ID number is assigned to each packet on the MACsec-secured link. When replay protection is enabled, the receiving interface checks the ID number of all packets that have traversed the MACsec-secured link. If a packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

You can require that all packets arrive in order by setting the replay window size to 0.

To enable replay protection with a window size of five on connectivity association ca-dynamic1:

```
[edit security macsec connectivity-association ca-dynamic1]  
user@switch# set replay-protect replay-window-size 5
```

10. (Optional) Exclude a protocol from MACsec:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set exclude-protocol protocol-name
```

For instance, if you did not want Link Level Discovery Protocol (LLDP) to be secured using MACsec:

```
[edit security macsec connectivity-association ca-dynamic1]
user@switch# set exclude-protocol lldp
```

When this option is enabled, MACsec is disabled for all packets of the specified protocol—in this case, LLDP—that are sent or received on the link.

11. Assign the connectivity association to an interface:

```
[edit security macsec]
user@switch# set interfaces interface-names connectivity-association
connectivity-association-name
```

Assigning the connectivity association to an interface is the final configuration step to enabling MACsec on an interface. For instance, to assign connectivity association `ca-dynamic1` to interface `xe-0/0/1`:

```
[edit security macsec]
user@switch# set interfaces xe-0/1/0 connectivity-association ca-dynamic1
```

Configuring MACsec Using Static Secure Association Key Security Mode to Secure a Switch-to-Switch Link

When you enable MACsec using static secure association key (SAK) security mode, one of up to two manually configured security keys is used to secure the point-to-point Ethernet link between the switches. All security key names and values are configured by the user; there is no key server or other tool that creates security keys. Security is maintained on the point-to-point Ethernet link by periodically rotating the security keys. Each security key name and value must have a corresponding matching value on the interface at the other end of the point-to-point Ethernet link to maintain MACsec on the link.

You configure static SAKs within secure channels when you are enabling MACsec using static SAK security mode. You configure secure channels within connectivity associations. A typical connectivity association for MACsec using static SAK security mode contains two secure channels—one for inbound traffic and one for outbound traffic—that have each been configured with two static SAKs. You must attach the connectivity association with the secure channel configurations to an interface to enable MACsec using static SAK security mode.

To configure MACsec on a switch-to-switch Ethernet link using static SAK security mode:

1. Create a connectivity association. You can skip this step if you are configuring an existing connectivity association.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name
```

For instance, to create a connectivity association named `ca1`, enter:

```
[edit security macsec]
user@switch# set connectivity-association ca1
```

2. Configure the MACsec security mode as **static-sak** for the connectivity association:

```
[edit security macsec]
```

```
user@switch# set connectivity-association connectivity-association-name security-mode static-sak
```

For instance, to configure the MACsec security mode to **static-sak** on connectivity association **ca1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 security-mode static-sak
```

3. Create a secure channel within the connectivity association. You can skip this step if you are configuring an existing secure channel.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel secure-channel-name
```

For instance, to create secure channel **sc1** in connectivity association **ca1**, enter:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1
```

4. Define the security associations and the static SAKs for the secure channel:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel secure-channel-name security-association number key key-string
```

where the **security-association number** is a number between 0 and 3, and the **key-string** is a 32-digit key defined statically by the network administrator.

The key string is a 32-digit hexadecimal number. The key string and the security association must match on both sides of an Ethernet connection to secure traffic using MACsec.

A secure channel must have at least two security associations with unique key strings. MACsec uses a security associations to establish a secure communications link, and periodically rotates to a new security association to keep the link secure. MACsec, therefore, must have at least one backup security association and key at all times.

To create one secure channel with two security associations and keys, for example:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 security-association 0 key d183c4002fa6fe3d2d9a852c20ab8412
user@switch# set connectivity-association ca1 secure-channel sc1 security-association 1 key b976c7494ab6fe2f2d4c432a90fd90a8
```

5. Specify whether the secure channel should be applied to traffic entering or leaving the switch:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel secure-channel-name direction [inbound | outbound]
```

where **inbound** applies the secure channel to traffic entering the switch, and **outbound** applies the secure channel to traffic leaving the switch.



NOTE: A secure channel can only be applied to traffic entering (inbound) or leaving (outbound) an interface on the switch.

If you need to configure MACsec using SAKs on inbound and outbound traffic on the same interface, you must configure a connectivity association with one secure channel for inbound traffic and a second secure channel for outbound traffic. The connectivity association is assigned to an interface later in this process.

For instance, to configure secure channel **sc1** to apply MACsec to incoming traffic:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 direction inbound
```

To configure secure channel **sc2** to apply MACsec to outgoing traffic:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc2 direction outbound
```

6. Specify a MAC address:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name id mac-address mac-address
```

If you are configuring a MAC address on a secure channel in the outbound direction, you should specify the MAC address of the interface as the **mac-address**.

If you are configuring a MAC address on a secure channel in the inbound direction, you should specify the MAC address of the interface at the other end of the link as the **mac-address**.

The **mac-address** variables must match on the sending and receiving secure channel on each side of a link to enable MACsec using static SAK security mode.



NOTE: You can see the MAC address of an interface in the **show interfaces** output.

To configure MACsec to accept frames from MAC address **12:34:56:ab:cd:ef** on secure channel **sc1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 id mac-address
12:34:56:ab:cd:ef
```

7. Specify a port:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name id port-id port-id-number
```

The **port-id-number** variables must match on a sending and receiving secure channel on each side of a link to enable MACsec.



NOTE: The only requirement for port numbers in this implementation of MACsec is that they match on the sending and receiving ends of an Ethernet link. When the port numbers match, MACsec is enabled for all traffic on the connection.

To specify port ID 4 on secure channel **sc1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 id port-id 4
```

8. (Optional) Enable encryption:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name encryption
```

You can enable MACsec without enabling encryption. If a secure channel is configured on an interface without encryption, traffic is forwarded across the Ethernet link in clear text, and you will be able to view unencrypted data in the Ethernet frame traversing the link when you are monitoring it. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic on the link does not represent a security threat.

Encryption is disabled by default when you are enabling MACsec using static SAK security mode. To ensure all traffic traversing secure-channel **sc1** is encrypted:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 encryption
```

9. (Optional) Set an offset to send the first 30 or 50 octets in unencrypted plain text when encryption is enabled.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name offset [0 | 30 | 50]
```

When the offset is set to 30, the IPv4 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic. When the offset is set to 50, the IPv6 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

The default offset is 0, so all traffic on the link is encrypted when the **encryption** option is enabled and an **offset** is not set.

To change the offset to 30 for secure channel **sc1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 offset 30
```

10. Assign the connectivity association to an interface:

```
[edit security macsec]
user@switch# set interfaces interface-names connectivity-association
connectivity-association-name
```

Assigning the connectivity association to an interface is the final configuration step to enabling MACsec on an interface.

For instance, to assign connectivity association ca1 to interface xe-0/0/1:

```
[edit security macsec]  
user@switch# set interfaces xe-0/1/0 connectivity-association ca1
```

MACsec using static SAK security mode is not enabled until a connectivity association on the opposite end of the link is also configured, and the configuration match on both ends of the link.

**Related
Documentation**

- [Understanding Media Access Control Security \(MACsec\) on page 139](#)

PART 6

Configuration Statements and Operational Commands

- [Configuration Statements for Firewall Filters on page 167](#)
- [Configuration Statements for Policers on page 177](#)
- [Configuration Statements for Port Security on page 197](#)
- [Configuration Statements for Port Security \(ELS CLI Only\) on page 211](#)
- [Configuration Statements for Device Security on page 233](#)
- [Configuration Statements for Device Security \(ELS CLI Only\) on page 241](#)
- [Configuration Statements for Media Access Control Security \(MACsec\) on page 249](#)
- [Operational Commands for Firewall Filters on page 277](#)
- [Operational Commands for Media Access Control Security \(MACsec\) on page 287](#)
- [Operational Commands for Port Security on page 299](#)

CHAPTER 6

Configuration Statements for Firewall Filters

- [family on page 168](#)
- [filter on page 169](#)
- [filter \(Layer 2 and Layer 3 Interfaces\) on page 170](#)
- [filter \(VLANs\) on page 171](#)
- [firewall on page 172](#)
- [from on page 173](#)
- [interface-specific on page 174](#)
- [term on page 175](#)
- [then \(Filters\) on page 176](#)

family

```
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 on page 13](#)
- [Configuring Firewall Filters on page 35](#)
- [Overview of Firewall Filters on page 3](#)

filter

Syntax	<pre> filter <i>filter-name</i> { <i>interface-specific</i>; term <i>term-name</i> { from { <i>match-conditions</i>; } then { <i>action</i>; <i>action-modifiers</i>; } } } </pre>
Hierarchy Level	[edit firewall family <i>family-name</i>]
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 firewall filters.
Options	<p><i>filter-name</i>—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>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions on page 13 • Configuring Firewall Filters on page 35 • Overview of Firewall Filters on page 3

filter (Layer 2 and Layer 3 Interfaces)

Syntax	<code>filter (input output) <i>filter-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Apply a firewall filter to traffic transiting a port or Layer 3 interface.
Default	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.
Options	<p><i>filter-name</i>—Name of a firewall filter defined at the [edit firewall family <i>family-name</i> filter] hierarchy level.</p> <p>input—Apply a firewall filter to traffic entering the port or Layer 3 interface.</p> <p>output—Apply a firewall filter to traffic exiting the port or Layer 3 interface.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Gigabit Ethernet Interfaces (CLI Procedure)• Configuring Firewall Filters on page 35• Overview of Firewall Filters on page 3

filter (VLANs)

Syntax	<code>filter (input output) <i>filter-name</i>;</code>
Hierarchy Level	<code>[edit vlans <i>vlan-name</i>],</code> <code>[edit vlans <i>vlan-name</i> forwarding-options]</code>
Release Information	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.
Description	Apply a firewall filter to traffic entering or exiting a VLAN.
Default	All incoming traffic is accepted unmodified to a VLAN, and all outgoing traffic is sent unmodified from a VLAN.
Options	<i>filter-name</i> —Name of a firewall filter defined at the <code>[edit firewall family <i>family-name</i> filter]</code> hierarchy level. input —Apply a firewall filter to VLAN ingress traffic. output —Apply a firewall filter to VLAN egress traffic.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Firewall Filters on page 35• Overview of Firewall Filters on page 3

firewall

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

Description Configure firewall filters and policers.

The remaining statements are explained separately.

Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions on page 13 • Configuring Firewall Filters on page 35 • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77 • Overview of Firewall Filters on page 3

from

Syntax	<pre>from { match-conditions; }</pre>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Match packet fields to values specified in a match condition. If the from statement is not included in a firewall filter configuration, all packets are considered to match and the actions and action modifiers in the then statement are implemented.
Options	match-conditions —Conditions that define the values or fields that the incoming or outgoing packets must contain for a match. You can specify one or more match conditions. If you specify more than one, they all must match for a match to occur and for the action in the then statement to be implemented.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions on page 13 • Configuring Firewall Filters on page 35 • Understanding Firewall Filter Match Conditions on page 9

interface-specific

Syntax	interface-specific;
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure separate counters for each interface to which a filter is applied.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Firewall Filter Match Conditions and Actions on page 13• Configuring Firewall Filters on page 35• Overview of Firewall Filters on page 3

term

Syntax	<pre>term <i>term-name</i> { from { <i>match-conditions</i>; } then { <i>action</i>; <i>action-modifiers</i>; } }</pre>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i>]
Release Information	Statement introduced in Junos OS Release 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	<p><i>term-name</i>—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions on page 13 • Configuring Firewall Filters on page 35 • Overview of Firewall Filters on page 3

then (Filters)

Syntax	<pre>then { action; action-modifiers; }</pre>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i>]
Release Information	Statement introduced in Junos OS Release 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	<p>action—Actions to accept, discard, or forward packets that match all conditions specified in a filter term.</p> <p>action-modifiers—Additional actions to analyze, classify, count, or police packets that match all conditions specified in a filter term.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Firewall Filter Match Conditions and Actions on page 13• Configuring Firewall Filters on page 35• Understanding Firewall Filter Match Conditions on page 9

CHAPTER 7

Configuration Statements for Policers

- [action on page 178](#)
- [bandwidth-limit on page 178](#)
- [burst-size-limit on page 179](#)
- [color-aware on page 180](#)
- [color-blind on page 181](#)
- [committed-burst-size on page 182](#)
- [committed-information-rate on page 183](#)
- [excess-burst-size on page 184](#)
- [filter-specific on page 185](#)
- [firewall on page 186](#)
- [if-exceeding on page 187](#)
- [loss-priority high then discard \(Three-Color Policer\) on page 188](#)
- [peak-burst-size on page 189](#)
- [peak-information-rate on page 190](#)
- [policer on page 191](#)
- [single-rate on page 192](#)
- [then \(Policers\) on page 193](#)
- [three-color-policer on page 194](#)
- [two-rate on page 195](#)

action

Syntax	<code>action { loss-priority high then discard; }</code>
Hierarchy Level	[edit <code>firewall three-color-policer name</code>]
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	Discard traffic on a logical interface using tricolor marking policing.
Options	The statements are explained separately.
Required Privilege Level	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.

bandwidth-limit

Syntax	<code>bandwidth-limit bps;</code>
Hierarchy Level	[edit <code>firewall policer policer-name if-exceeding</code>]
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 traffic rate in bits per second.
Options	<code>bps</code> —Traffic rate in bits per second. Specify <code>bps</code> as a decimal value or as a decimal number followed by one of the abbreviation <code>k</code> (1000), <code>m</code> (1,000,000), or <code>g</code> (1,000,000,000). Range: 32000 bps (32 Kbps) through 10,000,000,000 bps (10 Gbps)
Required Privilege Level	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77• Overview of Policers on page 59

burst-size-limit

Syntax	<code>burst-size-limit bytes;</code>
Hierarchy Level	[edit <code>firewall policer policer-name if-exceeding</code>]
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	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77• Overview of Policers on page 59

color-aware

Syntax	color-aware;
Hierarchy Level	[edit firewall three-color-policer <i>policer-name</i> single-rate], [edit firewall three-color-policer <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 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.
Default	If you omit the color-aware statement, the default behavior is color-aware mode.
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">• Overview of Policers on page 59• Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 66• Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 68• color-blind on page 181

color-blind

Syntax	color-blind;
Hierarchy Level	[edit firewall three-color-policer <i>policer-name</i> single-rate], [edit firewall three-color-policer <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 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.
Default	If you omit the color-blind statement, the default behavior is color-aware mode.
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"> • Overview of Policers on page 59 • Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 65 • Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 68 • Configuring Color-Blind Egress Policers for Medium-Low PLP on page 76 • color-aware on page 180

committed-burst-size


Syntax	<code>committed-burst-size bytes;</code>
Hierarchy Level	[edit <code>firewall three-color-policer policer-name</code> single-rate], [edit <code>firewall three-color-policer policer-name</code> 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).




NOTE: When you include the `committed-burst-size` statement in the configuration, you must also include the `committed-information-rate` statement at the same hierarchy level.

Options	bytes —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 k (1000), m (1,000,000), or g (1,000,000,000). Range: 512 bytes through 268435456 bytes (268 MB)
Required Privilege Level	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77• Overview of Policers on page 59

committed-information-rate

Syntax	<code>committed-information-rate <i>bits-per-second</i>;</code>
Hierarchy Level	[edit <code>firewall three-color-policer <i>policer-name</i></code> single-rate], [edit <code>firewall three-color-policer <i>policer-name</i></code> 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 guaranteed bandwidth under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green).
<div>  <p>NOTE: 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>	
Options	<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 k (1000), m (1,000,000), or g (1,000,000,000).</p> <p>Range: 32,000 bps through 10,000,000,000 bps (10 gbps)</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"> • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77 • Overview of Policers on page 59

excess-burst-size

Syntax	<code>excess-burst-size bytes;</code>
Hierarchy Level	[edit <code>firewall three-color-policer policer-name</code> 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> NOTE: 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	bytes —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 k (1000), m (1,000,000), or g (1,000,000,000). Range: 512 bytes through 268435456 bytes (268 MB)
Required Privilege Level	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77• Overview of Policers on page 59

filter-specific

Syntax	filter-specific;
Hierarchy Level	[edit firewall policer <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 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"> • 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. • The implicit counter counts all the packets are that 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.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration
Related Documentation	<ul style="list-style-type: none"> • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77 • Overview of Policers on page 59

firewall

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

Description Configure firewall filters and policers.

The remaining statements are explained separately.

Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Firewall Filter Match Conditions and Actions on page 13 • Configuring Firewall Filters on page 35 • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77 • Overview of Firewall Filters on page 3


if-exceeding

Syntax	<pre>if-exceeding { bandwidth-limit <i>bps</i>; burst-size-limit <i>bytes</i>; }</pre>
Hierarchy Level	[edit firewall policer <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	Configure policer rate limits. The remaining statements are explained separately.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77 • Overview of Policers on page 59


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

Syntax	loss-priority high then discard;
Hierarchy Level	[edit firewall three-color-policer <i>policer-name</i> action]
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>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>
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">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77• Overview of Policers on page 59

peak-burst-size

Syntax	<code>peak-burst-size bytes;</code>
Hierarchy Level	[edit <code>firewall three-color-policer policer-name two-rate</code>]
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 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>  NOTE: 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. </div>	
Options	bytes —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 k (1000), m (1,000,000), or g (1,000,000,000). Range: 1500 bytes through 100,000,000,000 bytes (100 GB)
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"> • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77 • Overview of Policers on page 59

peak-information-rate

Syntax	<code>peak-information-rate <i>bits-per-second</i>;</code>
Hierarchy Level	[edit <code>firewall three-color-policer <i>policer-name</i> two-rate</code>]
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> NOTE: 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 k (1000), m (1,000,000), or g (1,000,000,000). Range: 32,000 bps through 10,000,000,000 bps (10 gbps)
Required Privilege Level	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77• Overview of Policers on page 59


policer

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 firewall]
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 policer action modifier in the then 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"> • Configure a unique policer for each term. • Configure only one policer, but use a unique, explicit counter in each term.
Options	<p><i>policer-name</i>—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"> • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77 • Configuring Firewall Filters on page 35 • Overview of Policers on page 59

single-rate

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 firewall three-color-policer <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	<i>policer-name</i> —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">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77• Overview of Policers on page 59

then (Policers)

Syntax	then { <i>policer-action</i> ; }
Hierarchy Level	[edit firewall policer <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	Configure a policer action.
Options	<i>policer-action</i> —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.
<div>  NOTE: If you specify a policer in an egress firewall filter, the only supported action is discard. </div>	
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"> • Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77 • Configuring Firewall Filters on page 35 • Overview of Policers on page 59

three-color-policer

Syntax	<pre>three-color-policer <i>policer-name</i> { action { loss-priority high then discard; } single-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; excess-burst-size <i>bytes</i>; } two-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; peak-information-rate <i>bps</i>; peak-burst-size <i>bytes</i>; } }</pre>
Hierarchy Level	[edit firewall], [edit logical-systems <i>logical-system-name</i> 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	<p><i>policer-name</i>—Name of the three-color policer. Use this name when you apply the policer to an interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 77• Overview of Policers on page 59

two-rate


Syntax	<pre>two-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; peak-information-rate <i>bps</i>; peak-burst-size <i>bytes</i>; }</pre>
Hierarchy Level	[edit firewall three-color-policer <i>policer-name</i>]
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 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).</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 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).</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>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>

CHAPTER 8

Configuration Statements for Port Security

- `circuit-id` on page 198
- `dhcp-snooping-file` on page 199
- `fc-map` on page 200
- `fcoe-trusted` on page 202
- `mac-move-limit` on page 203
- `no-allowed-mac-log` on page 204
- `no-gratuitous-arp-request` on page 205
- `persistent-learning` on page 205
- `port-error-disable` on page 206
- `vendor-id` on page 208
- `write-interval` on page 209

circuit-id

Syntax	<pre> circuit-id { prefix { host-name; logical-system-name; routing-instance-name; } use-interface-description (device logical); use-vlan-id; } </pre>
Hierarchy Level	<ul style="list-style-type: none"> For platforms with Enhanced Layer 2 Software (ELS): [edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82] For platforms without ELS: [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82], [edit forwarding-options helpers bootp dhcp-option82] , [edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82] For MX Series platforms: [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82]
Release Information	<p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 14.1 for the MX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure the circuit-id suboption (suboption 1) of DHCP option 82 (the DHCP relay agent information option) in DHCP packets destined for a DHCP server. This suboption identifies the circuit (the interface, the VLAN, or both) on which the DHCP request arrived.</p> <p>The remaining statements are explained separately.</p>
Default	<p>If DHCP option 82 is enabled on the switch, the circuit ID is supplied by default in the format <i>interface-name:vlan-name</i> or, on a Layer 3 interface, just <i>interface-name</i>.</p>
<div>  <p>NOTE: When you configure circuit-id, remote-id is also enabled, even if you do not explicitly configure remote-id .</p> </div>	
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

Related Documentation	<ul style="list-style-type: none"> • <i>Configuring DHCP Option 82 to Help Protect the Switching Devices Against Attacks (CLI Procedure)</i> • <i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i> • <i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i> • <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i> • <i>Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)</i> • <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i> • RFC 3046, DHCP Relay Agent Information Option, at http://tools.ietf.org/html/rfc3046
------------------------------	--

dhcp-snooping-file

Syntax	<pre>dhcp-snooping-file { location <i>local_pathname</i> <i>remote_URL</i>; timeout <i>seconds</i>; write-interval <i>seconds</i>; }</pre>
Hierarchy Level	<p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options secure-access-port]</pre> <p>For platforms with ELS:</p> <pre>[edit system processes] dhcp-service]</pre>
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Specify a local pathname or remote URL for the DHCP snooping database file to maintain persistence of IP-MAC bindings.</p> <p>The remaining statements are explained separately.</p>
Default	The IP-MAC bindings in the DHCP snooping database file are not persistent. If the switch is rebooted, the bindings are lost.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92

fc-map

Syntax `fc-map fc-map-value;`

Hierarchy Level Original CLI

[edit ethernet-switching options secure-access-port vlan (all | *vlan-name*) examine-fip]

ELS CLI for Platforms that Support FCoE

[edit vlans *vlan-name* forwarding-options fip-security]



NOTE: The `fc-map` configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.

QFX Series that Support FCoE-FC Gateway Configuration

[edit fc-fabrics *fc-fabric-name* protocols fip]

Release Information Statement introduced in Junos OS Release 10.4 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

Description Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN).

You can configure the FC-MAP value or use the default value. The default FC-MAP value is different for VN_Port to VF_Port (VN2VF_Port) FIP snooping (0x0EFC00) than for VN_Port to VN_Port (VN2VN_Port) FIP snooping.

The FC switch provides the FC-MAP value to FCoE nodes (ENodes) in the FIP discovery advertisement message. If the EX Series switch or the QFX Series FCoE VLAN FC-MAP value does not match the FC switch FC-MAP value, neither device discovers the FC switch on that VLAN, and the ENodes on that VLAN cannot access the FC switch. The FC switch accepts only FCoE traffic that uses the correct FC-MAP value as part of the VN_Port MAC address.

When the QFX Series acts as an FCoE-FC gateway, the FC-MAP value for the gateway and the FCoE devices must match the FC switch FC-MAP value in order to communicate with the FC switch.



NOTE: Changing the FC-MAP value causes all logins to drop and forces the ENodes to log in again.

Options `fc-map-value`—FC-MAP value, hexadecimal value preceded by "0x".

Range: 0x0EFC00 through 0x0EFCFF


Default: 0x0EFC00 for VN2VF_Port FIP snooping 0x0EFD00 for VN2VN_Port FIP snooping

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


Related Documentation

- *examine-fip*
- *show fip snooping*
- *Example: Configuring an FCoE Transit Switch*
- *Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch*

fcoe-trusted

Syntax	fcoe-trusted;
Hierarchy Level	Original CLI [edit ethernet-switching-options secure-access-port interface <i>interface-name</i>] ELS CLI for Platforms that Support FCoE [edit vlans <i>vlan-name</i> forwarding-options fip-security interface <i>interface-name</i>]
	<div>  <p>NOTE: The fcoe-trusted configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.</p> </div>
	<p>QFX Series that Support FCoE-FC Gateway Configuration</p> <p>[edit fc-fabrics <i>fc-fabric-name</i> protocols fip]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced for the FC fabric in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.</p>
Description	<p>Configure the specified 10-Gigabit Ethernet interface to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.</p> <p>(QFX Series FCoE-FC gateway) Configure the specified local Fibre Channel fabric to trust FCoE traffic on all ports in the fabric. Changing the fabric ports from untrusted to trusted removes any existing FIP snooping filters from the ports. Changing the fabric ports from trusted to untrusted by removing the fcoe-trusted configuration from the fabric forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN_Ports log in again, the switch can build the appropriate FIP snooping filters.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>show fip snooping</i> • <i>Example: Configuring an FCoE Transit Switch</i> • <i>Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch</i> • <i>Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch</i>

mac-move-limit

Syntax	<code>mac-move-limit <i>limit</i> <fabric-limit <i>limit</i>> action <i>action</i>;</code>
Hierarchy Level	<p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options secure-access-port (all <i>vlan-name</i>)]</pre> <p>For platforms with ELS:</p> <pre>[edit vlans <i>vlan-name</i> switch-options],</pre>
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the number of times a MAC address can move to a new interface (port) in 1 second and the action to be taken by the switch if the MAC address move limit is exceeded.
	<div>  <p>CAUTION: Mac move limiting does not work properly on a QFX5100 switch used as a Node device in a QFabric system. Do not use this feature on a QFX5100 switch in a QFabric system.</p> </div>
Default	The default move limit is unlimited. The default action is drop .
Options	<p>fabric-limit—Specify the maximum number of moves in a QFabric system. If you do not specify a fabric limit, the value for mac-move-limit applies to the QFabric system.</p> <p>limit—Maximum number of moves to a new interface per second.</p> <p>action <i>action</i>—(Optional) Action to take when the MAC address move limit is reached:</p> <ul style="list-style-type: none"> • drop—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. This is the default. • log—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry. • none—No action. • shutdown—Logically disable the interface and generate a system log entry. If you have configured the switch with the port-error-disable statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the clear-ethernet-switch-port command.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- | | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring Basic Port Security Features</i>• <i>Configuring MAC Move Limiting (CLI Procedure)</i>• <i>Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)</i> |
|------------------------------|---|

no-allowed-mac-log

- | | |
|---------------------------------|---|
| Syntax | no-allowed-mac-log; |
| Hierarchy Level | <ul style="list-style-type: none">• For platforms without ELS:
[edit ethernet-switching-options secure-access-port interface (all <i>interface-name</i>)]• For platforms with ELS:
[edit switch-options interface <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | Specify that the switch should not log messages when it receives packets from invalid MAC addresses on an interface that has been configured for allowed MAC addresses. |
| Default | The switch logs messages when it receives packets from invalid MAC addresses on an interface that has been configured for particular allowed (specific) MAC addresses. |
| Required Privilege Level | routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Understanding MAC Limiting and MAC Move Limiting for Port Security on page 100• <i>Configuring MAC Limiting</i>• <i>mac-limit</i> |


no-gratuitous-arp-request

Syntax	no-gratuitous-arp-request;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces interface-range <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on both Layer 2 Ethernet switching interfaces and routed VLAN interfaces (RVIs).
Default	Gratuitous ARP responses are enabled on all Ethernet switching interfaces and RVIs.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring IRB Interfaces</i>

persistent-learning

Syntax	persistent-learning;
Hierarchy Level	<ul style="list-style-type: none"> • For platforms without ELS: [edit ethernet-switching-options secure-access-port interface (all <i>interface-name</i>)] • For platforms with ELS: [edit switch-options interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.4 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Hierarchy level [edit switch-options interface <i>interface-name</i>] introduced in Junos OS Release 13.2X50-D10
Description	Specify that learned MAC addresses persist on the specified interfaces across restarts of the switch and link-down conditions. This feature is also known as sticky MAC.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Basic Port Security Features</i> • <i>Configuring Persistent MAC Learning (CLI Procedure)</i> • <i>Configuring Persistent MAC Learning (CLI Procedure)</i>

port-error-disable

Syntax	<pre>port-error-disable { (disable-timeout <i>seconds</i> <i>recovery-timeout seconds</i>); }</pre>
Hierarchy Level	<ul style="list-style-type: none">For platforms without ELS: [edit ethernet-switching-options]For platforms with ELS: [edit switch-options]
Release Information	Statement introduced in Junos OS Release 11.1 on the QFX Series.
Description	Disable rather than block an interface when enforcing MAC limiting, MAC move limiting, and storm control, and allow the interface to recover automatically from the error condition after a specified period of time:
	<div>NOTE: The port-error-disable configuration does not apply to preexisting error conditions. It affects only error conditions that are detected after you enable and commit the port-error-disable statement. To clear a preexisting error condition and restore the interface to service, use the clear ethernet-switching port-error command.</div> <ul style="list-style-type: none">If you enable the <i>mac-limit</i> statement with the shutdown option and also enable the port-error-disable statement, the switch disables (rather than shuts down) the interface when the MAC address limit is reached.If you have enabled the mac-move-limit statement with the shutdown option and you enable the port-error-disable statement, the switch disables (rather than shuts down) the interface when the maximum number of moves to a new interface is reached.If you enable the storm-control statement with the action-shutdown option and you also enable port-error-disable, the switch disables (rather than shuts down) the interface when broadcast traffic and unknown unicast traffic exceed the specified levels.
Default	Not enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing—control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Understanding MAC Limiting and MAC Move Limiting for Port Security on page 100Understanding Storm Control on page 119Example: Configuring Storm Control to Prevent Network Outages

- *Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure)*
- [action-shutdown on page 234](#)
- *disable-timeout*
- [clear ethernet-switching port-error on page 302](#)

vendor-id

Syntax	<code>vendor-id <string>;</code>
For Platforms with Enhanced Layer 2 Software (ELS)	<code>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82]</code>
For Platforms Without ELS	<code>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82],</code> <code>[edit forwarding-options helpers bootp dhcp-option82],</code> <code>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]</code>
For MX Series Platforms	<code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82]</code>
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Hierarchy level <code>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security]</code> introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.) Hierarchy level <code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security]</code> introduced in Junos OS Release 14.1 for the MX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Insert a vendor ID in the DHCP option 82 information in a DHCP request packet header before forwarding or relaying the request to a DHCP server.
Default	If vendor-id is not explicitly configured for DHCP option 82, then no vendor ID is set.
Options	string —(Optional) A single string that designates the vendor ID. Range: 1–255 characters Default: If you specify vendor-id with no string value, then the default vendor ID Juniper Networks is configured.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Configuring DHCP Option 82 to Help Protect the Switching Devices Against Attacks (CLI Procedure)</i> <i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i> <i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i> <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i>

write-interval



Syntax	<code>write-interval <i>seconds</i>;</code>
Hierarchy Level	For platforms without ELS: [edit ethernet-switching-options secure-access-port dhcp-snooping-file] For platforms with ELS: [edit system processes] dhcp-service dhcp-snooping-file]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify how frequently the switch writes the database entries from memory into the specified DHCP snooping database file.
Default	None
Options	<i>seconds</i> —Value in seconds. Range: 60 through 86400
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92

CHAPTER 9

Configuration Statements for Port Security (ELS CLI Only)



- [accept-source-mac on page 212](#)
- [arp-inspection on page 214](#)
- [dhcp-security on page 216](#)
- [dhcp-service on page 219](#)
- [group \(DHCP Security\) on page 220](#)
- [interface \(DHCP Security\) on page 221](#)
- [interface-mac-limit on page 222](#)
- [no-dhcp-snooping on page 224](#)
- [no-option82 on page 225](#)
- [option-82 on page 226](#)
- [overrides \(DHCP Security\) on page 227](#)
- [recovery-timeout on page 228](#)
- [static-ip on page 230](#)
- [switch-options on page 231](#)
- [trusted on page 232](#)
- [untrusted on page 232](#)

accept-source-mac

Syntax	<pre> accept-source-mac { mac-address <i>mac-address</i> { policer { input <i>cos-policer-name</i>; output <i>cos-policer-name</i>; } } } </pre>
Hierarchy Level	<pre> [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>
Description	<p>For Gigabit Ethernet intelligent queuing (IQ) interfaces only, accept traffic from and to the specified remote media access control (MAC) address.</p> <p>The accept-source-mac statement is equivalent to the source-address-filter statement, which is valid for aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only. To allow the interface to receive packets from specific MAC addresses, include the accept-source-mac statement.</p> <p>On untagged Gigabit Ethernet interfaces, you should not configure the source-address-filter statement and the accept-source-mac statement simultaneously. On tagged Gigabit Ethernet interfaces, you should not configure the source-address-filter statement and the accept-source-mac statement with an identical MAC address specified in both filters.</p> <p>The statements are explained separately.</p>
	<p> NOTE: The policer statement is not supported on PTX Series Packet Transport Routers.</p>
	<p> NOTE: On QFX platforms, if you configure source MAC addresses for an interface using the static-mac or persistent-learning statements and later configure a different MAC address for the same interface using the accept-source-mac statement, the MAC addresses that you previously configured for the interface remain in the ethernet-switching table and can still be used to send packets to the interface.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Configuring MAC Address Filtering*
 - *Configuring MAC Address Filtering on PTX Series Packet Transport Routers*
 - *source-filtering*

arp-inspection

Syntax	<pre>arp-inspection { forwarding-class <i>class-name</i>; }</pre>
Hierarchy Level	<ul style="list-style-type: none"> For platforms with ELS: <ul style="list-style-type: none"> [edit vlans <i>vlan-name</i> forwarding-options dhcp-security], [edit forwarding-options dhcp-relay] For platforms without ELS: <ul style="list-style-type: none"> [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)], [edit forwarding-options dhcp-relay]
Release Information	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p>
Description	<p>Perform dynamic ARP inspection (DAI) on all VLANs or on the specified VLAN.</p> <p>When DAI is enabled, the switch logs invalid ARP packets that it receives on each interface, along with the sender's IP and MAC addresses. ARP probe packets, which have the sender IP address 0.0.0.0, are validated by DAI.</p>
<div>  <p>NOTE: If you configure DAI at the [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] hierarchy level:</p> <ul style="list-style-type: none"> DAI can be configured only for a specific VLAN, not for a list or a range of VLAN IDs. DHCP snooping is automatically enabled on the specified VLAN. The forwarding-class statement is not available at the [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] hierarchy level. <p>See <i>Enabling Dynamic ARP Inspection (CLI Procedure)</i> for more information about this configuration.</p> </div>	
<div>  <p>NOTE: On EX9200 switches, DAI is not supported in an MC-LAG scenario.</p> </div> <p>The remaining statement is explained separately.</p>	
Default	Disabled.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch</i>• <i>Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks</i>• <i>Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing</i>• <i>Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic</i>• <i>Enabling Dynamic ARP Inspection (CLI Procedure)</i>• <i>Enabling Dynamic ARP Inspection (J-Web Procedure)</i>

dhcp-security

```
Syntax  dhcp-security {
        arp-inspection;
        dhcpv6-options {
            option-16 {
                use-string string;
            }
            option-18 {
                prefix {
                    host-name;
                    logical-system-name;
                    routing-instance-name;
                    vlan-id;
                    vlan-name;
                }
                use-interface-mac;
                use-interface-index (device | logical);
                use-interface-description (device | logical);
                use-interface-name (device | logical);
                use-string string;
            }
            option-37 {
                prefix {
                    host-name;
                    logical-system-name;
                    routing-instance-name;
                    vlan-id;
                    vlan-name;
                }
                use-interface-mac;
                use-interface-index (device | logical);
                use-interface-description (device | logical);
                use-interface-name (device | logical);
                use-string string;
            }
        }
        group group-name {
            interface interface-name {
                static-ip ip-address {
                    mac mac-address;
                }
                static-ipv6 ip-address {
                    mac mac-address;
                }
            }
        }
        overrides {
            no-dhcpv6-options;
            no-option16;
            no-option18;
            no-option37;
            no-option82;
            trusted;
            untrusted;
        }
    }
```

```

    }
  }
  ip-source-guard;
  ipv6-source-guard;
  neighbor-discovery-inspection;
  no-dhcp-snooping;
  no-dhcpv6-snooping;
  option-82 {
    circuit-id {
      prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
      }
      use-interface-description (device | logical);
      use-vlan-id;
    }
    remote-id {
      host-name hostname;
      use-interface-description (device | logical);
      mac (Option 82);
      use-string string;
    }
    vendor-id {
      use-string string;
    }
  }
}

```

Hierarchy Level [edit vlans *vlan-name* forwarding-options]

Release Information Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. Support for **static-ipv6**, **neighbor-discovery-inspection**, **ipv6-source-guard**, **no-dhcpv6-snooping**, and **no-option37** introduced in Junos OS Release 13.2X51-D20 for EX Series switches. Support for **dhcpv6-options**, **option-16**, **option-18**, **option-37**, **no-dhcpv6-options**, **no-option16**, **no-option18**, and **no-option37** introduced in Junos OS Release 14.2 for EX Series switches.

Description Configure port security features on the switch. DHCP snooping is enabled automatically if you configure any of the following port security features within this hierarchy:

- Dynamic ARP inspection (DAI)
- IP source guard
- DHCP option 82
- Static IP

For switches that support DHCPv6, both DHCP snooping and DHCPv6 snooping are enabled automatically if you configure any of the afore-mentioned features or any of the following IPv6 features:

- IPv6 neighbor discovery inspection
- IPv6 source guard
- Static IPv6



NOTE: On EX9200 switches, DHCP Snooping, DHCPv6 Snooping and Port Security features are not supported in MC-LAG scenario.

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

- *Enabling Dynamic ARP Inspection (CLI Procedure)*
- *Configuring IP Source Guard (CLI Procedure)*
- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure)*

dhcp-service

Syntax	<pre>dhcp-service { dhcp-snooping-file (<i>local_pathname</i> <i>remote_URL</i>); write-interval <i>interval</i>; }</pre>
Hierarchy Level	[edit system processes]
Release Information	<p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p>
Description	<p>Enable DHCP services on the device. DHCP services automate network-parameter assignment to network devices. The DHCP service process is enabled by default. However, by default, IP-MAC bindings in the DHCP snooping database do not persist through device reboots. You can improve performance after rebooting by configuring the IP-MAC bindings to persist, by configuring a storage location for the DHCP database file. When specifying the location for the DHCP database, you must also specify how frequently the switch writes the database entries into the DHCP snooping database file.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance (CLI Procedure)</i>

group (DHCP Security)

Syntax

```
group group-name {  
  interface interface-name {  
    static-ip ip-address {  
      mac mac-address;  
    }  
    static-ipv6 ip-address {  
      mac mac-address;  
    }  
  }  
  overrides {  
    no-dhcpv6-options;  
    no-option16;  
    no-option18;  
    no-option37;  
    no-option82;  
    trusted;  
    untrusted;  
  }  
}
```

Hierarchy Level [edit vlans *vlan-name* forwarding-options **dhcp-security**]

Release Information Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. Support for the **static-ipv6** and **no-option37** statements introduced in Junos OS Release 13.2X51-D20 for EX Series switches.

Description Specify the name of a group of access interfaces that you want to configure for DHCP security attributes that are different from the attributes set for other interfaces in the VLAN. A group must contain at least one interface.

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

- *Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure)*
- *Enabling a Trusted DHCP Server (CLI Procedure)*
- [Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92](#)

interface (DHCP Security)

Syntax	<pre> interface <i>interface-name</i> { static-ip <i>ip-address</i> { mac <i>mac-address</i>; } static-ipv6 <i>ip-address</i> { mac <i>mac-address</i>; } } </pre>
Hierarchy Level	[edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Support for the static-ipv6 statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches.</p>
Description	<p>Configure an interface for a static IPv4 or IPv6 address to MAC address binding (IP-MAC binding) or configure an interface to belong to a group within the VLAN that has DHCP security attributes that are different from the attributes of other interfaces in the VLAN.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure)</i> • <i>Enabling a Trusted DHCP Server (CLI Procedure)</i> • <i>Configuring Port Security Features</i>

interface-mac-limit

Syntax	<pre>interface-mac-limit <i>limit</i> { packet-action drop; }</pre>
Hierarchy Level	<pre>[edit bridge-domains <i>bridge-domain-name</i> bridge-options], [edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> switch-options], [edit logical-systems <i>logical-system-name</i> switch-options interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> switch-options], [edit routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>], [edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the switch-options statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the virtual-switch type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.</p>



NOTE: For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the `interface-mac-limit` statement or changing the `interface-mac-limit` configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the `interface-mac-limit` statement or use the `commit at` configuration statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers or switches by using the `clear bridge mac-table` command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

Default For an access port, the default MAC limit is 1024 MAC addresses. For a trunk port, the default MAC limit is 8192 MAC addresses.

Options *limit*—Maximum number of MAC addresses learned from an interface.

Range: 1 through 524287 MAC addresses per interface

The remaining statement is explained separately.

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

Related Documentation

- *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
- *Layer 2 Learning and Forwarding for VLANs Overview*
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

no-dhcp-snooping

Syntax	no-dhcp-snooping;
Hierarchy Level (EX Series, QFX Series)	[edit vlans <i>vlan-name</i> forwarding-options dhcp-security]
Hierarchy Level (MX Series)	[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security]
Release Information	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. Hierarchy level [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 14.1 for the MX Series.
Description	Disable DHCP snooping for the specified VLAN or bridge domain.



NOTE: Explicitly disabling DHCP snooping also disables any other port security features that you have enabled under [edit vlans *vlan-name* forwarding-options **dhcp-security**], including dynamic ARP inspection (DAI) and IP source guard for the specified VLAN or bridge domain.

There is no configuration statement that explicitly enables DHCP snooping.

Default DHCP snooping is not enabled.



NOTE: Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style and MX Series routers do not have a configuration statement that explicitly enables DHCP snooping.

DHCP snooping is enabled automatically by Junos OS if any of the following is configured at the [edit vlans *vlan-name* forwarding-options **dhcp-security**] hierarchy level for EX Series switches or at the [edit bridge-domains *bridge-domain-name* forwarding-options **dhcp-security**] for MX Series routers:

- DAI
- IP source guard
- Static IP
- DHCP option 82

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

- Related Documentation**
- [Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92](#)

no-option82

Syntax	no-option82;
Hierarchy Level (EX Series, QFX Series)	[edit vlans <i>vlan-name</i> forwarding-options dhcp-security group group-name overrides]
Hierarchy Level (MX Series)	[edit bridge-domains <i>bridge-domain-name</i> forwarding-options group group <i>group-name</i> overrides]
Release Information	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. Statement introduced in Junos OS Release 14.1 for the MX Series.
Description	Configure a specific group of one or more access interfaces within the VLAN or bridge domain <i>not</i> to transmit DHCP option 82 information, even if the VLAN or bridge domain is configured to use option 82.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • option-82 on page 226 • Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) • Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks • Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92

option-82


Syntax	<pre> option-82 { circuit-id { prefix (host-name routing-instance-name); use-interface-description; use-vlan-id; } remote-id { host-name; mac (Option 82); use-interface-description; use-string string; } vendor-id { use-string string; } } </pre>
Hierarchy Level (EX Series, QFX Series)	[edit vlans <i>vlan-name</i> forwarding-options dhcp-security]
Hierarchy Level (MX Series)	[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security]
Release Information	<p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Juos OS Release 14.1 for the MX Series.</p>
Description	<p>Have the device insert DHCP option 82 (also known as the DHCP relay agent information option) information in the DHCP request packet header of a DHCP request that it receives from a DHCP client connected to one of its interfaces before it forwards or relays that DHCP request to a DHCP server. The server uses the option 82 information, which provides details about the circuit and host the request came from. However, in formulating the reply, the server does not make any changes to the option 82 information in the packet header. The device receives the reply and then removes the DHCP option 82 information before forwarding the reply to the client.</p> <p>The remaining statements are explained separately.</p>
Default	Insertion of DHCP option 82 information is not enabled.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i> • no-option82 on page 225 • <i>Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks</i>

- [Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92](#)
- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>

overrides (DHCP Security)

Syntax	<pre> overrides { no-dhcpv6-options; no-option16; no-option18; no-option37; no-option82; trusted; untrusted; } </pre>
Hierarchy Level	[edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Support for the no-option37 option introduced in Junos OS Release 13.2X51-D20 for EX Series switches.</p> <p>Support for the no-dhcpv6-options, no-option16 and no-option18 options introduced in Junos OS Release 14.2 for EX Series switches.</p>
Description	Modify selected DHCP attributes for a group of interfaces that is configured within a specified VLAN.
Options	<p>trusted—The interface specified in this group is trusted. DHCP snooping and DHCPv6 snooping do not apply to the trusted interface. Likewise, DAI, IP source guard, IPv6 source guard, and IPv6 neighbor discovery inspection—even if they are enabled for the VLAN—do not apply to the interface that is configured with the overrides and the trusted options. Access interfaces are untrusted by default.</p> <p>untrusted—(Only for EX9200) The interface specified in this group is untrusted. Trunk interface are trusted by default. Access interfaces are untrusted by default.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling a Trusted DHCP Server (CLI Procedure) • Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92 • Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks

recovery-timeout

Syntax	<code>recovery-timeout seconds;</code>
Hierarchy Level (EX Series and QFX Series)	[edit interfaces <i>interface-name</i> unit 0 family ethernet-switching]
Hierarchy Level (MX Series)	[edit interfaces <i>interface-name</i> unit 0 family bridge]
Release Information	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1 for the MX Series routers.
Description	<p>Configure an interface to be temporarily disabled when MAC limiting, MAC move limiting, or rate-limiting is in effect with the action shutdown. This enables the affected interface to recover automatically from the error condition after the specified period of time:</p> <ul style="list-style-type: none"> • If you configure MAC limiting with the shutdown option and you enable recovery-timeout, the interface is temporarily disabled when the MAC address limit is reached. The interface will recover automatically after the number of seconds specified. • If you enable MAC move limiting with the shutdown option and you enable recovery-timeout, the interface is temporarily disabled when the maximum number of MAC address moves is reached. The interface will recover automatically after the number of seconds specified. • If you enable MAC move limiting with the vlan-member-shutdown option and you enable recovery-timeout, the interface is temporarily disabled when the maximum number of MAC address moves is reached. The interface will recover automatically after the number of seconds specified. If the recovery timeout is not configured, the interface will recover automatically after 180 seconds. • If you enable storm control with the action-shutdown option and you enable recovery-timeout, the interface is temporarily disabled when applicable traffic exceeds the specified levels. Depending upon the configuration, applicable traffic can include broadcast, unknown unicast, and multicast traffic.
	<p> NOTE: The recovery-timeout configuration does not apply to preexisting error conditions. It impacts only error conditions that are detected after the recovery-timeout statement is configured and committed. To clear a preexisting error condition and restore the interface to service, use the operational mode commands <code>clear ethernet-switching recovery-timeout</code> for EX Series and QFX Series and <code>clear bridge recovery-timeout</code> for MX Series routers.</p>
Default	The interface does not automatically recover from an error condition.



NOTE: On EX9200 switches, if a MAC move limit is configured with the action `vlan-member-shutdown`, the interface automatically recovers from the disabled condition after 180 seconds by default.

Options *seconds*— Number of seconds that the interface remains in a disabled state due to a port error prior to automatic recovery.

Range: 10 through 3600


Required Privilege `system`—To view this statement in the configuration.

Level `system-control`—To add this statement to the configuration.

**Related
Documentation**

- *action-shutdown*
- *Configuring MAC Limiting (CLI Procedure)*
- *Configuring MAC Move Limiting (CLI Procedure)*
- *Configuring or Disabling Storm Control (CLI Procedure)*

static-ip

Syntax	static-ip <i>ip-addresses</i> { vlan <i>vlan-name</i> ; mac <i>mac-address</i> ; }
Hierarchy Level	<ul style="list-style-type: none"> For platforms with ELS: [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> interface <i>interface-name</i>] For platforms without ELS: [edit ethernet-switching-options secure-access-port interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches. Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)
Description	Configure a static IP address to MAC address (IP-MAC) binding to be added to the DHCP snooping database.
<div>  <p>NOTE: The VLAN is specified at the higher hierarchy level when static-ip is configured at [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> interface <i>interface-name</i>].</p> </div>	
Options	<p><i>ip-address</i>—Static IP address assigned to a device connected on the specified interface.</p> <p><i>macmac-address</i>—Static MAC address assigned to a device connected on the specified interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure)</i> <i>Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure)</i>

switch-options

Syntax	<pre> switch-options { interface <i>interface-name</i> { interface-mac-limit <i>limit</i> { packet-action drop; } no-mac-learning; static-mac <i>static-mac-address</i> { vlan-id <i>number</i>; } } interface-mac-limit <i>limit</i> { packet-action drop; } mac-statistics; mac-table-size <i>limit</i> { packet-action drop; } no-mac-learning; service-id <i>number</i>; vtep-source-interface } </pre>
Hierarchy Level	<pre> [edit <i>number</i>], [edit vlans <i>vlan--name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i>], [edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i>] </pre>
Release Information	<p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>
Description	<p>Configure Layer 2 learning and forwarding properties for a VLAN or a virtual switch.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

trusted

Syntax	trusted;
Hierarchy Level	[edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name overrides</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series.
Description	Allow DHCP responses from the specified interface. The interface is not subject to DHCP snooping, even if the VLAN is enabled for DHCP snooping.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Enabling a Trusted DHCP Server (CLI Procedure)</i>• Understanding Trusted DHCP Servers for Port Security on page 110• Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92

untrusted

Syntax	untrusted;
Hierarchy Level	[edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name overrides</i>]
Release Information	Statement introduced in Junos OS Release 13.2 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series.
Description	Override the default behavior of a trunk interface from trusted to untrusted.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Enabling a Trusted DHCP Server (CLI Procedure)</i>• Understanding Trusted DHCP Servers for Port Security on page 110• Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices on page 92

CHAPTER 10

Configuration Statements for Device Security

- [action-shutdown](#) on page 234
- [interface \(Unknown Unicast Forwarding\)](#) on page 235
- [no-broadcast](#) on page 236
- [no-multicast](#) on page 237
- [no-unknown-unicast](#) on page 238
- [rpf-check](#) on page 239
- [unknown-unicast-forwarding](#) on page 240

action-shutdown

Syntax	action-shutdown;
Hierarchy Level	For platforms without ELS: [edit ethernet-switching-options storm-control] For platforms with ELS: [edit forwarding-options storm-control-profiles]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Shut down or disable interfaces when the storm control level is exceeded, as follows:</p> <ul style="list-style-type: none">• If you set both the action-shutdown and the port-error-disable statements, the affected interfaces are disabled temporarily and recover automatically when the disable timeout expires.• If you set the action-shutdown statement and do not set the port-error-disable statement, the affected interfaces are shut down when the storm control level is exceeded, and they do not recover automatically. You must issue the clear ethernet-switching port-error command to clear the port error and restore the interfaces to service.
Default	The action-shutdown feature is disabled. If the storm control level is exceeded, the switch drops broadcast and unknown unicast messages on the specified interfaces.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Storm Control on page 119• <i>Example: Configuring Storm Control to Prevent Network Outages</i>• port-error-disable on page 206• clear ethernet-switching port-error on page 302

interface (Unknown Unicast Forwarding)

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	<ul style="list-style-type: none"> For platforms with ELS: [edit switch-options unknown-unicast-forwarding vlan <i>vlan-name</i>] For platforms without ELS: [edit ethernet-switching-options unknown-unicast-forwarding vlan <i>vlan-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p>
Description	Specify the interface to which unknown unicast packets will be forwarded.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <code>show vlans</code> <code>show ethernet-switching table</code> <i>Configuring Unknown Unicast Forwarding (CLI Procedure)</i> Understanding Unknown Unicast Forwarding on page 133

no-broadcast

Syntax	no-broadcast;
Hierarchy Level	For platforms without ELS: [edit ethernet-switching-options storm-control interface (all <i>interface-name</i>)] For platforms with ELS: [edit forwarding-options storm-control-profiles]
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	For interfaces configured for storm control, disable broadcast traffic storm control on the interface.
Default	When storm control is enabled on an interface, it is enabled for broadcast traffic (as well as multicast and unknown unicast traffic).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Storm Control on page 119• <i>Example: Configuring Storm Control to Prevent Network Outages</i>

no-multicast

Syntax	no-multicast;
Hierarchy Level	For platforms without ELS: [edit ethernet-switching-options storm-control interface (all <i>interface-name</i>)] For platforms with ELS: [edit forwarding-options storm-control-profiles]
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	Disable storm control for all multicast traffic (both registered multicast and unregistered multicast) for the specified interface or for all interfaces.
Default	Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Storm Control on page 119• <i>Example: Configuring Storm Control to Prevent Network Outages</i>

no-unknown-unicast

Syntax	no-unknown-unicast;
Hierarchy Level	For platforms without ELS: [edit ethernet-switching-options storm-control interface (all <i>interface-name</i>)] For platforms with ELS: [edit forwarding-options storm-control-profiles]
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	For interfaces configured for storm control, disable unknown unicast traffic storm control on the interface.
Default	When storm control is enabled on an interface, it is enabled for both unknown unicast traffic and broadcast traffic.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Storm Control on page 119• <i>Example: Configuring Storm Control to Prevent Network Outages</i>

rpf-check

Syntax	rpf-check;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>On EX3200 and EX4200 switches, enable a reverse-path forwarding (RPF) check on unicast traffic (except ECMP packets) on all ingress interfaces.</p> <p>On EX4300 switches, enable a reverse-path forwarding (RPF) check on unicast traffic, including ECMP packets, on all ingress interfaces.</p> <p>On EX8200 and EX6200 switches, enable an RPF check on unicast traffic, including ECMP packets, on the selected ingress interfaces.</p> <p>On QFX Series switches, enable an RPF check on unicast traffic (except ECMP packets) on the selected ingress interfaces.</p>
Default	Unicast RPF is disabled on all interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Unicast RPF on an EX Series Switch</i> • Configuring Unicast RPF (CLI Procedure) on page 128 • Disabling Unicast RPF (CLI Procedure) on page 132 • Understanding Unicast RPF on page 124

unknown-unicast-forwarding

Syntax	<pre>unknown-unicast-forwarding { vlan (Unknown Unicast Forwarding) (all <i>vlan-name</i>){ interface (Unknown Unicast Forwarding) <i>interface-name</i>; } }</pre>
Hierarchy Level	<ul style="list-style-type: none">For platforms with ELS: [edit switch-options]For platforms without ELS: [edit ethernet-switching-options]
Release Information	<p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the switch to forward all unknown unicast packets in a VLAN or on all VLANs to a particular interface.



NOTE: Before you can configure unknown unicast forwarding within a VLAN, you must first configure that VLAN.

The remaining statements are explained separately.


Default	Unknown unicast packets are flooded to all interfaces that belong to the same VLAN.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><i>show vlans</i><i>show ethernet-switching table</i><i>Configuring Unknown Unicast Forwarding (CLI Procedure)</i>Understanding Unknown Unicast Forwarding on page 133

CHAPTER 11


Configuration Statements for Device Security (ELS CLI Only)

- [bandwidth-level on page 242](#)
- [bandwidth-percentage on page 243](#)
- [no-registered-multicast on page 244](#)
- [no-unregistered-multicast on page 245](#)
- [storm-control on page 246](#)
- [storm-control-profiles on page 247](#)

bandwidth-level

Syntax	<code>bandwidth-level <i>kbps</i>;</code>
Hierarchy Level	[edit forwarding-options storm-control-profiles <i>profile-name</i> all]
Release Information	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1 for MX Series routers.
Description	Configure the storm control level as the bandwidth in kilobits per second of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams.
<div>  <p>NOTE: When you configure storm control level on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control level of 15,000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30,000 Kbps of combined broadcast, multicast, and unknown unicast traffic.</p> </div>	
Default	<p>On EX4300 switches—If you do not specify the storm control level using either the bandwidth-level or the bandwidth-percentage statements, the storm control level defaults to 80 percent of the available bandwidth used by the combined broadcast, unknown unicast, and multicast traffic streams.</p> <p>On EX9200 switches—Storm control is not enabled by default.</p> <p>On MX Series routers—Storm control is not enabled by default.</p>
Options	<p>bandwidth-level <i>kbps</i>—Traffic rate in kilobits per second of the combined broadcast, multicast, and unknown unicast traffic streams.</p> <p>Range: 100 through 10,000,000</p> <p>Range: 100 through 100,000,000 on QFX10000 Series switches</p> <p>Default: None</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • bandwidth-percentage on page 243 • <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i> • <i>Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers</i> • <i>Configuring or Disabling Storm Control (CLI Procedure)</i>

bandwidth-percentage

Syntax	<code>bandwidth-percentage <i>percentage</i>;</code>
Hierarchy Level	[edit forwarding-options storm-control-profiles <i>profile-name</i> all]
Release Information	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. Statement introduced in Junos OS Release 14.1 for MX Series routers.
Description	Configure the storm control level as the percentage of available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams on an interface. The storm control level is configured as part of the storm control profile.
<div>  <p>NOTE: When you configure storm control level on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control level of 15,000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30,000 Kbps of combined broadcast, multicast, and unknown unicast traffic.</p> </div>	
Default	<p>On EX4300 switches—The storm control level is 80 percent of the available bandwidth used by the combined broadcast, unknown unicast, and multicast traffic streams.</p> <p>On EX9200 switches—Storm control is not enabled by default.</p> <p>On MX Series routers—Storm control is not enabled by default.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • bandwidth-level on page 242 • <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i> • <i>Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers</i> • <i>Configuring or Disabling Storm Control (CLI Procedure)</i>


no-registered-multicast

Syntax	no-registered-multicast;
Hierarchy Level	<ul style="list-style-type: none">For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers): [edit forwarding-options storm-control-profiles <i>profile-name</i> all]For platforms without ELS: [edit ethernet-switching-options storm-control interface (all <i>interface-name</i>)]
Release Information	<p>Statement introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Hierarchy level [edit forwarding-options] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>(EX8200 switches only) Disable storm control for registered multicast traffic for the specified interface or for all interfaces.</p> <p>(EX4300 and EX9200 switches only) Exclude storm control for registered multicast traffic from a storm control profile.</p> <p>(MX Series routers only) Exclude storm control for registered multicast traffic from a storm control profile.</p>
Default	<p>EX4300 and EX8200 switches—Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic. The default storm control level is 80 percent of the available bandwidth used by the combined applicable traffic streams.</p> <p>EX9200 switches—Storm control is not enabled by default.</p> <p>MX Series routers—Storm control is not enabled by default.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><i>no-multicast</i>no-unregistered-multicast on page 245<i>Understanding Storm Control on EX Series Switches</i><i>Understanding Storm Control for Managing Traffic Levels on Switching Devices</i>

no-unregistered-multicast

Syntax	no-unregistered-multicast;
Hierarchy Level	<ul style="list-style-type: none"> For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers): [edit forwarding-options storm-control-profiles <i>profile-name</i> all] For platforms without ELS: [edit ethernet-switching-options storm-control interface (all <i>interface-name</i>)]
Release Information	<p>Statement introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Hierarchy level [edit forwarding-options] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>(EX8200 switches only) Disable storm control for unregistered multicast traffic for the specified interface or for all interfaces.</p> <p>(EX4300 and EX9200 switches only) Exclude storm control for unregistered multicast traffic from a storm control profile.</p> <p>(MX Series routers) Exclude storm control for unregistered multicast traffic from a storm control profile.</p>
Default	<p>EX4300 and EX8200 switches—Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic. The default storm control level is 80 percent of the available bandwidth used by the combined applicable traffic streams.</p> <p>EX9200 switches—Storm control is not enabled by default.</p> <p>MX Series routers—Storm control is not enabled by default.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>no-multicast</i> no-registered-multicast on page 244 <i>Understanding Storm Control on EX Series Switches</i> <i>Understanding Storm Control for Managing Traffic Levels on Switching Devices</i>

storm-control

Syntax	<code>storm-control storm-control-profile;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>number</i> family ethernet-switching], [edit interfaces <i>interface-name</i> unit <i>number</i> family bridge] [edit interfaces <i>interface-name</i> ether-options ethernet-switch-profile]
Release Information	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. Statement introduced in Junos OS Release 14.1 for the MX Series routers.
Description	<p>Bind a storm control profile to a logical interface.</p> <p>On switches running ELS software, storm control is enabled by default on all switch interfaces at a level of 80 percent of the combined broadcast and unknown unicast streams. (For the equivalent statement for platforms running non-ELS software, see <i>storm-control</i>.)</p>
<div> 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.</div>	
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i>• <i>Understanding Storm Control for Managing Traffic Levels on Switching Devices</i>

storm-control-profiles

Syntax `storm-control-profiles profile-name {
 action-shutdown;
 all {
 bandwidth-level;
 bandwidth-percentage;
 no-broadcast;
 no-multicast;
 no-registered-multicast;
 no-unknown-unicast;
 no-unregistered-multicast;
 }
 }`

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
 Statement introduced in Junos OS Release 13.2 for the QFX Series.
 Statement introduced in Junos OS Release 14.1 for MX Series routers.

Description Configure a storm control profile on a switch or router. Storm control is used to prevent network outages that are caused by broadcast traffic storms. Storm control enables the switching device to monitor traffic levels and to drop broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level or storm control bandwidth—is exceeded, thus preventing packets from proliferating and degrading the LAN.



NOTE: The name of the storm control profile can contain no more than 127 characters.

The remaining statements are explained separately.

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

Related Documentation

- *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*
- *Understanding Storm Control for Managing Traffic Levels on Switching Devices*

CHAPTER 12

Configuration Statements for Media Access Control Security (MACsec)

- cak on page 250
- ckn on page 251
- connectivity-association on page 252
- connectivity-association (MACsec Interfaces) on page 253
- direction on page 254
- encryption (MACsec) on page 255
- exclude-protocol on page 256
- id on page 257
- include-sci on page 258
- interfaces (MACsec) on page 259
- key (MACsec) on page 260
- key-server-priority (MACsec) on page 261
- mac-address (MACsec) on page 262
- macsec on page 263
- mka on page 264
- must-secure on page 265
- no-encryption (MACsec) on page 266
- offset on page 267
- port-id on page 268
- pre-shared-key on page 269
- replay-protect on page 270
- replay-window-size on page 271
- secure-channel on page 272
- security-association on page 273
- security-mode on page 274
- transmit-interval (MACsec) on page 275

cak

Syntax	<code>ckn <i>hexadecimal-number</i>;</code>
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i> pre-shared-key]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Specifies the connectivity association key (CAK) for a pre-shared key.</p> <p>A pre-shared key includes a connectivity association key name (CKN) and a CAK. A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security keys. The MACsec Key Agreement (MKA) protocol is enabled once the pre-shared keys are successfully exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link</p>
Default	No CAK exists, by default.
Options	<p><i>hexadecimal-number</i>—The key name, in hexadecimal format.</p> <p>The key name is 32 hexadecimal characters in length. If you enter a key name that is less than 32 characters long, the remaining characters are set to 0.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 147

ckn

Syntax	<code>ckn <i>hexadecimal-number</i>;</code>
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i> pre-shared-key]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Specifies the connectivity association key name (CKN) for a pre-shared key.</p> <p>A pre-shared key includes a CKN and a connectivity association key (CAK). A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security keys. The MACsec Key Agreement (MKA) protocol is enabled once the pre-shared keys are successfully exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link</p>
Default	No CKN exists, by default.
Options	<p><i>hexadecimal-number</i>—The key name, in hexadecimal format.</p> <p>The key name is 32 hexadecimal characters in length. If you enter a key name that is less than 32 characters long, the remaining characters are set to 0.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Media Access Control Security (MACsec) on page 147

connectivity-association

Syntax `connectivity-association connectivity-association-name {
 exclude-protocol protocol-name;
 include-sci;
 mka {
 must-secure;
 key-server-priority priority-number;
 transmit-interval interval;
 }
 no-encryption;
 offset (0|30|50);
 pre-shared-key {
 cak hexadecimal-number;
 ckn hexadecimal-number;
 }
 replay-protect {
 replay-window-size number-of-packets;
 }
 secure-channel secure-channel-name {
 direction (inbound | outbound);
 encryption (MACsec);
 id {
 mac-address mac-address;
 port-id port-id-number;
 }
 offset (0|30|50);
 security-association security-association-number {
 key key-string;
 }
 }
 security-mode security-mode;
 }`

Hierarchy Level [edit security [macsec](#)]

Release Information Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.

Description Create or configure a MACsec connectivity association.

A connectivity association is not applying MACsec to traffic until it is associated with an interface. MACsec connectivity associations are associated with interfaces using the [interfaces](#) statement in the [edit security macsec] hierarchy.

Default No connectivity associations are present, by default.

Options The remaining statements are explained separately.

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

Related Documentation • [Configuring Media Access Control Security \(MACsec\) on page 147](#)

connectivity-association (MACsec Interfaces)

Syntax	<code>connectivity-association <i>connectivity-association-name</i>;</code>
Hierarchy Level	[edit security macsec interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	Applies a connectivity association to an interface, which enables Media Access Control Security (MACsec) on that interface.
Default	No connectivity associations are associated with any interfaces.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	• Configuring Media Access Control Security (MACsec) on page 147

direction

Syntax	direction (inbound outbound);
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Configure whether the secure channel applies MACsec security to traffic entering or leaving an interface.</p> <p>If you need to apply MACsec on traffic entering and leaving an interface, you need to create one secure channel to apply MACsec on incoming traffic and another secure channel to apply MACsec on outgoing traffic within the same connectivity association. When you associate the connectivity association with an interface, MACsec is applied on traffic entering and leaving that interface.</p> <p>You only use this configuration option when you are configuring MACsec using static secure association keys (SAK) security mode. When you are configuring MACsec using static connectivity association keys (CAK) security mode, two secure channels that are not user-configurable—one inbound secure channel and one outbound secure channel—are automatically created within the connectivity association.</p>
Default	<p>This statement does not have a default value.</p> <p>If you have configured a secure channel to enable MACsec using static SAK security mode, you must specify whether the secure channel applies MACsec to traffic entering or leaving an interface. A candidate configuration that contains a secure channel that has not configured a direction cannot be committed.</p>
Options	<p>inbound—Enable MACsec security on traffic entering the interface that has applied the secure channel.</p> <p>outbound—Enable MACsec security on traffic leaving the interface that has applied the secure channel.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 147

encryption (MACsec)

Syntax	encryption;
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Enable MACsec encryption within a secure channel.</p> <p>You can enable MACsec without enabling encryption. If a connectivity association with a secure channel that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic has not been tampered with and does not represent a security threat.</p> <p>Traffic traversing a MAC-enabled point-to-point Ethernet link traverses the link at the same speed regardless of whether encryption is enabled or disabled. You cannot increase the speed of traffic traversing a MACsec-enabled Ethernet link by disabling encryption.</p> <p>This command is used to enable encryption when MACsec is configured using secure association key (SAK) security mode only. When MACsec is configuring using static connectivity association key (CAK) security mode, the encryption setting is configured outside of the secure channel using the no-encryption configuration statement.</p>
Default	MACsec encryption is disabled when MACsec is configured using static SAK security mode, by default.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 147

exclude-protocol

Syntax	<code>exclude-protocol <i>protocol-name</i>;</code>
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Specifies protocols whose packets are not secured using Media Access Control Security (MACsec) when MACsec is enabled on a link using static connectivity association key (CAK) security mode.</p> <p>When this option is enabled in a connectivity association that is attached to an interface, MACsec is not enabled for all packets of the specified protocols that are sent and received on the link.</p>
Default	<p>Disabled.</p> <p>All packets are secured on a link when MACsec is enabled, with the exception of all types of Spanning Tree Protocol (STP) packets.</p>
Options	<p><i>protocol-name</i>—Specifies the name of the protocol that should not be MACsec-secured. Options include:</p> <ul style="list-style-type: none">• cdp—Cisco Discovery Protocol.• lcp—Link Aggregation Control Protocol.• lldp—Link Level Discovery Protocol.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 147

id

Syntax	<pre>id { mac-address <i>mac-address</i>; port-id <i>port-id-number</i>; }</pre>
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	Specify a MAC address and a port that traffic on the link must be from to be accepted by the interface when MACsec is enabled using static secure association key (SAK) security mode.
Options	The remaining statements are explained separately.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 147

include-sci

Syntax	include-sci;
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.
Description	<p>Specifies that the SCI tag should be appended to each packet on a link that has enabled MACsec.</p> <p>You must enable SCI tagging on a switch that is enabling MACsec on an Ethernet link connecting to an EX4300 switch.</p> <p>SCI tags are automatically appended to packets leaving a MACsec-enabled interface on an EX4300 switch. This option is, therefore, not available on EX4300 switches.</p> <p>You should only use this option when connecting a switch to an EX4300 switch, or to a host device that requires SCI tagging. SCI tags are eight octets long, so appending an SCI tag to all traffic on the link adds a significant amount of unneeded overhead.</p>
Default	<p>SCI tagging is enabled on EX4300 switches that have enabled MACsec using static connectivity association key (CAK) security mode, by default.</p> <p>SCI tagging is disabled on all other interfaces, by default.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 147

interfaces (MACsec)

Syntax	<pre> interfaces <i>interface-name</i> { connectivity-association <i>connectivity-association-name</i>; } </pre>
Hierarchy Level	[edit security macsec]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Applies the specified connectivity association to the specified interface to enable MACsec.</p> <p>One connectivity association can be applied to multiple interfaces.</p> <p>You must always use this statement to apply a connectivity association to an interface to enable MACsec. You must complete this configuration step regardless of whether MACsec is enabled using static connectivity association key (CAK) security mode or static secure association key (SAK) security mode.</p> <p>If you are enabling MACsec using static SAK security mode and need to configure MACsec on inbound and outbound traffic on the same interface, you must configure a connectivity association with one secure channel for inbound traffic and a second secure channel for outbound traffic. The connectivity association is then applied to the interface using this statement to enable MACsec for traffic entering and leaving the interface.</p>
Default	Interfaces are not associated with any connectivity associations, by default.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 147

key (MACsec)

Syntax	<code>key <i>key-string</i>;</code>
Hierarchy Level	[edit security macsec connectivity-association connectivity-association-name secure-channel secure-channel-name security-association security-association-number]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Specifies the static security key to exchange to enable MACsec using static secure association key (SAK) security mode.</p> <p>The key string is a 32-digit hexadecimal number. The key string and the security association must match on both sides of an Ethernet connection to secure traffic using MACsec when enabling MACsec using SAK security mode.</p> <p>You must configure at least two security associations with unique security association numbers and key strings to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.</p>
Default	This statement does not have a default value.
Options	<i>key-string</i> —Specifies the key to exchange with the other end of the link on the secure channel. The <i>key-string</i> is a 32-digit hexadecimal string that is created by the user.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 147

key-server-priority (MACsec)

Syntax	<code>key-server-priority <i>priority-number</i>;</code>
Hierarchy Level	[edit security <code>macsec connectivity-association connectivity-association-name mka</code>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Specifies the key server priority used by the MACsec Key Agreement (MKA) protocol to select the key server when MACsec is enabled using static connectivity association key (CAK) security mode.</p> <p>The switch with the lower <i>priority-number</i> is selected as the key server.</p> <p>If the <i>priority-number</i> is identical on both sides of a point-to-point link, the MKA protocol selects the device with the lower MAC address as the key server.</p>
Default	The default key server priority number is 16.
Options	<p><i>priority-number</i>—Specifies the MKA server election priority number.</p> <p>The <i>priority-number</i> can be any number between 0 and 255. The lower the number, the higher the priority.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 147

mac-address (MACsec)

Syntax	<code>mac-address <i>mac-address</i>;</code>
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i> id]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Specify a MAC address to enable MACsec using static secure association key (SAK) security mode. The mac-address variables must match on the sending and receiving ends of a link to enable MACsec using static SAK security mode.</p> <p>If you are configuring a MAC address on a secure channel in the outbound direction, you should specify the MAC address of the interface as the mac-address.</p> <p>If you are configuring a MAC address on a secure channel in the inbound direction, you should specify the MAC address of the interface at the other end of the link as the mac-address.</p> <p>You only use this configuration option when you are configuring MACsec using static SAK security mode. This option does not need to be specified when you are enabling MACsec using static connectivity association key (CAK) security mode.</p>
Default	No MAC address is specified in the secure channel, by default.
Options	mac-address —The MAC address, in six groups of two hexadecimal digits.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 147

macsec

```
Syntax  macsec {
        connectivity-association connectivity-association-name {
            exclude-protocol protocol-name;
            include-sci;
            mka {
                must-secure;
                key-server-priority priority-number;
                transmit-interval interval;
            }
            no-encryption;
            offset (0|30|50);
            pre-shared-key {
                cak hexadecimal-number;
                ckn hexadecimal-number;
            }
            replay-protect {
                replay-window-size number-of-packets;
            }
            secure-channel secure-channel-name {
                direction (inbound | outbound);
                encryption (MACsec);
                id {
                    mac-address mac-address;
                    port-id port-id-number;
                }
                offset (0|30|50);
                security-association security-association-number {
                    key key-string;
                }
            }
            security-mode security-mode;
        }
        interfaces interface-name {
            connectivity-association connectivity-association-name;
        }
    }
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.

Description Configure Media Access Control Security (MACsec)..

Options The remaining statements are explained separately.

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

Related Documentation

- [Configuring Media Access Control Security \(MACsec\) on page 147](#)

mka

Syntax	<pre>mka { must-secure; key-server-priority <i>priority-number</i>; transmit-interval <i>interval</i>; }</pre>
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15.
Description	Specify parameters for the MACsec Key Agreement (MKA) protocol.
Options	The remaining statements are explained separately.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 147


must-secure

Syntax	<code>must-secure;</code>
Hierarchy Level	[edit security <code>macsec connectivity-association connectivity-association-name mka</code>]
Release Information	Statement introduced in Junos OS Release 14.1X53-D10.
Description	<p>Specifies that all traffic travelling on the MACsec-secured link must be MACsec-secured to be forwarded onward.</p> <p>When the must-secure option is enabled, all traffic that is not MACsec-secured that is received on the interface is dropped.</p> <p>When the must-secure option is disabled, all traffic from devices that support MACsec is MACsec-secured while traffic received from devices that do not support MACsec is forwarded through the network.</p> <p>The must-secure option is particularly useful in scenarios where multiple devices, such as a phone and a PC, are accessing the network through the same Ethernet interface. If one of the devices supports MACsec while the other device does not support MACsec, the device that doesn't support MACsec can continue to send and receive traffic over the network—provided the must-secure option is disabled—while traffic to and from the device that supports MACsec is MACsec-secured. In this scenario, traffic to the device that is not MACsec-secured must be VLAN-tagged.</p>
Default	The must-secure option is disabled.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 147

no-encryption (MACsec)

Syntax	no-encryption;
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Disables MACsec encryption for a connectivity association that is configured to enable MACsec using static connectivity association key (CAK) or dynamic security mode.</p> <p>You can enable MACsec without enabling encryption. If a connectivity association that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the packet, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic does not represent a security threat.</p> <p>This command is used to disable encryption when MACsec is configured using static CAK or dynamic security mode only. When MACsec is configuring using static secure association key (SAK) security mode, the encryption setting is managed in the secure channel using the encryption configuration statement.</p>
Default	MACsec encryption is enabled if MACsec is enabled using static CAK or dynamic security mode.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 147

offset

Syntax	offset (0 30 50);
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i>] [edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Specifies the number of octets in an Ethernet frame that are sent in unencrypted plain-text when encryption is enabled for MACsec.</p> <p>Setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the remaining traffic. Setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the remaining traffic.</p> <p>You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.</p> <p>You configure the offset in the [edit security macsec connectivity-association <i>connectivity-association-name</i>] hierarchy when you are enabling MACsec using static connectivity association key (CAK) or dynamic security mode.</p> <p>You configure the offset in the [edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i>] hierarchy when you are enabling MACsec using static secure association key (SAK) security mode.</p>
Default	0
Options	<p>0—Specifies that no octets are unencrypted. When you set the offset to 0, all traffic on the interface where the connectivity association or secure channel is applied is encrypted.</p> <p>30—Specifies that the first 30 octets of each Ethernet frame are unencrypted.</p>
	<div>  <p>NOTE: In IPv4 traffic, setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the rest of the traffic. An offset of 30, therefore, is typically used when a feature needs this information to perform a task on IPv4 traffic.</p> </div>
	<p>50—Specified that the first 50 octets of each Ethernet frame are unencrypted.</p>



NOTE: In IPv6 traffic, setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the rest of the traffic. An offset of 50, therefore, is typically used when a feature needs this information to perform a task on IPv6 traffic.

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

Related Documentation

- [Configuring Media Access Control Security \(MACsec\) on page 147](#)

port-id

Syntax `port-id port-id-number;`

Hierarchy Level [edit security [macsec connectivity-association](#) *connectivity-association-name* [secure-channel](#) *secure-channel-name* **id**]

Release Information Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.

Description Specify a port ID in a secure channel when enabling MACsec using static secure association key (SAK) security mode. The port IDs must match on a sending and receiving secure channel on each side of a link to enable MACsec.

Once the port numbers match, MACsec is enabled for all traffic on the connection.

You only use this configuration option when you are configuring MACsec using static SAK security mode. This option does not need to be specified when you are enabling MACsec using static connectivity association key (CAK) security mode.

Default No port ID is specified.

Options *port-id-number*—The port ID number.

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

Related Documentation

- [Configuring Media Access Control Security \(MACsec\) on page 147](#)

pre-shared-key

Syntax	<pre>pre-shared-key { cak hexadecimal-number; ckn hexadecimal-number; }</pre>
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Specifies the pre-shared key used to enable MACsec using static connectivity association key (CAK) security mode.</p> <p>A pre-shared key includes a connectivity association key name (CKN) and a connectivity association key (CAK). A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using static CAK security mode. The MACsec Key Agreement (MKA) protocol is enabled after the pre-shared keys are successfully verified and exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link.</p>
Default	No pre-shared keys exist, by default.
Options	The remaining statements are explained separately.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 147

replay-protect

Syntax	<pre>replay-protect { replay-window-size <i>number-of-packets</i>; }</pre>
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Enable replay protection for MACsec.</p> <p>A replay window size specified using the replay-window-size <i>number-of-packets</i> statement must be specified to enable replay protection.</p>
Options	The remaining statements are explained separately.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 147

replay-window-size

Syntax	<code>replay-window-size <i>number-of-packets</i>;</code>
Hierarchy Level	[edit security <code>macsec connectivity-association</code> <i>connectivity-association-name</i> replay-protect]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Specifies the size of the replay protection window.</p> <p>This statement has to be configured to enable replay protection.</p> <p>When MACsec is enabled on an Ethernet link, an ID number is assigned to each packet entering the link. The ID number of the packet is checked by the receiving interface after the packet has traversed the MACsec-enabled link.</p> <p>When replay protection is enabled, the sequence of the ID number of received packets are checked. If the packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.</p> <p>Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.</p> <p>Replay protection should not be enabled in cases where packets are expected to arrive out of order.</p>
Default	Replay protection is disabled.
Options	<p><i>number-of-packets</i>—Specifies the size of the replay protection window, in packets.</p> <p>When this variable is set to 0, all packets that arrive out-of-order are dropped.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Media Access Control Security (MACsec) on page 147

secure-channel

Syntax `secure-channel secure-channel-name {
 direction (inbound | outbound);
 encryption (MACsec);
 id {
 mac-address mac-address;
 port-id port-id-number;
 }
 offset (0|30|50);
 security-association security-association-number {
 key key-string;
 }
 }`

Hierarchy Level [edit security *macsec connectivity-association connectivity-association-name*]

Release Information Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.

Description Create and configure a secure channel to enable and configure MACsec when MACsec is enabled using static secure association key (SAK) security mode.

You do not need to use this option to enable MACsec using static connectivity association key (CAK) security mode. All configuration for MACsec using static CAK security mode is done inside of the connectivity association but outside of the secure channel. When MACsec is enabled using static CAK security mode, an inbound and an outbound secure channel—neither of which is user-configurable—is automatically created within the connectivity association.

Options The remaining statements are explained separately.

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

Related Documentation • [Configuring Media Access Control Security \(MACsec\) on page 147](#)

security-association

Syntax	<code>security-association <i>security-association-number</i> { key <i>key-string</i>; }</code>
Hierarchy Level	[edit security <i>macsec connectivity-association connectivity-association-name secure-channel secure-channel-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Specifies the number of one of the security associations in the secure channel when MACsec is enabled using static secure association key (SAK) security mode. Because SAKs are created by the key server when MACsec is enabled using static connectivity association key (CAK) security mode, the security-association statement is not used when enabling MACsec using static CAK security mode.</p> <p>You must configure at least two security associations to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.</p>
Default	No security keys are configured, by default.
Options	<p><i>security-association-number</i>—Specifies the security association number and creates the SAK.</p> <p>The security association number is a whole number between 0 and 3. You can configure two security associations in a secure channel when enabling MACsec using static security keys.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 147

security-mode

Syntax	<code>security-mode security-mode;</code>
Hierarchy Level	[edit security macsec connectivity-association <i>connectivity-association-name</i>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15. The dynamic security mode option was introduced in Junos OS Release 14.1X53-D10.
Description	<p>Configure the MACsec security mode for the connectivity association.</p> <p>We recommend enabling MACsec on switch-to-switch Ethernet links using static connectivity association key (CAK) security mode. Static CAK security mode ensures security by frequently refreshing to a new random secure association key (SAK) and by only sharing the SAK between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available when you enable MACsec using static CAK security mode.</p>
Options	<p>security-mode—Specifies the MACsec security mode. Options include:</p> <ul style="list-style-type: none">• dynamic—Dynamic mode. <p>Dynamic security mode is used to enable MACsec on switch-to-host Ethernet links. In dynamic mode, a master key is retrieved from a RADIUS server by a switch and a host as part of the AAA handshake in separate transactions. The MKA protocol is enabled when the master key is exchanged between the switch and the host.</p> <ul style="list-style-type: none">• static-cak—Static connectivity association key (CAK) mode. <p>Static CAK security mode is used to enable MACsec on switch-to-switch Ethernet links. In static-cak mode, the switch at one end of the point-to-point link acts as the key server and regularly transmits a randomized key using a process that does not transmit any traffic outside of the MACsec-secured point-to-point link.</p> <ul style="list-style-type: none">• static-sak—Static secure association key (SAK) mode. <p>Static SAK security mode is used to enable MACsec on switch-to-switch Ethernet links. In static-sak mode, one of two user-configured security keys is used to secure the point-to-point link. The two security keys are regularly rotated.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 147

transmit-interval (MACsec)

Syntax	<code>transmit-interval <i>interval</i>;</code>
Hierarchy Level	[edit security <code>macsec connectivity-association connectivity-association-name mka</code>]
Release Information	Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	<p>Specifies the transmit interval for MACsec Key Agreement (MKA) protocol data units (PDUs).</p> <p>The MKA transmit interval setting sets the frequency for how often the MKA PDU is sent to the directly connected device to maintain MACsec on a point-to-point Ethernet link. A lower <i>interval</i> increases bandwidth overhead on the link; a higher <i>interval</i> optimizes the MKA protocol data unit exchange process.</p> <p>The transmit interval settings must be identical on both ends of the link when MACsec using static connectivity association key (CAK) security mode is enabled.</p> <p>We recommend increasing the interval to 6000 ms in high-traffic load environments.</p>
Default	The default transmit interval is 2000 milliseconds.
Options	<i>interval</i> —Specifies the transmit interval, in milliseconds.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 147

CHAPTER 13

Operational Commands for Firewall Filters

- `clear firewall`
- `show firewall`
- `show firewall policer`
- `show interfaces filters`

clear firewall

Syntax	<code>clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i>)</code>
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	<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>
Options	<p>all—Clear the packet and byte counts for all firewall filter counters and clear the packet counts for all policer counters.</p> <p>counter <i>counter-name</i>—Clear the packet and byte counts for the specified firewall filter counter.</p> <p>filter <i>filter-name</i>—Clear the packet and byte counts for the specified firewall filter.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• Verifying That Firewall Filters Are Operational on page 38• Verifying That Two-Color Policers Are Operational on page 79• Overview of Firewall Filters on page 3• Overview of Policers on page 59

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

Syntax	show firewall <counter <i>counter-name</i> > <filter <i>filter-name</i> > <log <detail interface <i>interface-name</i> >> <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	<p>counter <i>counter-name</i>—(Optional) Display statistics about a particular firewall filter counter.</p> <p>filter <i>filter-name</i>—(Optional) Display statistics about a particular firewall filter.</p> <p>log—(Optional) Display log entries for all firewall filter activity.</p> <p>terse—(Optional) Display firewall filter names only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Verifying That Firewall Filters Are Operational on page 38 • Verifying That Two-Color Policers Are Operational on page 79 • Overview of Firewall Filters on page 3 • Overview of Policers on page 59
List of Sample Output	show firewall on page 280 show firewall filter <i>filter-name</i> on page 281 show firewall counter <i>counter-name</i> on page 281 show firewall log on page 281 show firewall log detail on page 281
Output Fields	Table 18 on page 279 lists the output fields for the show firewall command. Output fields are listed in the approximate order in which they appear.

Table 18: show firewall Output Fields

Field Name	Field Description	Level of Output
Filter	Name of the filter that is configured at the [edit firewall family <i>family-name</i> filter] hierarchy level.	All levels

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

Field Name	Field Description	Level of Output
Counters	Display filter counter information: <ul style="list-style-type: none"> • Name—Name of a filter counter that has been configured with the count firewall filter action modifier. • Bytes—Number of bytes that match the filter term where the count action modifier was specified. • Packets—Number of packets that matched the filter term where the count action modifier was specified. 	All levels
Policers	Display policer information: <ul style="list-style-type: none"> • Name—Name of the policer that is configured at the [edit firewall policer] hierarchy level. • Packets—Number of packets that matched the filter term where the policer action modifier was specified. This is the number of packets that exceeded the rate limits that the policer specifies. 	All levels
Action	Filter action: <ul style="list-style-type: none"> • A—Accept • D—Discard 	All levels
Interface	Interface on which the firewall filter is applied.	All levels
Protocol	Name of the packet protocol.	All levels
Packet Length	Length of the packet.	All levels
Src Addr	Source address of the packet.	All levels
Dest Addr	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

Syntax	<code>show firewall policer</code> <code><policer-name></code>
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	<ul style="list-style-type: none"> • Verifying That Firewall Filters Are Operational on page 38 • Verifying That Two-Color Policers Are Operational on page 79 • Overview of Firewall Filters on page 3 • Overview of Policers on page 59
List of Sample Output	show firewall policer on page 283 show firewall policer policer-name on page 284
Output Fields	Table 19 on page 283 lists the output fields for the show firewall policer command. Output fields are listed in the approximate order in which they appear.

Table 19: 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"> • Filter—Name of filter that specifies the policer action modifier. • Name—Name of policer. • Packets—Number of packets that matched the filter term in which the policer action modifier is specified. This is the number of packets that exceed the rate limits that the policer specifies. 	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

Syntax	<code>show interfaces filters</code> <code><interface-name></code>
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 firewall filters that are configured on each interface in a switch.
Options	none —Display firewall filter information about all interfaces. interface-name —(Optional) Display firewall filter information about a particular interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show firewall on page 279
List of Sample Output	show interfaces filters on page 285 show interfaces filters interface-name on page 286
Output Fields	Table 20 on page 285 lists the output fields for the show interfaces filters command. Output fields are listed in the approximate order in which they appear.

Table 20: show interfaces filters Output Fields

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

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		

CHAPTER 14

Operational Commands for Media Access Control Security (MACsec)

- `clear security mka statistics`
- `show security macsec connections`
- `show security macsec statistics`
- `show security mka sessions`
- `show security mka statistics`

clear security mka statistics

Syntax	clear security mka statistics <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches. Command introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.
Description	<p>Clear—reset to zero (0)—all MACsec Key Agreement (MKA) protocol statistics.</p> <p>You are clearing the statistics that are viewed using the show security mka statistics when you enter this command.</p>
Options	<p>none—Clear all MKA counters for all interfaces on the switch.</p> <p>interface <i>interface-name</i>—(Optional) Clear MKA traffic counters for the specified interface only.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show security mka statistics on page 297• show security mka sessions on page 295• Understanding Media Access Control Security (MACsec) on page 139

Sample Output

clear security mka statistics

```
user@switch> clear security mka statistics
```

show security macsec connections

Syntax	show security macsec connections <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	Display the status of the active MACsec connections on the switch. This command does not display output when MACsec is enabled using static secure association key (SAK) security mode.
Options	none —Display MACsec connection information for all interfaces on the switch. interface <i>interface-name</i> —(Optional) Display MACsec connection information for the specified interface only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show security macsec statistics on page 291
List of Sample Output	show security macsec connections on page 290
Output Fields	Table 21 on page 289 lists the output fields for the show security macsec connections command. Output fields are listed in the approximate order in which they appear.

Table 21: show security macsec connections Output Fields

Field Name	Field Description
Fields for Interface	
Interface name	Name of the interface.
CA name	<p>Name of the connectivity association.</p> <p>A connectivity association is named using the connectivity-association statement when you are enabling MACsec.</p>
Cipher suite	Name of the cipher suite used for encryption.
Encryption	<p>Encryption setting. Encryption is enabled when this output is on and disabled when this output is off.</p> <p>The encryption setting is set using the no-encryption statement in the connectivity association when using static connectivity association key (CAK) security mode and is set using the encryption statement in the secure channel when using static secure association key (SAK) or dynamic security mode.</p>

Table 21: show security macsec connections Output Fields (*continued*)

Field Name	Field Description
Key server offset	<p>Offset setting.</p> <p>The offset is set using the offset statement when configuring the connectivity association when using static connectivity association key (CAK) or dynamic security mode or the secure channel when using static secure association key (SAK) security mode.</p>
Include SCI	<p>SCI tagging. The SCI tag is included on packets in a secure channel when this output is yes, and not included on packets in a secure channel when this output is no.</p> <p>You can enable SCI tagging using the include-sci statement in the connectivity association.</p> <p>NOTE: SCI tags are automatically appended to packets leaving a MACsec-enabled interface on an EX4300 switch. The include-sci option is, therefore, not available on EX4300 switches. The output for the Include SCI field is yes.</p>
Replay protect	<p>Replay protection setting. Replay protection is enabled when this output is on and disabled when this output is off.</p> <p>You can enable replay protection using the replay-protect statement in the connectivity association.</p>
Replay window	<p>Replay protection window setting. This output is set to 0 when replay protection is disabled, and is the size of the replay window, in number of packets, when replay protection is enabled.</p> <p>The size of the replay window is configured using the replay-window-size statement in the connectivity association.</p>

Sample Output

show security macsec connections

```

user@host> show security macsec connections
Interface name: xe-0/1/0
  CA name: CA1
  Cipher suite: GCM-AES-128   Encryption: on
  Key server offset: 0        Include SCI: no
  Replay protect: off         Replay window: 0

```

show security macsec statistics

Syntax show security macsec statistics
<brief | detail>
<interface *interface-name*>

Release Information Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches.

Description Display Media Access Control Security (MACsec) statistics.

This command does not display output when MACsec is enabled using static secure association key (SAK) security mode.

Options **none**—Display MACsec statistics in brief form for all interfaces on the switch.

brief | detail—(Optional) Display the specified level of output. Using the **brief** option is equivalent to entering the command with no options (the default). The **detail** option displays additional fields that are not visible in the **brief** output.



NOTE: The field names that only appear in this command output when you enter the **detail** option are mostly useful for debugging purposes by Juniper Networks support personnel.

interface *interface-name*—(Optional) Display MACsec statistics for the specified interface only.

Required Privilege Level view

Related Documentation • [show security macsec connections on page 289](#)

List of Sample Output [show security macsec statistics interface xe-0/1/0 detail on page 293](#)

Output Fields [Table 22 on page 291](#) lists the output fields for the **show security macsec statistics** command. Output fields are listed in the approximate order in which they appear.

The field names that appear in this command output only when you enter the **detail** option are mostly useful for debugging purposes by Juniper Networks support personnel. Those field names are, therefore, not included in this table.

Table 22: show security macsec statistics Output Fields

Field Name	Field Description	Level of Output
Interface name	Name of the interface.	All levels

Fields for Secure Channel transmitted

Table 22: show security macsec statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Encrypted packets	<p>Total number of packets transmitted out of the interface in the secure channel that were secured and encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p>	All levels
Encrypted bytes	<p>Total number of bytes transmitted out of the interface in the secure channel that were secured and encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p>	All levels
Protected packets	<p>Total number of packets transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p>	All levels
Protected bytes	<p>Total number of bytes transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p>	All levels
Fields for Secure Association transmitted		
Encrypted packets	<p>Total number of packets transmitted out of the interface in the connectivity association that were secured and encrypted using MACsec.</p> <p>The total includes the data packets transmitted in the secure channel and secured using a SAK and the control packets secured using a connectivity association key (CAK).</p>	All levels
Protected packets	<p>Total number of packets transmitted out of the interface in the connectivity association that were secured but not encrypted using MACsec.</p> <p>The total includes the data packets transmitted in the secure channel and secured using a SAK and the control packets secured using a connectivity association key (CAK).</p>	All levels
Fields for Secure Channel received		
Accepted packets	<p>The number of received packets that have been accepted by the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link.</p> <p>A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check.</p> <p>This counter increments for traffic that is and is not encrypted using MACsec.</p>	All levels

Table 22: show security macsec statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Validated bytes	<p>The number of bytes that have been validated by the MACsec integrity check and received on the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link.</p> <p>This counter does not increment when MACsec encryption is disabled.</p>	All levels
Decrypted bytes	<p>The number of bytes received in the secure channel on the interface that have been decrypted. The secure channel is used to send all data plane traffic on a MACsec-enabled link.</p> <p>An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.</p>	All levels
Fields for Secure Association received		
Accepted packets	<p>The number of received packets that have been accepted in the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface.</p> <p>A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check.</p>	All levels
Validated bytes	<p>The number of bytes that have been validated by the MACsec integrity check and received on the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface.</p> <p>This counter does not increment when MACsec encryption is disabled.</p>	All levels
Decrypted bytes	<p>The number of bytes received in the connectivity association on the interface that have been decrypted. The counter includes all control and data plane traffic accepted on the interface.</p> <p>An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.</p>	All levels

Sample Output

show security macsec statistics interface xe-0/1/0 detail

```
user@host> show security macsec statistics interface xe-0/1/0 detail
```

```
Interface name: xe-0/1/0
Secure Channel transmitted
  Encrypted packets: 123858
  Encrypted bytes:   32190903
  Protected packets: 0
  Protected bytes:   0
Secure Association transmitted
```

```
    Encrypted packets: 123858
    Protected packets: 0
Secure Channel received
    Accepted packets: 123877
    Validated bytes: 0
    Decrypted bytes: 32196238
Secure Association received
    Accepted packets: 123877
    Validated bytes: 0
    Decrypted bytes: 32196238
Error and debug
Secure Channel transmitted packets
    Untagged: 0, Too long: 0
Secure Channel received packets
    Control: 0, Tagged miss: 3202804
    Untagged hit: 0, Untagged: 0
    No tag: 0, Bad tag: 0
    Unknown SCI: 0, No SCI: 0
    Control pass: 0, Control drop: 0
    Uncontrol pass: 123877, Uncontrol drop: 0
    Hit dropped: 0, Invalid accept: 0
    Late drop: 0, Delayed accept: 0
    Unchecked: 0, Not valid drop: 0
    Not using SA drop: 0, Unused SA accept: 0
```

show security mka sessions

Syntax	show security mka sessions <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	Display MACsec Key Agreement (MKA) session information.
Options	<ul style="list-style-type: none"> interface <i>interface-name</i>—(Optional) Display the MKA session information for the specified interface only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show security mka statistics on page 297 show security macsec connections on page 289 show security macsec statistics on page 291
List of Sample Output	show security mka sessions on page 296
Output Fields	Table 23 on page 295 lists the output fields for the show security mka sessions command. Output fields are listed in the approximate order in which they appear.

Table 23: show security mka sessions Output Fields

Field Name	Field Description
Interface name	Name of the interface.
Member identifier	Name of the member identifier.
CAK name	Name of the Connectivity Association Key (CAK). The CAK is configured using the cak keyword when configuring the pre-shared key.
Transmit interval	The transmit interval.
Outbound SCI	Name of the outbound secure channel identifier.
Message number	Number of the last data message.
Key number	Key number.
Key server	Key server status. The switch is the key server when this output is yes . The switch is not the key server when this output is no .

Table 23: show security mka sessions Output Fields (*continued*)

Field Name	Field Description
Key server priority	The key server priority. The key server priority can be set using the key-server-priority statement.
Latest SAK AN	Name of the latest secure association key (SAK) association number.
Latest SAK KI	Name of the latest secure association key (SAK) key identifier.
Fields for Peer list	
Member identifier	Name of the member identifier.
Hold time	Hold time, in seconds.
Message number	Number of the last data message
SCI	Name of the secure channel identifier.
Lowest acceptable PN	Number of the lowest acceptable packet number (PN).

Sample Output

show security mka sessions

```
user@host> show security mka sessions
```

```
Interface name: xe-0/1/0
Member identifier: 0CCBEE42F8778300F8D0C1DC
CAK name: 1234567890
Transmit interval: 2000(ms)
Outbound SCI: 2C:6B:F5:9D:4B:1B/1
Message number: 1526465    Key number: 0
Key server: no            Key server priority: 15
Latest SAK AN: 0          Latest SAK KI: 4F18CE25228178FD15976E4C/1
Previous SAK AN: 0        Previous SAK KI: 000000000000000000000000/0
Peer list
1. Member identifier: 4F18CE25228178FD15976E4C (live)
   Message number: 1526484 Hold time: 14500 (ms)
   SCI: 2C:6B:F5:9D:3A:1B/1
   Lowest acceptable PN: 121198
```

show security mka statistics

Syntax	show security mka statistics <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Description	Display MACsec Key Agreement (MKA) protocol statistics. The output for this command does not include statistics for MACsec data traffic. For MACsec data traffic statistics, see show security macsec statistics .
Options	<ul style="list-style-type: none"> interface <i>interface-name</i>—(Optional) Display the MKA information for the specified interface only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show security mka sessions on page 295 show security macsec statistics on page 291 show security macsec connections on page 289
List of Sample Output	show security mka statistics on page 298
Output Fields	Table 24 on page 297 lists the output fields for the show security mka statistics command. Output fields are listed in the approximate order in which they appear.

Table 24: show security mka statistics Output Fields

Field Name	Field Description
Received packets	<p>Number of received MKA control packets.</p> <p>This counter increments for received MKA control packets only. This counter does not increment when data packets are received.</p>
Transmitted packets	<p>Number of transmitted MKA packets</p> <p>This counter increments for transmitted MKA control packets only. This counter does not increment when data packets are transmitted.</p>
Version mismatch packets	Number of version mismatch packets.
CAK mismatch packets	<p>Number of Connectivity Association Key (CAK) mismatch packets.</p> <p>This counter increments when the connectivity association key (CAK) and connectivity association key name (CKN), which are user-configured values that have to match to enable MACsec, do not match for an MKA control packet.</p>

Table 24: show security mka statistics Output Fields (*continued*)

Field Name	Field Description
ICV mismatch packets	Number of ICV mismatched packets. This counter increments when the connectivity association key (CAK) value does not match on both ends of a MACsec-secured Ethernet link.
Duplicate message identifier packets	Number of duplicate message identifier packets.
Duplicate message number packets	Number of duplicate message number packets.
Duplicate address packets	Number of duplicate source MAC address packets.
Invalid destination address packets	Number of invalid destination MAC address packets.
Formatting error packets	Number of formatting error packets.
Old Replayed message number packets	Number of old replayed message number packets.

Sample Output

show security mka statistics

```
user@host> show security mka statistics
```

```

Received packets:          1525844
Transmitted packets:      1525841
Version mismatch packets: 0
CAK mismatch packets:     0
ICV mismatch packets:     0
Duplicate message identifier packets: 0
Duplicate message number packets: 0
Duplicate address packets: 0
Invalid destination address packets: 0
Formatting error packets:  0
Old Replayed message number packets: 0
```

CHAPTER 15

Operational Commands for Port Security

- clear arp inspection statistics
- clear dhcp snooping binding
- clear ethernet-switching port-error
- show arp inspection statistics
- show dhcp snooping binding

clear arp inspection statistics

Syntax	clear arp inspection statistics <interface <i>interface</i> >
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear ARP inspection statistics.
Options	none —Clears ARP statistics on all interfaces. interface <i>interface-names</i> —(Optional) Clear ARP statistics on one or more interfaces.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show arp inspection statistics on page 303• <i>Example: Configuring Basic Port Security Features</i>• Verifying That DAI Is Working Correctly on page 109
List of Sample Output	clear arp inspection statistics on page 300
Output Fields	This command produces no output.

Sample Output

clear arp inspection statistics

```
user@switch> clear arp inspection statistics
```


clear dhcp snooping binding

Syntax	clear dhcp snooping binding <mac (all <i>mac-address</i>)> <vlan (all <i>vlan-name</i>)> <vlan (all <i>vlan-name</i>) mac (all <i>mac-address</i>)>
Release Information	Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear the DHCP snooping database information.
Options	<p>mac (all <i>mac-address</i>)—(Optional) Clear DHCP snooping information for the specified MAC address or all MAC addresses.</p> <p>vlan (all <i>vlan-name</i>)—(Optional) Clear DHCP snooping information for the specified VLAN or all VLANs.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Basic Port Security Features</i> • show dhcp snooping binding on page 304
List of Sample Output	clear dhcp snooping binding on page 301
Output Fields	This command produces no output.

Sample Output

clear dhcp snooping binding

```
user@switch> clear dhcp snooping binding
```

clear ethernet-switching port-error

Syntax	clear ethernet-switching port-error <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch or from the specified interface, and restore the interfaces or the specified interface to service.
Options	none —Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch and restore the interfaces to service. interface <i>interface-name</i> —(Optional) Clear all MAC limiting, MAC move limiting, and storm control errors from the specified interface and restore the interface to service.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• <i>Configuring MAC Limiting</i>• <i>Example: Configuring Storm Control to Prevent Network Outages</i>• <i>Configuring Port Security (CLI Procedure)</i>• port-error-disable on page 206• <i>Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure)</i>
Output Fields	This command produces no output.

show arp inspection statistics

Syntax	show arp inspection statistics
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display ARP inspection statistics.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear arp inspection statistics on page 300 • <i>Example: Configuring Basic Port Security Features</i> • Verifying That DAI Is Working Correctly on page 109
List of Sample Output	show arp inspection statistics on page 303
Output Fields	Table 25 on page 303 lists the output fields for the show arp inspection statistics command. Output fields are listed in the approximate order in which they appear.

Table 25: show arp inspection statistics Output Fields

Field Name	Field Description	Level of Output
Interface	Interface on which ARP inspection has been applied.	All levels
Packets received	Total number of packets total that underwent ARP inspection.	All levels
ARP inspection pass	Total number of packets that passed ARP inspection.	All levels
ARP inspection failed	Total number of packets that failed ARP inspection.	All levels

Sample Output

show arp inspection statistics

```
user@switch> show arp inspection statistics
```

Interface	Packets received	ARP inspection pass	ARP inspection failed
-----	-----	-----	-----
ge-0/0/0	0	0	0
ge-0/0/1	0	0	0
ge-0/0/2	0	0	0
ge-0/0/3	0	0	0
ge-0/0/4	0	0	0
ge-0/0/5	0	0	0
ge-0/0/6	0	0	0
ge-0/0/7	703	701	2

show dhcp snooping binding

Syntax	show dhcp snooping binding <interface <i>interface-name</i>> <vlan <i>vlan-name</i>>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display the DHCP snooping database information.
Options	interface <i>interface-name</i> —(Optional) Display the DHCP snooping database information for an interface. vlan <i>vlan-name</i> —(Optional) Display the DHCP snooping database information for a VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear dhcp snooping binding • Example: Configuring Basic Port Security Features • Verifying That DHCP Snooping Is Working Correctly on page 99
List of Sample Output	show dhcp snooping binding on page 304
Output Fields	Table 26 on page 304 lists the output fields for the show dhcp snooping binding command. Output fields are listed in the approximate order in which they appear.

Table 26: show dhcp snooping binding Output Fields

Field Name	Field Description	Level of Output
MAC Address	MAC address of the network device; bound to the IP address.	All levels
IP Address	IP address of the network device; bound to the MAC address.	All levels
Lease	Lease granted to the IP address.	All levels
Type	How the MAC address was acquired.	All levels
VLAN	VLAN name of the network device whose MAC address is shown.	All levels
Interface	Interface address (port).	All levels

Sample Output

show dhcp snooping binding

```
user@switch> show dhcp snooping binding
```

DHCP Snooping Information:

MAC Address	IP Address	Lease	Type	VLAN	Interface
00:00:01:00:00:03	192.0.2.0	640	dynamic	guest	ge-0/0/12.0
00:00:01:00:00:04	192.0.2.1	720	dynamic	guest	ge-0/0/12.0
00:00:01:00:00:05	192.0.2.5	800	dynamic	guest	ge-0/0/13.0

