

# FastIron Ethernet Switch Traffic Management Configuration Guide

Supporting FastIron Software Release 8.0.30h

© 2016, Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, Brocade Assurance, the B-wing symbol, ClearLink, DCX, Fabric OS, HyperEdge, ICX, MLX, MyBrocade, OpenScript, VCS, VDX, Vplane, and Vyatta are registered trademarks, and Fabric Vision is a trademark of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned may be trademarks of others.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

The authors and Brocade Communications Systems, Inc. assume no liability or responsibility to any person or entity with respect to the accuracy of this document or any loss, cost, liability, or damages arising from the information contained herein or the computer programs that accompany it.

The product described by this document may contain open source software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

# Contents

---

<b>Preface</b> .....	<b>7</b>
Document conventions.....	7
Text formatting conventions.....	7
Command syntax conventions.....	7
Notes, cautions, and warnings.....	8
Brocade resources.....	8
Contacting Brocade Technical Support.....	8
Brocade customers.....	8
Brocade OEM customers.....	9
Document feedback.....	9
<b>About This Document</b> .....	<b>11</b>
Supported hardware.....	11
Supported Hardware.....	11
What's new in this document .....	11
How command information is presented in this guide.....	12
<b>Quality of Service</b> .....	<b>13</b>
QoS overview.....	13
Processing of classified traffic.....	13
QoS for Brocade stackable devices.....	20
QoS profile restrictions in a traditional stack.....	21
QoS behavior for trusting Layer 2 (802.1p) in a traditional stack.....	21
QoS behavior for trusting Layer 3 (DSCP) in a traditional stack .....	21
QoS behavior on port priority and VLAN priority in a traditional stack.....	21
QoS behavior for 802.1p marking in a traditional stack.....	21
QoS queues.....	21
Queues for the SX-FI48GPP interface module.....	22
Queues for the SX-FI-8XG interface module.....	22
Queues for the ICX 6430 switch.....	23
User-configurable scheduler profile.....	25
QoS priorities-to-traffic assignment.....	28
Changing a port priority.....	28
Assigning static MAC entries to priority queues.....	28
Buffer allocation and threshold for QoS queues.....	29
802.1p priority override.....	29
Configuration notes and feature limitations .....	29
Enabling 802.1p priority override.....	29
Marking.....	30
DSCP and CoS global remarking.....	30
Configuration considerations and limitations.....	30
Enabling DSCP marking.....	31
Example: DSCP marking.....	31
Enabling CoS marking.....	31
Example: CoS marking.....	31
DSCP-based QoS configuration.....	31
Application notes for DSCP-based QoS.....	32

Using ACLs to honor DSCP-based QoS.....	32
Trust DSCP for the SX-FI48GPP, SX-FI-24GPP, SX-FI-24HF, SX-FI-2XG, and SX-FI-8XG modules.....	33
Configuring QoS mapping configuration.....	33
Default DSCP to internal forwarding priority mappings.....	33
Changing the DSCP to internal forwarding priority mappings.....	35
Changing the VLAN priority 802.1p to hardwareforwarding queue mappings.....	36
Default scheduling configuration for the SX-FI48GPP module.....	36
Default scheduling configuration for the ICX 6430.....	37
Scheduling QoS information.....	37
Scheduling for the SX-FI48GPP module.....	37
QoS queuing methods.....	38
Selecting the QoS queuing method.....	39
Configuring the QoS queues.....	39
Viewing QoS settings.....	41
Viewing DSCP-based QoS settings.....	42
IPv6 Quality of Service .....	44
<b>Rate Limiting and Rate Shaping on FastIron X Series and FCX and ICX Series Switches.....</b>	<b>45</b>
Rate limiting overview.....	45
Rate limiting in hardware.....	45
How Fixed rate limiting works.....	46
Configuration notes for rate limiting.....	46
Configuring a port-based rate limiting policy.....	47
Configuring an ACL-based rate limiting policy.....	47
Displaying the fixed rate limiting configuration.....	48
Rate shaping overview.....	48
Configuration notes for rate shaping.....	49
Configuring outbound rate shaping for a port.....	50
Configuring outbound rate shaping for a specific priority.....	50
Configuring outbound rate shaping for a LAG port.....	50
Displaying rate shaping configurations.....	51
<b>Limiting broadcast, unknown unicast, and multicast traffic.....</b>	<b>53</b>
Configuration notes and feature limitations.....	53
Command syntax for packet-based limiting .....	53
On FastIron X-Series devices.....	53
On Brocade FCX Series, ICX 6430 , and ICX 6650 devices.....	54
On Brocade ICX 6610 and ICX 6450 devices.....	54
Command syntax for byte-based limiting.....	55
Viewing broadcast, multicast, and unknown unicast limits.....	55
BUM suppression port dampening.....	58
BUM suppression port dampening restrictions and limitations.....	58
Enabling BUM suppression port dampening.....	58
<b>Traffic Policies.....</b>	<b>61</b>
Traffic policies overview.....	61
Configuration notes and feature limitations for traffic policies.....	61
Maximum number of traffic policies supported on a device.....	62
Setting the maximum number of traffic policies supported on a Layer 3 device.....	63
ACL-based rate limiting using traffic policies.....	63
Support for fixed rate limiting and adaptive rate limiting.....	64
Configuring ACL-based fixed rate limiting.....	64

ACL-based adaptive rate limiting configuration.....	65
Specifying the action to be taken for packets that are over the limit.....	67
ACL statistics and rate limit counting.....	69
Enabling ACL statistics.....	69
Enabling ACL statistics with rate limiting traffic policies.....	70
Viewing ACL and rate limit counters.....	71
Clearing ACL and rate limit counters.....	72
Viewing traffic policies.....	72
CPU rate limiting.....	74
<b>ICX 7750 Flow Control and Buffer Management.....</b>	<b>75</b>
Priority flow control .....	75
Configuring priority flow control.....	76
Configuring priority flow control on a port.....	76
Configuring priority flow control on a range of ports.....	77
Packet buffer management.....	77



# Preface

---

- Document conventions..... 7
- Brocade resources..... 8
- Contacting Brocade Technical Support..... 8
- Document feedback..... 9

## Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Brocade technical documentation.

## Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used in the flow of the text to highlight specific words or phrases.

Format	Description
<b>bold text</b>	Identifies command names Identifies keywords and operands Identifies the names of user-manipulated GUI elements
<i>italic text</i>	Identifies text to enter at the GUI Identifies emphasis Identifies variables
Courier font	Identifies document titles Identifies CLI output Identifies command syntax examples

## Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
<b>bold text</b>	Identifies command names, keywords, and command options.
<i>italic text</i>	Identifies a variable.
value	In Fibre Channel products, a fixed value provided as input to a command option is printed in plain text, for example, <b>--show</b> WWN.
[ ]	Syntax components displayed within square brackets are optional.
{ x   y   z }	Default responses to system prompts are enclosed in square brackets. A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x   y	In Fibre Channel products, square brackets may be used instead for this purpose. A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.

Convention	Description
...	Repeat the previous element, for example, <i>member{member...}</i> .
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

## Notes, cautions, and warnings

Notes, cautions, and warning statements may be used in this document. They are listed in the order of increasing severity of potential hazards.

### NOTE

A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

### ATTENTION

An Attention statement indicates a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.



### CAUTION

A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.



### DANGER

*A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.*

## Brocade resources

Visit the Brocade website to locate related documentation for your product and additional Brocade resources.

You can download additional publications supporting your product at [www.brocade.com](http://www.brocade.com). Select the Brocade Products tab to locate your product, then click the Brocade product name or image to open the individual product page. The user manuals are available in the resources module at the bottom of the page under the Documentation category.

To get up-to-the-minute information on Brocade products and resources, go to [MyBrocade](http://MyBrocade). You can register at no cost to obtain a user ID and password.

Release notes are available on [MyBrocade](http://MyBrocade) under Product Downloads.

White papers, online demonstrations, and data sheets are available through the [Brocade website](http://Brocade website).

## Contacting Brocade Technical Support

As a Brocade customer, you can contact Brocade Technical Support 24x7 online, by telephone, or by e-mail. Brocade OEM customers contact their OEM/Solutions provider.

## Brocade customers

For product support information and the latest information on contacting the Technical Assistance Center, go to <http://www.brocade.com/services-support/index.html>.

If you have purchased Brocade product support directly from Brocade, use one of the following methods to contact the Brocade Technical Assistance Center 24x7.



Online	Telephone	E-mail
Preferred method of contact for non-urgent issues: <ul style="list-style-type: none"> <li>• <a href="#">My Cases</a> through MyBrocade</li> <li>• <a href="#">Software downloads</a> and licensing tools</li> <li>• <a href="#">Knowledge Base</a></li> </ul>	Required for Sev 1-Critical and Sev 2-High issues: <ul style="list-style-type: none"> <li>• Continental US: 1-800-752-8061</li> <li>• Europe, Middle East, Africa, and Asia Pacific: +800-AT FIBREE (+800 28 34 27 33)</li> <li>• For areas unable to access toll free number: +1-408-333-6061</li> <li>• <a href="#">Toll-free numbers</a> are available in many countries.</li> </ul>	<a href="mailto:support@brocade.com">support@brocade.com</a> Please include: <ul style="list-style-type: none"> <li>• Problem summary</li> <li>• Serial number</li> <li>• Installation details</li> <li>• Environment description</li> </ul>

## Brocade OEM customers

If you have purchased Brocade product support from a Brocade OEM/Solution Provider, contact your OEM/Solution Provider for all of your product support needs.

- OEM/Solution Providers are trained and certified by Brocade to support Brocade® products.
- Brocade provides backline support for issues that cannot be resolved by the OEM/Solution Provider.
- Brocade Supplemental Support augments your existing OEM support contract, providing direct access to Brocade expertise. For more information, contact Brocade or your OEM.
- For questions regarding service levels and response times, contact your OEM/Solution Provider.

## Document feedback

To send feedback and report errors in the documentation you can use the feedback form posted with the document or you can e-mail the documentation team.

Quality is our first concern at Brocade and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. You can provide feedback in two ways:

- Through the online feedback form in the HTML documents posted on [www.brocade.com](http://www.brocade.com).
- By sending your feedback to [documentation@brocade.com](mailto:documentation@brocade.com).

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.



# About This Document

- Supported hardware..... 11
- What's new in this document ..... 11
- How command information is presented in this guide..... 12

## Supported hardware

### Supported Hardware

This guide supports the following product families from Brocade:

- FCX Series
- FastIron X Series (FSX 800 and FSX 1600)
- ICX 6610 Series
- ICX 6430 Series (ICX 6430, ICX 6430-C12)
- ICX 6450 Series (ICX 6450, ICX 6450-C12-PD)
- ICX 6650 Series
- ICX 7750 Series
- ICX 7450 Series
- ICX 7250 Series

For information about the specific models and modules supported in a product family, refer to the hardware installation guide for that product family.

#### NOTE

The Brocade ICX 6430-C switch supports the same feature set as the Brocade ICX 6430 switch unless otherwise noted.

#### NOTE

The Brocade ICX 6450-C12-PD switch supports the same feature set as the Brocade ICX 6450 switch unless otherwise noted.

## What's new in this document

**TABLE 1** Summary of enhancements in FastIron release 08.0.30h

Feature	Description	Described in
BUM suppression port dampening	Rate limiting of BUM traffic is used to protect a switch, router node, or network from Denial of Service (DoS) attacks or unintentional excess traffic conditions. If the ingress BUM traffic exceeds the configured rate limit value, the excess traffic is dropped. If the traffic drop count exceeds a set number within a set time interval, the port is shutdown (dampened) for a user configured period.	<a href="#">BUM suppression port dampening</a> on page 58

## How command information is presented in this guide

For all new content supported in FastIron Release 08.0.20 and later, command information is documented in a standalone command reference guide.

In an effort to provide consistent command line interface (CLI) documentation for all products, Brocade is in the process of completing a standalone command reference for the FastIron platforms. This process involves separating command syntax and parameter descriptions from configuration tasks. Until this process is completed, command information is presented in two ways:

- For all new content supported in FastIron Release 08.0.20 and later, the CLI is documented in separate command pages included in the *FastIron Command Reference*. Command pages are compiled in alphabetical order and follow a standard format to present syntax, parameters, usage guidelines, examples, and command history.

### NOTE

Many commands from previous FastIron releases are also included in the command reference.

- Legacy content in configuration guides continues to include command syntax and parameter descriptions in the chapters where the features are documented.

If you do not find command syntax information embedded in a configuration task, refer to the *FastIron Command Reference*.

# Quality of Service

---

• QoS overview.....	13
• QoS for Brocade stackable devices.....	20
• QoS queues.....	21
• QoS priorities-to-traffic assignment.....	28
• 802.1p priority override.....	29
• Marking.....	30
• DSCP and CoS global remarking.....	30
• DSCP-based QoS configuration.....	31
• Configuring QoS mapping configuration.....	33
• Scheduling QoS information.....	37
• Viewing QoS settings.....	41
• Viewing DSCP-based QoS settings.....	42
• IPv6 Quality of Service .....	44

## QoS overview

Quality of Service (QoS) features are used to prioritize the use of bandwidth in a switch. When QoS features are enabled, traffic is classified as it arrives at the switch, and processed through on the basis of configured priorities. Traffic can be dropped, prioritized for guaranteed delivery, or subject to limited delivery options as configured by a number of different mechanisms.

This chapter describes how QoS is implemented and configured in FastIron devices.

*Classification* is the process of selecting packets on which to perform QoS, reading the QoS information, and assigning a priority to the packets. The classification process assigns a priority to packets as they enter the switch. These priorities can be determined on the basis of information contained within the packet or assigned to the packet as it arrives at the switch. Once a packet or traffic flow is classified, it is mapped to a forwarding priority queue.

Packets on Brocade devices are classified in up to eight traffic classes with values from 0 to 7. Packets with higher priority classifications are given a precedence for forwarding.

## Processing of classified traffic

The *trust level* in effect on an interface determines the type of QoS information the device uses for performing QoS. The Brocade device establishes the trust level based on the configuration of various features and whether the traffic is switched or routed. The trust level can be one of the following:

- Ingress port default priority.
- Static MAC address.
- Layer 2 Class of Service (CoS) value - This is the 802.1p priority value in the Ethernet frame. It can be a value from 0 through 7. The 802.1p priority is also called the *Class of Service*.
- Layer 3 Differentiated Services Code Point (DSCP) - This is the value in the six most significant bits of the IP packet header 8-bit DSCP field. It can be a value from 0 through 63. These values are described in RFCs 2472 and 2475. The DSCP value is sometimes called the *DiffServ value*. The device automatically maps the DSCP value of a packet to a hardware forwarding queue. Refer to [Viewing QoS settings](#) on page 41.
- ACL keyword - An ACL can also prioritize traffic and mark it before sending it along to the next hop. This is described under "QoS options for IP ACLs" section in the *FastIron Ethernet Switch Security Configuration Guide*.

Given the variety of different criteria, there are many possibilities for traffic classification within a stream of network traffic. For this reason, the priority of packets must be resolved based on which criteria takes precedence. Precedence follows the schemes illustrated in the *Determining a packet trust level - FSX devices* through *Determining a packet trust level - FCX, and ICX devices* figures.

## Determining the trust level of a packet

Packet trust level is determined differently on FSX devices than on FCX, and ICX series devices.

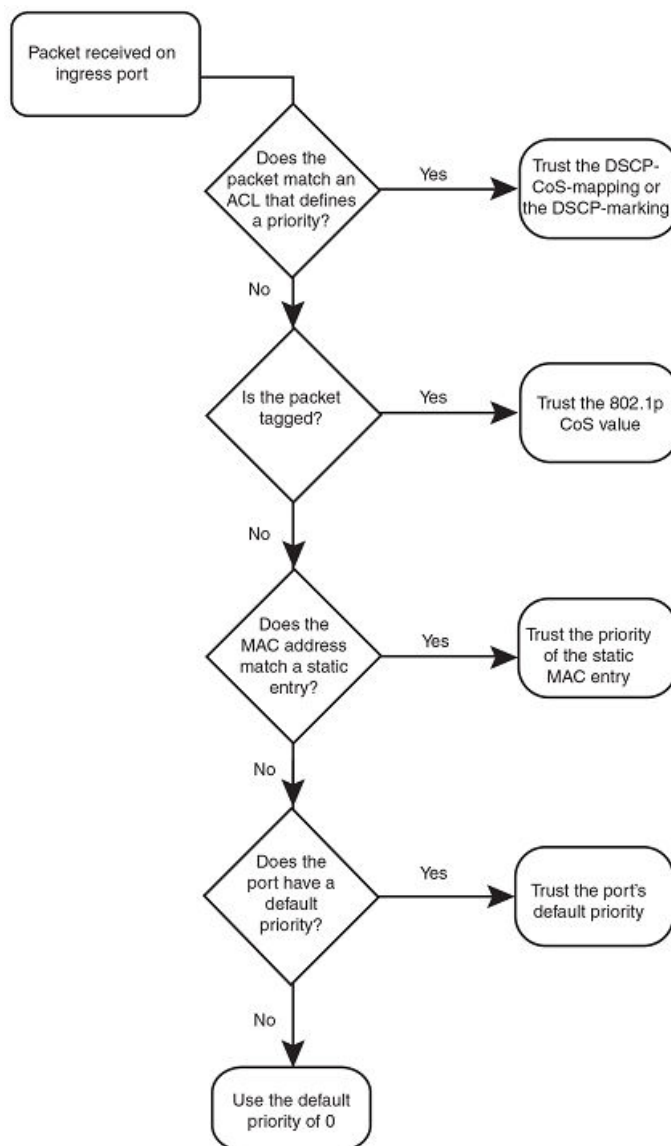
The following figure illustrates how FSX devices determine the trust level of a packet.

### NOTE

The *Determining a packet trust level - FSX devices* figure is not applicable to the third generation FSX interface modules. To determine the trust level of a packet for the SX-FI48GPP, SX-FI-24GPP, SX-FI-24HF, SX-FI-2XG, and SX-FI-8XG modules, refer to the *Determining a packet trust level - SX-FI48GPP, SX-FI-24GPP, SX-FI-24HF, SX-FI-2XG, and SX-FI-8XG modules* figure.

As shown in the flowchart, the first criteria considered is whether the packet matches on an ACL that defines a priority. Next, it checks if trust DSCP is enabled on the port. If this is not the case, the packet is next classified based on the static MAC address. If this is not true and the packet is tagged, the packet is classified with the 802.1p CoS value. If none of these is true, the packet is next classified based on the ingress port default priority or the default priority of zero (0).

FIGURE 1 Determining a packet trust level - FSX devices

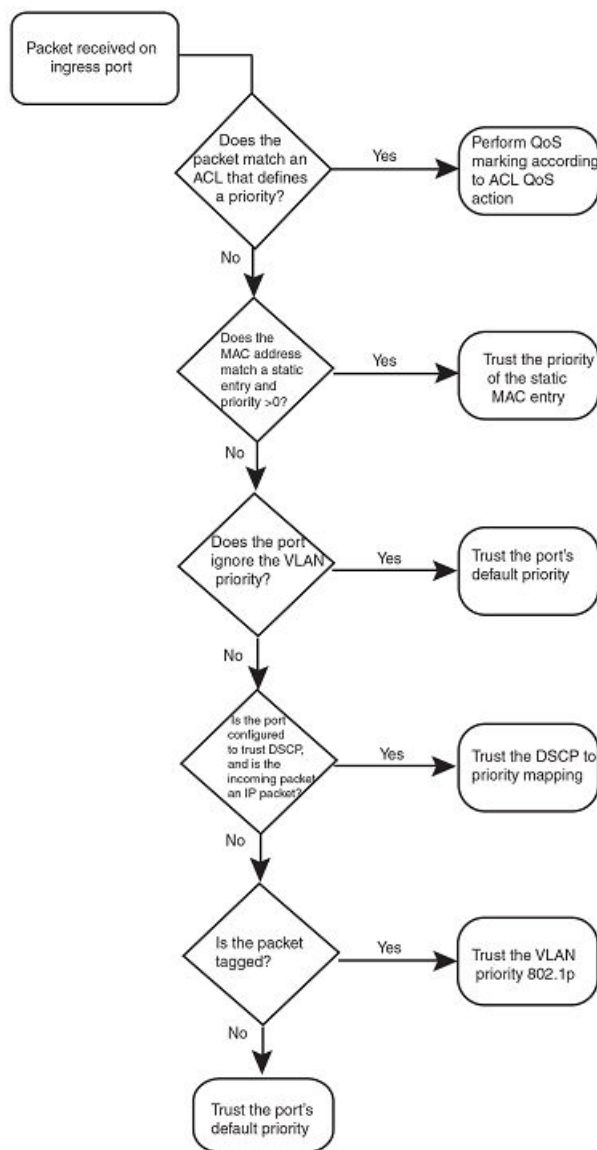


The *Determining a packet trust level - SX-FI48GPP, SX-FI-24GPP, SX-FI-24HF, SX-FI-2XG, and SX-FI-8XG modules* figure illustrates how the SX-FI48GPP, SX-FI-24GPP, SX-FI-24HF, SX-FI-2XG, and SX-FI-8XG modules determine the trust level of a packet. The marking process for these modules is similar to the marking process for other FastIron SX modules. However, there are major differences between these modules and other FastIron SX modules.

- For the SX-FI48GPP, SX-FI-24GPP, SX-FI-24HF, SX-FI-2XG, and SX-FI-8XG modules, static MAC priority takes higher precedence than VLAN priority. For other FastIron SX modules, VLAN priority takes higher precedence over static MAC priority.
- For other FastIron SX modules, the priority of the dynamically learned MAC address is inherited from the default port priority. For the SX-FI48GPP, SX-FI-24GPP, SX-FI-24HF, SX-FI-2XG, and SX-FI-8XG modules, the priority of the dynamically learned MAC address is not inherited from the default port priority because it is not desirable to allow the port priority to take

precedence over the VLAN priority. All dynamically learned MAC addresses are assigned a priority of 0 in the SX-FI48GPP, SX-FI-24GPP, SX-FI-24HF, SX-FI-2XG, and SX-FI-8XG modules. Therefore, configuring a static MAC with a priority of 0 has no effect on QoS marking.

**FIGURE 2** Determining a packet trust level - SX-FI48GPP, SX-FI-24GPP, SX-FI-24HF, SX-FI-2XG, and SX-FI-8XG modules



The following figure illustrates how FCX, and ICX series devices determine the trust level of a packet. As shown in the flowchart, the first criteria considered is whether the packet matches on an ACL that defines a priority. If this is not the case and the MAC address of the packet matches a static entry, the packet is classified with the priority of the static MAC entry. If neither of these is true, the packet is next classified with the ingress port default priority, then DSCP/ToS value, then 802.1p CoS value, and finally the default priority of zero (0).



FIGURE 3 Determining a packet trust level - FCX, and ICX devices

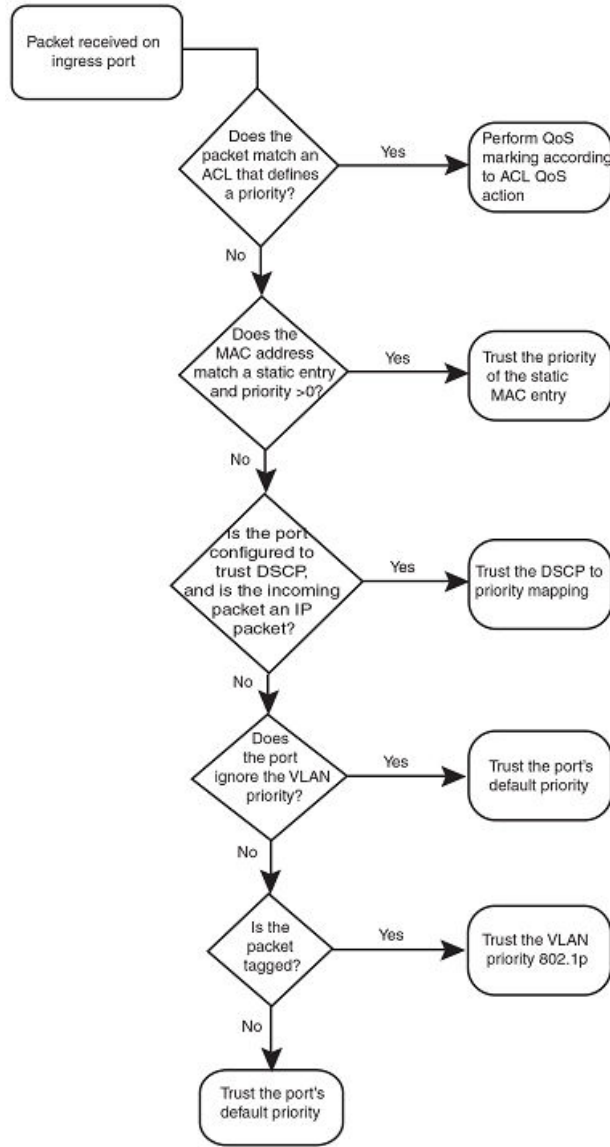
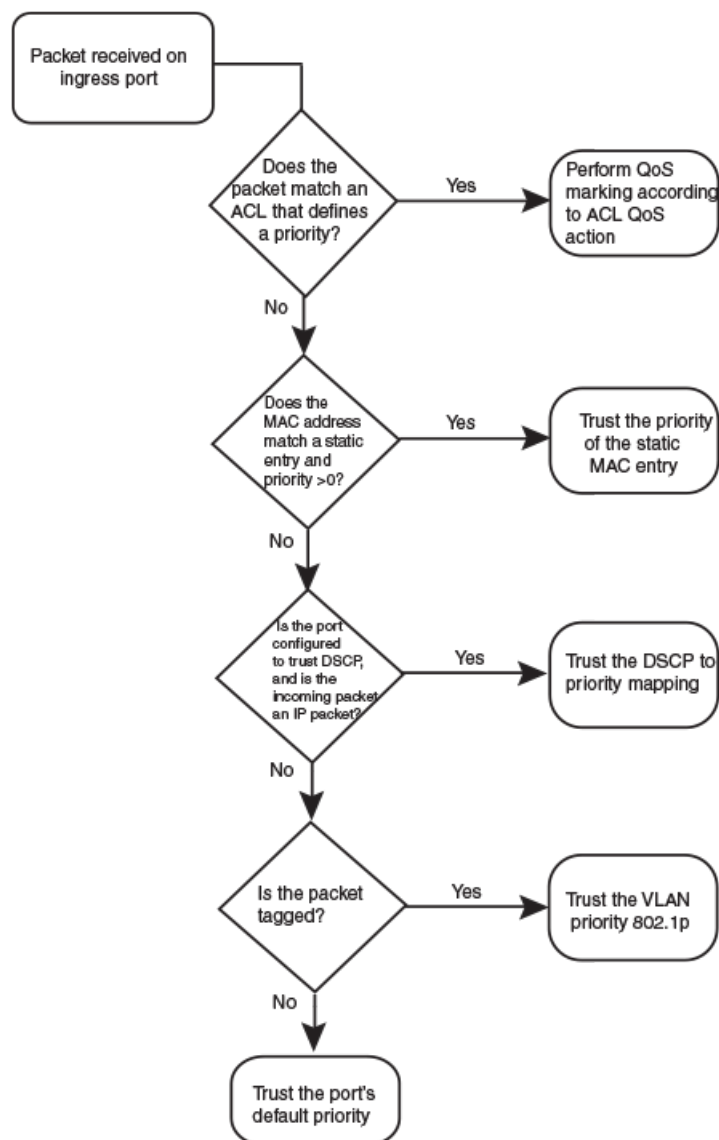


FIGURE 4 Determining the trust level of a packet - ICX 7750 devices



Once a packet is classified, it is mapped to a forwarding queue. For all products except the SX-FI48GPP interface module and ICX 6430 switch, there are eight queues designated from 0 through 7. The internal forwarding priority maps to one of these eight queues. For the SX-FI48GPP interface module and ICX 6430 switch, internal forwarding priority maps to four forwarding queues. The mapping between the internal priority and the forwarding queue cannot be changed.

The following tables show the default QoS mappings for FCX platforms that are used if the trust level for CoS or DSCP is enabled. For information on the SX-FI48GPP interface module, refer to [Queues for the SX-FI48GPP interface module](#) on page 22. For information on default QoS mappings for the ICX 6430 switch, refer to [Queues for the ICX 6430 switch](#) on page 23.

**TABLE 2** Default QoS mappings for FCX platforms, columns 0 to 15

DSCP value	0	1	2	3	4	5	6	7	8	9	10	11	12	12	14	15
802.1p (CoS) value	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
DSCP value	0	1	2	3	4	5	6	7	8	9	10	11	12	12	14	15
Internal forwarding priority	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
Forwarding queue	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1

**TABLE 3** Default QoS mappings for FCX platforms, columns 16 to 31

DSCP value	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
802.1p (CoS) value	2	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3
DSCP value	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Internal forwarding priority	2	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3
Forwarding queue	2	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3

**TABLE 4** Default QoS mappings for FCX platforms, columns 32 to 47

DSCP value	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
802.1p (CoS) value	4	4	4	4	4	4	4	4	5	5	5	5	5	5	5	5
DSCP value	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
Internal forwarding priority	4	4	4	4	4	4	4	4	5	5	5	5	5	5	5	5

**TABLE 4** Default QoS mappings for FCX platforms, columns 32 to 47 (continued)

DSCP value	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
Forwarding queue	4	4	4	4	4	4	4	4	5	5	5	5	5	5	5	5

**TABLE 5** Default QoS mappings for FCX platforms, columns 48 to 63

DSCP value	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
802.1p (CoS) value	6	6	6	6	6	6	6	6	7	7	7	7	7	7	7	7
DSCP value	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
Internal forwarding priority	6	6	6	6	6	6	6	6	7	7	7	7	7	7	7	7
Forwarding queue	6	6	6	6	6	6	6	6	7	7	7	7	7	7	7	7

Mapping between the DSCP value and forwarding queue cannot be changed. However, mapping between DSCP values and other properties can be changed as follows:

- DSCP to internal forwarding priority mapping - You can change the mapping between the DSCP value and the internal forwarding priority value from the default values shown in the above tables. This mapping is used for CoS marking and determining the internal priority when the trust level is DSCP. Refer to *Changing the DSCP to internal forwarding priority mappings*.
- VLAN priority (802.1p) to hardware forwarding queue - You can change the mapping between the 802.1p value and hardware forwarding queue from the default value. Refer to *Changing the VLAN priority 802.1p to hardware forwarding queue mappings*.

## QoS for Brocade stackable devices

Brocade FastIron units in a traditional stack support QoS. Units in a stack communicate the stack topology information and other proprietary control information through the stacking links. For more information about stacking links and traditional stack technology, refer to the *FastIron Ethernet Switch Stacking Configuration Guide*.

In addition to control information, the stacking links also carry user network data packets. In a traditional stack topology, the priority of stacking-specific control packets is elevated above that of data path packets, preventing loss of control packets, and timed retries that affect performance. This prioritization also prevents stack topology changes that may occur if enough stack topology information packets are lost.

Traditional stack technology reserves one QoS profile to provide a higher priority for stack topology and control traffic.

On ICX 7450 stacking devices, Priority7 multicast traffic is not treated as Strict Priority. Multicast queues in ICX 7450 are limited, so Priority 6 and Priority 7 Multicast traffic is mapped to Multicast Queue7. Therefore, even if you configure Priority 7 as Strict Priority, and

Priority 6 as non-Strict, scheduling weight `sched_6_wt+sched_7_wt` is applied on Multicast Queue7 so that Priority 7 traffic is not scheduled as Strict. See the "User-configurable scheduler profile configuration" section for more information on scheduling weights.

## QoS profile restrictions in a traditional stack

In a stacking topology, because CoS level 7 is reserved for stacking, quality profiles for `qosp7` cannot be configured. If an attempt is made to configure a profile for `qosp7`, the system ignores the configuration.

### NOTE

This applies only when the device is operating in stacking mode. It does not apply to standalone devices.

## QoS behavior for trusting Layer 2 (802.1p) in a traditional stack

By default, Layer 2 trust is enabled. Because priority 7 is reserved for stacking control packets, any ingress data traffic with priority 7 is mapped to internal hardware queue 6. All other priorities are mapped to their corresponding queues.

## QoS behavior for trusting Layer 3 (DSCP) in a traditional stack

When the **trust dscp mode** is enabled, packets arriving with DSCP values 56 to 63 are mapped to internal hardware queue 6. All other DSCP values are mapped to their corresponding internal hardware queues.

## QoS behavior on port priority and VLAN priority in a traditional stack

Port priority has a higher precedence than the 802.1p priority examination. If port priority is set to 7, all incoming traffic is mapped to internal hardware queue 6.

When stacking is not enabled on a device, all priorities are mapped to their corresponding queues without restrictions.

## QoS behavior for 802.1p marking in a traditional stack

By default, 802.1p marking is not enabled in a traditional stack. Outgoing tagged traffic will not be marked based on the hardware queue into which ingress traffic was classified. 802.1p marking can be achieved using ACL. For configuration syntax, rules, and examples of QoS marking, refer to the "QoS options for IP ACLs" section in the *FastIron Ethernet Switch Security Configuration Guide*.

## QoS queues

Brocade devices support the eight QoS queues (`qosp0` through `qosp7`) listed in the following table.

**TABLE 6** QoS queues

QoS priority level	QoS queue
0	<code>qosp0</code> (lowest priority queue)
1	<code>qosp1</code>
2	<code>qosp2</code>
3	<code>qosp3</code>
4	<code>qosp4</code>
5	<code>qosp5</code>
6	<code>qosp6</code>
7	<code>qosp7</code> (highest priority queue)

**NOTE**

On ICX 7450 devices, both Priority 6 and Priority 7 traffic maps to one multicast queue (MCQ:7).

The queue names listed in the table are the default names. If desired, you can rename the queues as shown in the “Renaming the queues” section.

Packets are classified and assigned to specific queues based on the criteria shown in the figures described in the “Determining the trust level of a packet” section.

For FCX and ICX devices, ingress packets are classified into the eight priorities, which map to eight hardware queues or traffic classes (TCs) based on the priority. Exceptions to this model are the SX-FI48GPP and SX-FI-8XG interface modules and the ICX 6430 switch as explained in the following sections.

## Queues for the SX-FI48GPP interface module

The SX-FI48GPP interface module consists of two separate hardware Network Processors (NPs). The front-end NP supports four hardware queues, and the back-end NP supports eight hardware queues. Ingress packets are classified into eight priorities mapped into four hardware queues. In the egress, traffic is destined to two adjacent network ports (for example, ports 1/1/1 and 1/1/2), and aggregated into one 1-GbE port in the back-end NP. The two network ports share the same hardware queues, and therefore they have the same buffer and descriptor limits and scheduling algorithm for transmission.

Ingress packets are classified into eight QoS priority levels at the front-end NP of the SX-FI48GPP module. The eight priorities are mapped into four hardware queues based on the priority queue configuration in the following table. QoS priority 7 is the highest priority, and QoS 0 is the lowest priority.

**TABLE 7** Priority queues for the SX-FI48GPP

QoS priority level	Hardware queues (traffic classes)
0	0
1	0
2	1
3	1
4	2
5	2
6	3
7	3

QoS classification occurs in two iterations; initially in the front-end NP, followed by the back-end NP. The back-end NP has the same classification and marking capabilities of existing FastIron SX interface modules, but the front-end NP does not support ACL and static MAC priority. The front-end NP supports basic QoS features, such as port priority, QoS-ToS mapping, 802.1p to priority mapping, 802.1p override, and trust DSCP mode.

The default scheduling configuration for Weighted Round Robin (WRR), Hybrid WRR and Strict Priority (SP), and SP mode for the eight QoS priority queues mapped to the four hardware queues is described under [Default scheduling configuration for the SX-FI48GPP module](#) on page 36.

## Queues for the SX-FI-8XG interface module

The SX-FI-8XG interface module consists of two separate hardware Network Processors (NP). The front-end NP supports 8 hardware queues, and the back-end NP supports eight hardware queues. In the egress, traffic is destined to four adjacent ports (for example, ports

1/1 to 1/4), and aggregated into one 10GbE port in the back-end NP. The four network ports share the same hardware queues; therefore, they have the same buffer and descriptor limits and scheduling algorithm for transmission.

QoS classification occurs in two iterations; initially in the front-end NP, followed by the back-end NP. The back-end NP has the same classification and marking capabilities of existing FSX interface modules, however, the front -end NP does not support ACL and static MAC priority. The front-end NP supports basic QoS features, such as port priority, qos-tos mapping, 802.1p to priority mapping, 802.1p override, and trust-dscp mode.

## Queues for the ICX 6430 switch

For the ICX 6430 switch, ingress packets are classified into eight QoS priority levels. These are mapped internally to four hardware forwarding queues or traffic classes as shown in the following table. QoS priority 7 is the highest priority, and QoS 0 is the lowest QoS priority (qosp) level.

**TABLE 8** Priority queues for the ICX 6430

QoS priority level	Hardware queues (Traffic classes)
0	0
1	0
2	1
3	1
4	1
5	2
6	2
7	3

For the ICX 6430 switch, internal forwarding priority maps to hardware forwarding queues 0 through 3. The mapping between the internal priority and hardware forwarding queue cannot be changed. The following tables show the default QoS mappings that are used if the trust level for CoS or DSCP is enabled. Mappings are the same for standalone and stacking systems.

**TABLE 9** Default QoS mappings for ICX 6430, columns 0 to 15

DSCP value	0	1	2	3	4	5	6	7	8	9	10	11	12	12	14	15
802.1p (CoS) value	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
DSCP value	0	1	2	3	4	5	6	7	8	9	10	11	12	12	14	15
Internal forwarding priority	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
Forwarding queue	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

**TABLE 10** Default QoS mappings for ICX 6430, columns 16 to 31

DSCP value	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
802.1p (CoS) value	2	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3
DSCP value	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Internal forwarding priority	2	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3
Forwarding queue	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

**TABLE 11** Default QoS mappings for ICX 6430, columns 32 to 47

DSCP value	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
802.1p (CoS) value	4	4	4	4	4	4	4	4	5	5	5	5	5	5	5	5
DSCP value	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
Internal forwarding priority	4	4	4	4	4	4	4	4	5	5	5	5	5	5	5	5
Forwarding queue	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2

**TABLE 12** Default QoS mappings for ICX 6430, columns 48 to 63

DSCP value	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
802.1p (CoS) value	6	6	6	6	6	6	6	6	7	7	7	7	7	7	7	7
DSCP value	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
Internal forwarding priority	6	6	6	6	6	6	6	6	7	7	7	7	7	7	7	7



**TABLE 12** Default QoS mappings for ICX 6430, columns 48 to 63 (continued)

DSCP value	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
Forwarding queue	2	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3

Mapping between DSCP value and forwarding queue cannot be changed. However, mapping between DSCP values and other properties can be changed as follows:

- DSCP to internal forwarding priority mapping - You can change the mapping between the DSCP value and the internal forwarding priority value from the default values shown in the above tables. This mapping is used for CoS marking and determining the internal priority when the trust level is DSCP. See the “Changing the DSCP to internal forwarding priority mappings” section.
- VLAN priority (802.1p) to hardware forwarding queue - You can change the mapping between the 802.1p value and hardware forwarding queue from the default value. See the “Changing the VLAN priority 802.1p to hardware forwarding queue mappings” section.

## User-configurable scheduler profile

The user-configurable scheduler profile is a template that defines either the scheduling mechanism or scheduling profile (weights assigned to the queues) or both for the egress queues. A configured user-configurable scheduler profile for egress queues can be applied to any hardware device. The default QoS is applicable to the entire system. If the scheduler profile is configured using the **qos mech strict** command, all devices in the system will be configured with the strict priority. The user-configurable scheduler profile is applicable only to the specific devices, leaving the remaining devices running default QoS. On any device, the user-configurable scheduler profile has high priority over the default QoS. On any device, user-configurable scheduler profile has high priority over the default QoS. The user-configurable scheduler profile should be in line with default QoS commands in both stacking and standalone systems.

On Brocade ICX 7750 devices, scheduler profiles are applied at the port, rather than at the device (port region), level. See the description of the **scheduler-profile** command for more information.

## User-configurable scheduler profile configuration

Configuring a user-configurable scheduler profile involves, selecting a proper mechanism and appropriate weights for the traffic classes (TCs) corresponding to that mechanism. It is highly recommended that you let the system use the default scheduling mechanism unless user knows what parameters you intend to modify and for what reasons.

There are two ways of creating a user-configurable scheduler profile. The scheduler-profile can be created either by specifying a mechanism (WRR, Strict, or Mixed) or by specifying weights.

The user-configurable scheduler profile can be created by specifying a mechanism. There are three available mechanisms:

- Strict Priority (SP)
- Weighted Round Robin (WRR)
- Mixed (combination of SP and WRR)

Following is the command format for creating a profile while specifying a mechanism.

**Syntax:** **qos scheduler-profile** *user\_profile\_name* **mechanism** *scheduling\_mechanism*

The *user\_profile\_name* variable is the name of the profile you are creating.

The *scheduling\_mechanism* variable is SP, WRR, or Mixed.

The user-configurable scheduler profile can be created by specifying weights from **qosp0** through **qosp7**, as shown in the following command format.

**Syntax:** `qos scheduler-profile user_profile_name qosp0 w0 qosp1 w1 qosp2 w2 qosp3 w3 qosp4 w4 qosp5 w5 qosp6 w6 qosp7 w7`

The `user_profile_name` variable is the name of the profile you are creating.

Profile **qosp0** through **qosp7** are the default queue names.

The `w0` through `w7` variables are the assigned weights.

If you create a profile specifying only the weights (**qosp0** through **qosp7**) without specifying the mechanism, the default mechanism is used. The default mechanism for stacking systems is *Mixed*, and *WRR* for standalone systems.

If you change the profile mechanism, the weights also get changed according to the mechanism. The weights can be modified according to the following requirements:

- If the mechanism is changed to *WRR*, the default system weights get assigned
- If the mechanism is changed to *Mixed*, the default mix weights get assigned
- If the mechanism is changed to *Strict*, the weights are ignored and remain untouched.

Scheduler-profile modifications take effect dynamically on an active profile. The operational defaults for all scheduling types for stacking and standalone systems are listed in the "Default values for scheduling type for stacking and standalone systems (for FCX and ICX 6450 platforms)" table.

## Displaying the user-configurable scheduler profile configuration

To display the specified user-configurable scheduler profile configuration, use the **show scheduler-profile** `user_profile_name` command.

The `user_profile_name` variable is the name of the profile you are creating.

To display all the scheduler profiles configured in the runtime configuration for the system, use the **show scheduler-profile all** command.

### NOTE

On the Brocade ICX 7450 and ICX 7750, the **show scheduler-profile** command does not work; you must use the **show qos scheduler-profile** command on these platforms.

## FCX and ICX 6450 platforms

The following tables show the default values for the scheduling type for stacking and standalone FCX and ICX 6450 platforms.

**TABLE 13** Default values for scheduling type for stacking systems (for FCX and ICX 6450 platforms)

Traffic Class	SP	SP Jumbo	WRR	WRR Jumbo	Mixed	Mixed Jumbo
TC0	SP	SP	3	8	15	15
TC1	SP	SP	3	8	15	15
TC2	SP	SP	3	8	15	15
TC3	SP	SP	3	8	15	15
TC4	SP	SP	3	8	15	15
TC5	SP	SP	10	16	25	25
TC6	SP	SP	75	44	SP	SP
TC7	SP	SP	SP	SP	SP	SP

**TABLE 14** Default values for scheduling type for standalone systems (for FCX and ICX 6450 platforms)

	SP	SP Jumbo	WRR	WRR Jumbo	Mixed	Mixed Jumbo
TC0	SP	SP	3	8	15	15
TC1	SP	SP	3	8	15	15
TC2	SP	SP	3	8	15	15
TC3	SP	SP	3	8	15	15
TC4	SP	SP	3	8	15	15
TC5	SP	SP	3	8	25	25
TC6	SP	SP	7	8	SP	SP
TC7	SP	SP	75	44	SP	SP

### ICX 6650 platforms

The following tables show the default values for the scheduling type for ICX 6650 platforms.

**TABLE 15** Default values for scheduling type for ICX 6650 platforms

	SP	SP Jumbo	WRR	WRR Jumbo	Mixed	Mixed Jumbo
TC0	SP	SP	3	8	15	15
TC1	SP	SP	3	8	15	15
TC2	SP	SP	3	8	15	15
TC3	SP	SP	3	8	15	15
TC4	SP	SP	3	8	15	15
TC5	SP	SP	3	8	25	25
TC6	SP	SP	7	8	SP	SP
TC7	SP	SP	75	44	SP	SP

### ICX 6430 platforms

The following table shows the default values for scheduling type for stacking and standalone ICX 6430 platforms. The lowest weighted priority is for qosp0, while the highest is for qosp7.

Note that values are provided for QoS priority (QSP) levels. The weights applied to the traffic class (TC) are the sum of the weights of the QSP levels that map to that TC. For example, QSP0 and QSP1 map to TC0. If the weight for QSP0 is 6 and the weight for QSP1 is 6, then the weight for TC0 is 12. Refer to the "Priority queues for the ICX 6430" table for QoS priority to traffic class mapping.

**TABLE 16** Default values for scheduling type for stacking systems (for ICX 6430 platforms)

QSP Level	SP	SP Jumbo	WRR	WRR Jumbo	Mixed	Mixed Jumbo
QSP0	SP	SP	3	8	15	15
QSP1	SP	SP	3	8	15	15
QSP2	SP	SP	3	8	15	15
QSP3	SP	SP	3	8	15	15
QSP4	SP	SP	3	8	40	40
QSP5	SP	SP	10	16	SP	SP
QSP6	SP	SP	75	44	SP	SP
QSP7	SP	SP	SP	SP	SP	SP

**TABLE 17** Default values for scheduling type for standalone systems (for ICX 6430 platforms)

	SP	SP Jumbo	WRR	WRR Jumbo	Mixed	Mixed Jumbo
QSP0	SP	SP	3	8	15	15
QSP1	SP	SP	3	8	15	15
QSP2	SP	SP	3	8	15	15
QSP3	SP	SP	3	8	15	15
QSP4	SP	SP	3	8	40	40
QSP5	SP	SP	3	8	SP	SP
QSP6	SP	SP	7	8	SP	SP
QSP7	SP	SP	75	44	SP	SP

## QoS priorities-to-traffic assignment

By default, all traffic is in the best-effort queue (qosp0) and is honored on tagged ports on all FastIron switches. You can assign traffic to a higher queue based on the following:

- Incoming port (sometimes called the *ingress port*)
- Static MAC entry

When you change the priority, you specify a number from 0 through 7. The priority number specifies the IEEE 802.1 equivalent to one of the eight QoS queues on Brocade devices. The numbers correspond to the queues as shown in the *QoS queues* table.

Although it is possible for a packet to qualify for an adjusted QoS priority based on more than one of the criteria, the system always gives a packet the highest priority for which it qualifies. Thus, if a packet is entitled to the premium queue because of its IP source and destination addresses, but is entitled only to the high queue because of its incoming port, the system places the packet in the premium queue on the outgoing port.

## Changing a port priority

To change the QoS priority of port 1/1/1 to the premium queue (qosp7), enter the following commands.

```
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)#priority 7
```

The device assigns priority 7 to untagged switched traffic received on port 1/1/1.

Use the following command to assign priority levels.

**Syntax:** [no] priority *num*

The *num* variable can be from 0 through 7 and specifies the IEEE 802.1 equivalent to one of the eight QoS queues listed in the *QoS queues* table.

## Assigning static MAC entries to priority queues

By default, all MAC entries are in the best-effort queue. When you configure a static MAC entry, you can assign the entry to a higher QoS level.

To configure a static MAC entry and assign the entry to the premium queue, enter commands such as the following.

```
Device(config)#vlan 9
Device(config-vlan-9)#static-mac-address 0000.0063.67FF ethernet 1/1/1 priority 7
Device(config-vlan-9)#write memory
```

Use the following command to configure a MAC entry and assign the entry to a priority queue.

**Syntax:** `[no] static-mac-address mac-addr ethernet unit/slot/port [ priority num ]`

The `mac-addr` variable is the MAC address.

The `priority num` variable can be from 0 through 7 and specifies the IEEE 802.1 equivalent to one of the eight QoS queues.

## Buffer allocation and threshold for QoS queues

By default, Brocade IronWare software allocates a certain number of buffers to the outbound transport queue for each port based on QoS priority. The buffers control the total number of packets permitted in the outbound queue for the port. If desired, you can increase or decrease the maximum number of outbound transmit buffers allocated to all QoS queues, or to specific QoS queues on a port or group of ports. For more information, refer to the *FastIron Ethernet Switch Platform and Layer 2 Switching Configuration Guide*.

### NOTE

On ICX 6650 devices, you cannot increase or decrease the maximum number of outbound transmit buffers allocated to all QoS queues, or to specific QoS queues on a port or group of ports.

## 802.1p priority override

You can configure a port to ignore the 802.1p priority for traffic classification for an incoming packet. When this feature is enabled, packets will be classified as follows:

- If the packet matches an ACL that defines the priority, then ACL priority will be used.
- If the packet source or destination MAC address matches a configured static MAC address with priority, then static MAC priority will be used.
- If the ingress port has a configured priority, then port priority will be used.
- If the other situations do not apply, the configured or default port priority (0) will be used.

Note that the original 802.1p priority in the packet will be retained. This feature does not re-mark the 802.1p value.

## Configuration notes and feature limitations

- 802.1p priority override is supported on physical ports and trunk ports. When applied to the primary port of a trunk group, the configuration applies to all members of the trunk group.
- This feature is not supported together with the `trust dscp` command.

## Enabling 802.1p priority override

To enable 802.1p priority override, enter the following command at the interface level of the CLI.

```
device(config-if-e1000-2)#priority ignore-8021p
```

**Syntax:** `[no] priority ignore-8021p`

Use the following command to show whether 802.1p priority override is enabled on a port.

```
device# show run interface ethernet 1
interface ethernet 1
priority ignore-8021p
```

**Syntax:** `show run interface ethernet port`

## Marking

*Marking* is the process of changing the packet QoS information (the 802.1p and DSCP information in a packet) for the next hop. For example, for traffic coming from a device that does not support Differentiated Services (DiffServ), you can change the packet IP precedence value into a DSCP value before forwarding the packet.

You can mark a packet's Layer 2 CoS value, its Layer 3 DSCP value, or both values. The Layer 2 CoS or DSCP value the device marks in the packet is the same value that results from mapping the packet QoS value into a Layer 2 CoS or DSCP value.

Marking is optional and is disabled by default. In releases prior to IronWare 8.0, marking is performed only using ACLs. For configuration syntax, rules, and examples of QoS marking, refer to "QoS options for IP ACLs" section in the *FastIron Ethernet Switch Security Configuration Guide*.

## DSCP and CoS global remarking

### NOTE

DSCP and CoS global marking is not supported on the ICX 6650.

When marking is not used, the device performs the mappings listed for scheduling the packet, but leaves the packet QoS values unchanged when the device forwards the packet. For more information, refer to the "QoS overview" section. When marking is not enabled using ACLs, a rogue host that wants preferential treatment for all its traffic can mark the DSCP field as per its requirements and send the traffic to the device.

Prior to 08.0.00, the only way to prevent such threats was to mark all packets using ACLs. Beginning with 08.0.00, the internal forwarding priority can be set using an ACL only for flows that require preferential QoS treatment. For all other flows, you can remark DSCP and CoS fields globally. Traffic marked by the ACL method always has a higher priority than the global marking.

When DSCP marking is configured on a given port, the DSCP field of any IPv4 packet received on the port is remarked to the configured value.

When CoS marking is configured, the PCP bit value in the VLAN header is remarked to the desired value for all tagged packets. CoS marking can be configured on a port. When configured on a port, the PCP bit in the VLAN header for all packets that egress the port is remarked to the configured value.

Both DSCP and CoS global marking can be configured on the ports of the modules that are configured but not physically present. When the modules are hot-swapped, the marking is automatically applied or removed.

The DSCP and CoS remarking can be configured through the command line interface (CLI) at the global level and the interface level. The global DSCP and CoS marking can coexist with other security features configured on the same port. The coexistence rules are the same as those for IPv4 ACLs.

## Configuration considerations and limitations

- When an ACL is configured on a port without remarking and global DSCP remarking is enabled, the global DSCP remarking is enabled for the permitted traffic.

- DSCP and CoS global remarking are supported on the same interface together.
- DSCP and CoS global remarking cannot coexist with MAC filters and MAC-based VLANs.

The following table summarizes the behavior when the remarking is set.

**TABLE 18** DSCP and PCP remarking

DSCP	Remarking set	Remarking set	Not set
CoS	Remarking set	Not set	Remarking set
DSCP action	Remark DSCP at the ingress	Remark DSCP at the ingress	N/A
PCP action	Remark PCP at the egress	N/A	Remark PCP at the egress
Traffic class	Apply the TC equivalent to DSCP	Apply the TC equivalent to DSCP	Apply the TC equivalent to PCP

## Enabling DSCP marking

To enable DSCP remarking globally, use the **ip dscp-remark** command in global configuration mode.

To enable DSCP remarking on a port, use the **ip dscp-remark** command in interface configuration mode.

### Example: DSCP marking

The following example shows how to set the DSCP value to 3 for all IP packets:

```
device(config)# ip dscp-remark 3
```

The following example shows how to set the DSCP value to 4 of all IP packets on a specific port:

```
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# ip dscp-remark 4
```

## Enabling CoS marking

To enable CoS marking globally, use the **ip pcp-remark** command in global configuration mode.

To enable CoS marking on a port, use the **ip pcp-remark** command in interface configuration mode.

### Example: CoS marking

The following example shows how to set the PCP value to 3 for all VLAN tagged packets:

```
device(config)# ip pcp-remark 3
```

The following example shows how to set the PCP value to 4 of all IP packets on a specific port:

```
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# ip pcp-remark 4
```

## DSCP-based QoS configuration

Brocade IronWare releases support basic DSCP-based QoS (also called Type of Service (ToS)-based QoS) as described in this chapter. However, the FastIron family of switches does not support other advanced DSCP-based QoS features.

Brocade IronWare releases also support marking of the DSCP value. The software can read Layer 3 Quality of Service (QoS) information in an IP packet and select a forwarding queue for the packet based on the information. The software interprets the value in

the six most significant bits of the IP packet header 8-bit ToS field as a DSCP value, and maps that value to an internal forwarding priority.

#### NOTE

MAC filter and DSCP marking cannot be configured on the same port.

The internal forwarding priorities are mapped to one of the eight forwarding queues (qosp0 through qosp7) on the Brocade device. During a forwarding cycle, the device gives more preference to the higher-numbered queues, so that more packets are forwarded from these queues. For example, queue qosp7 receives the highest preference, while queue qosp0, the best-effort queue, receives the lowest preference.

## Application notes for DSCP-based QoS

- DSCP-based QoS is not automatically honored for routed and switched traffic. The default is 802.1p to CoS mapping. To honor DSCP-based QoS, you must either use an ACL or enable trust DSCP. Refer to [Using ACLs to honor DSCP-based QoS](#) on page 32.
- When DSCP marking is enabled, the device changes the contents of the inbound packet ToS field to match the DSCP-based QoS value.

## Using ACLs to honor DSCP-based QoS

This section shows how to configure Brocade devices to honor DSCP-based QoS for routed and switched traffic.

### FCX and ICX devices

Brocade FCX, ICX 6430, ICX 6450, ICX 6610, ICX 6650, and ICX 7750 devices support DSCP-based QoS on a per-port basis. DSCP-based QoS is not automatically honored for switched traffic. The default is 802.1p to CoS mapping. To honor DSCP-based QoS, enter the **trust dscp** command at the interface level of the CLI.

#### Syntax: [no] trust dscp

To disable the configuration, use the **no** form of the command.

When the **trust dscp** command is enabled, the interface honors the Layer 3 DSCP value. By default, the interface honors the Layer 2 CoS value.

#### NOTE

The **trust dscp** command is not supported with 802.1p priority override.

#### NOTE

On ICX 7750 devices, configuring the **trust dscp** command to honor DSCP-based QoS classification on the ingress port works on all traffic except GRE tunnels; classification on these remains based on Layer 2 (802.1p) trust. For GRE tunnels you can use ACLs to configure classification based on the DSCP value.

### FSX Series devices

FSX devices require the use of an ACL to honor DSCP-based QoS for routed traffic in the Layer 3 image, or for switched traffic in the Layer 2 image. To enable DSCP-based QoS on these devices, apply an ACL entry such as the following.

```
device(config)#access-list 101 permit ip any any dscp-cos-mapping
```



**NOTE**

Use the **bridged-routed** keyword in the ACL to honor DSCP for switched traffic in the Layer 3 image. Refer to "Enabling ACL support for switched traffic in the router image" section in the *FastIron Ethernet Switch Security Configuration Guide*.

**NOTE**

The **access-list 101 permit ip any any dscp-cos-mapping** command is supported on the SX-FI48GPP interface module. For more information on QoS queues for the SX-FI48GPP interface module, refer to the "Queues for the SX-FI48GPP interface" section.

## Trust DSCP for the SX-FI48GPP, SX-FI-24GPP, SX-FI-24HF, SX-FI-2XG, and SX-FI-8XG modules

On the following modules, trust DSCP can be enabled on a per-port basis:

- SX-FI48GPP
- SX-FI-24GPP
- SX-FI-24HF
- SX-FI-2XG
- SX-FI-8XG

Each port on the supported modules corresponds to a front-end panel port. By default, trust VLAN priority is enabled.

**NOTE**

For all ports in the other FastIron SX modules, ACL should be used to implement the trust DSCP mode.

For example, to enable trust DSCP on interface ethernet 1/48 on the SX-FI48GPP module, enter the **trust dscp** command.

**Syntax:** **[no] trust dscp**

To disable the configuration, use the **no** form of the command.

## Configuring QoS mapping configuration

You can optionally change the following QoS mappings:

- DSCP to internal forwarding priority
- VLAN priority (802.1p) to hardware forwarding queue

The mappings are globally configurable and apply to all interfaces.

### Default DSCP to internal forwarding priority mappings

The DSCP values are described in RFCs 2474 and 2475. The following table lists the default mappings of DSCP values to internal forwarding priority values.

**TABLE 19** Default DSCP to internal forwarding priority mappings

Internal forwarding priority	DSCP value
0 (lowest priority queue)	0 - 7
1	8 - 15
2	16 - 23
3	24 - 31

**TABLE 19** Default DSCP to internal forwarding priority mappings (continued)

Internal forwarding priority	DSCP value
4	32 - 39
5	40 - 47
6	48 - 55
7 (highest priority queue)	56 - 63

Notice that DSCP values range from 0 through 63, whereas the internal forwarding priority values range from 0 through 7. Any DSCP value within a given range is mapped to the same internal forwarding priority value. For example, any DSCP value from 8 through 15 maps to priority 1.

After performing this mapping, the device maps the internal forwarding priority value to one of the hardware forwarding queues.

On FCX and ICX devices, you can use QoS queue 1 for priority traffic, even when sFlow is enabled on the port. This differs from the FastIron X Series devices, which support seven priorities for user data instead of eight when sFlow is enabled. QoS queue 1 is reserved for sFlow and not used by other packets. Any non-sFlow packets assigned to QoS queue 1 will be directed to QoS queue 0. Note that the ICX 6430 does not support sFlow.

The following table lists the default mappings of internal forwarding priority values to the hardware forwarding queues for the ICX 6430.

**TABLE 20** Default mappings of internal forwarding priority values for the ICX 6430

Internal forwarding priority	Forwarding queues
0 (lowest priority queue)	qosp0
1	qosp0
2	qosp1
3	qosp1
4	qosp1
5	qosp2
6	qosp2
7 (highest priority queue)	qosp3

You can change the DSCP to internal forwarding mappings. You also can change the internal forwarding priority to hardware forwarding queue mappings.

## Changing the DSCP to internal forwarding priority mappings

To change the DSCP to internal forwarding priority mappings for all the DSCP ranges, enter commands such as the following at the global CONFIG level of the CLI.

```
Device(config)#qos-tos map dscp-priority 0 2 3 4 to 1
Device(config)#qos-tos map dscp-priority 8 to 5
Device(config)#qos-tos map dscp-priority 16 to 4
Device(config)#qos-tos map dscp-priority 24 to 2
Device(config)#qos-tos map dscp-priority 32 to 0
Device(config)#qos-tos map dscp-priority 40 to 7
Device(config)#qos-tos map dscp-priority 48 to 3
Device(config)#qos-tos map dscp-priority 56 to 6
Device(config)#qos-tos map dscp-priority 56 t
```

Use the following command to map priority levels to DSCP values.

**Syntax:** `[no] qos-tos map dscp-priority dscp-value1 dscp-value2 dscp-value3 dscp-value4 dscp-value5 dscp-value6 dscp-value7 dscp-value8 to priority`

The *dscp-value* variable specifies the DSCP value ranges you are remapping.

You can specify up to eight DSCP values in the same command, to map to the same forwarding priority.

The *priority* variable specifies the internal forwarding priority.

This is an example of using this command:

```
qos-tos map dscp-priority 1 2 3 4 5 6 7 8 to 6
```

This is output displayed from using the **show qos-tos** command as a result of issuing the preceding command:

```
Device#show qos-tos
Portions of table omitted for simplicity.
DSCP-Priority map: (dscp = d1d2)
  d2| 0  1  2  3  4  5  6  7  8  9
  d1|
-----+-----
  0 | 1
  0 | 1  1  1
  0 | 0  0  5
  1
  1 | 6  1  1  1  1  1  4
  2 | 2  2
  2 | 2  2  2  2  2
  3 | 3  3  3  3
  3 | 3  3  0
  4 | 4  4  4  4  4  4
  4 | 7
  5 | 5  5  5  5  5  5  3
  6
  5 | 6  6  6  6  6  6  6
  7 | 7  7
  6 | 7  7  7  7
```

This output displays mappings in the DSCP to forwarding priority portion of the QoS information display. To read this part of the display, select the first part of the DSCP value from the d1 column and select the second part of the DSCP value from the d2 row. For example, to read the DSCP to forwarding priority mapping for DSCP value 24, select 2 from the d1 column and select 4 from the d2 row. The mappings that are changed by the example **qos-tos map dscp-priority** command are shown in bold type.

## Changing the VLAN priority 802.1p to hardwareforwarding queue mappings

To map a VLAN priority to a different hardware forwarding queue, enter commands such as the following at the global CONFIG level of the CLI.

**Syntax:** `[no] qos tagged-priority num queue`

The *num* variable can be from 0 through 7 and specifies the VLAN priority.

The *queue* variable specifies the hardware forwarding queue to which you are reassigning the priority.

The default queue names are as follows:

- qosp7
- qosp6
- qosp5
- qosp4
- qosp3
- qosp2
- qosp1
- qosp0

Following is an example of using this command.

```
device(config)#qos tagged-priority 2 qosp0
```

## Default scheduling configuration for the SX-FI48GPP module

The default scheduling configuration for Weighted Round Robin (WRR), Strict Priority (SP), and mixed WRR and SP mode for the eight QoS priority (qosp) queues mapped to the four hardware queues is described in the following table.

**TABLE 21** Default configuration for 8 to 4 queues for the SX-FI48GPP module

Hardware Queue	Weighted Round Robin (WRR) mode	Mixed WRR and SP	Strict Priority (SP) mode
3	Weight 82%	Strict Priority	Strict Priority
2	Weight 6%	Weight 40%	Strict Priority
1	Weight 6%	Weight 30%	Strict Priority
0	Weight 6%	Weight 30%	Strict Priority

Note that the above table includes values for default, non-jumbo mode WRR. The hardware queues are calculated using default qosp values from the *Default values for scheduling type for stacking and standalone systems (for FCX and ICX 6450 platforms)* table as follows:

- Front end queue 3 = 75% (qosp7) + 7% (qosp6) = 82%
- Front end queue 2 = 3% (qosp4) + 3% (qosp5) = 6%
- Front end queue 1 = 3% (qosp2) + 3% (qosp3) = 6%
- Front end queue 0 = 3% (qosp0) + 3% (qosp1) = 6%

The hardware queues for mixed WRR and SP mode are calculated as follows:

- Front end queue 3 is Strict Priority as default values for qosp7 and qosp6 are SP
- Front end queue 2 = 25% (qosp4) + 15% (qosp5) = 40%

- Front end queue 1 = 15% (qosp2) + 15% (qosp3) = 30%
- Front end queue 0 = 15% (qosp0) + 15% (qosp1) = 30%

## Default scheduling configuration for the ICX 6430

The default scheduling configuration for Weighted Round Robin (WRR), Strict Priority (SP), and mixed WRR and SP mode for the eight QoS priority (qosp) queues mapped to the four hardware queues for an ICX 6430 is described in the following table.

**TABLE 22** Default configuration for 8 to 4 queues (standalone system)

Hardware queue	Weighted Round Robin (WRR) mode	Mixed WRR and SP	Strict Priority (SP) mode
3	Weight 75%	Strict Priority	Strict Priority
2	Weight 10%	Strict Priority	Strict Priority
1	Weight 9%	Weight 70%	Strict Priority
0	Weight 6%	Weight 30%	Strict Priority

The above table includes values for default, non-jumbo mode WRR for a standalone system. The hardware queues are calculated using default qosp values from the *Default values for scheduling type for stacking and standalone systems (for ICX 6430 platforms)* table as follows:

- Queue 3 = 75% (qosp7)
- Queue 2 = 3% (qosp5) + 7% (qosp6) = 10%
- Queue 1 = 3% (qosp2) + 3% (qosp3) + 3% (qosp4) = 9%
- Queue 0 = 3% (qosp0) + 3% (qosp1) = 6%

The hardware queues for mixed WRR and SP mode are calculated as follows:

- Queue 3 is Strict Priority as the default value for qosp7 is SP
- Queue 2 is Strict Priority as default values for qosp5 and qosp6 are SP
- Queue 1 = 15% (qosp2) + 15% (qosp3) + 40% (qosp4) = 70%
- Queue 0 = 15% (qosp0) + 15% (qosp1) = 30%

### NOTE

If any qosp value is SP, then the weight of the hardware queue is SP.

## Scheduling QoS information

*Scheduling* is the process of mapping a packet to an internal forwarding queue based on its QoS information, and servicing the queues according to a mechanism.

### Scheduling for the SX-FI48GPP module

The SX-FI48GPP module supports scheduling at the front-end and back-end NP. If egress congestion occurs at the front-end NP of the SX-FI48GPP module, scheduling is based on four queues instead of eight. For more information on default configuration for 8 to 4 queue mapping, refer to the *Priority queues for the SX-FI48GPP* table. If egress congestion occurs at the back-end of the SX-FI48GPP module, then scheduling is based on eight queues. When SX-FI48GPP ports are running at a reduced speed (100 Mbps or 10 Mbps), egress congestion usually occurs at the front-end NP.

## QoS queuing methods

The following QoS queuing methods are supported in all IronWare releases for the FastIron devices:

- Weighted Round Robin (WRR) - WRR ensures that all queues are serviced during each cycle. A WRR algorithm is used to rotate service among the eight queues on the FastIron devices. The rotation is based on the weights you assign to each queue. This method rotates service among the queues, forwarding a specific number of packets in one queue before moving on to the next one.

### NOTE

In stacking mode, the qosp7 queue is reserved as Strict Priority (SP) under weighted queuing. Attempts to change the qosp7 setting will be ignored.

WRR is the default queuing method and uses a default set of queue weights.

The number of packets serviced during each visit to a queue depends on the percentages you configure for the queues. The software automatically converts the percentages you specify into weights for the queues.

### NOTE

Queue cycles on the Brocade FastIron devices (except on ICX 6650) are based on bytes. These devices service a given number of bytes (based on weight) in each queue cycle. The QoS WRR on ICX 6650 is configured to operate in packet count mode.

- Strict Priority (SP) - SP ensures service for high-priority traffic. The software assigns the maximum weights to each queue, to cause the queuing mechanism to serve as many packets in one queue as possible before moving to a lower queue. This method biases the queuing mechanism to favor the higher queues over the lower queues.

For example, strict queuing processes as many packets as possible in qosp3 before processing any packets in qosp2, then processes as many packets as possible in qosp2 before processing any packets in qosp1, and so on.

- Hybrid WRR and SP - This configurable queueing mechanism combines both the SP and WRR mechanisms. The combined method enables the Brocade device to give strict priority to delay-sensitive traffic such as VoIP traffic, and weighted round robin priority to other traffic types.

By default, when you select the combined SP and WRR queueing method, the Brocade device assigns strict priority to traffic in qosp7 and qosp6, and weighted round robin priority to traffic in qosp0 through qosp5. Thus, the Brocade device schedules traffic in queue 7 and queue 6 first, based on the strict priority queueing method. When there is no traffic in queue 7 and queue 6, the device schedules the other queues in round-robin fashion from the highest priority queue to the lowest priority queue.

### NOTE

Brocade stackable devices that are operating as members of a stack reserve queue 7 for stacking functions. For more information, refer to the "QoS for Brocade stackable devices" section.

By default, when you specify the combined SP and WRR queuing method, the system balances the traffic among the queues as shown in the following table. If desired, you can change the default bandwidth values as shown in the "Bandwidth allocations of the hybrid WRR and SP queues" section.

**TABLE 23** Default bandwidth for combined SP and WRR queueing methods

Queue	Default bandwidth
qosp7	Strict Priority (highest priority)
qosp6	Strict Priority
qosp5	25%
qosp4	15%
qosp3	15%

**TABLE 23** Default bandwidth for combined SP and WRR queuing methods (continued)

Queue	Default bandwidth
qosp2	15%
qosp1	15%
qosp0	15% (lowest priority)

## Selecting the QoS queuing method

By default, Brocade devices use the WRR method of packet prioritization. To change the method to SP, enter the **qos mechanism strict** command at the global CONFIG level of the CLI.

```
device(config)#qos mechanism strict
```

To change the method back to WRR, enter the **qos mechanism weighted** command.

```
device(config)#qos mechanism weighted
```

To change the queuing mechanism to the combined SP and WRR method, enter the **qos mechanism mixed-sp-wrr** command at the global CONFIG level of the CLI.

```
device(config)#qos mechanism mixed-sp-wrr
```

**Syntax:** [no] qos mechanism { strict | weighted | mixed-sp-wrr }

## Configuring the QoS queues

Each of the queues has the following configurable parameters:

- The queue name
- The minimum percentage of a port outbound bandwidth guaranteed to the queue

## Renaming the queues

The default queue names are qosp7, qosp6, qosp5, qosp4, qosp3, qosp2, qosp1, and qosp0. You can change one or more of the names if desired.

To rename queue " qosp3 " to " 92-octane ", enter the following command.

```
device(config)#qos name qosp3 92-octane
```

**Syntax:** qos name *old-name new-name*

The *old-name* variable specifies the name of the queue before the change.

The *new-name* variable specifies the new name of the queue. You can specify an alphanumeric string up to 32 characters long.

## Changing the minimum bandwidth percentages of the WRR queues

If you are using the weighted round robin mechanism instead of the strict priority mechanism, you can change the weights for each queue by changing the minimum percentage of bandwidth you want each queue to guarantee for its traffic.

### NOTE

On the SX-FI48GPP interface module, the bandwidth percentages for 8 to 4 queue mapping for WRR queues are different from other Brocade SX modules. For more information on 8 to 4 queue mapping on the SX-FI48GPP interface module, refer to the "Default scheduling configuration for the SX-FI48GPP" section.

By default, the eight QoS queues on FastIron devices receive the minimum guaranteed percentages of a port's total bandwidth, as shown in the following table. Note that the defaults differ when jumbo frames are enabled.

**TABLE 24** Default minimum bandwidth percentages on Brocade devices

Queue	Default minimum percentage of bandwidth	
	Without jumbo frames	With jumbo frames
qosp7	75%	44%
qosp6	7%	8%
qosp5	3%	8%
qosp4	3%	8%
qosp3	3%	8%
qosp2	3%	8%
qosp1	3%	8%
qosp0	3%	8%

When the queuing method is WRR, the software internally translates the percentages into weights. The weight associated with each queue controls how many packets are processed for the queue at a given stage of a cycle through the weighted round robin algorithm.

#### NOTE

Queue cycles on the Brocade FastIron devices are based on bytes. These devices service a given number of bytes (based on the weight) in each queue cycle.

The bandwidth allocated to each queue is based on the relative weights of the queues. You can change the bandwidth percentages allocated to the queues by changing the queue weights.

There is no minimum bandwidth requirement for a given queue. For example, queue qosp3 is not required to have at least 50 percent of the bandwidth.

To change the bandwidth percentages for the queues, enter commands such as the following. Note that this example uses the default queue names.

```
Device(config)#qos profile qosp7 25 qosp6 15 qosp5 12 qosp4 12 qosp3 10 qosp2
10 qosp1 10 qosp0 6
Profile qosp7      : Priority7  bandwidth requested 25% calculated 25%
Profile qosp6      : Priority6  bandwidth requested 15% calculated 15%
Profile qosp5      : Priority5  bandwidth requested 12% calculated 12%
Profile qosp4      : Priority4  bandwidth requested 12% calculated 12%
Profile qosp3      : Priority3  bandwidth requested 10% calculated 10%
Profile qosp2      : Priority2  bandwidth requested 10% calculated 10%
Profile qosp1      : Priority1  bandwidth requested 10% calculated 10%
Profile qosp0      : Priority0  bandwidth requested 6%  calculated 6%
```

**Syntax:** `[no] qos profile name7{ sp | percentage } name6{ sp | percentage } name5{ sp | percentage } name4{ sp | percentage } name3{ sp | percentage } name2{ sp | percentage } name1{ sp | percentage } name0{ sp | percentage }`

Each *name* variable specifies the name of a queue. You can specify the queues in any order on the command line, but you must specify each queue.

The *percentage* variable specifies a number for the percentage of the device outbound bandwidth that is allocated to the queue. Brocade QoS queues require a minimum bandwidth percentage of 3 percent for each priority. When jumbo frames are enabled, the minimum bandwidth requirement is 8 percent. If these minimum values are not met, QoS may not be accurate.



## Configuration notes for changing the bandwidth

- The total of the percentages you enter must equal 100.
- FastIron devices do not adjust the bandwidth percentages you enter.

## Bandwidth allocations of the hybrid WRR and SP queues

### NOTE

On the SX-FI48GPP interface module, the bandwidth percentages for 8 to 4 queue mapping for hybrid WRR and SP queues are different from other Brocade SX modules. For more information on 8 to 4 queue mapping on the SX-FI48GPP interface module, refer to the “Default scheduling configuration for the SX-FI48GPP module” section.

To change the default bandwidth percentages for the queues when the device is configured to use the combined SP and WRR queuing mechanism, enter commands such as the following. Note that this example uses the default queue names.

```
device(config)#qos profile qosp7 sp qosp6 sp qosp5 20 qosp4 16 qosp3 16 qosp2 16 qosp1 32 qosp0 sp qosp1 16
qosp0 16
```

**Syntax:** `[no] qos profile name7{ sp | percentage } name6{ sp | percentage } name5{ sp | percentage } name4{ sp | percentage } name3{ sp | percentage } name2{ sp | percentage } name1{ sp | percentage } name0{ sp | percentage }`

Each *name* variable specifies the name of a queue, such as 7, 6, 5, 4, 3, 2, 1, and 0. You can specify the queues in any order on the command line, but you must specify each queue. Note that queue 7 supports Strict Priority (**sp**) only, queue 6 supports both SP and WRR queuing mechanisms (**sp** |), and queues 0 through 5 support the WRR queuing mechanism only.

### NOTE

Brocade stackable devices that are operating as members of a stack reserve queue 7 for stacking functions.

The **sp** parameter configures strict priority as the queuing mechanism. Note that only queue 7 and queue 6 support this method.

The *percentage* variable configures WRR as the queuing mechanism and specifies the percentage of the device outbound bandwidth allocated to the queue. The queues require a minimum bandwidth percentage of 3 percent for each priority. When jumbo frames are enabled, the minimum bandwidth requirement is 8 percent. If these minimum values are not met, QoS may not be accurate.

### NOTE

The percentages must add up to 100. The Brocade FastIron devices do not adjust the bandwidth percentages you enter.

## Viewing QoS settings

To display the QoS settings for all of the queues, enter the **show qos-profiles** command.

The following example shows the output on an FSX device.

```
Device# show qos-profiles all
bandwidth scheduling mechanism: weighted priority
Profile qosp7      : Priority7  bandwidth requested  25% calculated  25%
Profile qosp6      : Priority6  bandwidth requested  15% calculated  15%
Profile qosp5      : Priority5  bandwidth requested  12% calculated  12%
Profile qosp4      : Priority4  bandwidth requested  12% calculated  12%
Profile qosp3      : Priority3  bandwidth requested  10% calculated  10%
Profile qosp2      : Priority2  bandwidth requested  10% calculated  10%
Profile qosp1      : Priority1  bandwidth requested  10% calculated  10%
Profile qosp0      : Priority0  bandwidth requested   6% calculated   6%
```

The following example displays information, including multicast queue weights, for all the queues on an ICX 7450 device.

```
Device#show qos-profiles all
bandwidth scheduling mechanism: mixed weighted priority with strict priority
```

```

Unicast Traffic
Profile qosp7      : Priority7(Highest) Set as strict priority
Profile qosp6      : Priority6          Set as strict priority
Profile qosp5      : Priority5          bandwidth requested 25% calculated 25%
Profile qosp4      : Priority4          bandwidth requested 15% calculated 15%
Profile qosp3      : Priority3          bandwidth requested 15% calculated 15%
Profile qosp2      : Priority2          bandwidth requested 15% calculated 15%
Profile qosp1      : Priority1          bandwidth requested 15% calculated 15%
Profile qosp0      : Priority0(Lowest)  bandwidth requested 15% calculated 15%
Multicast Traffic
Profile qosp7+qosp6 : Priority7(Highest),6      Set as strict priority
Profile qosp5       : Priority5          bandwidth requested 25% calculated 25%
Profile qosp4+qosp3+qosp2 : Priority4,3,2      bandwidth requested 45% calculated 45%
Profile qosp1+qosp0 : Priority1,0(Lowest)    bandwidth requested 30% calculated 30%

```

**Syntax:** `show qos-profiles { all | name }`

The **all** parameter displays the settings for all eight queues.

The **name** variable displays the settings for the specified queue.

## Viewing DSCP-based QoS settings

To display configuration information for DSCP-based QoS, enter the **show qos-tos** command at any level of the CLI.

```

device#show qos-tos
DSCP-->Traffic-Class map: (DSCP = d1d2: 00, 01...63)
  d2| 0  1  2  3  4  5  6  7  8  9
  d1|
-----+-----
  0 | 0  0  0  0  0  0  0  0  0  1  1
  1 | 1  1  1  1  1  1  2  2  2  2  2
  2 | 2  2  2  2  3  3  3  3  3  3  3
  3 | 3  3  3  4  4  4  4  4  4  4  4
  4 | 5  5  5  5  5  5  5  5  6  6  6
  5 | 6  6  6  6  6  6  7  7  7  7  7
  6 | 7  7  7  7  7  7  7  7  7  7  7

Traffic-Class-->802.1p-Priority map (use to derive DSCP--802.1p-Priority):
Traffic | 802.1p
Class   | Priority
-----+-----
  0     | 0
  1     | 1
  2     | 2
  3     | 3
  4     | 4
  5     | 5
  6     | 6
  7     | 7
-----+-----

```

**Syntax:** `show qos-tos`

The following table shows the output information for the **show qos-tos** command.

**TABLE 25** DSCP-based QoS configuration information

Field	Description
DSCP to traffic class map	
d1 and d2	The DSCP to forwarding priority mappings that are currently in effect.

**TABLE 25** DSCP-based QoS configuration information (continued)

Field	Description
	<p><b>NOTE</b></p> <p>The example shows the default mappings. If you change the mappings, the command displays the changed mappings</p>
Traffic class to 802.1p priority map	
Traffic Class and 802.1p Priority	<p>The traffic class to 802.1p priority mappings that are currently in effect.</p> <p><b>NOTE</b></p> <p>The example shows the default mappings. If you change the mappings, the command displays the changed mappings.</p>

The **show qos-tos** command can also be used to display configuration information for 8 to 4 queue mapping. The following example displays an 8 to 4 queue mapping configuration.

```

device#show qos-tos
DSCP-->Traffic-Class map: (DSCP = d1d2: 00, 01...63)
  d2| 0  1  2  3  4  5  6  7  8  9
  d1|
-----+-----
  0 | 0  0  0  0  0  0  0  0  0  1  1
  1 | 1  1  1  1  1  1  2  2  2  2  2
  2 | 2  2  2  2  3  3  3  3  3  3  3
  3 | 3  3  4  4  4  4  4  4  4  4  4
  4 | 5  5  5  5  5  5  5  5  5  6  6
  5 | 6  6  6  6  6  6  7  7  7  7  7
  6 | 7  7  7  7  7
Traffic-Class-->802.1p-Priority map (use to derive DSCP--802.1p-Priority):
Traffic | 802.1p
Class   | Priority
-----+-----
  0 | 0
  1 | 1
  2 | 2
  3 | 3
  4 | 4
  5 | 5
  6 | 6
  7 | 7
-----+-----
8to4 queue mapping:
Priority| Hardware Queue
-----+-----
  0 | 0
  1 | 0
  2 | 1
  3 | 1
  4 | 2
  5 | 2
  6 | 3
  7 | 3
-----+-----

```

The following table shows the output information for 8 to 4 queue mapping for the **show qos-tos** command.

**TABLE 26** 8 to 4 queue mapping configuration information

Field	Description
8 to 4 queue mapping	
Priority and Hardware Queue	The priority to hardware queues that are currently in effect for 8 to 4 queue mapping. QoS priority 7 is the highest priority, and QoS 0 is the lowest priority

## IPv6 Quality of Service

QoS features supported for IPv6 environments include packet classification, queueing, traffic shaping, packet marking, and policing of IPv6 packets. These features are available for all FastIron products. The feature set is identical to that in IPv4.

To implement QoS in networks running IPv6, follow the same steps as those to implement QoS in networks running only IPv4. The recommended steps are as follows:

- Identify applications in your network and understand the characteristics of the applications so that you can make decisions about what QoS features to apply.
- Depending on network topology, link-layer header sizes are affected by changes and forwarding.
- Decide the method of classification, marking, and rate limiting. If the same network is carrying IPv4 and IPv6 traffic, decide if you want to treat both the same or differently, and specify match criteria accordingly. If you want to treat them the same, use match statements such as **match dscp** and **set dscp**. If you want to treat them differently, add match criteria such as **match protocol ip** and **match protocol ipv6** in the match criteria.

### NOTE

The command syntax for IPv6 access control lists (ACLs) is different from the syntax for IPv4. See the "Configuring an IPv6 ACL" section in the *FastIron Ethernet Switch Security Configuration Guide*.

# Rate Limiting and Rate Shaping on FastIron X Series and FCX and ICX Series Switches

- [Rate limiting overview.....](#)45
- [Rate limiting in hardware.....](#)45
- [Rate shaping overview.....](#)48

## Rate limiting overview

### NOTE

Rate limiting is packet-based on ICX 6650 devices in contrast with the other platforms on which it is byte-based.

Port-based fixed rate limiting is supported on inbound ports. This feature allows you to specify the maximum number of bytes (packets in the case of ICX 6650 devices) a given port can receive. The port drops bytes that exceed the limit you specify. You can configure a Fixed rate limiting policy on a port inbound direction only. Fixed rate limiting applies to all traffic on the rate limited port.

Fixed rate limiting is at line rate and occurs in hardware. Refer to the "Rate limiting in hardware" section.

When you specify the maximum number of bytes, you specify it in kilobits per second (kbps). On ICX 6650 devices, you specify the maximum number of packets. The Fixed rate limiting policy applies to one-second intervals and allows the port to receive the number of bytes (packets in the case of ICX 6650 devices) you specify in the policy, but drops additional bytes (packets in the case of ICX 6650 devices). Unused bandwidth is not carried over from one interval to the next.

### NOTE

For ICX 6450, ICX 6610, and ICX 6650 devices only, Port based Rate Limiting affects only known-unicast traffic. Broadcast, Multicast and Unknown-unicast (BUM) is not affected by this rate. To rate limit the BUM traffic, use BUM rate limiting as described in chapter BUM Rate Limiting.

### NOTE

Brocade recommends that you do not use Fixed rate limiting on ports that receive route control traffic or Spanning Tree Protocol (STP) control traffic. If the port drops control packets due to the Fixed rate limiting policy, routing or STP can be disrupted.

Rate limiting applies to inbound ports and rate shaping applies to outbound ports.

## Rate limiting in hardware

Each Brocade device supports line-rate rate limiting in hardware. The device creates entries in Content Addressable Memory (CAM) for the rate limiting policies. The CAM entries enable the device to perform the rate limiting in hardware instead of sending the traffic to the CPU. The device sends the first packet in a given traffic flow to the CPU, which creates a CAM entry for the traffic flow. A CAM entry consists of the source and destination addresses of the traffic. The device uses the CAM entry for rate limiting all the traffic within the same flow. A rate limiting CAM entry remains in the CAM for two minutes before aging out.

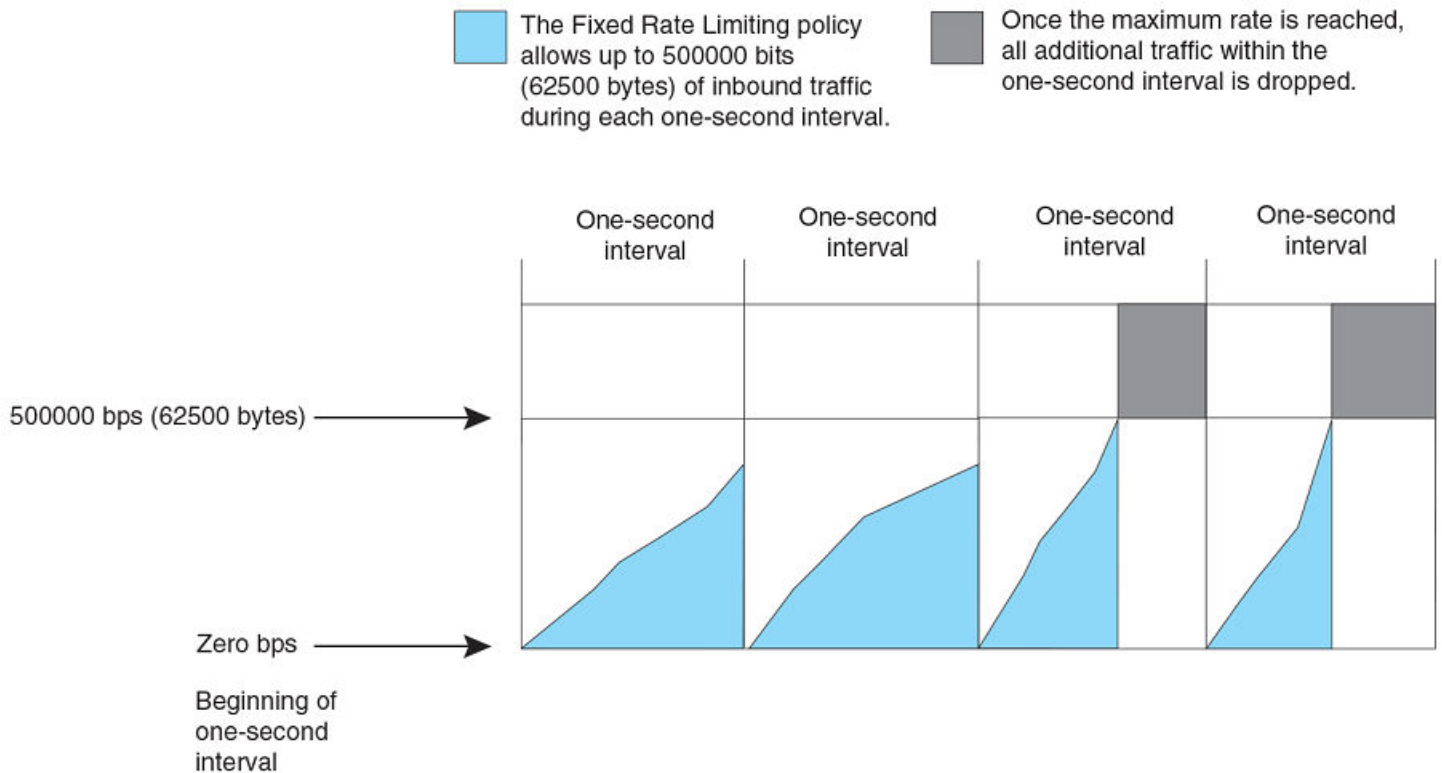
## How Fixed rate limiting works

Fixed rate limiting counts the number of bytes (packets in ICX 6650 devices) that a port receives, in one second intervals. If the number exceeds the maximum number you specify when you configure the rate, the port drops all further inbound packets for the duration of the one-second interval.

Once the one-second interval is complete, the port clears the counter and re-enables traffic.

The following figure shows an example of how Fixed rate limiting works. In this example, a Fixed rate limiting policy is applied to a port to limit the inbound traffic to 500000 bits (62500 bytes) a second. During the first two one-second intervals, the port receives less than 500000 bits (packets in ICX 6650 devices) in each interval. However, the port receives more than 500000 bits (packets in ICX 6650 devices) during the third and fourth one-second intervals, and consequently drops the excess traffic.

FIGURE 5 Fixed rate limiting



### NOTE

The software counts the bytes (packets in ICX 6650) by polling statistics counters for the port every 100 milliseconds, which provides 10 readings each second. Due to the polling interval, the Fixed Rate Limiting policy has an accuracy of within 10% of the port's line rate. It is therefore possible for the policy to sometimes allow more traffic than the limit you specify, but the extra traffic is never more than 10% of the port's line rate.

## Configuration notes for rate limiting

- Rate limiting is available only on inbound ports.

- Port based Rate limiting is not supported on the SX-FI62XG and SX-FI42XG modules of the FastIron X Series devices.
- FastIron X Series devices do not support fixed rate limiting on tagged ports in the base Layer 3 and full Layer 3 images.
- The rate limit on IPv6 hardware takes several seconds to take effect at higher configured rate limit values. For example, if the configured rate limit is 750 Mbps (1500000 packets/second, in ICX 6650 devices), line-rate limiting could take up to 43 seconds to take effect.
- You can enable rate limiting on Static LAG only. You cannot enable rate limiting on other types of LAG.
- You can configure rate limiting on individual ports of the LAG. You cannot configure rate limiting on the LAG itself.

## Configuring a port-based rate limiting policy

To configure rate limiting on a port, enter commands such as the following.

```
Device(config)#interface ethernet 1/2/4
Device(config-if-e1000-1/2/4)#rate input fixed 500
```

These commands configure a fixed rate limiting policy that allows port 24 to receive a maximum of 500 kbits (500 packets in ICX 6650 devices) per second. If the port receives additional bytes (packets in ICX 6650 devices) during a given one-second interval, the port drops all inbound packets on the port until the next one-second interval starts.

**Syntax:** `[no] rate-limit input fixed average-rate`

For FastIron devices, the *average-rate* parameter specifies the maximum number of kilobits per second (kbps) (packets per second (pkts/s) in ICX 6650 devices) the port can receive. The minimum rate that can be configured is 64 kbps (125 pkts/s in ICX 6650 devices).

On Brocade ICX 7750 devices, the syntax is:

**Syntax:** `[no] rate-limit input fixed average-rate [ burst burst-size ]`

The optional burst parameter specifies the burst size in kilobits.

By default, the least burst size is set to accommodate five jumbo frames. The upper limit depends on the burstiness of traffic in a deployment scenario and should be user configured.

### NOTE

When traffic reaches the rate limiting threshold, TD2 sends traditional pause or PFC frames, depending on the flow-control configuration.

When PFC is enabled, TD2 transmits PFC for all priorities mapped to the lossless priority group that reaches the XOFF limit (TD2 chip limitation).

## Configuring an ACL-based rate limiting policy

IP ACL-based rate limiting of inbound traffic provides the facility to limit the rate for IP traffic that matches the permit conditions in extended IP ACLs. This feature is available in the Layer 2 and Layer 3 code.

To configure ACL-based rate limiting on a Brocade device, you create individual *traffic policies*, then reference the traffic policies in one or more ACL entries (also called clauses or statements). The traffic policies become effective on ports to which the ACLs are bound.

For configuration procedures for ACL-based rate limiting, refer to *Traffic Policies* chapter.

## Displaying the fixed rate limiting configuration

To display the fixed rate limiting configuration on the device, enter the **show rate-limit input** command.

```
device#show rate-limit input
Total rate-limited interface count: 5.
  Port      Configured Input Rate      Actual Input Rate
  1/1/1      65000                       65000
  1/1/2      195000                      195000
  1/1/6      1950                        1950
  1/5/2      230432                      230000
  1/5/6      234113                      234000
```

### Syntax: show rate-limit input

The command lists the ports on which fixed rate limiting is configured, and provides the information listed in the following table for each of the ports.

**TABLE 27** CLI display of fixed rate limiting information

Field	Description
Total rate-limited interface count	The total number of ports that are configured for fixed rate limiting.
Port	The port number in unit/slot/port format.
Configured Input Rate	The maximum rate requested for inbound traffic. The rate is measured in kilobits per second (kbps).
Actual Input Rate	The actual maximum rate provided by the hardware. The rate is measured in bps.

To display the fixed rate limiting configuration on the ICX 6650 device, enter the **show rate-limit input** command.

```
device#show rate-limit input
Total rate-limited interface count: 5.
  Port      Configured Input Rate      Actual Input Rate
  1/1/1      65000 pkts/s              65000 pkts/s
  1/1/2      195000 pkts/s             195000 pkts/s
  1/1/6      1950 pkts/s               1950 pkts/s
  1/5/2      230432 pkts/s             230000 pkts/s
  1/5/6      234113 pkts/s             234000 pkts/s
```

### Syntax: show rate-limit input

The command lists the ports on which fixed rate limiting is configured, and provides the information listed in the following table for each of the ports.

**TABLE 28** CLI display of fixed rate limiting information

Field	Description
Total rate-limited interface count	The total number of ports that are configured for fixed rate limiting.
Port	The port number in unit/slot/port format.
Configured Input Rate	The maximum rate requested for inbound traffic. The rate is measured in packets per second (pkts/s).
Actual Input Rate	The actual maximum rate provided by the hardware. The rate is measured in pkts/s.

## Rate shaping overview

Outbound Rate Shaping is a port-level feature that is used to shape the rate and control the bandwidth of outbound traffic on a port. This feature smooths out excess and bursty traffic to the configured maximum limit before it is sent out on a port. Packets are stored in



available buffers and then forwarded at a rate no greater than the configured limit. This process provides for better control over the inbound traffic of neighboring devices.

The device has one global rate shaper for a port and one rate shaper for each port priority queue. Rate shaping is done on a single-token basis, where each token is defined to be 1 byte (1 packet for ICX 6650 devices).

## Configuration notes for rate shaping

The following rules apply when configuring outbound rate shapers:

- Outbound rate shapers can be configured *only* on physical ports, not on virtual or loopback ports.
- When outbound rate shaping is enabled on a port on an IPv4 device, the port QoS queuing method (*qos mechanism*) will be strict mode. This applies to IPv4 devices only. On IPv6 devices, the QoS mechanism is whatever method is configured on the port, even when outbound rate shaping is enabled.
- You can configure a rate shaper for a port and for the individual priority queues of that port. However, if a port rate shaper is configured, that value overrides the rate shaper value of a priority queue if the priority queue rate shaper is greater than the rate shaper for the port.
- You can configure rate shaping on individual ports of the link aggregation group (LAG). You cannot configure rate shaping on the LAG itself.
- You can enable rate shaping on the individual ports for static and dynamic LAG. You cannot enable rate shaping on other types of LAG (for example, keep-alive).
- You cannot configure rate shaping on devices where the packet-forwarding method is cut-through. On ICX 7750 devices, the default packet-forwarding method is cut-through so you must first configure the **store-and-forward** command to change the method to store-and-forward. On ICX 7450 devices, store-and-forward is the default method.

### NOTE

You must save the configuration and reload for the change to take effect. See the description of the **store-and-forward** command for more information.

- When configuring rate shaping on dynamic LAG for ICX 7750 devices, you should configure the queues where Link Aggregation Control Protocol (LACP) packets are not forwarded. If you configure rate shaping on dynamic LAG ports (either on the port or on the queue), LACP packets can be dropped and cause a dynamic LAG failure

For more information on LAG configuration, refer to the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide*.

The configured rate shaper values are rounded up to the nearest multiples of minimum values supported on the platform. The following table shows the minimum and maximum values for output rate shaping on various devices. Values are in Kbps for all the platforms except those for ICX 6650 devices, which are in pkts/s.

**TABLE 29** Output rate shaping on FastIron devices

Device	Module	Minimum	Maximum
ICX 6610	1 Gbps ports	3	999750
ICX 6610	10 Gbps ports	3388	9996513
ICX 6430	All	55	890232
ICX 6450	All	73	999868
ICX 6650	40 Gbps ports	20	80000000
ICX 6650	10 Gbps ports	20	20000000
ICX 7750	10 Gbps ports	128 kbps	10,000,000 kbps
ICX 7750	40 Gbps ports	8 kbps	40,000,000 kbps

**TABLE 29** Output rate shaping on FastIron devices (continued)

Device	Module	Minimum	Maximum
ICX 7450	1 Gbps ports	8	999936
ICX 7450	10 Gbps ports	128 kbps	10,000,000 kbps
ICX 7450	40 Gbps ports	8 kbps	40,000,000 kbps
ICX 7250	1 Gbps ports	8	999936
ICX 7250	10 Gbps ports	128 kbps	10,000,000 kbps
FCX	1 Gbps ports	89	999666
FCX	10 Gbps ports	3388	9996513
FSX	SX-FI48GPP module	51	999485
FSX	SX-FI-24GPP module	51	999485
FSX	SX-FI424P module	651	999936
FSX	SX-FI-2XG module	3	9999000
FSX	SX-FI-8XG module	3	9999000

## Configuring outbound rate shaping for a port

To configure the maximum rate at which outbound traffic is sent out on a port, enter commands such as the following.

```
Device(config)#interface ethernet 1/1/2
Device(config-if-e1000-1/1/2)#rate-limit output shaping 1300
```

The configured 1300 Kbps outbound rate shaping on port 1/1/2 is rounded up to the nearest multiple of the minimum value on the device.

On ICX 6650 devices, the *value* is specified in pkts/s.

**Syntax:** `[no] rate-limit output shaping value`

## Configuring outbound rate shaping for a specific priority

To configure the maximum rate at which outbound traffic is sent out on a port priority queue, enter commands such as the following.

```
Device(config)#interface ethernet 1/1/2
device(config-if-e1000-1/1/2)#rate-limit output shaping 500 priority 7
```

The configured 500 Kbps limit for outbound traffic on Priority queue 7 on port 1/1/2 is rounded up to the nearest multiple of 651 Kbps.

On ICX 6650 devices, the *value* specified is in pkts/s.

**Syntax:** `[no] rate-limit output shaping value priority priority-queue`

Specify 0-7 for *priority-queue*.

## Configuring outbound rate shaping for a LAG port

This feature is supported on individual ports of a LAG group.

To configure the maximum rate at which outbound traffic is sent out on a LAG port, enter the following on each LAG port where outbound traffic will be shaped.

```
device(config)#lag lag1 static
device(config-lag-lag1)#rate-limit output shaping ethe 1/1/5 651
```

**Syntax:** `[no] rate-limit output shaping ethernet unit/slot/port ratelimit`

## Displaying rate shaping configurations

To display the configured outbound rate shaper on a device, enter the `show rate-limit output-shaping` command.

The following is a sample output from an ICX 6650 device. The display lists the ports on a device, the configured outbound rate shaper on a port and for a priority for a port.

```
Device#show rate-limit output-shaping
Outbound Rate Shaping Limits in packets/s:
Port  PortMax  Prio0  Prio1  Prio2  Prio3  Prio4  Prio5  Prio6  Prio7
1/1/1  -        -      -      -      -      -      -      -      651
1/1/2  1302    -      -      -      -      -      -      -      -
1/1/3  651     -      -      -      -      -      -      -      -
```



# Limiting broadcast, unknown unicast, and multicast traffic

---

- Configuration notes and feature limitations..... 53
- Command syntax for packet-based limiting .....53
- Command syntax for byte-based limiting..... 55
- Viewing broadcast, multicast, and unknown unicast limits.....55
- BUM suppression port dampening..... 58

## Configuration notes and feature limitations

Brocade devices can forward all flooded traffic at wire speed within a VLAN. However, some third-party networking devices cannot handle high rates of broadcast, multicast, or unknown-unicast traffic. If high rates of traffic are being received by the Brocade device on a given port of that VLAN, you can limit the number of broadcast, multicast, or unknown-unicast packets or bytes received each second on that port. This can help to control the number of such packets or bytes that are flooded on the VLAN to other devices.

Brocade ICX 7250, ICX 7450, and ICX 7750 devices support byte-based and packet-based rate-limiting.

## Command syntax for packet-based limiting

### On FastIron X-Series devices

To enable broadcast limiting on a group of ports by counting the number of packets received, enter commands such as the following.

```
device(config)# interface ethernet 1 to 8
device(config-mif-e1000-1-8)# broadcast limit 65536
```

These commands configure packet-based broadcast limiting on ports 1 - 8. On each port, the maximum number of broadcast packets per second cannot exceed 65,536 packets per second.

To include multicasts in the 65536 packets per second limit on each of the ports, enter the **multicast limit** command after enabling broadcast limiting.

```
device(config-mif-e1000-1-8)# multicast limit
```

To enable unknown unicast limiting by counting the number of packets received, enter commands such as the following.

```
device(config)# interface ethernet 1

device(config-if-e1000-1)# unknown-unicast limit 65536
The combined number of inbound Unknown Unicast packets permitted
for ports 1 to 12 is now set to 65536
device(config-if-e1000-1)#
```

#### NOTE

On the SX-FI48GPP module, multicast and unknown-unicast limiting use the value defined in broadcast limiting. You cannot set a separate limit for unknown-unicast limiting and multicast limiting.

**Syntax:** `[no] broadcast limit num`

**Syntax:** `[no] multicast limit`

**Syntax:** `[no] unknown-unicast num`

The *num* variable specifies the maximum number of packets per second. It can be any number that is a multiple of 8192, up to a maximum value of 2147418112. If you enter the **multicast limit** or **unknown-unicast limit** command, multicast packets or unknown-unicast limit are included in the corresponding limit. If you specify 0, limiting is disabled. If you specify a number that is not a multiple of 8192, the software rounds the number to the next multiple of 8192. Limiting is disabled by default.

## On Brocade FCX Series, ICX 6430 , and ICX 6650 devices

To enable broadcast limiting on a group of ports by counting the number of packets received, enter commands such as the following.

```
device(config)# interface ethernet 1/1/1 to 1/1/8
device(config-mif-e1000-1/1/1-1/1/8)# broadcast limit 65536
```

To include unknown unicast limiting by counting the number of packets received, enter commands such as the following.

```
device(config-mif-e1000-1/1/1-1/1/8)# unknown-unicast limit
```

To include multicasts limiting, enter the **multicast limit** command after enabling broadcast limiting.

```
device(config-mif-e1000-1-8)# multicast limit
```

**Syntax:** `[no] broadcast limit num`

**Syntax:** `[no] multicast limit`

**Syntax:** `[no] unknown-unicast limit`

The *num* variable specifies the maximum number of packets per second. It can be any number that is a multiple of 65536, up to a maximum value of 2147418112. If you enter the **multicast limit** command, multicast packets are included in the corresponding limit. If you specify 0, limiting is disabled. If you specify a number that is not a multiple of 65536, the software rounds the number to the next multiple of 65536. Limiting is disabled by default.

## On Brocade ICX 6610 and ICX 6450 devices

To enable broadcast limiting on a group of ports by counting the number of bits received, enter commands such as the following.

```
Device(config)# interface ethernet 1/1/1 to 1/1/8
Device(config-mif-e1000-1/1/1-1/1/8)# broadcast limit 8192
```

To include unknown-unicast limiting, enter the **unknown-unicast limit** command after enabling broadcast limiting.

```
Device(config-mif-e1000-1/1/1-1/1/8)# unknown-unicast limit
```

To include multicasts limiting, enter the **multicast limit** command after enabling broadcast limiting.

```
Device(config-mif-e1000-1-8)# multicast limit
```

**Syntax:** `[no] broadcast limit num`

**Syntax:** `[no] multicast limit`

**Syntax:** `[no] unknown-unicast limit`

The *num* variable specifies the maximum number of Kilobits per second. It can be any number that is a multiple of 8192, up to a maximum value of 2147418112. If you enter the **multicast limit** or **unknown-unicast limit** command, multicast or unknown-unicast packets are included in the corresponding limit. If you specify 0, limiting is disabled. If you specify a number that is not a multiple of 8192, the software rounds the number to the next multiple of 8192. Limiting is disabled by default.

## Command syntax for byte-based limiting

### NOTE

Byte-based limiting is supported only on FSX devices with SXR 800 and SXR 1600 generation 2 line cards. Byte-based limiting is not supported on the Brocade FCX Series, ICX 6430 devices and the SX-FI8XG, SX-FI-2XG, SX-FI-24HF, SX-FI-24GPP, and SX-FI-48GPP modules.

### NOTE

Refer to the "On Brocade ICX 6610 and ICX 6450 devices" section for information on the command syntax for byte-based limiting on ICX 6610 and ICX 6450 devices.

*Byte-based limiting* provides the ability to rate limit traffic based on byte count. When the byte mode is enabled, packets will be received on a port as long as the number of bytes received per second is less than the corresponding limit. Once the limit is reached, further packets will be dropped.

To enable broadcast limiting on a group of ports by counting the number of bytes received, enter commands such as the following.

```
Device(config)# interface ethernet 9 to 10
Device(config-mif-e1000-9-10)# broadcast limit 131072 bytes
```

These commands configure byte-based broadcast limiting on ports 9 and 10. On each port, the total number of bytes received from broadcast packets cannot exceed 131,072 per second.

To include multicasts in the 131072 bytes per second limit on each of the ports, enter the **multicast limit** command after enabling broadcast limiting.

```
Device(config-mif-e1000-1-8)# multicast limit
```

To enable unknown unicast limiting, enter commands such as the following.

```
Device# config terminal
Device(config)# interface ethernet 13
Device(config-if-e1000-13)# unknown-unicast limit 65536 bytes
The combined number of bytes of inbound Unknown Unicast packets
permitted for ports 13 to 24 is now set to 65536
Device(config-if-e1000-13)#
```

**Syntax:** `[no] broadcast limit num bytes`

**Syntax:** `[no] multicast limit`

**Syntax:** `[no] unknown-unicast limit num bytes`

The *num* variable can be any number that is a multiple of 65536, up to a maximum value of 2147418112. If you enter the **multicast limit** command, multicast packets are included in the limit you specify. If you specify 0, limiting is disabled. If you specify a number that is not a multiple of 65536, the software rounds the number to the next multiple of 65536. Limiting is disabled by default. The **unknown-unicast limit *num* bytes** command is supported on FSX devices.

## Viewing broadcast, multicast, and unknown unicast limits

You can use the **show run interface** command to display the broadcast, multicast, and unknown-unicast limits configured on the device.

You can use the following commands, in addition to the **show run interface** command, to display the broadcast, multicast, and unknown-unicast limits configured on the device:

- **show rate-limit unknown-unicast**
- **show rate-limit broadcast**

**NOTE**

The **show rate-limit unknown-unicast** command is supported only on FSX and the ICX 7750 devices.

Use the **show run interface** command to view the broadcast, multicast, and unknown-unicast limit configured on each port.

```
Device# show run interface
interface ethernet 4
broadcast limit 1245184 bytes
multicast limit
!
interface ethernet 5
broadcast limit 1245184 bytes
multicast limit
!
interface ethernet 12
unknown-unicast limit 524288
!
interface ethernet 13
unknown-unicast limit 65536 bytes
!
interface ethernet 14
broadcast limit 65536
!
interface ethernet 23
broadcast limit 131072
multicast limit
!
```

**Sample output - ICX 7750 devices**

```
Device# show run interface

interface management 1
no ip dhcp-client enable
ip address 10.21.113.8 255.255.255.128
!
interface ethernet 1/1/30
broadcast limit 96
multicast limit 96
unknown-unicast limit 96
```

**Sample output - ICX 6650 devices**

```
Device# show run interface
interface management 1
ip address 10.21.113.7 255.255.248.0
!
interface ethernet 1/1/1
broadcast limit 65536
multicast limit
unknown-unicast limit
!
interface ethernet 1/1/2
broadcast limit 65536
multicast limit
unknown-unicast limit
!
interface ethernet 1/1/3
broadcast limit 65536
multicast limit
unknown-unicast limit
!
interface ethernet 1/1/4
broadcast limit 65536
multicast limit
unknown-unicast limit
!
interface ethernet 1/1/5
broadcast limit 65536
multicast limit
```



```

unknown-unicast limit
!
interface ethernet 1/1/6
broadcast limit 65536
multicast limit
unknown-unicast limit
!
interface ethernet 1/1/7
broadcast limit 65536
multicast limit
unknown-unicast limit
!
interface ethernet 1/1/8
broadcast limit 65536
multicast limit
unknown-unicast limit
!
interface ethernet 1/2/1
optical-monitor 0
!
!
interface tunnel 1
!

```

**Syntax: show run interface**

Use **show rate-limit multicast** command on the ICX 7750 to display the multicast limit for each port to which it applies.

**NOTE**

This command applies only to ICX 7750 devices.

```
Device# show rate-limit multicast
```

```
Multicast Rate Limit Settings:
Port      Limit  Packets/Kbps
1/1/30    96     Packets
```

Use the **show rate-limit unknown-unicast** command to display the unknown unicast limit for each port region to which it applies.

```
Device# show rate-limit unknown-unicast
Unknown Unicast Limit Settings:
Port Region Combined Limit Packets/Bytes
1/1/1 - 12      524288      Packets
1/1/3 - 24      65536       Bytes
```

**Syntax: show rate-limit unknown-unicast**

Use the **show rate-limit broadcast** command to display the broadcast limit or broadcast and multicast limit for each port to which it applies.

**NOTE**

For ICX 7750 devices, this command shows only the broadcast limit.

Sample output - ICX 6650 devices

```
Device# show rate-limit broadcast
Broadcast/Multicast/Unknown Unicast Limit Settings:
Port  Limit  Packets/Bytes  Packet Type(s)
1/1/1  65536  Packets        Broadcast + Multicast + Unknown Unicast
1/1/2  65536  Packets        Broadcast + Multicast + Unknown Unicast
1/1/3  65536  Packets        Broadcast + Multicast + Unknown Unicast
1/1/4  65536  Packets        Broadcast + Multicast + Unknown Unicast
1/1/5  65536  Packets        Broadcast + Multicast + Unknown Unicast
1/1/6  65536  Packets        Broadcast + Multicast + Unknown Unicast
1/1/7  65536  Packets        Broadcast + Multicast + Unknown Unicast
1/1/8  65536  Packets        Broadcast + Multicast + Unknown Unicast
```

Sample output - ICX 7750 devices

```
Device# show rate-limit broadcast

Broadcast Limit Settings:
Port          Limit  Packets/Bytes
1/1/30       96    Packets
```

**Syntax:** `show rate-limit broadcast`

## BUM suppression port dampening

The BUM suppression port dampening feature allows you to monitor BUM traffic drops in a port for a configured time span.

Rate limiting of BUM traffic is used to protect a switch, router node, or network from Denial of Service (DoS) attacks or unintentional excess traffic conditions. If the ingress BUM traffic exceeds the configured rate limit value, the excess traffic is dropped. If the traffic drop count exceeds a set number within a set time interval, the port is shutdown (dampened) for a user configured period.

## BUM suppression port dampening restrictions and limitations

Consider these restrictions and limitations when configuring BUM suppression port dampening.

- MAC-filter cannot be configured along with this feature.
- This feature does not work when configured with the route-only feature.
- This feature is not supported for BUM rate limit configured in pps mode.
- This feature cannot be enabled if hardware resources are insufficient for configuring a BUM ACL.
- The BUM packet drop is read for every **rate-limit-log** interval. Unlike the BUM suppression syslog feature, configuring **rate-limit-log 0** does not disable this feature instead the command would set the polling period of BUM suppression - port dampening feature to the default interval of 300 seconds.
- If the port shutdown period is configured as 0, the port is left in ERR-DISABLED state until you manually re-enable it.
- If this feature is enabled in primary port of a LAG, when the threshold is exceeded, the operationally UP ports in the LAG are placed in ERR-DISABLED state.
- Considering CPU load and resource usage, a 30 second timer is used for monitoring the ports in an ERR-DISABLED state. As a result, there is a delay of up to a maximum 30 seconds while recovering the port from an ERR-DISABLED state.
- If **errdisable recovery cause all** and **errdisable recovery interval time** are configured along with this feature, the ERR-DISABLED port is recovered after the minimum of **port shutdown period** and **errdisable recovery interval** time setting.
- The SNMP trap for BUM suppression - port dampening is not generated by default. The following command must be enabled in global configuration mode to obtain the syslog notification trap: **snmp-server enable traps syslog**.

## Enabling BUM suppression port dampening

Follow these steps to enable BUM suppression port dampening.

1. Enter global configuration mode.

```
device# configure terminal
```

2. If you want to receive SNMP traps, enable the SNMP Syslog notification traps.

```
device(config)# snmp-server enable traps syslog
```

The SNMP trap for the BUM suppression port dampening feature is generated using Syslog notification traps, these must be enabled as they are not generated by default.

3. Globally configure the log interval.

```
device(config)# rate-limit-log 6
```

If this command is not configured, a default polling interval of 5 minutes is used.

4. Enter interface configuration mode

```
device(config)# interface ethernet 1/1/17
```

5. Enable BUM suppression port dampening limits.

- a) Enable broadcast suppression port dampening.

```
device(config-if-e10000-1/1/17)# broadcast limit 100 kbps threshold 1000 action port-shutdown
```

As no value is set for the port-shutdown parameter, the default of 300 seconds is applied.

- b) Enable multicast suppression port dampening.

```
device(config-if-e10000-1/1/17)# multicast limit 1000 kbps threshold 1000 action port-shutdown 600
```

- c) Enable unknown unicast suppression port dampening.

```
device(config-if-e10000-1/1/17)# unknown-unicast limit 1000 kbps threshold 1000 action port-shutdown 240
```

6. Verify the configuration.

```
device(config-if-e10000-1/1/17)# show running-config | begin 1/1/17
interface ethernet 1/1/17
  broadcast limit 100 kbps threshold 1000 action port-shutdown
  multicast limit 1000 kbps threshold 1000 action port-shutdown 600
  unknown-unicast limit 1000 kbps threshold 1000 action port-shutdown 240
...
```

7. Display information for a port that is in an ERR-DISABLED state.

```
device(config-if-e10000-1/1/17)# show internet ethernet 1/1/17
GigabitEthernet1/1/17 is ERR-DISABLED (broadcast traffic threshold exceeded), line protocol is down
Port down for 40 second(s)
Hardware is GigabitEthernet, address is cc4e.248c.6848 (bia cc4e.248c.6848)
Configured speed auto, actual unknown, configured duplex fdx, actual unknown
Configured mdi mode AUTO, actual unknown
EEE Feature Disabled
```

## BUM suppression – port dampening configuration example

```
device# configure terminal
device(config)# snmp-server enable traps syslog
device(config)# interface ethernet 1/1/17
device(config-if-e10000-1/1/17)# broadcast limit 100 kbps threshold 1000 action port-shutdown
device(config-if-e10000-1/1/17)# multicast limit 1000 kbps threshold 1000 action port-shutdown 600
device(config-if-e10000-1/1/17)# unknown-unicast limit 1000 kbps threshold 1000 action port-shutdown 240
device(config-if-e10000-1/1/17)# show running-config | begin 1/1/17
```



# Traffic Policies

---

• Traffic policies overview.....	61
• Maximum number of traffic policies supported on a device.....	62
• ACL-based rate limiting using traffic policies.....	63
• ACL statistics and rate limit counting.....	69
• Viewing traffic policies.....	72
• CPU rate limiting.....	74

## Traffic policies overview

This chapter describes how traffic policies are implemented and configured on the FastIron devices.

Traffic policies are rules that define rate limits on packets permitted by ACLs. As traffic policies apply rate limits on specific interfaces using ACLs, this method is also called ACL-based rate limiting. The process for applying a traffic policy to an interface involves:

1. Creating a traffic policy
2. Adding a reference to the traffic policy in an ACL entry
3. Binding the ACL associated with this ACL entry to an interface

Brocade devices use traffic policies for the following reasons:

- To rate limit inbound traffic
- To count the packets and bytes per packet to which ACL permit or deny clauses are applied

Traffic policies consist of policy names and policy definitions:

- Traffic policy name - A string of up to eight alphanumeric characters that identifies individual traffic policy definitions.
- Traffic policy definition (TPD) - The command filter associated with a traffic policy name. A TPD can define any one of the following:
  - Rate limiting policy
  - ACL counting policy
  - Combined rate limiting and ACL counting policy (not applicable on ICX 6650 devices)

The maximum number of supported active TPDs is a system-wide parameter and depends on the device you are configuring. The total number of active TPDs cannot exceed the system maximum. Refer to "Maximum number of traffic policies supported on a device."

When you apply a traffic policy to an interface, you do so by adding a reference to the traffic policy in an ACL entry, instead of applying the individual traffic policy to the interface. The traffic policy becomes an active traffic policy or active TPD when you bind its associated ACL to an interface.

To configure traffic policies for ACL-based rate limiting, refer to "Configuring ACL-based fixed rate limiting" and "ACL-based adaptive rate limiting configuration".

To configure traffic policies for ACL counting, refer to "Enabling ACL statistics."

## Configuration notes and feature limitations for traffic policies

Note the following when configuring traffic policies:

- Traffic policies applies to IP ACLs only.

- Traffic policies are supported on FastIron X Series devices, but not on the 10 Gbps Ethernet interfaces of the SX-FI62XG and SX-FI42XG modules.
- The maximum number of supported active TPDs is a system-wide parameter and depends on the device you are configuring. The total number of active TPDs cannot exceed the system maximum. Refer to the “Maximum number of traffic policies supported on a device” section.
- You can reference the same traffic policy in more than one ACL entry within an ACL. For example, two or more ACL statements in ACL 101 can reference a TPD named TPD1.
- You can reference the same traffic policy in more than one ACL. For example, ACLs 101 and 102 could both reference a TPD named TPD1.
- Rate limits and ACL counting are applied at the traffic policy level, and are cumulative across ACLs and ACL entries on which they are applied. However, they are not cumulative across port regions.
- For all types of rate limiting on Brocade ICX 6650 devices (ACL-based; Port-based; and Broadcast, unknown Unicast, and Multicast rate limiting), the minimum value is 125 packets and can be increased in steps of 125 packets.
- To modify or delete an active traffic policy, you must first unbind the ACL that references the traffic policy.
- When you define a TPD (when you enter the **traffic-policy** command), explicit marking of CoS parameters, such as traffic class and 802.1p priority, are not available on the device. In the case of a TPD defining rate limiting, the device re-marks CoS parameters based on the DSCP value in the packet header and the determined conformance level of the rate limited traffic, as shown in the following table.

**TABLE 30** CoS parameters for packets that use rate limiting traffic policies

Packet conformance level	Packet DSCP value	Traffic class and 802.1p priority
0 (Green) or 1 (Yellow)	0 - 7	0 (lowest priority queue)
	8 - 15	1
	16 - 23	2
	24 - 31	3
	32 - 39	4
	40 - 47	5
	48 - 55	6
	56 - 63	7 (highest priority queue)
2 (Red)	N/A	0 (lowest priority queue)

- When you define a TPD, reference the TPD in an ACL entry, and then apply the ACL to a VE in the Layer 3 router code, the rate limit policy is accumulative for all of the ports in the port region. If the VE or VLAN contains ports that are in different port regions, the rate limit policy is applied per port region.

For example, TPD1 has a rate limit policy of 600M and is referenced in ACL 101. ACL 101 is applied to VE 1, which contains ethernet ports 1/1/1 to 1/1/4. Because ethernet ports 1/1/1 and 1/1/2 are in a different port region than ports 1/1/3 and 1/1/4, the rate limit policy will be 600M for ports 1/1/1 and 1/1/2, and 600M for ports 1/1/3 and 1/1/4.

## Maximum number of traffic policies supported on a device

The maximum number of supported active traffic policies is a system-wide parameter and depends on the device you are configuring, as follows:

- By default, up to 1024 active traffic policies are supported on Layer 2 switches. This value is fixed on Layer 2 switches and cannot be modified.

- For FastIron devices other than the FCX, the number of active traffic policies supported on Layer 3 switches varies depending on the configuration and the available system memory. The default value and also the maximum number of traffic policies supported on Layer 3 switches is 50.
- On FCX devices, up to 1024 active traffic policies are supported on Layer 3 switches. This is the default value as well as the maximum value.
- The maximum number of active TPDs (traffic policy definitions) supported by Brocade ICX 6650 is 896.

#### NOTE

On FCX devices, by default 992 of the maximum of 1024 active traffic policies are applied. The other 32 are reserved and the **show traffic** command returns zero references/bindings beyond the 992 traffic policies. The **show default values** command displays the maximum number of traffic conditioners that can be applied, in the `hw-traffic-conditioner` section of the results. The configurable tables and their defaults and maximum values can be obtained using the **show default** command.

## Setting the maximum number of traffic policies supported on a Layer 3 device

#### NOTE

This configuration is supported on FastIron devices with the exception of the FCX and ICX 6650 platforms. Setting the `system-max` for traffic policies is not required on FCX platforms as the default number of traffic policies is also the maximum number.

If desired, you can adjust the maximum number of active traffic policies that a Layer 3 device will support. To do so, enter commands such as the following at the global CONFIG level of the CLI.

```
Device(config)#system-max hw-traffic-conditioner 25
Device(config)#write memory
Device(config)#reload
```

#### NOTE

You must save the configuration and reload the software to place the change into effect.

**Syntax:** `[no] system-max hw-traffic-conditioner num`

The `num` variable is a value from 0 through `n`, where 0 disables hardware resources for traffic policies, and `n` is a number up to 50. The maximum number you can configure depends on the configuration and available memory on your device. If the configuration you enter causes the device to exceed the available memory, the device rejects the configuration and displays a warning message on the console.

#### NOTE

Brocade does not recommend setting the system maximum for traffic policies to 0 (zero), because this renders traffic policies ineffective.

## ACL-based rate limiting using traffic policies

ACL-based rate limiting provides the facility to limit the rate for IP traffic that matches the permit conditions in extended IP ACLs. This feature is available in the Layer 2 and Layer 3 code.

To configure ACL-based rate limiting, you create individual traffic policies, and then reference the traffic policies in one or more ACL entries (also called clauses or statements). The traffic policies become effective on ports to which the ACLs are bound.

When you configure a traffic policy for rate limiting, the device automatically enables rate limit counting, similar to the two-rate three-color marker (trTCM) mechanism described in RFC 2698 for adaptive rate limiting, and the single-rate three-color marker (srTCM)

mechanism described in RFC 2697 for fixed rate limiting. This feature counts the number of bytes and trTCM or srTCM conformance level per packet to which rate limiting traffic policies are applied. Refer to [ACL statistics and rate limit counting](#) on page 69.

You can configure ACL-based rate limiting on the following interface types:

- Physical Ethernet interfaces
- Virtual interfaces
- Trunk ports
- Specific VLAN members on a port (refer to "Applying an IPv4 ACL to specific VLAN members on a port (Layer 2 devices only)" section in the *FastIron Ethernet Switch Security Configuration Guide*).
- A subset of ports on a virtual interface (refer to "Applying an IPv4 ACL to a subset of ports on a virtual interface (Layer 3 devices only)" section in the *FastIron Ethernet Switch Security Configuration Guide*).

## Support for fixed rate limiting and adaptive rate limiting

FastIron devices support the following types of ACL-based rate limiting:

- Fixed rate limiting - Enforces a strict bandwidth limit. The device forwards traffic that is within the limit but either drops all traffic that exceeds the limit, or forwards all traffic that exceeds the limit at the lowest priority level, according to the action specified in the traffic policy.
- Adaptive rate limiting - Enforces a flexible bandwidth limit that allows for bursts above the limit. You can configure adaptive rate limiting to forward traffic, modify the IP precedence of and forward traffic, or drop traffic based on whether the traffic is within the limit or exceeds the limit.

## Configuring ACL-based fixed rate limiting

Use the procedures in this section to configure ACL-based fixed rate limiting. Before configuring this feature, see what to consider in the "Configuration notes and feature limitations for traffic policies" section.

Fixed rate limiting enforces a strict bandwidth limit. The port forwards traffic that is within the limit. If the port receives more than the specified number of fragments in a one-second interval, the device either drops or forwards subsequent fragments in hardware, depending on the action you specify.

To implement the ACL-based fixed rate limiting feature, first create a traffic policy, and then reference the policy in an extended ACL statement. Lastly, bind the ACL to an interface. Complete the following steps.

1. Create a traffic policy. Enter a command such as the following.

```
Device(config)#traffic-policy TPD1 rate-limit fixed 100 exceed-action drop
```

2. Create an extended ACL entry or modify an existing extended ACL entry that references the traffic policy. Enter a command such as the following.

```
Device(config)#access-list 101 permit ip host 10.10.12.2 any traffic-policy TPD1
```

3. Bind the ACL to an interface. Enter commands such as the following.

```
Device(config)#interface ethernet 1/1/5
Device(config-if-31000-1/1/5)#ip access-group 101 in
Device(config-if-e1000-1/1/5)#exit
```



The previous commands configure a fixed rate limiting policy that allows port 1/1/5 to receive a maximum traffic rate of 100 kbps (100 pkts/s for ICX 6650 devices). If the port receives additional bits during a given one-second interval, the port drops the additional inbound packets that are received within that one-second interval.

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit fixed cir-value exceed-action drop [ count ]`

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit fixed cir-value exceed-action permit-at-low-pri [ count | remark-cos [ count ] ]`

**Syntax:** `access-list num {permit | deny... } traffic policy TPDname`

**Syntax:** `[no] ip access-group num in`

#### NOTE

For brevity, some parameters were omitted from the **access-list** syntax.

The software allows you to add a reference to a non-existent TPD in an ACL statement and to bind that ACL to an interface. The software does not issue a warning or error message for non-existent TPDs.

Use the **no** form of the command to delete a traffic policy definition. Note that you cannot delete a traffic policy definition if it is currently in use on a port. To delete a traffic policy, first unbind the associated ACL.

The *TPDname* variable is the name of the traffic policy definition. This value can be eight or fewer alphanumeric characters.

The **rate-limit fixed** *cir-value* parameter specifies that the traffic policy will enforce a strict bandwidth. The *cirvalue* variable is the committed information rate in kbps. This value can be from 64 through 1,000,000 Kbps.

#### NOTE

For ICX 6650 devices, the *cir-value* variable is the committed information rate in packets per second. This value can be from 125 through 15,000,000 packets per second.

The **exceed-action** *action* parameter specifies the action to be taken when packets exceed the configured committed information rate (CIR) value. Refer to the "Specifying the action to be taken for packets that are over the limit" section.

The **remark-cos** parameter sets 802.1p priority of the dropped packet to 0. If the option is specified, then the COS/PCP field value set to 0 for the low priority traffic for any packet exceeding the rate-limit set by the traffic policy. If the option is not used, there is no change in the existing behavior.

The **count** parameter is optional and enables ACL counting. Refer to the "ACL statistics and rate limit counting" section.

## ACL-based adaptive rate limiting configuration

Adaptive rate limiting enforces a flexible bandwidth limit. The port forwards traffic that is within the limit. If the port receives more than the specified number of fragments in a one-second interval, the device either drops or forwards subsequent fragments in hardware, depending on the exceed action you specify.

Use the procedures in this section to configure ACL-based adaptive rate limiting. Before configuring this feature, see what to consider in the "Configuration notes and feature limitations for traffic policies" section.

The following table lists the configurable parameters for ACL-based adaptive rate limiting.

**TABLE 31** ACL based adaptive rate limiting parameters

Parameter	Definition
Committed Information Rate (CIR)	The guaranteed kilobit (packets per second in ICX 6650 devices) rate of inbound traffic that is allowed on a port.
Committed Burst Size (CBS)	The number of bytes per second (packets per second in ICX 6650 devices) allowed in a burst before some packets will exceed the

TABLE 31 ACL based adaptive rate limiting parameters (continued)

Parameter	Definition
	committed information rate. Larger bursts are more likely to exceed the rate limit. The CBS must be a value greater than zero (0). Brocade recommends that this value be equal to or greater than the size of the largest possible IP packet in a stream.
Peak Information Rate (PIR)	The maximum kilobit (packets per second in ICX 6650 devices) rate for inbound traffic on a port. The PIR must be equal to or greater than the CIR.
Peak Burst Size (PBS)	The number of bytes per second (packets per second in ICX 6650 devices) allowed in a burst before all packets will exceed the peak information rate. The PBS must be a value greater than zero (0). Brocade recommends that this value be equal to or greater than the size of the largest possible IP packet in the stream.

If a port receives more than the configured bit or byte (packets per second in ICX 6650 devices) rate in a one-second interval, the port will either drop or forward subsequent data in hardware, depending on the action you specify.

## Configuring ACL-based adaptive rate limiting

To implement the ACL-based adaptive rate limiting feature, first create a traffic policy, and then reference the policy in an extended ACL statement. Lastly, bind the ACL to an interface. Complete the following steps.

1. Create a traffic policy. Enter a command such as the following.

```
Device(config)#traffic-policy TPDfour rate-limit adaptive cir 10000 cbs 1600 pir 20000 pbs 4000
exceed-action drop
```

2. Create a new extended ACL entry or modify an existing extended ACL entry that references the traffic policy. Enter a command such as the following.

```
Device(config)#access-list 104 permit ip host 10.10.12.2 any traffic-policy TPDfour
```

3. Bind the ACL to an interface. Enter commands such as the following.

```
Device(config)#interface ethernet 1/1/7
Device(config-if-e1000-1/1/7)#ip access-group 104 in
Device(config-if-e1000-1/1/7)#exit
```

The previous commands configure an adaptive rate limiting policy that enforces a guaranteed committed rate of 10000 kbps on port 1/1/7 and allows bursts of up to 1600 bytes. It also enforces a peak rate of 20000 kbps and allows bursts of 4000 bytes above the PIR limit. If the port receives additional bits during a given one-second interval, the port drops all packets on the port until the next one-second interval starts.

### NOTE

On ICX 6650 devices, rate-limiting is packet-based.

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit adaptive cir cir-value cbs cbs-value pir pir-value pbs pbs-value exceed-action drop [ count ]`

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit adaptive cir cir-value cbs cbs-value pir pir-value pbs pbs-value exceed-action permit-at-low-pri [ count | remark-cos [ count ] ]`

**Syntax:** `access-list num { permit | deny... } traffic policy TPDname`

**Syntax:** `[no] ip access-group num in`

**NOTE**

For brevity, some parameters were omitted from the **access-list** syntax.

The software allows you to add a reference to a non-existent TPD in an ACL statement and to bind that ACL to an interface. The software does not issue a warning or error message for non-existent TPDs.

Use the **no** form of the command to delete a traffic policy definition. Note that you cannot delete a traffic policy definition if it is currently in use on a port. To delete a traffic policy, first unbind the associated ACL.

The *TPDname* variable is the name of the traffic policy definition. This value can be eight or fewer alphanumeric characters.

The **rate-limit adaptive cir** parameter specifies that the policy will enforce a flexible bandwidth limit that allows for bursts above the limit. The *cir-value* variable is the committed information rate in kbps on all the other supported platform except on ICX 6650 devices, where it is packets/second. Refer to the *ACL based adaptive rate limiting parameters* table.

The **cbs cbs-value** parameter is the committed burst size in bytes on all supported platforms except on ICX 6650 devices, where it is in packets.

The **pir pir-value** parameter is the peak information rate in kbps on all supported platforms except on ICX 6650 devices, where it is packets/second.

The **pbs pbs-value** parameter is the peak burst size in bytes on all supported platforms except on ICX 6650 devices, where it is in packets.

The **exceed-action** parameter specifies the action to be taken when packets exceed the configured values. Refer to the “Specifying the action to be taken for packets that are over the limit” section.

The **count** parameter is optional and enables ACL counting. Refer to the “ACL statistics and rate limit counting” section.

## Inspecting the 802.1p bit in the ACL for adaptive rate limiting

You can configure the Brocade device to rate limit traffic for a specified 802.1p priority value. To do so, complete the following configuration steps.

1. Create an adaptive rate limiting traffic policy. Enter a command such as the following:

```
Device(config)#traffic-policy adap rate-limit adaptive cir 1000 cbs 1000 pir
2000 pbs 10000 exceed-action drop
```

2. Create an IPv4 extended ACL or IPv6 ACL that includes the traffic policy and 802.1p priority matching value. Enter a command such as the following:

```
Device(config)#access-list 136 permit ip any any 802.1p-priority matching 3
traffic-policy adap
```

3. Bind the ACL to an interface. Enter commands such as the following.

```
Device(config)#interface ethernet 1/1/7
Device(config-if-1/1/7)#ip access-group 136 in
Device(config-if-1/1/7)#exit
```

Use the **show access-list accounting** command to view accounting statistics. For more information, refer to the “Viewing ACL and rate limit counters” section.

## Specifying the action to be taken for packets that are over the limit

You can specify the action to be taken when packets exceed the configured CIR value for fixed rate limiting, or the CIR, CBS, PIR, and PBS values for adaptive rate limiting. You can specify one of the following actions:

- Drop packets that exceed the limit.
- Permit packets that exceed the limit and forward them at the lowest priority level.

## Dropping packets that exceed the limit

The ultimate action that a device can take on a packet is to drop the packet. You can apply the drop action on packets that exceed the rate limit in both fixed rate limiting and adaptive rate limiting traffic policies. In fixed rate limiting policies, a packet is dropped only when the packet rate exceeds the CIR limit. Whereas, in adaptive rate limiting policies, a packet is dropped only when the packet rate exceeds PIR limit + PBS within one second.

This section shows some example configurations and provides the CLI syntax for configuring a port to drop packets that exceed the configured limits for rate limiting.

The following example shows a fixed rate limiting configuration:

```
Device(config)#traffic-policy TPD1 rate-limit fixed 10000 exceed-action drop
```

The command sets the fragment threshold at 10,000 packets per second. If the port receives more than 10,000 packets in a one-second interval, the device drops the excess fragments.

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit fixed cir-value exceed-action drop [ count ]`

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit fixed cir-value exceed-action permit-at-low-pri [ count | remark-cos [ count ] ]`

The following example shows an adaptive rate limiting configuration.

```
Device(config)#traffic-policy TPD4 rate-limit adaptive cir 10000 cbs 1600 pir 20000 pbs 4000 exceed-action drop
```

The command configures an adaptive rate limiting policy that enforces a guaranteed committed rate of 10000 kbps (10000 pkts/s in ICX 6650 devices) and allows bursts of up to 1600 bytes (1600 packets in ICX 6650 devices). It also enforces a peak rate of 20000 kbps (20000 pkts/s in ICX 6650 devices) and allows bursts of 4000 bytes (4000 packets in ICX 6650 devices) above the PIR limit. If the port receives additional bits during a given one-second interval, the port drops all packets on the port until the next one-second interval starts.

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit adaptive cir cir-value cbs cbs-value pir pir-value pbs pbs-value exceed-action drop [ count ]`

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit adaptive cir cir-value cbs cbs-value pir pir-value pbs pbs-value exceed-action permit-at-low-pri [ count | remark-cos [ count ] ]`

## Permitting packets that exceed the limit

This section shows some example configurations and provides the CLI syntax for configuring a port to permit packets that exceed the configured limit for rate limiting.

The following example shows a fixed rate limiting configuration.

```
Device(config)#traffic-policy TPD1 rate-limit fixed 10000 exceed-action permit-at-low-pri
```

The command sets the fragment threshold at 10,000 packets per second. If the port receives more than 10,000 packets in a one-second interval, the device takes the specified action. The action specified with this command is to permit excess fragments and forward them at the lowest priority level.

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit fixed cir-value exceed-action drop [ count ]`

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit fixed cir-value exceed-action permit-at-low-pri [ count | remark-cos [ count ] ]`

The following example shows an adaptive rate limiting configuration.

```
Device(config)#traffic-policy TPD4 rate-limit adaptive cir 10000 cbs 1600
pir 20000 pbs 4000 exceed-action permit-at-low-pri
```

The command configures an adaptive rate limiting policy that enforces a guaranteed committed rate of 10000 kbps (10000 pkts/s in ICX 6650 devices) and allows bursts of up to 1600 bytes (1600 packets in ICX 6650 devices). It also enforces a peak rate of 20000 kbps (20000 pkts/s in ICX 6650 devices) and allows bursts of 4000 bytes (4000 packets in ICX 6650 devices) above the PIR limit. If the port receives additional bits during a given one-second interval, the port permits all packets on the port and forwards the packets at the lowest priority level.

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit adaptive cir cir-value cbs cbs-value pir pir-value pbs pbs-value exceed-action drop [ count ]`

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit adaptive cir cir-value cbs cbs-value pir pir-value pbs pbs-value exceed-action permit-at-low-pri [ count | remark-cos [ count ] ]`

## ACL statistics and rate limit counting

ACL statistics, also called *ACL counting*, enables the Brocade device to count the number of packets and the number of bytes per packet to which ACL filters are applied.

Rate limit counting counts the number of bytes and the conformance level per packet to which rate limiting traffic policies are applied. The device uses the counting method similar to the two-rate three-color marker (trTCM) mechanism described in RFC 2698 for adaptive rate limiting, and the single-rate three-color marker (srTCM) mechanism described in RFC 2697 for fixed rate limiting. Rate limit counting is automatically enabled when a traffic policy is enforced (active). You can view these counters using the **show** commands listed in [Viewing traffic policies](#) on page 72.

## Enabling ACL statistics

### NOTE

*ACL statistics* and *ACL counting* are used interchangeably throughout this chapter and mean the same thing.

Use the procedures in this section to configure ACL statistics. Before configuring ACL statistics, see what to consider in [Configuration notes and feature limitations for traffic policies](#) on page 61.

To enable ACL statistics on a device, first create a traffic policy, and then reference the traffic policy in an extended ACL entry. Lastly, bind the ACL to an interface. The ACL counting policy becomes effective on ports to which the ACLs are bound.

You also can enable ACL statistics when you create a traffic policy for rate limiting. Refer to [Enabling ACL statistics with rate limiting traffic policies](#) on page 70.

Complete the following steps to implement the ACL statistics feature.

1. Create a traffic policy. Enter a command such as the following.

```
device(config)#traffic-policy TPD5 count
```

2. Create an extended ACL entry or modify an existing extended ACL entry that references the traffic policy definition. Enter a command such as the following.

```
device(config)#access-list 101 permit ip host 10.10.12.2 any traffic-policy TPD5
```

3. Bind the ACL to an interface. Enter commands such as the following.

```
device(config)#interface ethernet 1/1/4
device(config-if-e1000-1/1/4)#ip access-group 101 in
device(config-if-e1000-1/1/4)#exit
```

The previous commands configure an ACL counting policy and apply it to port 1/1/4. Port 1/1/4 counts the number of packets and the number of bytes on the port that were permitted or denied by ACL filters.

**Syntax:** `[no] traffic-policy TPDname count`

**Syntax:** `ip access-list extended 101`

Device (config-std-nacl)# `{ permit | deny } protocol source destination traffic-policy TPDname`

**Syntax:** `[no] ip access-group num in`

#### NOTE

For brevity, some parameters were omitted from the `access-list` syntax.

The software allows you to add a reference to a non-existent TPD in an ACL statement and to bind that ACL to an interface. The software does not issue a warning or error message for non-existent TPDs.

Use the `no` form of the command to delete a traffic policy definition. Note that you cannot delete a traffic policy definition if it is currently in use on a port. To delete a traffic policy, first unbind the associated ACL.

The `TPDname` variable is the name of the traffic policy definition. This value can be eight alphanumeric characters or fewer.

## Enabling ACL statistics with rate limiting traffic policies

The configuration example in the section "Enabling ACL statistics" shows how to enable ACL counting without having to configure parameters for rate limiting. You also can enable ACL counting while defining a rate limiting traffic policy, as illustrated in the following configuration examples.

To enable ACL counting while defining traffic policies for fixed rate limiting, enter commands such as the following at the global config level of the CLI.

```
Device(config)#traffic-policy TPD1 rate-limit fixed 1000 count
Device(config)#traffic-policy TPD2 rate-limit fixed 10000 exceed-action drop count
```

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit fixed cir-value exceed-action drop [ count ]`

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit fixed cir-value exceed-action permit-at-low-pri [ count | remark-cos [ count ] ]`

See the *FastIron Command Reference* for a description of the `traffic-policy rate-limit fixed` command.

To enable ACL counting while defining traffic policies for adaptive rate limiting, enter commands such as the following at the global config level of the CLI.

```
Device(config)#traffic-policy TPDA4 rate-limit adaptive cir 10000 cbs 1600 pir 20000 pbs 4000 count
Device(config)#traffic-policy TPDA5 rate-limit adaptive cir 10000 cbs 1600 pir 20000 pbs 4000 exceed-action
permit-at-low-pri count
```

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit adaptive cir cir-value cbs cbs-value pir pir-value pbs pbs-value exceed-action drop [ count ]`

**Syntax:** `[no] traffic-policy traffic-policy-def rate-limit adaptive cir cir-value cbs cbs-value pir pir-value pbs pbs-value exceed-action permit-at-low-pri [ count | remark-cos [ count ] ]`

See the *FastIron Command Reference* for a description of the **traffic-policy rate-limit adaptive** command.

## Viewing ACL and rate limit counters

When ACL counting is enabled on the Brocade device, you can use **show** commands to display the total packet count and byte count of the traffic filtered by ACL statements. The output of the **show** commands also displays the rate limiting traffic counters, which are automatically enabled for active rate limiting traffic policies.

Use either the **show access-list accounting traffic-policy** command or the **show statistics traffic-policy** command to display ACL and traffic policy counters. The output of these commands is identical.

### NOTE

In the SX-FI48GPP module only, the output of these commands is identical with one exception. When ACL counting is shown by **show statistics traffic-policy**, the Packet Count is not supported and displays "N/A".

The following example shows the output from the **show access-list accounting** command.

```
Device#show access-list accounting traffic-policy g_voip
Traffic Policy - g_voip:
General Counters:
Port Region#           Byte Count           Packet Count
-----
7 (4/1 - 4/12)         85367040             776064
All port regions       84367040             776064
Rate Limiting Counters:
Port Region#           Green Conformance    Yellow Conformance    Red Conformance
-----
7 (4/1 - 4/12)         329114195612139520  37533986897781760    0
All port regions       329114195612139520  37533986897781760    0
```

The following example shows the output on an ICX 6650 device.

```
Device#show access-list accounting traffic-policy tf125c
Traffic Policy tf125c:
Port Regions:
-----
0 : 1/1/1-1/1/56, 1/3/1-1/3/8, 1/2/1-1/2/4
General Counters:
Port Region# Byte Count Packet Count
-----
7 (4/1 - 4/12) 85367040 776064
All port regions 84367040 776064
Rate Limiting Counters (in Packets):
Port Region# Green Conformance Yellow Conformance Red Conformance
-----
7 (4/1 - 4/12) 551475 224589 0
All port regions 551475 224589 0
```

The following example shows the output on an ICX 7750 device.

```
Device#show statistics traffic-policy abc
Traffic Policy tf125c:

General Counters:
Port Region#           Byte Count           Packet Count
-----
0                       235400192            1839051
All port regions       235400192            1839051

Rate Limiting Counters (in bytes):
Port Region#           Green/Yellow Conformance    Red Conformance
-----
0                       225023872                10376320
All port regs          225023872                10376320
```

**Syntax:** `show access-list accounting traffic-policy [ TPDname ]`

or

**Syntax:** `show statistics traffic-policy [ TPDname ]`

The *TPDname* variable is the name of the traffic policy definition for which you want to display ACL and traffic policy counters.

The following table explains the output of the `show access-list accounting traffic-policy` and `show statistics traffic-policy` commands.

**TABLE 32** ACL and rate limit counting statistics

Parameter	Description
Traffic Policy	The name of the traffic policy.
General Counters	
Port Region #	The port region to which the active traffic policy applies.
Byte Count	The number of bytes (packets in ICX 6650 devices) that were filtered (matched ACL clauses).
Packet Count	The number of packets that were filtered (matched ACL clauses).
Rate Limiting Counters	
Port Region#	The port region to which the active traffic policy applies.
Green Conformance	The number of bytes (packets in ICX 6650 devices) that did not exceed the CIR packet rate.
Yellow Conformance	The number of bytes (packets in ICX 6650 devices) that exceeded the CIR packet rate.
Red Conformance	The number of bytes (packets in ICX 6650 devices) that exceeded the PIR packet rate.

## Clearing ACL and rate limit counters

The Brocade device keeps a running tally of the number of packets and the number of bytes per packet that are filtered by ACL statements and rate limiting traffic policies. You can clear these accumulated counters, essentially resetting them to zero. To do so, use either the `clear access-list accounting traffic-policy` command or the `clear statistics traffic-policy` command.

To clear the counters for ACL counting and rate limit counting, enter commands such as the following.

```
Device(config)#clear access-list accounting traffic-policy CountOne
Device(config)#clear statistics traffic-policy CountTwo
```

**Syntax:** `clear access-list accounting traffic-policy TPDname`

or

**Syntax:** `clear statistics traffic-policy TPDname`

The *TPDname* variable is the name of the traffic policy definition for which you want to clear traffic policy counters.

## Viewing traffic policies

To view traffic policies that are currently defined on the Brocade device, enter the `show traffic-policy` command. The following example shows displayed output. The following table explains the output of the `show traffic-policy` command.

```
Device#show traffic-policy t_voidp
Traffic Policy - t_voidp:
Metering Enabled, Parameters:
    Mode: Adaptive Rate-Limiting
```



```

        cir: 100 kbps,      cbs: 2000 bytes,      pir: 200 kbps,      pbs: 4000 bytes
Counting Not Enabled
Number of References/Bindings:1

```

Sample output from an ICX 6650 device:

```

Device#show traffic-policy t_voip
Traffic Policy - t_voip:
Metering Enabled, Parameters:
Mode: Adaptive Rate-Limiting
cir: 100 Pkts/s, cbs: 2000 Pkts, pir: 200 Pkts/s, pbs: 4000 Pkts
Counting Not Enabled

```

**Syntax:** `show traffic-policy TPDname`

To display all traffic policies, enter the `show traffic-policy` command without entering a TPD name.

**TABLE 33** Traffic policy information

Parameter	Description
Traffic Policy	The name of the traffic policy.
Metering	Shows whether or not rate limiting was configured as part of the traffic policy: <ul style="list-style-type: none"> <li>Enabled - The traffic policy includes a rate limiting configuration.</li> <li>Disabled - The traffic policy does not include a rate limiting configuration.</li> </ul>
Mode	If rate limiting is enabled, this field shows the type of metering enabled on the port: <ul style="list-style-type: none"> <li>Fixed Rate-Limiting</li> <li>Adaptive Rate-Limiting</li> </ul>
cir	The committed information rate, in kbps (pkts/s in ICX 6650 devices), for the adaptive rate limiting policy.
cbs	The committed burst size, in bytes (packets in ICX 6650 devices) per second, for the adaptive rate-imiting policy.
pir	The peak information rate, in kbps (pkts/s in ICX 6650 devices), for the adaptive rate limiting policy.
pbs	The peak burst size, in bytes (packets in ICX 6650 devices) per second, for the adaptive rate limiting policy.
Counting	Shows whether or not ACL counting was configured as part of the traffic policy: <ul style="list-style-type: none"> <li>Enabled - Traffic policy includes an ACL counting configuration.</li> <li>Not Enabled - Traffic policy does not include an ACL traffic counting configuration.</li> </ul>
Number of References/Bindings	<p><b>NOTE</b></p> <p>This field does not apply to FastIron X and ICX 6650 devices.</p> <p>The number of port regions to which this traffic policy applies. For example, if the traffic policy is applied to a trunk group that includes ethernet ports 1/9/9, 1/9/10, 1/9/11, and 1/9/12, the value in this field would be 2, because these four trunk ports are in two different port regions.</p>

## CPU rate limiting

Unnecessary traffic to the switch CPU lowers the efficiency of the CPU and delays handling of other traffic that requires processing. CPU rate limiting is a CPU protection scheme which limits certain traffic types.

CPU rate limiting identifies the traffic type and assigns a maximum rate limit to the traffic type. The traffic types which are subjected to rate limiting include broadcast ARP and other exceptions, such as TTL exceed, IP MTU failed, reverse path check failed, IP fragments, and unsupported tunneling. Each of these types is rate-limited individually.

The following table shows the rate limits for each rate-limited packet type and shows which platforms on which each rate limit applies. These rates cannot be configured by users currently.

**TABLE 34** CPU rate limits for packet type and applicable platforms

Packet type	Rate limit in packets per second	Applicable platforms
ARP	6000	All
IP TTL exceed, or Reverse path check failed	150	All
IP MTU exceed, IP tunnel-terminated packets which are fragmented or has options, or IP tunnel-terminated packets with unsupported GRE tunnel header	3000	All
IP Unicast packets mirrored to CPU due to ICMP redirect	100	All
Bridge packets forward to CPU	5000	FCX and ICX

All currently supported FastIron devices support the CPU rate-limiting feature. However, on the FSX devices, only the following modules support this feature:

- SX-FI-24GPP
- SX-FI-24HF
- SX-FI-2XG
- SX-FI-8XG
- SX-FI48GPP

# ICX 7750 Flow Control and Buffer Management

---

- Priority flow control .....75
- Configuring priority flow control.....76
- Configuring priority flow control on a port.....76
- Configuring priority flow control on a range of ports.....77
- Packet buffer management.....77

## Priority flow control

Brocade's priority flow control (PFC) feature prevents frame loss from congestion by pausing traffic based on the congested priority without affecting the traffic of uncongested priorities.

Flow control enables feedback from a receiver to its sender to communicate buffer availability. PFC implements IEEE 802.1Qbb PFC on Brocade ICX 7750 devices, supporting eight priorities and four priority groups (PGs) that can be subject to flow control independently. You can configure PGs for priority flow control and ingress buffer management.

Because multiple priorities can be mapped to a single PG, congestion on one priority in a PG might generate a pause, stopping transmission of all priorities in the PG. Therefore, it is important to create a custom priority-to-PG map to meet your application needs, using either PFC pause honoring or PFC pause transmission.

### PFC pause honoring

- The MAC decodes the class enable vector field to extract the priorities and pause timer value from the packet.
- The per priority XOFF/XON status is passed to the pausing logic to pause/resume packet scheduling to the corresponding queue of the egress port.

### PFC pause transmission

- Each priority can be mapped to a PG. The mapping is configurable.
- When buffer threshold of a PG exceeds XOFF value, a PFC pause frame is sent. The pause frame is encoded with all priorities that belong to the PG in class enable vector.

A receiver using PFC must predict the potential for buffer exhaustion for a PG and respond by generating an explicit PAUSE frame for that class when that condition arises. At any time the receiver must have enough ingress buffers available to store any packet that might be in flight while the PAUSE frame travels back to the sender and gets processed there. In ICX 7750 devices, the number of ingress buffers is set automatically according to the port speed when PFC is enabled.

### NOTE

Configuring PFC commands interrupts traffic temporarily.

You can configure the **qos priority-to-pg** command to change the default priority to PG mapping. See the description of the **qos priority-to-pg** command for more information on the default mapping.

By default, the Brocade ICX 7750 device boots up with tail-drop mode, which means that packets are dropped at the egress queues during congestion. By default, all ports honor IEEE 802.3x pause. However, when transmission of the 802.3x pause is disabled, PFC is also disabled. You can configure the **symmetrical-flow-control enable** command to enable the transmission of the 802.3x pause. See the description of the **symmetrical-flow-control enable** command for more information on symmetrical flow control (SFC).

**NOTE**

Enabling flow control on ports that have auto-neg enabled causes flap because the port pause capabilities must be advertised and negotiated again with peer.

Ports that have auto-neg disabled do not experience flap.

In FastIron Release 08.0.20 and later releases, the following conditions apply on ICX 7750 devices:

- PFC and SFC are not supported for ports across stack units.
- SFC is not supported for ports across master and slave packet-processor (pp) devices in ICX7450-48 units.

In FastIron Release 08.0.20 and later releases, the following conditions apply on ICX 7450 devices:

- PFC is not supported.
- SFC is not supported for ports across stack units, or for ports across master and slave packet-processor (pp) devices in ICX7450-48 units.

## Configuring priority flow control

Enables priority flow control (PFC) globally and for a priority group.

1. Enable PFC globally.

```
Device(config)# priority-flow-control enable
```

2. Enable PFC for priority group (PG) 0.

```
Device(config)# priority-flow-control 0
```

## Configuring priority flow control on a port

Enables priority flow control (PFC) for a priority group on a port.

PFC must be enabled on at least one priority group (PG) before you can configure the **priority-flow-control enable** command on an interface.

1. Enable PFC globally.

```
Device(config)# priority-flow-control enable
```

2. Enable PFC for PG 0.

```
Device(config)# priority-flow-control 0
```

3. Enable interface configuration mode.

```
Device(config)# interface ethernet 1/1/1
```

4. Enable PFC on the interface.

```
Device(config-if-e10000-1/1/1)# priority-flow-control enable
```

## Configuring priority flow control on a range of ports

Enables priority flow control (PFC) for a priority group on a range of ports.

PFC must be enabled on at least one priority group (PG) before you can configure the **priority-flow-control enable** command on an interface.

1. Enter configuration mode.

```
Device# configure terminal
```

2. Enter interface configuration mode for ports in the range you want to enable PFC

```
Device(config)#interface ethernet 1/1/42 ethernet 1/1/43
```

Enables PFC for the ports in the range specified.

3. Enable PFC globally as described in the "Configuring priority flow control" section.
4. Enable PFC on the required ports.

```
Device(config-mif-1/1/42-1/1/43)# priority-flow-control enable
```

Enables PFC on the specified ports.

5. Display the PFC status for the ports in the range specified.

```
Device(config)# show running interface ethernet 1/1/43 ethernet 1/1/42
interface ethernet 1/1/42
  priority-flow-control enable
!
interface ethernet 1/1/43
  load-interval 30
  priority-flow-control enable
!
```

Confirms that PFC is enabled on the range of ports specified.

## Packet buffer management

Ingress- and egress-buffer management as packet admission mechanisms.

On Brocade ICX 7750 devices, packet memory can support 960 Gbps bandwidth. The total packet memory is 12M bytes. ICX 7750 devices run in cut-through mode, which means that cut-through eligible packets are not buffered. If a packet needs to be buffered, it is buffered after Layer 2 and 3 lookup. Packet priority is classified before buffering.

There are two independent packet admission mechanisms: ingress buffer management and egress buffer management.

### Ingress buffer management

- The ingress buffer mechanism determines whether a packet should be admitted into memory-based on the state of available memory and the amount of buffer resources in use by the ingress port priority group.
- It aims to support fair access to buffering resources while also enabling loss-less operation across a network.
- The memory is logically divided into three sections: guaranteed, shared, and headroom for flow control in fly packets.
- Ingress buffer limits are automatically configured based on user configuration to support either loss-less or tail drop operation.
- You can configure the **qos ingress-buffer-profile** command to configure a share level, which determines the maximum number of buffers a PG can use as a fraction of the total sharing pool. For example, if PG 0 is at level4, it can use up to 1/9 of the total sharing buffers in the sharing pool. The actual number of buffers a PG can use depends on the number currently available in the system. See the description of the **qos ingress-buffer-profile** command for more information on ingress buffer profiles.

### Egress buffer management

- The egress buffer mechanism tracks buffer utilization on a per egress port-priority basis. As these accounting structures reach the limit, packets that are destined to the congested egress port-priority are tail-dropped.
- It aims to support fair access to the buffering resources among congested egress ports.
- Any incoming packet is counted only once per egress port regardless of whether it is unicast or multicast.
- Memory is logically divided into two sections: guaranteed and shared.
- You can configure the **qos egress-buffer-profile** command to configure a share level, which determines the maximum number of buffers an egress queue can use as a fraction of the total sharing pool. For example, if queue 4 is at level4, it can use up to 1/9 of the total sharing buffers in the sharing pool. You can configure eight levels of sharing. The actual number of buffers a queue can use depends on the number currently available in the system. See the description of the **qos egress-buffer-profile** command for more information on egress buffer profiles.