RUCKUS

COMMSCOPE

RUCKUS FastIron
QoS and Traffic Management
Configuration Guide, 10.0.20

Supporting FastIron Software Release 10.0.20

Export Restrictions

These products and associated technical data (in print or electronic form) may be subject to export control laws of the United States of America. It is your responsibility to determine the applicable regulations and to comply with them. The following notice is applicable for all products or technology subject to export control:

These items are controlled by the U.S. Government and authorized for export only to the country of ultimate destination for use by the ultimate consignee or end-user(s) herein identified. They may not be resold, transferred, or otherwise disposed of, to any other country or to any person other than the authorized ultimate consignee or end-user(s), either in their original form or after being incorporated into other items, without first obtaining approval from the U.S. government or as otherwise authorized by U.S. law and regulations.

Disclaimer

THIS CONTENT AND ASSOCIATED PRODUCTS OR SERVICES ("MATERIALS"), ARE PROVIDED "AS IS" AND WITHOUT WARRANTIES OF ANY KIND, WHETHER EXPRESS OR IMPLIED. TO THE FULLEST EXTENT PERMISSIBLE PURSUANT TO APPLICABLE LAW, COMMSCOPE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, TITLE, NON-INFRINGEMENT, FREEDOM FROM COMPUTER VIRUS, AND WARRANTIES ARISING FROM COURSE OF DEALING OR COURSE OF PERFORMANCE. CommScope does not represent or warrant that the functions described or contained in the Materials will be uninterrupted or error-free, that defects will be corrected, or are free of viruses or other harmful components. CommScope does not make any warranties or representations regarding the use of the Materials in terms of their completeness, correctness, accuracy, adequacy, usefulness, timeliness, reliability or otherwise. As a condition of your use of the Materials, you warrant to CommScope that you will not make use thereof for any purpose that is unlawful or prohibited by their associated terms of use.

Limitation of Liability

IN NO EVENT SHALL COMMSCOPE, COMMSCOPE AFFILIATES, OR THEIR OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS, LICENSORS AND THIRD PARTY PARTNERS, BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, EXEMPLARY OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER, EVEN IF COMMSCOPE HAS BEEN PREVIOUSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES, WHETHER IN AN ACTION UNDER CONTRACT, TORT, OR ANY OTHER THEORY ARISING FROM YOUR ACCESS TO, OR USE OF, THE MATERIALS. Because some jurisdictions do not allow limitations on how long an implied warranty lasts, or the exclusion or limitation of liability for consequential or incidental damages, some of the above limitations may not apply to you.

Trademarks

CommScope and the CommScope logo are registered trademarks of CommScope and/or its affiliates in the U.S. and other countries. For additional trademark information see https://www.commscope.com/trademarks. All product names, trademarks, and registered trademarks are the property of their respective owners.

Patent Marking Notice

For applicable patents, see www.cs-pat.com.

Contents

Contact Information, Resources, and Conventions

Contacting RUCKUS Customer Services and Support

The Customer Services and Support (CSS) organization is available to provide assistance to customers with active warranties on their RUCKUS products, and to customers and partners with active support contracts.

For product support information and details on contacting the Support Team, go directly to the RUCKUS Support Portal using https://support.ruckuswireless.com, or go to https://www.ruckusnetworks.com and select **Support**.

What Support Do I Need?

Technical issues are usually described in terms of priority (or severity). To determine if you need to call and open a case or access the self-service resources, use the following criteria:

- Priority 1 (P1)—Critical. Network or service is down and business is impacted. No known workaround. Go to the **Submit a Case** section.
- Priority 2 (P2)—High. Network or service is impacted, but not down. Business impact may be high. Workaround may be available. Go to the **Submit a Case** section.
- Priority 3 (P3)—Medium. Network or service is moderately impacted, but most business remains functional. Click the **CONTACT** tab at the top of the page and explore the **Self-Service Online Help** options.
- Priority 4 (P4)—Low. Requests for information, product documentation, or product enhancements. Click the **CONTACT** tab at the top of the page and explore the **Self-Service Online Help** options.

Open a Case

When your entire network is down (P1), or severely impacted (P2), call the appropriate telephone number listed below to get help:

- Continental United States: 1-855-782-5871
- Canada: 1-855-782-5871
- Europe, Middle East, Africa, Central and South America, and Asia Pacific, toll-free numbers are available at https://support.ruckuswireless.com/contact-us and Live Chat is also available.
- Worldwide toll number for our support organization. Phone charges will apply: +1-650-265-0903

We suggest that you keep a physical note of the appropriate support number in case you have an entire network outage.

Self-Service Resources

The RUCKUS Support Portal at https://support.ruckuswireless.com offers a number of tools to help you to research and resolve problems with your RUCKUS products, including:

- Technical Documentation—https://support.ruckuswireless.com/documents
- Community Forums—https://community.ruckuswireless.com
- Knowledge Base Articles—https://support.ruckuswireless.com/answers
- Software Downloads and Release Notes—https://support.ruckuswireless.com/#products_grid
- Security Bulletins—https://support.ruckuswireless.com/security

Using these resources will help you to resolve some issues, and will provide the Technical Assistance Center (TAC) with additional data from your troubleshooting analysis if you still require assistance through a support case or Return Merchandise Authorization (RMA). If you still require help, open and manage your case at https://support.ruckuswireless.com/case_management.

Document Feedback

RUCKUS is interested in improving its documentation and welcomes your comments and suggestions.

You can email your comments to RUCKUS at #Ruckus-Docs@commscope.com.

When contacting us, include the following information:

- Document title and release number
- Document part number (on the cover page)
- Page number (if appropriate)

For example:

- RUCKUS SmartZone Upgrade Guide, Release 5.0
- Part number: 800-71850-001 Rev A
- Page 7

RUCKUS Product Documentation Resources

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

Release Notes and other user documentation are available at https://support.ruckuswireless.com/documents. You can locate the documentation by product or perform a text search. Access to Release Notes requires an active support contract and a RUCKUS Support Portal user account. Other technical documentation content is available without logging in to the RUCKUS Support Portal.

White papers, data sheets, and other product documentation are available at https://www.ruckusnetworks.com.

Online Training Resources

To access a variety of online RUCKUS training modules, including free introductory courses to wireless networking essentials, site surveys, and products, visit the RUCKUS Training Portal at https://commscopeuniversity.myabsorb.com/. The registration is a two-step process described in this video. You create a CommScope account and then register for, and request access for, CommScope University.

Document Conventions

The following table lists the text conventions that are used throughout this guide.

TABLE 1 Text Conventions

Convention	Description	Example
monospace	Identifies command syntax examples	`device(config)# interface ethernet 1/1/6`
bold	User interface (UI) components such as screen or page names, keyboard keys, software buttons, and field names	On the **Start** menu, click **All Programs**.
italics	Publication titles	Refer to the *RUCKUS Small Cell Release Notes* for more information.

Notes, Cautions, and Safety 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 some information that you must read before continuing with the current action or task.

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.

Command Syntax Conventions

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

Convention	Description
bold text	Identifies command names, keywords, and command options.
italic text	Identifies a variable.
[ ]	Syntax components displayed within square brackets are optional.
	Default responses to system prompts are enclosed in square brackets.
{x| y| z}	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x|y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, *member*[*member*...].
\	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.

About This Document

New In This Document

The following table describes information added or modified in this guide for FastIron 10.0.20.

TABLE 2 Key Features and Enhancements in *FastIron 10.0.20a*

Feature	Description	Reference
Self generated packets	Options to configure DSCP and CoS values for self generated packets.	Configuring DSCP and CoS Marking for the Self-Generated Packets on page 30
Updates to address defects	**Updated**: Minor updates on content throughout to address defects.	All chapters
Minor editorial updates	**Updated**: Minor editorial updates were made throughout the Configuration Guide.	All chapters

TABLE 3 Key Features and Enhancements in *FastIron 10.0.20*

Feature	Description	Reference
Updates to address defects	**Updated**: Minor updates on content throughout to address defects.	All chapters
Minor editorial updates	**Updated**: Minor editorial updates were made throughout the Configuration Guide.	All chapters

Supported Hardware

This guide supports the following RUCKUS products:

- RUCKUS ICX 8200 Switches
- RUCKUS ICX 7850 Switches
- RUCKUS ICX 7650 Switches
- RUCKUS ICX 7550 Switches

For information about what models and modules these devices support, refer to the hardware installation guide for the specific product family.

Quality of Service

Quality of Service Overview

Quality of Service (QoS) provides preferential treatment to specific traffic.

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 based on configured priorities. Traffic can be dropped, prioritized for guaranteed delivery, or subject to the delivery options as configured by several different mechanisms.

Classification is the process of selecting packets on which to perform QoS, reading or ignoring 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 based on information contained within the packet or assigned to the packet as it arrives at the switch. Once a packet or traffic flow is identified and marked, it is mapped to a forwarding priority queue.

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

There are two traffic types in QoS:

- Data—These can be either network-to-network traffic or traffic from the CPU. QoS parameters can be assigned and modified for data traffic. The device also supports setting or modifying the IEEE 802.1p user priority or the IP header DSCP field.

- Control—Packets to and from the CPU is considered control traffic. The QoS parameters associated with the control traffic are preassigned and not configurable.

Classified Traffic Processing

The *trust level* in effect on an interface determines the type of QoS information the device uses for performing QoS.

A RUCKUS ICX 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—If the packet does not match on an ACL that defines a priority and the MAC address of the packet matches a static entry, the packet is classified with the priority of the static MAC entry.

- 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*.

- 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 *RUCKUS FastIron Security Configuration Guide*.

Given the variety of 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 figures in Packet Trust Level on page 12.

Packet Trust Level

The following figure illustrates how 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).

> **NOTE**
> ICX 8200 devices determine internal priority differently. In ICX 8200 devices, ACL matches are first considered, and DSCP/ToS priority is considered next, followed by the priority of the static MAC entry, then default ingress port priority, 802.1p CoS value, and finally the default priority of zero (0).

FIGURE 1 Determining the Trust Level of a Packet for Most ICX Devices
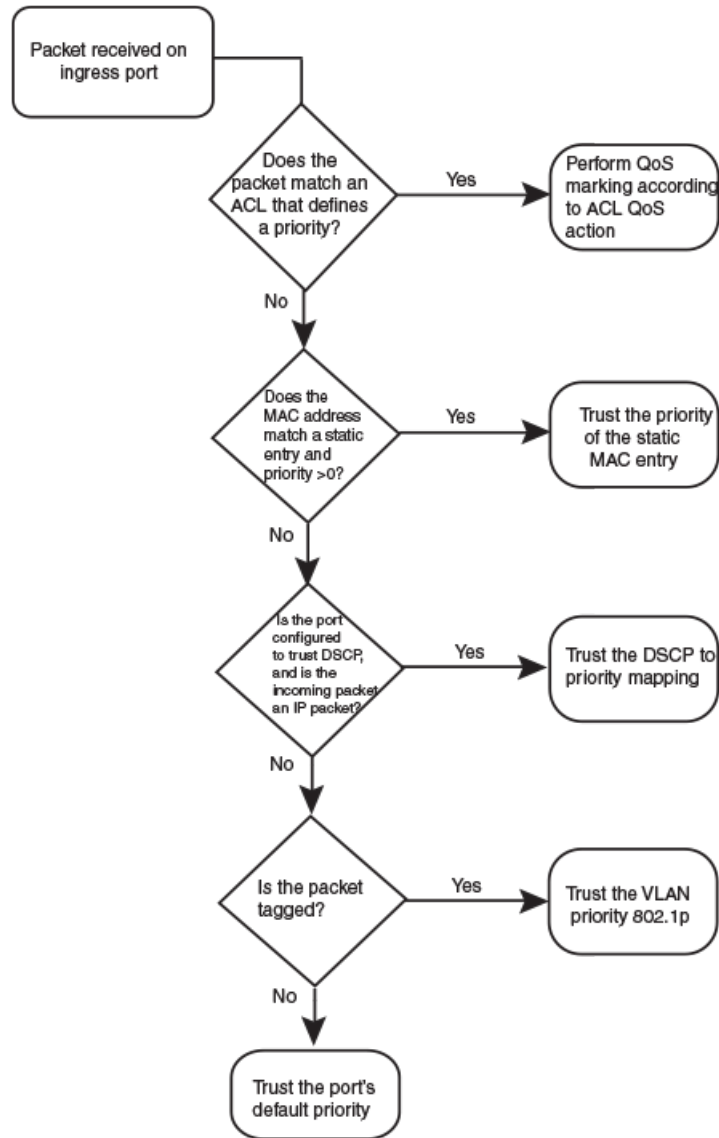
FIGURE 2 Determining the Trust Level of a Packet for ICX 8200 Devices
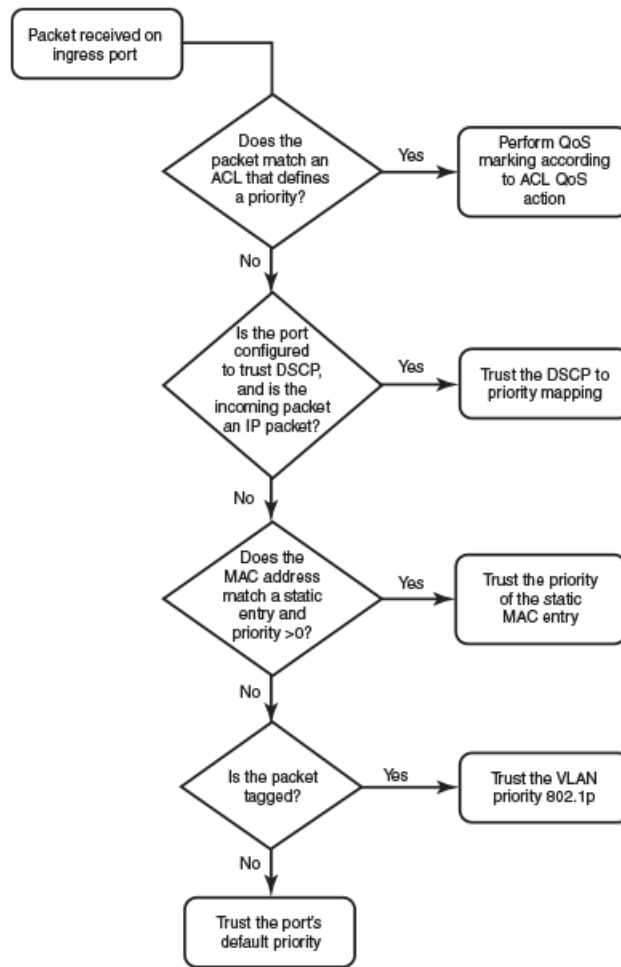
Once a packet is classified, it is mapped to a forwarding queue. There are eight queues designated from 0 through 7. The internal forwarding priority maps to one of these eight queues. The mapping between the internal priority and the forwarding queue cannot be changed.

The following tables show the default QoS mappings for ICX platforms that are used if the trust level for CoS or DSCP is enabled.

TABLE 4 Default QoS Mappings for ICX Platforms, Columns 0 to 15

DSCP value	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
802.1p (CoS) value	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
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 5 Default QoS Mappings for ICX 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
Internal forwarding priority	2	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3

TABLE 5 Default QoS Mappings for ICX Platforms, Columns 16 to 31 (continued)

DSCP value	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Forwarding queue	2	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3

TABLE 6 Default QoS Mappings for ICX 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
Internal forwarding priority	4	4	4	4	4	4	4	4	5	5	5	5	5	5	5	5
Forwarding queue	4	4	4	4	4	4	4	4	5	5	5	5	5	5	5	5

TABLE 7 Default QoS Mappings for ICX 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
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 on page 32.

- 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 on page 33.

QoS for RUCKUS ICX Stackable Devices

RUCKUS 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 *RUCKUS FastIron 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.

QoS Behaviors in a Stack

QoS Profile Restrictions

In a stacking topology, 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 switches.

QoS Behavior for Trusting Layer 2 (802.1p)

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)

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

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

By default, 802.1p marking is not enabled in a traditional stack. Outgoing tagged traffic is not 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 *RUCKUS FastIron Security Configuration Guide*.

QoS Queues

RUCKUS devices support the eight QoS queues (qosp0 through qosp7).

The supported queues are:

TABLE 8 QoS Queues

QoS Priority Level	QoS Queue
0	qosp0 (lowest priority queue)
1	qosp1
2	qosp2
3	qosp3
4	qosp4
5	qosp5
6	qosp6
7	qosp7 (highest priority queue)

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

Packets are classified and assigned to specific queues based on the criteria shown in the figures described in

For ICX devices, ingress packets are classified into the eight priorities, which map to eight hardware queues or traffic classes (TCs) based on the priority.

TABLE 9 QoS Queue Range of the RUCKUS Devices

Device	UC Queue (Unicast)	MC Queue (Multicast)
RUCKUS ICX 7550	queue 0-9	queue 0-9
RUCKUS ICX 7650	queue 0-9	queue 0-9
RUCKUS ICX 7850	queue 0-9	queue 0-9
RUCKUS ICX 8200	queue 0-7 (each queue is for both Unicast + Multicast)	

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 applies to the entire system. If the scheduler profile is configured using the **qos mech strict** command, all devices in the system are configured with the strict priority. The user-configurable scheduler profile applies only to the specific devices, leaving the remaining devices running default QoS. On any device, the user-configurable scheduler profile has priority over the default QoS. The user-configurable scheduler profile should be in line with the default QoS commands in both stacking and standalone systems.

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)

  NOTE
  Limited by the buffer capacity, RUCKUS ICX 8200 switches may drop frames before they are properly scheduled. Under this condition, configure the egress buffer profile to support WRR or Mixed mode and increase the port-share level of the profile.

If you create a profile specifying only the weights 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 following tables show the default values for the scheduling type for stacking and standalone ICX devices.

TABLE 10 Default Values for Scheduling Type for Stacking Systems

Traffic Class	SP	SP Jumbo	WRR	WRR Jumbo	Mixed	Mixed Jumbo
TC 0	SP	SP	3	8	15	15
TC 1	SP	SP	3	8	15	15
TC 2	SP	SP	3	8	15	15
TC 3	SP	SP	3	8	15	15
TC 4	SP	SP	3	8	15	15
TC 5	SP	SP	10	16	25	25
TC 6	SP	SP	75	44	SP	SP
TC 7	SP	SP	SP	SP	SP	SP

TABLE 11 Default Values for Scheduling Type for Standalone Systems

Traffic Class	SP	SP Jumbo	WRR	WRR Jumbo	Mixed	Mixed Jumbo
TC 0	SP	SP	3	8	15	15
TC 1	SP	SP	3	8	15	15
TC 2	SP	SP	3	8	15	15
TC 3	SP	SP	3	8	15	15
TC 4	SP	SP	3	8	15	15
TC 5	SP	SP	3	8	25	25
TC 6	SP	SP	7	8	SP	SP
TC 7	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, specify a number from 0 through 7. The priority number specifies the IEEE 802.1 equivalent to one of the eight QoS queues on RUCKUS devices. The numbers correspond to the queues as shown in the table in the QoS Queues on page 16 section.

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.

Buffer Allocation and Threshold for QoS Queues

By default, RUCKUS FastIron 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 *RUCKUS FastIron Layer 2 Switching Configuration Guide*.

QoS Marking

The term—QoS marking—is the process of initially changing the packet QoS information for the next hop.

Layer 2 (802.1p) and Layer 3 (Differentiated Services Code Point (DSCP)) information in a packet can be marked. As an example of marking traffic coming from a device that does not support DSCP, you can change the packet IP precedence value into a DSCP value before forwarding the packet.

Class of Service (CoS) is a 3-bit field within an Ethernet frame header known as the Priority Code Point (PCP) when using a 802.1 network. This field specifies a priority value between 0 and 7, inclusive, that can be used by Quality of Service (QoS) to differentiate traffic.

The Differentiated Services Code Point (DSCP) is a 6-bit field in an IP header for the classification of packets. Differentiated Services is a technique used to classify and manage network traffic and it helps to provide QoS for modern Internet networks. It can provide services to all kinds of networks.

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 that 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.

DSCP Remarking Overview

Differentiated Services Code Point (DSCP) remarking can be configured using three main types of configuration with different levels of QoS precedence.

There is a debate between using the terms "marking" or "remarking." Almost all devices initially mark the DSCP packets with a value. Every packet has one of 64 values, a decimal number from 0 to 63, in the DSCP field. Each of these values, including 0, is a legitimate DSCP. When the packet is processed by a DSCP marker, we can use the term "remarking" the packet, even though the DSCP may not change.

DSCP remarking is performed on ICX devices using different types of configuration:

- ACL—Traffic matching a specific pattern is remarked.
- Interface (Physical, LAG, VE)—Traffic entering a physical, LAG, or VE interface (except traffic matched by an ACL) is remarked with a configured value.

ACL Remarking

ACLs can be configured to match a specific pattern and remark DSCP values. When remarking is not enabled using ACLs, a rogue host that wants preferential treatment for all its traffic could mark the DSCP field for its requirements and send the traffic to the device.

For more information on QoS marking using ACLs, refer to "QoS options for IP ACLs" in the *RUCKUS FastIron Security Configuration Guide*.

Interface (Physical, LAG, VE) Interface Remarking

Packets entering a physical, LAG, or VE interface can be remarked with a configured DSCP value. Remarking at the interface or VLAN level can be referred to as Class of Service (CoS) remarking although the values set are DSCP values. Remember that DSCP remarking configuration at the ACL level takes precedence over the DSCP remarking configuration at the interface level.

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

For a configuration example of QoS remarking at the interface level, refer to

For information about the QoS remarking using physical, LAG, or VE interfaces in VXLANs, refer to the Quality of Service Support topic in the *RUCKUS FastIron Layer 2 Switching Configuration Guide*.

DSCP-Based QoS Configuration

RUCKUS FastIron releases support basic DSCP-based QoS (also called Type of Service [ToS]-based QoS). However, the FastIron family of switches does not support other advanced DSCP-based QoS features.

RUCKUS FastIron 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 ACLs 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 RUCKUS 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.
- 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 RUCKUS devices to honor DSCP-based QoS for routed and switched traffic.

DSCP-based QoS is not automatically honored for switched traffic. The default is 802.1p to CoS mapping. To honor DSCP-based QoS, you must enter the **trust dscp** command at the interface level of the CLI.

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**
> On ICX 7650 device, 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.

> **NOTE**
> On ICX 8200 devices, configuring both DSCP trust and Traffic Policy is not recommended. When dscp trust and Traffic Policy are configured on the same interface, trust DSCP will not be honored on Layer3 (routed) traffic, and all traffic including green conformance will be sent out on best effort delivery Queue 0.

Remarking Configuration Considerations and Limitations

Keep the following considerations in mind when configuring remarking.

For a router image:

- When DSCP remarking is enabled on a virtual interface (VE), an ingress ACL applied to the VLAN to which the VE belongs honors the DSCP remarking value configured on the VE.
- When PCP remarking is enabled on a VE, an egress ACL applied on the VLAN to which the VE belongs honors the PCP remarking value configured on the VE.

- DSCP/PCP remarking is not supported on tagged interfaces in the router image.
- When DSCP remarking is enabled on a Layer 3 physical or LAG interface, an ingress ACL applied to the physical or LAG interface honors the DSCP remarking value configured for the physical or LAG interface.

QoS Mapping Configuration

You can optionally change the following QoS mappings:

- DSCP to internal forwarding priority
- VLAN priority (802.1p) to hardware forwarding queue, as described in Changing the VLAN Priority 802.1p to Hardware Forwarding Queue Mappings on page 33

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 12 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
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 ICX devices, you can use QoS queue 1 for priority traffic, even when sFlow is enabled on the port.

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

QoS Scheduling and Queuing Methods

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 queuing method.

The following QoS queuing methods are supported for the FastIron devices:

- Weighted Round Robin (WRR)—This method 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 under weighted queuing. Attempts to change the qosp7 setting are 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 FastIron devices are based on bytes. These devices service a given number of bytes (based on weight) in each queue cycle.

- Strict Priority (SP)—This 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 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 device assigns strict priority to traffic in qosp7 and qosp6, and weighted round robin priority to traffic in qosp0 through qosp5. Thus, the 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**
  > Stackable devices that are operating as members of a stack reserve queue 7 for stacking functions. For more information, refer to QoS for RUCKUS ICX Stackable Devices on page 15.

  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.

TABLE 13 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 13 Default Bandwidth for Combined SP and WRR Queueing Methods (continued)

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

IPv6 QoS

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. Refer to IPv6 ACLs in the *RUCKUS FastIron Security Configuration Guide*.

Flow Control and Buffer Management

Using flow control and buffer management techniques, data packet transmission rates and buffer queue capacity can be managed to provide the preferred quality of service (QoS).

Flow control manages the rate of data transmission between two devices to avoid overloading the receiving device with data. Using a technique that allows the receiving device to control the data transmission speed, flow control can prevent data packets from being dropped.

Buffer management controls whether the data packets are channeled to buffer queues before processing or allowed to pass through the device. Packet buffer management uses priorities, and lower priority data traffic is routed to buffers which have a finite amount of memory. If the device buffers are full when a packet arrives, the packet may be dropped.

Priority Flow Control

The RUCKUS implementation of priority flow control (PFC) prevents frame loss from congestion by pausing traffic based on the congested priority without affecting the traffic of uncongested priorities.

  NOTE
  Priority flow control (PFC) is supported on per priority level for RUCKUS ICX 7550 and ICX 7850 devices only.

Flow control enables feedback from a receiver to its sender to communicate buffer exhaustion and buffer availability. A receiver generates a MAC control frame and sends a PAUSE request to a sender when a specified threshold (Xoff value) of the receiver buffer is filled. Upon receiving a PAUSE request, the sender stops transmission of any new packets until the receiver notifies the sender that it has sufficient buffer space (Xon value). The RUCKUS implementation of IEEE 802.1Qbb PFC supports eight priorities and four priority groups (PGs) that can be subjected to flow control independently. You can configure PGs for priority flow control and ingress buffer management.

> **NOTE**
> PFC is disabled by default. It can be enabled by using the **priority-flow-control enable** command.

> **NOTE**
> Priority flow control and 802.3x flow control are mutually exclusive; therefore, configuring the **priority-flow-control enable** command disables 802.3x in both the transmit and receive directions.

> **NOTE**
> PFC is always symmetric. Once it is enabled, it enables both PFC RX and TX.

Because multiple priorities can be mapped to a single PG, congestion on one priority in a PG may generate a pause, stopping transmission of all priorities in that 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.

The ICX 7550 and ICX 7850 devices support a flat queuing structure. Therefore, it is not possible to channel the traffic from different sources with the same priority to separate lossy and loss-less queues. A pause frame applied to this queue affects the traffic from both sources (lossy and loss-less). Therefore, PFC cannot be enabled at the interface level on ICX 7550 and ICX 7850 devices. Instead, the ICX 7550 and ICX 7850 devices support per-priority PFC, in which lossy and loss-less traffic is channeled to different egress queues based on the priorities mapped to the priority groups.

PFC Pause Honoring

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

PFC Pause Transmission

- Priorities 0 through 6 can be mapped to a PG; Priority 7 cannot be mapped.
- The mapping is configurable.
- When the buffer threshold of a PG exceeds the Xoff value, a PFC pause frame is sent. The pause frame is encoded with all priorities that belong to the PG in the 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 may be in flight while the pause frame travels back to the sender and gets processed there. In ICX 7550 and ICX 7850 devices, the number of ingress buffers is set automatically according to the port speed when PFC is enabled.

> **NOTE**
> Configuring PFC commands may temporarily interrupt traffic.

You can configure the **qos priority-to-pg** command to change the default priority to PG mapping.

By default, the ICX 7550 and ICX 7850 devices boot 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.

NOTE

Enabling flow control on ports that have auto-negotiation enabled causes flapping because the port pause capabilities must be advertised and negotiated again with the peer. Ports with auto-negotiation disabled do not experience flapping.

PFC Support on ICX Devices

RUCKUS ICX 7550 devices: Priority flow control (PFC) is supported on per priority level. It is not supported at port level.

RUCKUS ICX 7650 devices: PFC is not supported.

RUCKUS ICX 7850 devices: PFC is supported on per priority level. It is not supported at port level. SFC is not supported for 1-Gbps ports as well as for ports across stack units. Port flow control is not supported in 1-Gbps ports.

RUCKUS ICX 8200 devices: PFC is not supported.

Packet Buffer Management

The following table lists the packet memory bandwidth and the total packet memory on ICX devices.

TABLE 14 Packet Memory on ICX Devices

ICX Device	Total Bandwidth	Total Packet Memory
RUCKUS ICX 7550	348 Gbps (24/24P) 480 Gbps (48/48P/24ZP) 560 Gbps (24F/48F/48ZP)	8 MB
RUCKUS ICX 7650	480 Gbps 564 Gbps (IO 48ZP/48ZF) 328 Gbps (IO 48P)	8 MB
RUCKUS ICX 7850	2.0 Tbps 3.2 Tbps	32 MB
RUCKUS ICX 8200	300 Gbps	3 MB

ICX 7850 devices support cut-through mode although they run in store-and-forward mode by default. If enabled, the cut-through-eligible packets are not buffered. If a packet must be buffered, it is buffered after Layer 2 and Layer 3 lookup. The packet priority is classified before buffering. On ICX 7850 devices, 1-Gbps ports always operate in store-and-forward mode, irrespective of the global switch mode configuration. On ICX 7850-48C devices, cut-through mode is supported only on Module 2 (ports 1/2/1 through 1/2/8 only).

ICX 7850-48C devices do not support cut-through switching on either 1-Gbps or 10-Gbps ports. Furthermore, when cut-through mode is enabled, while 1-Gbps and 10-Gbps ports remain in store-and-forward mode (the default), flow control is disabled on these ports in both directions.

NOTE

The ICX 7650 and ICX 8200 devices support only store-and-forward mode for packet forwarding.

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

Ingress Buffer Management

On the ICX 8200, ingress buffer management tracks buffer utilization on a per-ingress-port basis.

- Loss-less behavior through symmetric flow control is supported.

- Buffers are reserved for high-priority traffic.

As these accounting structures reach their limits, incoming packets to the ingress port are dropped.

On the ICX 7650 and ICX 8200 devices, there is a default profile for ingress buffer management, but it is not configurable because PFC is not supported.

The ingress buffer profile is not supported on the ICX 7550 and ICX 7850 devices.

Egress Buffer Management

This mechanism tracks buffer utilization on a per-egress port and priority basis. As these accounting structures reach their limits, packets that are destined to the congested egress port-priority are tail-dropped. The aim of the mechanism is 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 (on a per-port-priority basis).

- On a per-port-priority basis for the ICX 7550, ICX 7650, ICX 7850, and ICX 8200 devices.

On the RUCKUS devices, sharing is a ratio of the remaining buffers. You can configure the share level to determine the maximum number of buffers that an egress queue can use as a fraction of the total sharing pool. For example, if queue 4 is at level 4, 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 that a queue can use depends on the number currently available in the system.

Configuring QoS

The configuration of QoS includes the following components:

- Port priority

- Static MAC entries to priority queues

- QoS marking

- DSCP to internal forwarding priority mappings

- VLAN priority 802.1p to hardware forwarding queue mappings

- Queuing method

- QoS queue naming and percentage of a port outbound bandwidth guaranteed to the queues

- Minimum bandwidth of WRR queues

- Bandwidth allocation for hybrid WRR and SP queues

- Priority flow control

- Ingress and egress buffer profile

Displaying User-Configurable Scheduler Profile Information

Follow these steps to display configurable scheduler profile information. The information shown may be different from your profile.

1. Display a specific profile.

```
device# show qos scheduler-profile test
User Scheduler Profile: test      Scheduling Option: Mixed-SP-WRR

Ports attached: (U1)    --
Ports attached: (U2)    --
Ports attached: (LAG)   --
Ports attached: (LAG)   -- Unicast per
Queue details:             Bandwidth%
Traffic Class 0                   15%
Traffic Class 1                   15%
Traffic Class 2                   15%
Traffic Class 3                   15%
Traffic Class 4                   15%
Traffic Class 5                   25%
Traffic Class 6                   sp
Traffic Class 7                   sp
Multicast per Queue details: Bandwidth%
Traffic Class 0                   15%
Traffic Class 1                   15%
Traffic Class 2                   15%
Traffic Class 3                   15%
Traffic Class 4                   15%
Traffic Class 5                   25%
Traffic Class 6                   sp
Traffic Class 7                   sp

Minimum Guaranteed Rate:
Unicast per Queue details:  Bandwidth%
Traffic Class 0                    0
Traffic Class 1                    0
Traffic Class 2                    0
Traffic Class 3                    0
Traffic Class 4                    0
Traffic Class 5                    0
Traffic Class 6                    0
Traffic Class 7                    0
```

2. Display all user profiles.

```
device# show scheduler-profile all

User Scheduler Profile: test    Scheduling Option: Mixed-SP-WRR

Ports attached: (U1)      --
Ports attached: (U2)      --
Ports attached: (LAG)     --
Ports attached: (LAG)     --
Unicast per Queue details:   Bandwidth%
Traffic Class 0                 15%
Traffic Class 1                 15%
Traffic Class 2                 15%
Traffic Class 3                 15%
Traffic Class 4                 15%
Traffic Class 5                 25%
Traffic Class 6                 sp
Traffic Class 7                 sp
Multicast per Queue details: Bandwidth%
Traffic Class 0                 15%
Traffic Class 1                 15%
Traffic Class 2                 15%
Traffic Class 3                 15%
Traffic Class 4                 15%
Traffic Class 5                 25%
Traffic Class 6                 sp
Traffic Class 7                 sp

Minimum Guaranteed Rate:    Bandwidth%
Unicast per Queue details:  Bandwidth%
Traffic Class 0                 0%
Traffic Class 1                 0%
Traffic Class 2                 0%
Traffic Class 3                 0%
Traffic Class 4                 0%
Traffic Class 5                 0%
Traffic Class 6                 0%
Traffic Class 7                 0%

User Scheduler Profile: test2    Scheduling Option: Weighted round-robin

Ports attached: (U1)    --
Ports attached: (U2)    --
Ports attached: (LAG)   --
Ports attached: (LAG)   --
Unicast per Queue details:   Bandwidth%
Traffic Class 0                 3%
Traffic Class 1                 3%
Traffic Class 2                 3%
Traffic Class 3                 3%
Traffic Class 4                 3%
Traffic Class 5                 3%
Traffic Class 6                 7%
Traffic Class 7                 75%
Multicast per Queue details: Bandwidth%
Traffic Class 0                 3%
Traffic Class 1                 3%
Traffic Class 2                 3%
Traffic Class 3                 3%
Traffic Class 4                 3%
Traffic Class 5                 3%
Traffic Class 6                 7%
Traffic Class 7                 75%

Minimum Guaranteed Rate:    Bandwidth%
Unicast per Queue details:  Bandwidth%
Traffic Class 0                 0%
Traffic Class 1                 0%
Traffic Class 2                 0%
Traffic Class 3                 0%
```

```
Traffic Class 4                        0%
Traffic Class 5                        0%
Traffic Class 6                        0%
Traffic Class 7                        0%
```

Changing a Port Priority

Follow these steps to change the QoS priority of a specific port.

1. Enter global configuration mode.

    ```
    device# configure terminal
    ```

2. Enter interface configuration mode.

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

3. Set the port priority.

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

    This step assigns priority 7 to untagged switched traffic received on port 1/1/1.

4. Return to privileged EXEC mode.

    ```
    device(config-if-e1000-1/1/1)# end
    ```

5. Verify the configuration.

    ```
    device# show interface brief
    Port       Link     State    Dupl Speed Trunk Tag Pvid Pri MAC            Name
    1/1/1      Down     None     None None  None  Yes 4000 7   cc4e.248b.b050 ERSPAN
    1/1/2      Down     None     None None  None  No  5    0   cc4e.248b.b050
    1/1/3      Down     None     None None  None  No  5    0   cc4e.248b.b052
    1/1/4      Down     None     None None  None  No  5    0   cc4e.248b.b053
    1/1/5      Down     None     None None  2     Yes N/A  0   cc4e.248b.b054
    ...
    ```

    The interface priority is listed under the heading `Pri`.

Changing a Port Priority Configuration Example

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# priority 7
device(config-if-e1000-1/1/1)# end
device# show interface brief
```

Assigning Static MAC Entries to Priority Queues

Follow these steps to configure a static MAC entry and assign the entry to the premium queue.

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.

1. Enter global configuration mode.

    ```
    device# configure terminal
    ```

2. Enter VLAN configuration mode.

```
device(config)# vlan 1
```

3. Assign the priority.

```
device(config-vlan-1)# static-mac-address 0000.0063.67FF ethernet 1/1/1 priority 7
```

4. Return to privileged EXEC mode.

```
device(config-vlan-1)# end
```

5. Verify the MAC address configuration.

```
device# show mac-address ethernet 1/1/1
Total static entries from port 1/1/1 = 1
MAC-Address      Port                 Type          VLAN
0000.0063.67ff   1/1/1                Static        1
```

6. Verify the priority configuration.

```
device(config-vlan-1)# show running-config vlan 1
vlan 1 by port
 static-mac-address 0000.0063.67ff ethernet 1/1/1 priority 7
!
```

7. Save the configuration.

```
device# write memory
```

Assign Static MAC Entries to Priority Queues Configuration Example

```
device# configure terminal
device(config)# vlan 1
device(config-vlan-1)# static-mac-address 0000.0063.67FF ethernet 1/1/1 priority 7
device(config-vlan-1)# end
device# show mac-address ethernet 1/1/1
device# show running-config interface ethernet 1/1/1
device# write memory
```

Configuring DSCP and CoS Marking for the Self-Generated Packets

Follow these steps to configure DSCP and CoS remarking for self-generated packets.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Configure the IEEE 802.1p priority marking for SSH traffic

```
device(config)# ip self-generate-traffic ssh 802-1p-priority-marking 5
```

3. Configure the DSCP marking for SSH traffic

```
device(config)# ip self-generate-traffic ssh dscp-marking 45
```

4. Configure additional options for the self-generated packets.

```
ICX8200-C08PF(config)#ip self-generate-traffic
  all         Change DSCP marking for all
  dns         Change DSCP marking for DNS
  icmp        Change DSCP marking for ICMP
  ntp         Change DSCP marking for NTP
  radius      Change DSCP marking for RADIUS
  snmp        Change DSCP marking for SNMP
  snmp-trap   Change DSCP marking for SNMP-TRAP
  ssh         Change DSCP marking for SSH
  syslog      Change DSCP marking for SYSLOG
  telnet      Change DSCP marking for TELNET
  tftp        Change DSCP marking for TFTP
```

5. Exit global configuration mode.

```
device(config)# exit
```

6. Save the configuration

```
device# write memory
```

```
device# configure terminal
device(config)# ip self-generate-traffic ssh 802-1p-priority-marking 5
device(config)# ip self-generate-traffic ssh dscp-marking 45
device(config)# exit
device# write memory
```

Configuring DSCP and CoS Remarking at the Interface Level

Follow these steps to configure DSCP and CoS remarking for specific ports. DSCP and CoS remarking are disabled by default.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter interface configuration mode to set DSCP and CoS values for a specific port.

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

3. Set the DSCP value to 4 of all IP packets on the port.

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

4. Set the PCP value to 4 of all IP packets on the port.

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

5. Return to privileged EXEC mode.

```
device(config-if-e1000-1/1/1)# end
```

6. Verify the configuration.

```
device# show running-config interface ethernet 1/1/1
interface ethernet 1/1/1
 port-name ERSPAN
 dual-mode
 ip dscp-remark 4
 ip pcp-remark 4
 rate-limit input fixed 40000 burst 120000
 mon profile 1 both
 priority 7
 speed-duplex 1000-full
 broadcast limit 96 kbps
 multicast limit 400 kbps
 unknown-unicast limit 96 kbps
 pvst-mode
 port security
 age 2 absolute
!
```

7. Save the configuration.

```
device# write memory
```

DSCP and CoS Remarking Configuration Example

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# ip dscp-remark 4
device(config-if-e1000-1/1/1)# ip pcp-remark 4
device(config-if-e1000-1/1/1)# end
device# show running-config interface ethernet 1/1/1
device# write memory
```

Changing the DSCP to Internal Forwarding Priority Mappings

Follow this example to change the DSCP to internal forwarding priority mappings for all the DSCP ranges.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Change the DSCP to internal forwarding priority mappings.

```
device(config)# qos-tos map dscp-priority 1 2 3 4 5 6 7 8 to 6
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
```

3. Return to privileged EXEC mode.

```
device(config)# exit
```

4. Verify the 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   |  1   6   6   6   6   6   6   6   6   1
   1   |  1   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

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
--------+---------
```

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.

5. Save the configuration.

```
device# write memory
```

Change the DSCP to Internal Forwarding Priority Mappings Configuration Example

```
device# configure terminal
device(config)# qos-tos map dscp-priority 1 2 3 4 5 6 7 8 to 6
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)# exit
device# show qos-tos
device# write memory
```

Changing the VLAN Priority 802.1p to Hardware Forwarding Queue Mappings

Follow this example to map a VLAN priority to a different hardware forwarding queue.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Map a VLAN priority.

```
device(config)# qos tagged-priority 2 qosp0
802.1p priority 2 mapped to qos profile qosp0
```

3. Return to privileged EXEC mode.

```
device(config)# exit
```

4. Verify the configuration.

```
device# show running-config | include qosp0
qos scheduler-profile voice profile qosp0 15 qosp1 15 qosp2 15 qosp3 15 qosp4 15 qosp5 25 qosp6 sp
qosp7 sp
...
qos tagged-priority 2 qosp0
```

5. Save the configuration.

```
device# write memory
```

Change the VLAN Priority 802.1p to Hardware Forwarding Queue Mappings Configuration Example

```
device# configure terminal
device(config)# qos tagged-priority 2 qosp0
device(config)# exit
device# show running-config | include qosp0
device# write memory
```

Selecting the QoS Queuing Method

Follow these steps to change the queuing method.

By default, RUCKUS devices use the weighted round robin (WRR) method of packet prioritization.

1. Enter global configuration mode.

```
device# configure terminal
```

2.   Select the QoS queuing method.

   ●   Change the queuing method to strict priority (SP).

```
device(config)# qos mechanism strict
bandwidth scheduling mechanism: strict priority
Qos profile bandwidth percentages are ignored
```

   ●   Change the queuing method to mixed SP and WRR.

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

   Observe that the verification step is not necessary with either of these choices.

3.   Return to privileged EXEC mode.

```
device(config)# exit
```

4.   Verify the configuration

```
device# show qos scheduler-profile all
User Scheduler Profile: test    Scheduling Option: Mixed-SP-WRR

Ports attached: (U1)        --
Ports attached: (U2)        --
Ports attached: (LAG)       --
Ports attached: (LAG)       --
Unicast per Queue details:       Bandwidth%
Traffic Class 0                    15%
Traffic Class 1                    15%
Traffic Class 2                    15%
Traffic Class 3                    15%
Traffic Class 4                    15%
Traffic Class 5                    25%
Traffic Class 6                    sp
Traffic Class 7                    sp
Multicast per Queue details:     Bandwidth%
Traffic Class 0                    15%
Traffic Class 1                    15%
Traffic Class 2                    15%
Traffic Class 3                    15%
Traffic Class 4                    15%
Traffic Class 5                    25%
Traffic Class 6                    sp
Traffic Class 7                    sp

Minimum Guaranteed Rate:         Bandwidth%
Unicast per Queue details:       Bandwidth%
Traffic Class 0                    0%
Traffic Class 1                    0%
Traffic Class 2                    0%
Traffic Class 3                    0%
Traffic Class 4                    0%
Traffic Class 5                    0%
Traffic Class 6                    0%
Traffic Class 7                    0%
```

5.   Save the configuration.

```
device# write memory
```

Select the QoS Queuing Method Configuration Example

```
device# configure terminal
device(config)# qos mechanism mixed-sp-wrr
device(config)# exit
device# show qos scheduler-profile all
device# write memory
```

Configuring the QoS Queue Name and Guaranteed Bandwidth

Follow these steps to change a queue name and the minimum percentage of a port outbound bandwidth guaranteed to the queue.

> **NOTE**
> Stackable devices that are operating as members of a stack reserve queue 7 for stacking functions.

NOTE
ICX 8200 devices do not support "qos guaranteed-rate" configuration.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Change minimum percentage of a port outbound bandwidth guaranteed to the queues.

```
device(config)# qos guaranteed-rate qosp0 10 qosp1 10 qosp2 15 qosp3 15 qosp4 10 qosp5 10 qosp6 10
qosp7 10
```

3. Change the name of QoS queue 3.

```
device(config)# qos name qosp3 r3d3
```

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.

4. Return to privileged EXEC mode.

```
device(config)# exit
```

5. Verify the configuration

```
device# show qos guaranteed-rate
Profile qosp7     : Minimum Guaranteed bandwidth 10
Profile qosp6     : Minimum Guaranteed bandwidth 10
Profile qosp5     : Minimum Guaranteed bandwidth 10
Profile qosp4     : Minimum Guaranteed bandwidth 10
Profile r3d3      : Minimum Guaranteed bandwidth 15
Profile qosp2     : Minimum Guaranteed bandwidth 15
Profile qosp1     : Minimum Guaranteed bandwidth 10
Profile qosp0     : Minimum Guaranteed bandwidth 10
```

6. Save the configuration.

```
device# write memory
```

QoS Queue Configuration Example

```
device# configure terminal
device(config)# qos guaranteed-rate qosp0 10 qosp1 10 qosp2 15 qosp3 15 qosp4 10 qosp5 10 qosp6 10 qosp7
10
device(config)# qos name qosp3 r3d3
device(config)# exit
device# show qos guaranteed-rate
device# write memory
```

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.

By default, the eight QoS queues on RUCKUS 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 15 Default Minimum Bandwidth Percentages on RUCKUS ICX 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 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.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Change the bandwidth percentages for the queues.

```
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%
```

The assigned bandwidths must total 100%. The configuration is immediately verified by the command output.

There is no minimum bandwidth requirement for a given queue.

> **NOTE**
> FastIron devices do not adjust the bandwidth percentages you enter.

3. Return to privileged EXEC mode.

```
device(config)# exit
```

4. Save the configuration.

```
device# write memory
```

Change the Minimum Bandwidth Percentages of the WRR Queues Example

```
device# configure terminal
device(config)# qos profile qosp7 25 qosp6 15 qosp5 12 qosp4 12 qosp3 10 qosp2
10 qosp1 10 qosp0 6
device(config)# exit
device# write memory
```

Allocating Bandwidth for Hybrid WRR and SP Queues

Follow these steps to change the default bandwidth percentages for the queues when the device is configured to use the combined SP and WRR queuing mechanism.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Change minimum percentage of a port outbound bandwidth guaranteed to the queues.

```
device(config)# qos profile qosp7 sp qosp6 sp qosp5 20 qosp4 15 qosp3 15 qosp2 20  qosp1 15 qosp0 15
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  20% calculated  20%
Profile qosp4: Priority4         bandwidth requested  15% calculated  15%
Profile qosp3: Priority3         bandwidth requested  15% calculated  15%
Profile qosp2: Priority2         bandwidth requested  20% calculated  20%
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+qosp4+qosp3+qosp2: Priority5,4,3,2       bandwidth requested  70% calculated  70%
Profile qosp1+qosp0            : Priority1,0(Lowest)  bandwidth requested  30% calculated  30%
```

The assigned bandwidths must total 100%. The configuration is immediately verified by the command output.

3. Return to privileged EXEC mode.

```
device(config)# exit
```

4. Display all QoS configuration settings.

```
device# show running-config | include qos
qos mechanism mixed-sp-wrr
qos profile qosp7 sp qosp6 sp qosp5 20 qosp4 15 qosp3 15 qosp2 20 qosp1 15 qosp0 15
qos scheduler-profile voice mechanism mixed-sp-wrr
qos scheduler-profile voice profile qosp0 15 qosp1 15 qosp2 15 qosp3 15 qosp4 15 qosp5 25 qosp6 sp
qosp7 sp
qos scheduler-profile voice guaranteed-rate qosp0 5 qosp1 5 qosp2 5 qosp3 5 qosp4 5 qosp5 25 qosp6 5
qosp7 5
qos priority-to-pg  qosp0 0 qosp1 0 qosp2 1 qosp3 1 qosp4 1 qosp5 2 qosp6 3 qosp7 4
qos guaranteed-rate qosp7 10 qosp6 10 qosp5 10 qosp4 10 qosp3 15 qosp2 15 qosp1 10 qosp0 10
qos tagged-priority 2 qosp0
qos-tos map dscp-priority 32 to 0
qos-tos map dscp-priority 0 to 1
qos-tos map dscp-priority 24 to 2
qos-tos map dscp-priority 48 to 3
qos-tos map dscp-priority 16 to 4
qos-tos map dscp-priority 1 2 3 4 5 6 7 8 to 6
qos-tos map dscp-priority 56 to 6
qos-tos map dscp-priority 40 to 7
```

5. Display information about QoS profiles.

```
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  20% calculated  20%
Profile qosp4: Priority4          bandwidth requested  15% calculated  15%
Profile qosp3: Priority3          bandwidth requested  15% calculated  15%
Profile qosp2: Priority2          bandwidth requested  20% calculated  20%
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+qosp4+qosp3+qosp2: Priority5,4,3,2          bandwidth requested  70% calculated  70%
Profile qosp1+qosp0            : Priority1,0(Lowest)      bandwidth requested  30% calculated  30%
```

6. Save the configuration.

```
device# write memory
```

Allocate Bandwidth for Hybrid WRR and SP Queues Configuration Example

```
device# configure terminal
device(config)# qos profile qosp7 sp qosp6 sp qosp5 20 qosp4 15 qosp3 15 qosp2 20  qosp1 15 qosp0 15
device(config)# exit
device# show running-config | include qos
device# write memory
```

Enabling Priority Flow Control Globally

Follow these steps to enable PFC globally.

> **NOTE**
> PFC is supported only on RUCKUS ICX 7550 and ICX 7850 devices.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable PFC globally.

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

3. Return to privileged EXEC mode.

```
device(config)# exit
```

4. Verify the configuration.

```
device# show priority-flow-control
Global PFC Status: Enabled
PFC Disabled on PG0
PFC Disabled on PG1
PFC Enabled on PG2
PFC Disabled on PG3
```

5. Save the configuration.

```
device# write memory
```

Enable Priority Flow Control Globally Configuration Example

```
device# configure terminal
device(config)# priority-flow-control enable
device(config)# exit
device# show priority-flow-control
device# write memory
```

Enabling Priority Flow Control for a Single Priority Group

Follow these steps to enable PFC for a single priority group (PG).

> **NOTE**
> PFC is supported only on RUCKUS ICX 7550 and ICX 7850 devices.

1. Enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Enable PFC for PG 1.

   ```
   device(config)# priority-flow-control 1
   ```

   There are four PGs numbered from 0 through 3.

3. Return to privileged EXEC mode.

   ```
   device(config)# exit
   ```

4. Verify the configuration.

   ```
   device# show priority-flow-control
   Global PFC Status: Enabled
   PFC Disabled on PG0
   PFC Enabled on PG1
   PFC Enabled on PG2
   PFC Disabled on PG3
   ```

   Observe that PFC for PG 1 is enabled.

5. Save the configuration.

   ```
   device# write memory
   ```

Enable Priority Flow Control for a Single Priority Group Configuration Example

```
device# configure terminal
device(config)# priority-flow-control 1
device(config)# exit
device# show priority-flow-control
device# write memory
```

Configuring the Share Queue Level for an Egress Buffer Profile

The share level is the maximum number of buffers that a priority group (PG) can use as a portion of the total sharing pool.

Follow these steps to configure the egress buffer share queue level.

1. Enter global configuration mode.

    ```
    device# configure terminal
    ```

2. Configure the share level.

    ```
    device(config)# qos egress-buffer-profile egress1 queue-share-level level3-1/16 7
    ```

3. Return to privileged EXEC mode.

    ```
    device(config)# exit
    ```

4. Verify the configuration

    ```
    device# show qos egress-buffer-profile egress1
    Egress Buffer Profile: egress1
    Ports attached:  --

     Per Queue Details:      Share Level:
    Queue 0                 level5-1/5
    Queue 1                 level4-1/9
    Queue 2                 level4-1/9
    Queue 3                 level4-1/9
    Queue 4                 level4-1/9
    Queue 5                 level4-1/9
    Queue 6                 level4-1/9
    Queue 7                 level3-1/16
    ```

5. Save the configuration.

    ```
    device# write memory
    ```

Share Queue Level for an Egress Buffer Profile Configuration Example

```
device# configure terminal
device(config)# qos egress-buffer-profile egress1 queue-share-level level3-1/16 7
device(config)# exit
device# show qos egress-buffer-profile egress1
device# write memory
```

Configuring a Port to the Egress Queue Drop Counters

Each port has its own set of drop counters.

To display the queue drop counters, use the **show interfaces ethernet** command for the port.

Rate Limiting and Rate Shaping

Rate Limiting

Non ACL-Based Rate Limiting

Port-Based Fixed Rate Limiting Configuration Notes

- Rate limiting is available only on inbound ports.
- 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, 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.
- The RUCKUS ICX 7650 applies rating limiting starting with the Layer 1 overhead. All other ICX devices begin counting the packet size from the Layer 2 header onwards.

Port-Based Fixed Rate Limiting

You can configure a fixed rate limiting policy on a port in the inbound direction only. This feature allows you to specify the maximum number of bytes in kilobits per second (kbps) a given port can receive. The port drops bytes that exceed the limit you specify.

Fixed rate limiting applies to all traffic on the rate-limited port. It counts the number of bytes that a port receives in one second intervals. If the number of bytes exceeds the maximum number you specified when you configured the rate, the port drops all further inbound packets for the duration of the one-second interval. Unused bandwidth is not carried over from one interval to the next. Once the one-second interval is complete, the port clears the counter and re-enables traffic.

> **NOTE**
> RUCKUS 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.

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

Unused bandwidth can be used to consume bursts in traffic. When a group of packets which has shorter inter-packet gaps comes in, a higher bandwidth is required to consume these bursts. This additional bandwidth is borrowed from unused bandwidth in an set interval.

Each device supports line-rate rate limiting in hardware. The device creates entries in Content Addressable Memory (CAM) for the rate limiting policies. 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.
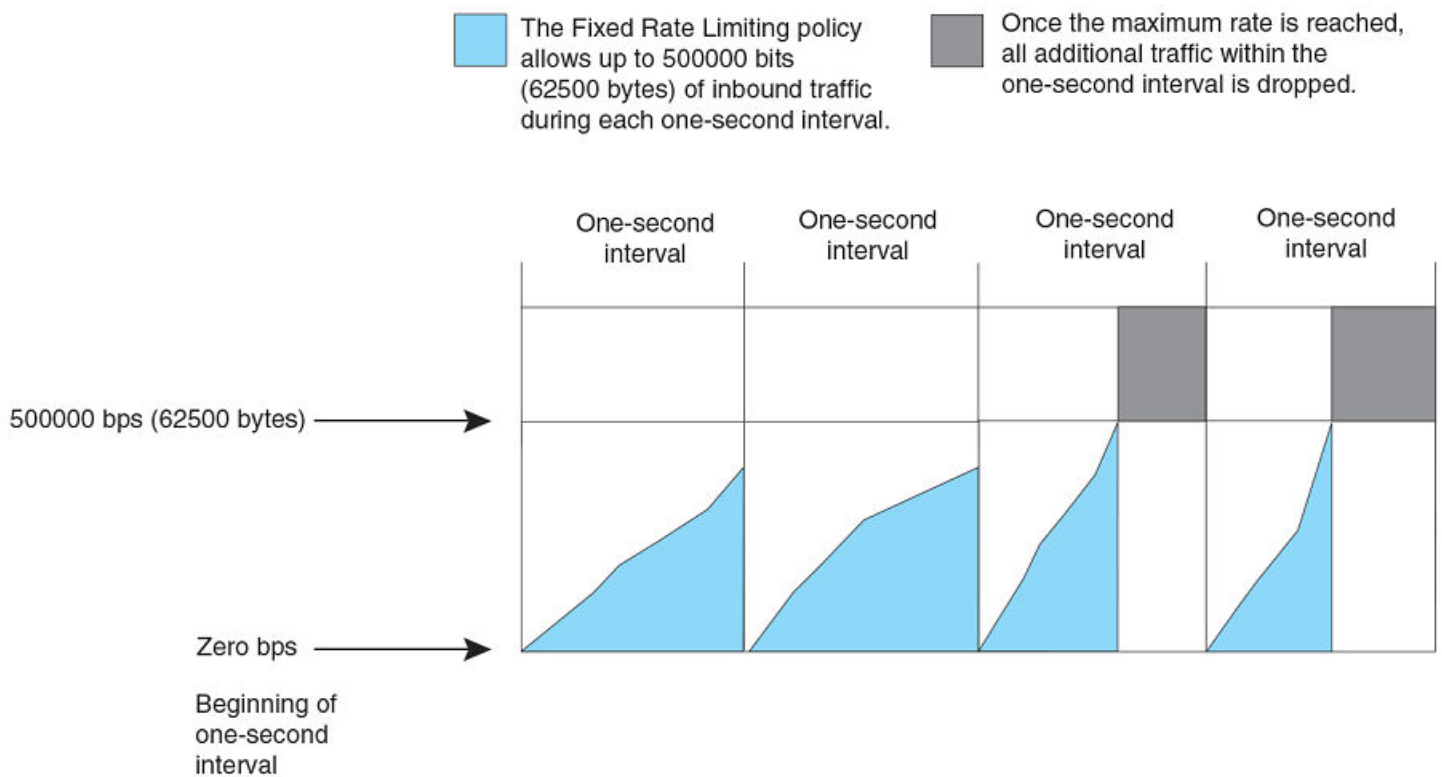
> **NOTE**
> CPU traffic can also be rate-limited by applying traffic policies using ACLs. Refer to Applying ACLs to Rate Limit Inbound CPU Traffic on page 59.

How Port-Based Fixed Rate Limiting Works

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 500,000 bits (62,500 bytes) a second. During the first two one-second intervals, the port receives less than 500,000 bits in each interval. However, the port receives more than 500,000 bits during the third and fourth one-second intervals and consequently drops the excess traffic.

FIGURE 3 Fixed Rate Limiting

> **NOTE**
> The software counts the bytes by polling statistics counters for the port every 100 milliseconds, which provides 10 readings each second. As such, the fixed rate limiting policy has an accuracy of within 10% of the port's line rate. It is possible, then, 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.

The following figure shows how adding a burst size to fixed rate limiting works.

FIGURE 4 Rate Limiting with Burst Option

An example of this is, traffic coming into port 1 and going out of port 2 is a 1G link. The rate is limited to 10Mb on port one. This means only 10Mb data is going out of port 2. If only 9MB of traffic is sent, it will all go through without even hitting the set rate-limit. This leaves 1Mb of unused bandwidth. This 1Mb of unused bandwidth can be borrowed in case if a burst of packets comes in. Now, the "rate-limit fixed 10000 **burst <burst size>"** can be used if you want to restrict the borrowing of 1M to say only 0.1M.

Default System Rate Limiting of Inbound CPU Traffic

Default CPU rate limiting is a CPU protection scheme that limits certain traffic types.

Unnecessary traffic to the switch CPU lowers the efficiency of the CPU and delays handling of other traffic that requires processing. Default CPU rate limiting identifies the traffic type and assigns a maximum rate limit to the traffic type. The traffic types that 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. You cannot configure these rates.

All currently supported FastIron devices support the default CPU rate limiting feature.

> **NOTE**
> It is possible to configure fixed rate limiting for inbound CPU traffic. Refer to for more information.

TABLE 16 Default CPU Rate Limits for Packet Type

Packet Type	Rate Limit in Packets per Second
ARP	6000

TABLE 16 Default CPU Rate Limits for Packet Type (continued)

Packet Type	Rate Limit in Packets per Second
IP TTL exceed	150
Reverse path check failed	
IP MTU failed	3000
IP tunnel-terminated packets that are fragmented or have options	
IP tunnel-terminated packets with unsupported GRE tunnel header	
IP Unicast packets mirrored to CPU due to ICMP redirect	100
Bridge packets forwarded to CPU	5000

Rate Limiting Broadcast, Unknown Unicast, and Multicast Traffic

RUCKUS ICX devices can forward all flooded traffic at wire speed within a VLAN. However, some third party networking devices cannot handle high rates of broadcast, unknown unicast, and multicast (BUM) traffic.

If high rates of traffic are being received by the device on a given port of that VLAN, you can limit the number of BUM 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.

RUCKUS ICX devices support byte-based and packet-based rate limiting.

For the ICX 8200, the packet per second rate limiting uses increments of 125.

Traffic Policy ACL-Based Rate Limiting

Traffic Policies for ACL-Based Rate Limit Configuration Notes

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.

Applying a traffic policy to an interface includes the following steps:

1. Creating a traffic policy

2. Adding a reference to the traffic policy in an ACL entry

3. Binding the ACL that contains the ACL entry to an interface

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

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 go into effect 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.

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

- Physical Ethernet interfaces

- Virtual interfaces

- LAG ports

- Specific VLAN members on a port

- A subset of ports on a virtual interface

For more information on ACLs, refer to the *RUCKUS FastIron Security Configuration Guide*.

For ICX 8200 devices, the traffic policy and DSCP/802.1p/internal; priority markings are supported together.

ACL-Based Fixed 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.

ACL-Based Adaptive Rate Limiting

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.

Traffic Policies for ACL-Based Rate Limit Restrictions and Limitations

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.

Note the following when configuring traffic policies:

- 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."

- 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.

- You can also reference the same traffic policy in more than one rule in the same ACL.

- On all ICX devices except ICX 8200 devices, when the same traffic policy is referenced by different rules in the same ACL, the rate limit is separate for each traffic flow to which the traffic policy applies. In contrast, on ICX 8200 devices, rate limiting is cumulative for all rules within the same ACL that apply the same traffic policy. For example, suppose a traffic policy sets a rate limit of 3000 KBPS, and two ACL rules, r1 and r2, apply the traffic policy to a traffic flow on a match. The total rate across all flows that match r1 and r2 will be 3000 KBPS. In contrast, all other ICX switches allow the full rate limit of 3000 KBPS individually for every traffic flow that matches either rule.

- 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.

- 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 17 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 VLAN or VLAN interfaces, the rate limit policy is cumulative for all of the ports in the port region. If the 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 VLAN 100, 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 from 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 that 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 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.

Configuring Rate Limiting

Configuring Port-Based Fixed Rate Limiting

Follow these steps to configure a rate limiting policy for a port.

These commands configure a fixed rate limiting policy that allows Ethernet port 1/1/1 to receive a maximum of 500 kbps. If the port receives additional bytes during a given one-second interval, all inbound packets on the port are dropped until the next one-second interval starts.

When traffic reaches the rate limiting threshold, Traffic Domain 2 (TD2) sends traditional pause or Priority Flow Control (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).

When **rate-limit input fixed** is configured, it limits the traffic coming into the port. The rate-limited traffic is seen on the egress port's statistics.

1. Enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Enter interface configuration mode.

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

3. Set the fixed rate limiting policy.

```
device(config-if-e1000-1/1/1)# rate-limit input fixed 500
device(config-if-e1000-1/1/1)# end
```

4. Verify the configuration.

```
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
```

Fixed Rate Limiting Configuration Example

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# rate-limit input fixed 500
device(config-if-e1000-1/1/1)# end
device# show rate-limit input
```

Configuring Port-Based Fixed Rate Limiting with Burst Option

Follow these steps to configure a rate limiting policy for a port with a burst size included.

These commands configure a fixed rate limiting policy with a burst option that allows Ethernet port 1/1/1 to receive a maximum of 10000 kbps (i.e. 10 Mbps and therefore forward a maximum of 10 Mb to the egress port. Now if only 9 Mb of traffic is sent, there is 1 Mb of unused bandwidth. This 1 Mb of unused bandwidth can be borrowed in case if a burst of packets comes in. Now, the burst size option can be used to restrict the borrowing of the 1 Mb to only 0.1Mb (i.e. 1000 kbps).

The burst option command is optional and is not part of the required command.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter interface configuration mode.

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

3. Set the fixed rate limiting policy.

```
device(config-if-e1000-1/1/1)# rate-limit input fixed 10000 burst 1000
device(config-if-e1000-1/1/1)# end
```

4. Verify the configuration.

```
device# show rate-limit input
Total rate-limited interface count: 5.
    Port        Configured Input Rate    Actual Input Rate
    1/1/1       10000                    10000end
    1/1/2       195000                   195000
    1/1/6       1950                     1950
    1/5/2       230432                   230000
    1/5/6       234113                   234000
```

Fixed Rate Limiting with Burst Configuration Example

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# rate-limit input fixed 10000 burst 10000
device(config-if-e1000-1/1/1)# end
device# show rate-limit input
```

Configuring Rate Limiting for BUM Traffic

If high rates of traffic are being received by the device on a given port of that VLAN, you can limit the number of BUM packets or kilobytes received each second on that port.

1. Enter global configuration mode.

    ```
    device# configure terminal
    ```

2. Enter interface configuration mode.

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

3. Enable BUM rate limits.

    The rate limit can be set in kilobits per second (kbps) or packets per second (pps). You must specify "kbps" to set the rate limit in kilobits per second. If the unit of measure is not mentioned explicitly, the rate limit is measured in packets per second (pps).

4. Set a broadcast rate limit by port speed in kbps.

    ```
    device(config-if-e40000-1/1/1)# broadcast limit 100 kbps
    ```

    To set the broadcast rate limit by packets per second, enter the following command.

    > **NOTE**
    > For the ICX 8200, the broadcase rate limit will be 125.

    ```
    device(config-if-e40000-1/1/1)# broadcast limit 100
    ```

5. Set a multicast rate limit by port speed in kbps.

    ```
    device(config-if-e40000-1/1/1)# multicast limit 400 kbps
    ```

    To set the multicast rate limit by packets per second, enter the following command.

    ```
    device(config-if-e40000-1/1/1)# multicast limit 400
    ```

6. Set an unknown unicast rate limit by port speed in kbps.

```
device(config-if-e40000-1/1/1)# unknown-unicast limit 100 kbps
```

To set the unknown unicast rate limit by packets per second, enter the following command.

```
device(config-if-e40000-1/1/1)# unknown-unicast limit 100
```

7. Verify the configuration.

```
device# show running-config interface ethernet 1/1/1 | i limit
broadcast limit 100
multicast limit 400
unknown-unicast limit 100

device# show running-config interface ethernet 1/1/1 | i limit
broadcast limit 100 kbps
multicast limit 400 kbps
unknown-unicast limit 100 kbps
```

Rate Limiting BUM Traffic Configuration Example

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e40000-1/1/1)# broadcast limit 100 kbps
device(config-if-e40000-1/1/1)# unknown-unicast limit 100 kbps
device(config-if-e40000-1/1/1)# multicast limit 400 kbps

device(config-if-e40000-1/1/1)# show running-config interface | begin  1/1/1
interface ethernet 1/1/1
broadcast limit 100 kbps
unknown-unicast limit 100 kbps
multicast limit 400 kbps


device# configure terminal
device(config)# interface ethernet 1/1/2
device(config-if-e40000-1/1/2)# broadcast limit 100
device(config-if-e40000-1/1/2)# unknown-unicast limit 100
device(config-if-e40000-1/1/2)# multicast limit 400

device(config-if-e40000-1/1/2)# show running-config interface | begin  1/1/2
interface ethernet 1/1/2
broadcast limit 100
unknown-unicast limit 100
multicast limit 400
```

BUM Suppression Port Dampening

The BUM suppression port dampening 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 shut down (dampened) for a user-configured period.

Enabling BUM Suppression Port Dampening

NOTE
You can set the rate limit interval in either kilobits per second (kbps) or packets per second (pps).

NOTE
Set the log timer interval when kbps rate limit is configured (broadcast, unknown-unicast, and multicast).

Complete the following steps to enable BUM suppression port dampening in kbps.

1. Enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Set the rate limit log interval.

   ```
   device(config)# rate-limit-log 3
   ```

   In this example, the kbps rate interval is set to 3 minutes. The value can be from 1 through 10 minutes. The default value is 5 minutes.

   At every rate-limit-log interval, the device checks whether the threshold configured is exceeded or not to take the port-shutdown action.

3. Enter interface configuration mode.

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

4. Enable broadcast suppression port dampening in kbps with a shut down interval.

   ```
   device(config-if-e10000-1/1/1)# broadcast limit 100 kbps threshold 2000 action port-shutdown 7
   ```

5. Enable multicast suppression port dampening in kbps with a shut down interval.

   ```
   device(config-if-e10000-1/1/1)# multicast limit 100 kbps threshold 2000 action port-shutdown 7
   ```

6. Enable unknown unicast suppression port dampening in kbps.

   ```
   device(config-if-e10000-1/1/1)# unknown-unicast limit 100 kbps threshold 2000 action port-shutdown
   ```

   Because no value is indicated for the port-shutdown parameter, the default of 5 minutes is applied.

7. Verify the configuration.

   ```
   device# show running-config interface ethernet 1/1/1 | i limit
   broadcast limit 100 kbps threshold 2000 action port-shutdown 7
   multicast limit 100 kbps threshold 2000 action port-shutdown
   unknown-unicast limit 100 kbps threshold 2000 action port-shutdown
   ```

Complete the following steps to enable BUM suppression port dampening in pps.

1. Enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Set the rate limit log interval.

   ```
   device(config)# rate-limit-pps-log 200
   ```

   In this example, the pps rate interval is set to 200 seconds. The value can be from 1 through 600 seconds. The default value is 300 seconds.

   At every rate-limit-log-pps interval, the device checks whether the threshold configured is exceeded or not to take the port-shutdown action.

3. Enter interface configuration mode.

   ```
   device(config)# interface ethernet 1/1/2
   ```

4. Enable broadcast suppression port dampening in pps with a shutdown interval.

   ```
   device(config-if-e10000-1/1/2)# broadcast limit 1500 pps threshold 2000 action port-shutdown 7
   ```

5. Enable multicast suppression port dampening in pps with a shutdown interval.

```
device(config-if-e10000-1/1/2)# multicast limit 100 pps threshold 2000 action port-shutdown 7
```

6. Enable unknown unicast suppression port dampening in pps.

```
device(config-if-e10000-1/1/2)# unknown-unicast limit 1500 pps threshold 2000 action port-shutdown
```

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

7. Verify the configuration.

```
device#show running-config interface ethernet 1/1/2 | i limit
broadcast limit 1500 pps threshold 2000 action port-shutdown 7
multicast limit 100 pps threshold 2000 action port-shutdown 7
unknown-unicast limit 1500 pps threshold 2000 action port-shutdown
```

BUM Suppression Port Dampening Configuration Example

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e40000-1/1/1)# broadcast limit 100 kbps threshold 2000 action port-shutdown 7
device(config-if-e40000-1/1/1)# unknown-unicast limit 100 kbps threshold 2000 action port-shutdown 7
device(config-if-e40000-1/1/1)# multicast limit 100 kbps threshold 2000 action port-shutdown 7

device# configure terminal
device(config)# interface ethernet 1/1/2
device(config-if-e40000-1/1/2)# broadcast limit 100 pps threshold 2000 action port-shutdown 7
device(config-if-e40000-1/1/2)# unknown-unicast limit 100 pps threshold 2000 action port-shutdown 7
device(config-if-e40000-1/1/2)# multicast limit 100 pps threshold 2000 action port-shutdown 7
```

Enabling BUM Suppression Logging

Complete the following steps to enable BUM suppression logging.

> **NOTE**
> Rate limiting must be enabled.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter Ethernet interface configuration mode.

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

3. Enable rate limiting.

```
device(config-if-e10000-1/1/1)# broadcast limit 8388607
```

Broadcast is used in this example. The command syntax is same for multicast or unknown unicast with the corresponding command (**multicast** or **unknown-unicast**).

4. Enable logging when the specified limit is exceeded.

```
device(config-if-e10000-1/1/1)# broadcast limit 100 kbps log
```

Broadcast is used in this example. The command syntax is same for multicast or unknown unicast with the corresponding command (**multicast** or **unknown-unicast**).

5. Use the **threshold** option for the broadcast limit to shut-down the port and generate the syslog.

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

    The following example shuts down the port for 300 seconds (default) when the packet drop threshold value exceeds 1000 Kbps.

    ```
    device(config)# interface ethernet 1/2/1
    device(config-if-e10000-1/2/1)# unknown-unicast limit 100 kbps threshold 1000 action port-shutdown
    ```

6. Globally configure the log interval.

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

7. Verify the logging interval.

    ```
    device(config)# show running-config | include rate-limit-log
     rate-limit-log 6
    ```

8. Verify the configuration.

    ```
    device# show logging | include 1/1/1
    Jan 13 12:02:12:I:Security: Interface ethernet 1/1/1 reached the Broadcast traffic limit and 1434 kB
    are dropped
    ```

Enabling BUM Suppression Logging Configuration Example

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e10000-1/1/1)# broadcast limit 8388607
device(config-if-e10000-1/1/1)# broadcast limit 100 kbps log
device(config-if-e10000-1/1/1)# broadcast limit 100 kbps threshold 1000 action port-shutdown
device(config)# rate-limit-log 6
device(config)# show running-config | include rate-limit-log
device(config)# exit
device# show logging | include 1/1/1
```

Viewing BUM Suppression Syslog Notifications

Use the following commands to display BUM suppression syslog notification information.

Use the **show logging** command to view the BUM suppression syslog notifications for all interfaces.

```
device# show logging
Jan 13 12:02:12:I:Security: Interface ethernet 1/1/1 reached the Broadcast traffic limit and 11620 kB are
dropped
Jan 13 12:14:23:I:Security: Interface ethernet 1/3/12 reached the Multicast traffic limit and 870 kB are
dropped
Jan 13 12:45:38:I:Security: Interface ethernet 3/2/14 reached the Unknown-Unicast traffic limit and 2321 kB
are dropped
```

The first section of the output is `mmm dd hh:mm:ss:Info:System`.

To view the BUM suppression syslog notifications for a specific interface, use the following command.

```
device# show logging | include 1/1/1
Jan 13 12:02:12:I:Security: Interface ethernet 1/1/1 reached the Broadcast traffic limit and 11620 kB are
dropped
```

Configuring ACL-Based Fixed Rate Limiting Using Traffic Policies

Use the procedure in this section to configure ACL-based fixed rate limiting.

> **NOTE**
> Before configuring this feature, see what to consider in "Configuration notes and feature limitations for traffic policies."

These commands:

- Set the maximum number of traffic policies.

- Create a fixed traffic policy that enables ACL statistics (counting).

- Create a new extended ACL entry and bind the ACL to an interface.

- Verify the configuration.

1. Enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Create a traffic policy and set parameters.

   - Create a policy that drops packets or bytes that exceed the CIR (committed information rate) limit. The following examples use a packet-based fixed CIR.

     ```
     device(config)# traffic-policy TPDF1 rate-limit packet-based fixed cir 10000 exceed-action drop
     count
     ```

   - Create a policy that permits packets or bytes that exceed the CIR limit.

     ```
     device(config)# traffic-policy TPDF1 rate-limit packet-based fixed cir 10000 exceed-action
     permit-at-low-pri count
     ```

   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. If the port receives additional bits during a given one-second interval, the port either drops all packets on the port until the next one-second interval starts or permits packets that exceed the limit.

   Use the keyword **byte-based** if you want to specify a CIR based on byte count rather than a packet count.

3. Create an extended ACL entry or modify an existing extended ACL entry that references the traffic policy.

   ```
   device(config)# ip access-list extended 101
   device(config-ext-ipacl-101)# permit ip host 10.10.12.2 any traffic-policy TPDF1
   ```

4. Bind the ACL to an interface.

   a) Enter interface configuration mode.

      ```
      device(config-ext-ipacl-101)# interface ethernet 1/1/5
      ```

   b) Bind the ACL to the interface.

      ```
      device(config-if-e1000-1/1/5)# ip access-group 101 in
      ```

   c) Exit interface configuration mode.

      ```
      device(config-if-e1000-1/1/5)# exit
      ```

   These commands allow port 1/1/5 to receive a maximum traffic rate of 100 kbps. 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.

   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.

5. Verify the configuration.

   ```
   device(config)# show traffic-policy TPDF1
   Traffic Policy - TPDF1:
   Metering Enabled, Parameters:
        Mode: Fixed Rate-Limiting
         cir: 100 kbps, cbs: 2000 bytes, pir: 200 kbps, pbs: 4000 bytes
   Counting Not Enabled
   Number of References/Bindings:1
   ```

6. View the ACL and rate limit counters.

   ```
   device(config)# show access-list accounting ethernet 1/1/5 in
   MAC Filters Accounting Information
       0: DA ANY SA 0000.0000.0001 - MASK FFFF.FFFF.FFFF
         action to take : DENY
        Hit Count:   (1Min)                     0    (5Sec)                    0
                     (PktCnt)                   0  (ByteCnt)                   0
        -----------------------------------------------------------
   65535: Implicit Rule deny any any
        Hit Count:   (1Min)           5028    (5Sec)               2129
                     (PktCnt)         5028  (ByteCnt)             643584
        -----------------------------------------------------------
   ```

7. Clear the ACL and rate limit counters.

   a) Clear the ACL counters.

   ```
   device(config)# clear access-list accounting all
   ```

   b) Clear the rate limit counters.

   ```
   device(config)# clear statistics traffic-policy TPDF1
   ```

8. Configure DSCP marking and 802.1p priority marking traffic policies.

   ```
   device(config)#ip access-list ip3
   device(config-ext-ipacl-ip3)# sequence 9 permit ip any any dscp-marching 48 dscp-marking 46 traffic-
   policy VOIP
   device(config-ext-ipacl-ip3)# sequence 10 permit ip any any dscp-matching 48 dscp-marking 46 802.1p-
   and-internal-marking 2 traffic-policy VOIP
   device(config-ext-ipacl-ip3)# sequence 11 permit ip any any dscp-marking 0 traffic-policy VOIP
   device(config-ext-ipacl-ip3)# sequence 1 permit ip any any 802.1p-priority-matching 5 802.1p-
   priority-marking 6 internal-priority-marking 3 traffic-policy VOIP
   device(config-ext-ipacl-ip3)#sequence 3 permit ip any any internal-priority-marking 4 traffic-policy
   VOIP
   ```

ACL-Based Fixed Rate Limiting Using Traffic Policies Configuration Example

```
device# configure terminal
device(config)# traffic-policy TPDF1 rate-limit fixed packet-based cir 10000 exceed-action drop
device(config)# ip access-list extended 101
device(config-ext-ipacl-101)# permit ip host 10.10.12.2 any traffic-policy TPDF1
device(config-ext-ipacl-101)# interface ethernet 1/1/5
device(config-if-e1000-1/1/5)# ip access-group 101 in
device(config-if-e1000-1/1/5)# exit
device(config)# show traffic-policy TPDF1
device(config)# clear access-list accounting all
device(config)# clear statistics traffic-policy TPDF1
device(config)#ip access-list ip3
device(config-ext-ipacl-ip3)# sequence 9 permit ip any any dscp-marching 48 dscp-marking 46 traffic-policy
VOIP
device(config-ext-ipacl-ip3)# sequence 10 permit ip any any dscp-matching 48 dscp-marking 46 802.1p-and-
internal-marking 2 traffic-policy VOIP
device(config-ext-ipacl-ip3)# sequence 11 permit ip any any dscp-marking 0 traffic-policy VOIP
device(config-ext-ipacl-ip3)# sequence 1 permit ip any any 802.1p-priority-matching 5 802.1p-priority-
marking 6 internal-priority-marking 3 traffic-policy VOIP
device(config-ext-ipacl-ip3)# sequence 3 permit ip any any internal-priority-marking 4 traffic-policy VOIP
device(config-ext-ipacl-ip3)# exit
```

Configuring ACL-Based Adaptive Rate Limiting Using Traffic Policies

You can configure adaptive rate limiting to forward traffic, modify the IP precedence of and then forward traffic, or drop traffic based on whether the traffic is within the limit or exceeds the set limit.

These commands:

- Set the maximum number of traffic policies.

- Create an adaptive traffic policy that enables ACL statistics (counting).

- Create a new extended ACL entry and bind the ACL to an interface.

- Verify the configuration.

1. Enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Create traffic policies and set parameters.

   a) Create a policy, TPDrop, that drops packets that exceed the limit.

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

   b) Create a policy, TPallow, that permits packets that exceed the limit.

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

   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. If the port receives additional bits during a given one-second interval, the port either drops all packets on the port until the next one-second interval starts or permits packets that exceed the limit.

3. Verify the traffic policy configuration.

    ```
    device(config)# show traffic-policy
    Traffic Policy - TPallow:

            Metering Enabled, Parameters:
                    Mode: Adaptive Rate-Limiting
                    cir: 400000 kbps,    cbs: 125000 kbits,    pir: 12000000 kbps,    pbs: 1250000000
    kbits

            Counting Enabled
            Number of References/Bindings: 1
    Traffic Policy - TPDrop:

            Metering Enabled, Parameters:
                    Mode: Adaptive Rate-Limiting
                    cir: 10000 kbps,    cbs: 1600 kbits,    pir: 20000 kbps,    pbs: 4000 kbits

            Counting Enabled
            Number of References/Bindings: 0
    ```

4. Create appropriate ACL entries.

    a) Create a new extended IPv4 ACL or modify an existing extended IPv4 ACL that references the traffic policy.

        ```
        device(config)# ip access-list extended 104
        device(config-ext-ipacl-104)# permit ip host 1.1.1.2 any traffic-policy TPallow
        device(config-ext-ipacl-104)# exit
        ```

        The previous example creates an IPv4 extended ACL that allows traffic from the specified IP host and applies the traffic policy created earlier to allow traffic at a low priority.

    b) Create or modify an IPv6 ACL that references the traffic policy.

        ```
        device(config)# ipv6 access-list acl105
        device(config-ipv6-access-list acl105)# permit ip any any 802.1p-priority matching 3 traffic-
        policy TPdrop
        device(config-ipv6-access-list acl105)# exit
        ```

        The previous example creates an IPv6 ACL that permits an IP traffic, and if it matches 802.1p priority 3, it applies the traffic policy TPdrop, created earlier to drop and count traffic that exceed the defined rate limit.

5. Verify the ACLs.

    ```
    device(config)# show access-list all
    ...

    Extended IP access list 104 : 1 entry
    permit ip host 1.1.1.2 any traffic-policy TPallow

    Extended IP access list 105 : 1 entry
    permit ip any any 802.1p-priority-matching 3 traffic-policy TPdrop
    ...
    ```

6. Bind the ACL to an interface.

    a) Enter interface configuration mode.

        ```
        device(config)# interface ethernet 1/1/6
        ```

    b) Bind the ACL to the interface.

        ```
        device(config-if-e1000-1/1/6)# ip access-group 104 in
        device(config-if-e1000-1/1/6)# exit

        device(config-if-e1000-1/1/6)# ipv6 access-group acl105 in
        device(config-if-e1000-1/1/6)# exit
        ```

7. Clear the ACL and rate limit counters.

    a) Clear the ACL counters.

        ```
        device(config)# clear access-list accounting all
        Traffic Policy TPallow: cleared
        ```

    b) Clear the rate limit counters.

        ```
        device(config)# clear statistics traffic-policy TPallow
        device(config)# clear statistics traffic-policy TPdrop
        ```

8. Configure DSCP marking and 802.1p priority marking traffic policies.

    ```
    device (config)# ipv6 access-list ipv6
    device(config-ipv6acl-ipv6)# sequence 2 permit ipv6 any any dscp-matching 48 dscp-marking 62
    internal-priority-marking 7
    device(config-ipv6acl-ipv6)# sequence 3 permit ipv6 any any dscp-matching 48 dscp-marking 62 802.1p-
    priority-marking 4 internal-priority-marking 4
    device(config-ipv6acl-ipv6)# sequence 8 permit ipv6 any any dscp-marking 0
    ```

ACL-Based Adaptive Rate Limiting Using Traffic Policies Configuration Example

```
device# configure terminal
device(config)# traffic-policy TPDrop rate-limit adaptive cir 10000 cbs 1600 pir 20000 pbs 4000 exceed-
action drop count
device(config)# traffic-policy TPallow rate-limit adaptive cir 10000 cbs 1600 pir 20000 pbs 4000 exceed-
action permit-at-low-pri count
device(config)# show traffic-policy
!
device(config)# ip access-list extended 104
device(config-ext-ipacl-104)# permit ip host 1.1.1.2 any traffic-policy TPallow
device(config-ext-ipacl-104)# exit
device(config)# ipv6 access-list acl105
device(config-ipv6-access-list acl105)# permit ip any any 802.1p-priority matching 3 traffic-policy TPdrop
device(config-ipv6-access-list acl105)# exit
device(config)# show access-list all
device(config)# interface ethernet 1/1/6
device(config-if-e1000-1/1/6)# ip access-group 104 in
device(config-if-e1000-1/1/6)# ipv6 access-group acl105 in
device(config-if-e1000-1/1/6)# exit
device(config)# clear access-list accounting all
device(config)# clear statistics traffic-policy TPDA4
device (config)# ipv6 access-list ipv6
device(config-ipv6acl-ipv6)# sequence 2 permit ipv6 any any dscp-matching 48 dscp-marking 62 internal-
priority-marking 7
device(config-ipv6acl-ipv6)# sequence 3 permit ipv6 any any dscp-matching 48 dscp-marking 62 802.1p-
priority-marking 4 internal-priority-marking 4
device(config-ipv6acl-ipv6)# sequence 8 permit ipv6 any any dscp-marking 0
```

Applying ACLs to Rate Limit Inbound CPU Traffic

Traffic policies can be included in ACLs and applied to incoming CPU traffic.

You can apply ACLs to the CPU on an ICX standalone unit or an ICX stack to filter or rate limit specific incoming traffic. For example, you can create
ACLs to perform any of the following actions:

- Rate limit DHCP packets sent to the CPU
- Permit ICMP packets
- Permit packets from a known subnet
- Deny Telnet and SSH connections

> **NOTE**
> On ICX 8200 devices, configuring both DSCP trust and Traffic Policy is not recommended. When dscp trust and Traffic Policy are configured on the same interface, Layer 3 traffic wont honor trust dscp and all traffic including green conformance will be egressing on best effort delivery Queue 0.

Constraints on ACLs Applied to Inbound CPU Traffic

The following restrictions pertain to ACLs applied to inbound CPU traffic:

- Only fixed rate limiting is supported for CPU ACLs.
- CPU ACLs cannot be applied on any other type of interface.
- Inbound traffic on the management port is also subject to the ACL applied to the CPU.
- In a stack, an ACL can be applied to the CPU of the active controller only. This does not create issues in a stack system during a failover because the ACL is applied to all stack units when it is bound to the active controller CPU.
- The maximum number of filters in an ACL applied to the CPU is 15.
- Standard IPv4 ACLs are not supported. IPv4 extended ACLs or IPv6 ACLs must be used.
- MAC ACLs are not supported.
- A DROP action is the only exceed-action allowed in a traffic policy applied to the CPU.
- Without a **permit ip any any** statement in an IPv4 ACL or a **permit ipv6 any any** statement in an IPv6 ACL, the ACL cannot be bound to a CPU port.
- The **deny any any** statement is not supported for ACLs applied to the CPU.
- ACLs applied to the CPU do not support implicit filters, such as **deny any any**, characteristic of other ACLs applied on ICX devices. In other words, it is not possible to block all traffic to the CPU.
- In TCP and UDP filters, only the equal option (eq) is supported.
- Accounting, mirroring, and logging are not supported in ACLs applied to the CPU.

Unsupported Protocols and Options

The following protocols and options are not supported in IPv4 ACLs applied to the CPU:

- esp
- gre
- 802.1p-and-internal-marking
- 802.1p-priority-marking
- 802.1p-priority-matching
- dscp-marking
- dscp-matching
- internal-priority-marking
- precedence
- tos

The following protocols and options are not supported in IPv6 ACLs applied to the CPU:

- ahp
- esp
- sctp

- 802.1p-priority-marking

- 802.1p-priority-matching

- dscp-marking

- dscp-matching

- fragments

- internal-priority-marking

- routing

- precedence, gt, lt, neq

Configuring Fixed Rate Limiting on the CPU

To apply fixed rate limiting on the CPU, you must complete the following actions:

- Create a traffic policy.

- Create an ACL containing the traffic policy, or add a statement containing the traffic policy to an existing ACL.

- Bind the ACL containing the policy to the CPU.

Perform the following steps to bind the ACL containing the policy to the CPU.

1. Enter global configuration mode.

   ```
   device# configure terminal
   device(config)#
   ```

2. Enter CPU configuration mode.

   ```
   device(config)# interface cpu active
   device(config-if-cpu-active)#
   ```

   The **active** keyword used in the command designates the active controller of a stack but is also used for a standalone unit, whether or not stacking is enabled.

3. Bind the ACL that was previously created with the desired traffic policy to the CPU.

   ```
   device(config-if-cpu-active)# ipv6 access-group ipv6_icmp in
   ```

   The previous example binds an IPv6 ACL (ipv6_icmp) to the CPU interface and applies the ACL to incoming traffic.

   > **NOTE**
   > IPv4 extended ACLs and IPv6 ACLs can be applied to the CPU interface. Use the **ip access-group** command followed by the **in** keyword to apply an IPv4 ACL. The following example binds an existing IPv4 extended ACL (block_telnet) to the CPU interface.

   ```
   device# configure terminal
   device(config)# interface cpu active
   device(config-if-cpu-active)# ip access-group block_telnet in
   device(config-if-cpu-active)# exit
   ```

4. (Optional) Enter the **show running-config interface cpu active** command to verify that the ACL has been applied.

   ```
   device(config-if-cpu-active)# show running-config interface cpu active
   interface cpu active
   ip access-group block_telnet in
   ```

5. When you are finished, exit CPU configuration mode.

   ```
   device(config-if-cpu-active)# exit
   device(config)#
   ```

The following example creates a traffic policy, adds it to an ACL (cpu_ipv4), applies the ACL to the CPU interface, and verifies the configuration.

```
device# configure terminal
device(config)# traffic-policy TPDF1 rate-limit packet-based fixed cir 10000 exceed-action drop
device(config)# ip access-list extended cpu_ipv4
device(config-ext-ipacl-cpu_ipv4)# permit ip host 10.10.12.2 any traffic-policy TPDF1
device(config-ext-ipacl-cpu_ipv4)# interface cpu active
device(config-if-cpu-active)# ip access-group cpu_ipv4 in
device(config-if-cpu-active)# show running-config interface cpu active
interface cpu active
ip access-group cpu_ipv4 in
device(config-if-cpu-active)# exit
device(config)#
```

Rate Shaping

Rate Shaping Configuration Notes

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.

Rate shaping smooths 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 following apply when configuring outbound rate shaping:

- Outbound rate shaping can be configured only on physical ports, not on virtual or loopback ports.

- RUCKUS ICX devices have 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.

- 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 any other type of LAG (for example, keepalive).

- You cannot configure rate shaping on devices where cut-through mode is configured. If cut-through mode is configured, you must first configure the **store-and-forward** command to change the method to store-and-forward.

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

- When rate-limit output shaping is configured on a port, the traffic going out of this port is rate-limited, so traffic statistics on this port will show the rate-limited traffic.

The configured rate shaping 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 kilobits per second (Kbps) for all the platforms.

TABLE 18 Output Rate Shaping on FastIron Devices

Device	Module	Minimum	Maximum
ICX 8200	10 Gbps ports	8 Kbps	10,000,000 Kbps
ICX 8200	1 Gbps ports	8 Kbps	999,936 Kbps
ICX 7850	25 Gbps ports	8 Kbps	25,000,000 Kbps
ICX 7850	40 Gbps ports	8 Kbps	40,000,000 Kbps
ICX 7850	100 Gbps ports	8 Kbps	100,000,000 Kbps
ICX 7850	10 Gbps ports	8 Kbps	10,000,000 Kbps
ICX 7650	1 Gbps ports	8 Kbps	999,936 Kbps
ICX 7650	10 Gbps ports	8 Kbps	10,000,000 Kbps
ICX 7650	40 Gbps ports	8 Kbps	40,000,000 Kbps
ICX 7650	100 Gbps ports	8 Kbps	100,000,000 Kbps
ICX 7550	1Gbps	8 Kbps	999,936 Kbps
ICX 7550	10 Gbps	8 Kbps	10,000,000 Kbps
ICX 7550	25 Gbps	8 Kbps	25,000,000 Kbps
ICX 7550	40 Gbps	8 Kbps	40,000,000 Kbps
ICX 7550	100 Gbps	8 Kbps	100,000,000 Kbps

Configuring Rate Shaping

Follow these steps to configure outbound rate shaping on an Ethernet or a link aggregation group (LAG) port.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter interface configuration mode.

```
device(config)# interface ethernet 1/1/2
```

3. Configure the maximum rate at which outbound traffic is sent out on a port.

```
device(config-if-e1000-1/1/2)# rate-limit output shaping 1300
Outbound Rate Shaping on Port 1/1/2 Config: 1300 Kbps, Actual: 1304 Kbps
```

4. Configure the maximum rate at which outbound traffic is sent out on a port priority queue.

```
device(config-if-e1000-1/1/2)# rate-limit output shaping 500 priority 7
Outbound Rate Shaping on Port 1/1/2 for Priority 7
        Config: 500 Kbps, Actual: 500 Kbps
device(config-if-e1000-1/1/2)# end
```

5. Verify the configuration.

```
device# show rate-limit output-shaping
Outbound Rate Shaping Limits in Kbps:
Port    PortMax Prio0 Prio1 Prio2 Prio3 Prio4 Prio5 Prio6 Prio7
1/1/2     1304    -     -     -     -     -     -     -    500
1/1/3     1302    -     -     -     -     -     -     -     -
1/1/4      651    -     -     -     -     -     -     -     -
```

Outbound Rate Shaping Configuration Example

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

Configuring Rate Shaping on 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, configure the following on each LAG port where outbound traffic is shaped.

1. Enter global configuration mode.

    ```
    device# configure terminal
    ```

2. Enter LAG configuration sub-mode for an existing LAG. In this example, the LAG is static LAG lag1.

    ```
    device(config)# lag lag1
    ```

3. Configure the maximum rate at which outbound traffic is sent out on the LAG port.

    ```
    device(config-lag-lag1)# rate-limit output shaping ethe 1/1/5 651
    Outbound Rate Shaping on Port 1/1/5 Config: 651 Kbps, Actual: 656 Kbps
    device(config-lag-lag1)# exit
    ```

    Be sure to use the abbreviated form of Ethernet - ethe.

4. Verify the configuration.

    ```
    device(config)# show rate-limit output-shaping
    Outbound Rate Shaping Limits in kbps:
         Port    PortMax    Prio0    Prio1    Prio2    Prio3    Prio4    Prio5    Prio6    Prio7
        1/1/5        656        -        -        -        -        -        -        -        -
    ```

Outbound Rate Shaping on a LAG Port Configuration Example

```
device# configure terminal
device(config)# lag lag1
device(config-lag-lag1)# rate-limit output shaping ethe 1/1/5 651
device(config-lag-lag1)# exit
device(config)# show rate-limit output-shaping
```