

AOS-CX 10.09 Quality of Service Guide

6200, 6300, 6400 Switch Series



a Hewlett Packard
Enterprise company

Copyright Information

© Copyright 2022 Hewlett Packard Enterprise Development LP.

Open Source Code

This product includes code licensed under the GNU General Public License, the GNU Lesser General Public License, and/or certain other open source licenses. A complete machine-readable copy of the source code corresponding to such code is available upon request. This offer is valid to anyone in receipt of this information and shall expire three years following the date of the final distribution of this product version by Hewlett Packard Enterprise Company. To obtain such source code, send a check or money order in the amount of US \$10.00 to:

Hewlett Packard Enterprise Company
6280 America Center Drive
San Jose, CA 95002
USA

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Optane™, Pentium®, Xeon®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the U.S. and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

All third-party marks are property of their respective owners.

Contents	3
About this document	5
Applicable products	5
Latest version available online	5
Command syntax notation conventions	5
About the examples	6
Identifying switch ports and interfaces	6
Identifying modular switch components	7
QoS overview	8
End-to-end QoS behavior	8
Best effort service	8
Class of Service	8
Differentiated services	9
QoS on the switch	10
QoS trust	12
Dynamic QoS trust mode	13
Port rate limiting	13
Queue profiles	13
Schedule profiles	15
Egress queue shaping	16
Egress port shaping	16
Virtual output queues	16
6200 and 6300 switch series	16
6400 switch series	17
Terms	17
QoS configuration	19
Configuring QoS	19
Configuring expedited forwarding for VoIP traffic	20
Configuring rate limiting	22
Configuring egress queue shaping	23
Configuring egress port shaping	23
Supporting Ethernet 802.1D Class of Service	24
Monitoring queue operation	25
QoS commands	27
apply qos	27
dwrr queue	29
map queue	30
min-bandwidth	31
name queue	32
qos cos	33
qos cos-map	34
qos dscp	35
qos dscp-map	37
qos queue-profile	38

qos schedule-profile	39
qos shape	40
qos trust	41
rate-limit	43
show interface queues	45
show interface qos	46
show qos cos-map	47
show qos dscp-map	48
show qos queue-profile	49
show qos schedule-profile	51
show qos trust	52
strict queue	53

Support and Other Resources 55

Accessing Aruba Support	55
Accessing Updates	56
Aruba Support Portal	56
My Networking	56
Warranty Information	56
Regulatory Information	56
Documentation Feedback	57

This document describes features of the AOS-CX network operating system. It is intended for administrators responsible for installing, configuring, and managing Aruba switches on a network.

Applicable products

This document applies to the following products:

- Aruba 6200 Switch Series (JL724A, JL725A, JL726A, JL727A, JL728A)
- Aruba 6300 Switch Series (JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL762A)
- Aruba 6400 Switch Series (JL741A, R0X26A, R0X27A, R0X29A, R0X30A)

Latest version available online

Updates to this document can occur after initial publication. For the latest versions of product documentation, see the links provided in [Support and Other Resources](#).

Command syntax notation conventions

Convention	Usage
<code>example-text</code>	Identifies commands and their options and operands, code examples, filenames, pathnames, and output displayed in a command window. Items that appear like the example text in the previous column are to be entered exactly as shown and are required unless enclosed in brackets ([]).
example-text	In code and screen examples, indicates text entered by a user.
Any of the following: <ul style="list-style-type: none">■ <code><example-text></code>■ <i>example-text</i>■ <code>example-text</code>■ <i>example-text</i>	Identifies a placeholder—such as a parameter or a variable—that you must substitute with an actual value in a command or in code: <ul style="list-style-type: none">■ For output formats where italic text cannot be displayed, variables are enclosed in angle brackets (< >). Substitute the text—including the enclosing angle brackets—with an actual value.■ For output formats where italic text can be displayed, variables might or might not be enclosed in angle brackets. Substitute the text including the enclosing angle brackets, if any, with an actual value.
	Vertical bar. A logical OR that separates multiple items from which you can choose only one. Any spaces that are on either side of the vertical bar are included for readability and are not a required part of the command syntax.

Convention	Usage
{ }	Braces. Indicates that at least one of the enclosed items is required.
[]	Brackets. Indicates that the enclosed item or items are optional.
... or ...	Ellipsis: <ul style="list-style-type: none"> ■ In code and screen examples, a vertical or horizontal ellipsis indicates an omission of information. ■ In syntax using brackets and braces, an ellipsis indicates items that can be repeated. When an item followed by ellipses is enclosed in brackets, zero or more items can be specified.

About the examples

Examples in this document are representative and might not match your particular switch or environment. The slot and port numbers in this document are for illustration only and might be unavailable on your switch.

Understanding the CLI prompts

When illustrating the prompts in the command line interface (CLI), this document uses the generic term `switch`, instead of the host name of the switch. For example:

```
switch>
```

The CLI prompt indicates the current command context. For example:

```
switch>
```

Indicates the operator command context.

```
switch#
```

Indicates the manager command context.

```
switch (CONTEXT-NAME)#
```

Indicates the configuration context for a feature. For example:

```
switch(config-if)#
```

Identifies the `interface` context.

Variable information in CLI prompts

In certain configuration contexts, the prompt may include variable information. For example, when in the VLAN configuration context, a VLAN number appears in the prompt:

```
switch(config-vlan-100)#
```

When referring to this context, this document uses the syntax:

```
switch(config-vlan-<VLAN-ID>)#
```

Where `<VLAN-ID>` is a variable representing the VLAN number.

Identifying switch ports and interfaces

Physical ports on the switch and their corresponding logical software interfaces are identified using the format:

```
member/slot/port
```

On the 6200 Switch Series

- *member*: Member number of the switch in a Virtual Switching Framework (VSF) stack. Range: 1 to 8. The primary switch is always member 1. If the switch is not a member of a VSF stack, then member is 1.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface 1/1/4 in software is associated with physical port 4 in slot 1 on member 1.

On the 6300 Switch Series

- *member*: Member number of the switch in a Virtual Switching Framework (VSF) stack. Range: 1 to 10. The primary switch is always member 1. If the switch is not a member of a VSF stack, then member is 1.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface 1/1/4 in software is associated with physical port 4 on member 1.

On the 6400 Switch Series

- *member*: Always 1. VSF is not supported on this switch.
- *slot*: Specifies physical location of a module in the switch chassis.
 - Management modules are on the front of the switch in slots 1/1 and 1/2.
 - Line modules are on the front of the switch starting in slot 1/3.
- *port*: Physical number of a port on a line module.

For example, the logical interface 1/3/4 in software is associated with physical port 4 in slot 3 on member 1.

Identifying modular switch components

- Power supplies are on the front of the switch behind the bezel above the management modules. Power supplies are labeled in software in the format: *member/power supply*:
 - *member*: 1.
 - *power supply*: 1 to 4.
- Fans are on the rear of the switch and are labeled in software as: *member/tray/fan*:
 - *member*: 1.
 - *tray*: 1 to 4.
 - *fan*: 1 to 4.
- Fabric modules are not labeled on the switch but are labeled in software in the format: *member/module*:
 - *member*: 1.
 - *member*: 1 or 2.
- The display module on the rear of the switch is not labeled with a member or slot number.

Quality of Service (QoS) enables network administrators to customize how different types of traffic are serviced on a network, taking into account the unique characteristics of each traffic type and its importance within an organization's infrastructure. QoS ensures uniform and efficient traffic handling, keeping the most important traffic moving at an acceptable speed, regardless of current bandwidth usage. It also provides methods for administrators to control the priority settings of inbound traffic arriving at each network device.

End-to-end QoS behavior

The QoS settings on each network device must be aligned to achieve the desired end-to-end QoS behavior for a network. Three service types can be used to categorize and prioritize network traffic:

- Best Effort Service
- Ethernet Class of Service (CoS)
- Internet Differentiated Services (DiffServ)

For a network as a whole, it is best to select one service type to use as the primary end-to-end behavior, and then use the other two service types as needed.

Best effort service

This is the simplest service type. All traffic is treated equally in a first-come, first-served manner. If the traffic load is low in relation to the capacity of the network links, then there is no need for the administrative complexity and costs of maintaining a more complex end-to-end policy. This is sometimes called over-provisioning, as all link speeds are much higher than peak loads on the network.

Class of Service

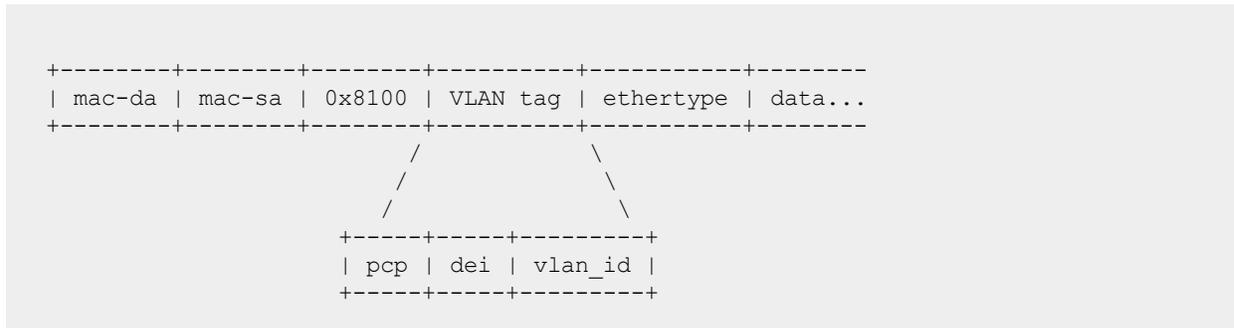
Class of Service (CoS) is a method for classifying network traffic at layer 2 by marking 802.1Q VLAN Ethernet frames with one of eight service classes.

CoS	Traffic type	Example protocols
7	Network Control	STP, PVST
6	Internetwork Control	BGP, OSPF, PIM
5	Voice (<10ms latency)	VoIP(UDP)
4	Video (<100ms latency)	RTP
3	Critical Applications	SQL RPC, SNMP
2	Excellent Effort	NFS, SMB

CoS	Traffic type	Example protocols
0	Best Effort	HTTP, TELNET
1	Background	SMTP, IMAP

CoS 1 is deliberately set as the lowest CoS. This enables a traffic service level below the default (best effort) traffic level to be specified.

The 3-bit Priority Code Point (PCP) field within the 16-bit Ethernet VLAN tag is used to mark the CoS.



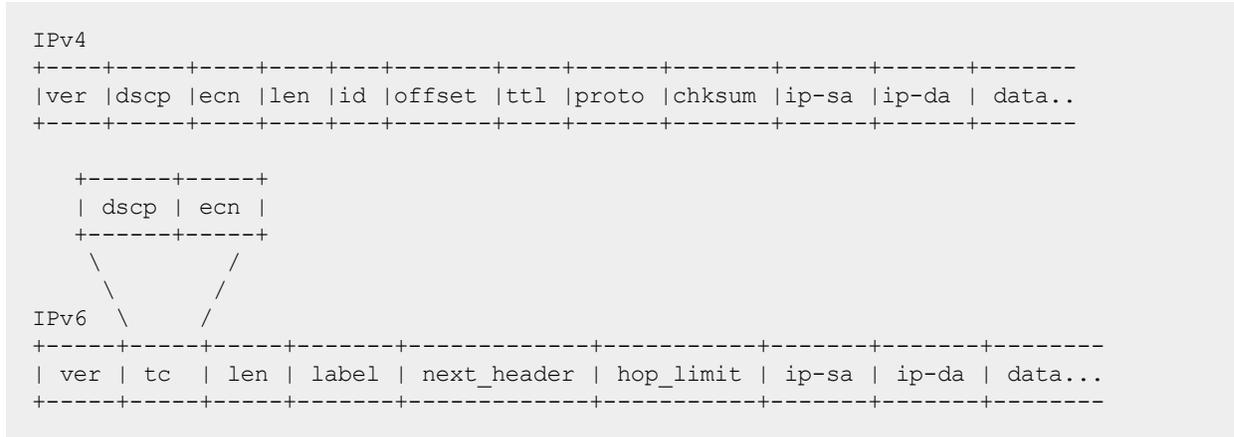
Differentiated services

Differentiated services (DiffServ) is a method for classifying network traffic at layer 3 by marking packets with one of 64 different service classes. Services classes are identified by the Differentiated services Code Point (DSCP) value. Some common DSCP values are:

DSCP	Name	Service class	RFC
48	CS6	Network Control	2474
46	EF	Telephony	3246
40	CS5	Signaling	2474
34, 36, 38	AF41, AF42, AF43	Multimedia Conferencing	2597
32	CS4	Real-Time Interactive	2474
26, 28, 30	AF31, AF32, AF33	Multimedia Streaming	2597
24	CS3	Broadcast Video	2474
18, 20, 22	AF21, AF22, AF23	Low-Latency Data	2597
16	CS2	OAM	2474
00	CS0, BE, DF	Best Effort	2474
10, 12, 14	AF11, AF12, AF13	Bulk Data	2597
08	CS1	Low-Priority Data	3662

DSCP CS1 (08) CoS 1 is deliberately set as the lowest priority. This enables a traffic service level below the standard (best effort or default forwarding) level to be specified.

The DSCP value is carried within the IPv4 DSCP field or the upper 6-bits of the 8-bit IPv6 Traffic Class (TC) field.



QoS on the switch

There are five key stages a packet passes through when traversing a switch: ingress, prioritization, destination determination, egress queuing, and transmission. The following table provides an overview of each stage, and lists the commands that can be used to configure QoS settings.

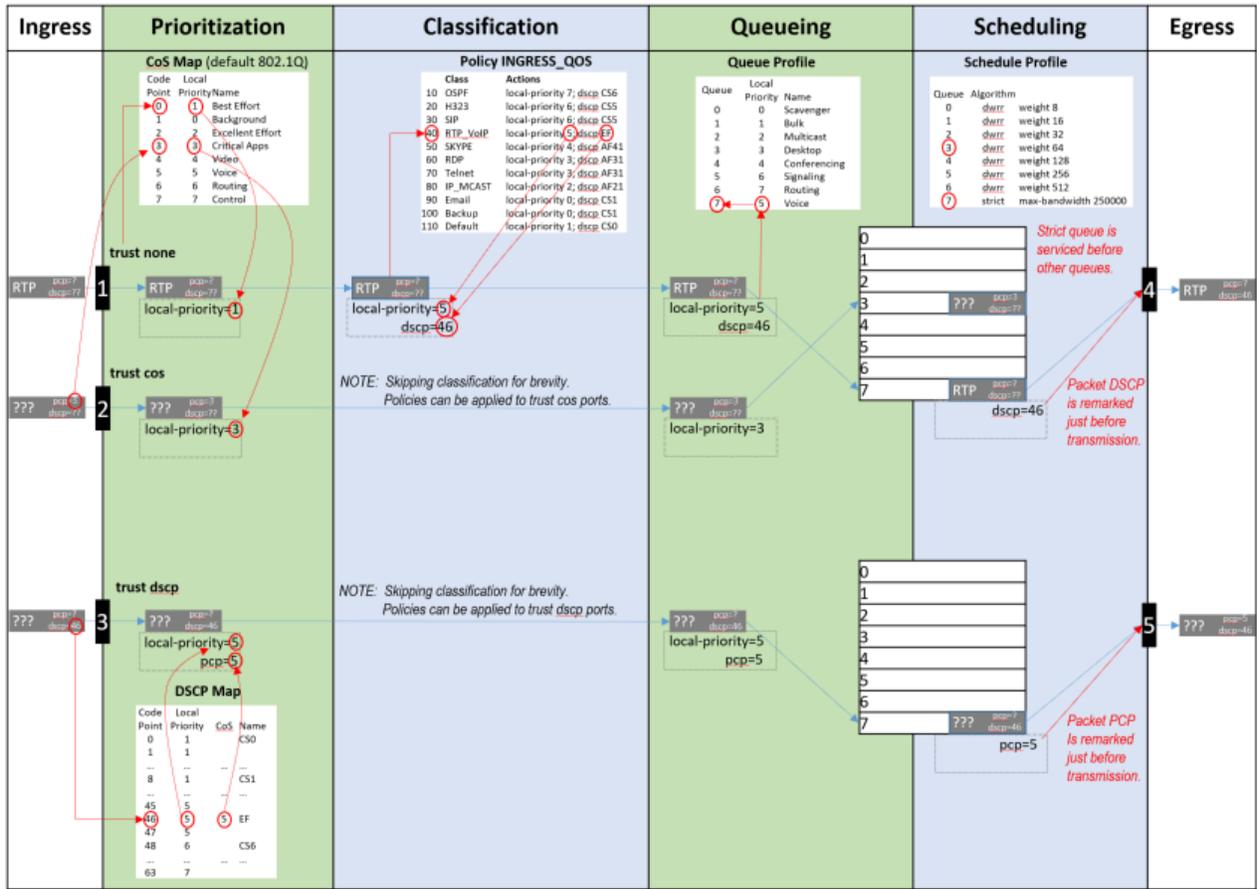


Switches with at least 52 ports will experience negative performance if a flood occurs where at least 42 ports are members of the same VLAN and all 52 ports have QoS rules applied to them.

Ingress	Prioritization	Classification	Queueing	Transmission
Packets arrive at switch interface (port).	An initial local-priority value is assigned to the packet based on VLAN CoS or IP DSCP.	Packets local-priority and code points can be remarked, and rates policed, based on many packet header fields. <i>(see ACL & Policy Guide)</i>	Packets are queued based on the destination interface and local-priority of the packet.	A scheduler defines the order in which packets are selected from queues to be transmitted from the interface.
rate-limit Rate limiting can control ingress flow by packet type: broadcast, multicast, unknown-unicast, or ICMP.	qos trust {cos dscp none} Assigns which packet values are used to determine initial local-priority either globally for all interfaces, or to specific interface(s). qos cos-map Defines how CoS values are mapped to local-priority values. qos dscp-map Defines how IP DSCP values are mapped to local-priority values. qos cos Remarks egress CoS value. qos dscp Remarks egress DSCP value.	class {ip ipv6 mac} NAME match TERMS [count] ignore TERMS [count] Creates an ordered list of rules to identify packets to match the class. policy NAME class NAME ACTIONS Creates an ordered list of classes to which traffic is evaluated, and the actions to be taken on matching packets: remarking, policing etc apply policy NAME Assigns the policy to evaluate all traffic inbound or outbound globally, or on specific interfaces or VLANs.	qos queue-profile Creates a profile that defines transmit queues. map queue Assigns a local-priority to a queue. Packets with a matching local-priority are placed in the queue. name queue Assigns a name to a queue. apply qos Assigns queue, schedule, or threshold profiles globally to all interfaces, or schedule or threshold profiles to a specific interface(s).	qos schedule-profile Creates a profile that defines the order in which packets are selected from the queues for transmission. strict queue Assigns the strict priority algorithm, with optional max-bandwidth, to a queue. dwrr queue Assigns the weighted DWRR algorithm to a queue. qos shape Assigns the maximum rate for traffic to be transmitted by the interface.

Ingress	Prioritization	Classification	Queueing	Transmission
Packets arrive at switch interface (port).	An initial local-priority value is assigned to the packet based on VLAN CoS or IP DSCP.	Packets local-priority and code points can be remarked, and rates policed, based on many packet header fields. <i>(see ACL & Policy Guide)</i>	Packets are queued based on the destination interface and local-priority of the packet.	A scheduler defines the order in which packets are selected from queues to be transmitted from the interface.
	<p>show int IFNAME qos Displays all the QoS settings configured on an interface:</p> <ul style="list-style-type: none"> • Profiles • Rate-limits • CoS or DSCP remarks • Egress shape • Trust <p>show qos cos-map Displays the contents of the CoS Map table.</p> <p>show qos dscp-map Displays the content of the DSCP Map table.</p> <p>show qos trust Display the current default trust mode.</p>	<p>show class {ip ipv6 mac} Displays the ordered list of the match and ignore statements in the class.</p> <p>show policy Displays the order list of classes in the policy.</p> <p>show policy hitcounts Displays the ordered list of classes in the policy. For each class the number of times that each match or ignore statement in the class matched a packet.</p>	<p>show int IFNAME queues Displays the per-queue counts of packets transmitted, packets dropped, and bytes transmitted.</p> <p>show qos queue-profile Displays the definition of queue profiles.</p>	<p>show qos schedule-profile Displays the definition of schedule profiles.</p>

The following diagram shows how different packets might traverse a switch. It also shows how QoS configuration settings apply at each stage.



QoS trust

When trust mode is set to CoS or DSCP, the switch translates the QoS settings in VLAN tags (for CoS), or the DS field in an IP header (for DSCP), to local priority values on the switch. Translation is controlled by the CoS map or DSCP map tables.

For example:

CoS Map

When QoS trust mode = CoS

Code Point	Local Priority	Name
0	1	Best Effort
1	0	Background
2	2	Excellent Effort
3	3	Critical Apps
4	4	Video
5	5	Voice
6	6	Routing
7	7	Control

Untagged packet → Packet local-priority = 1
 Tagged PCP=3 packet → Packet local-priority = 3 & PCP = 3

DSCP Map

When QoS trust mode = DSCP

Code Point	Local Priority	CoS	Name
0	1		CS0
1	1		
...
8	1		CS1
...
45	5		
46	5	5	EF
47	5		
48	6		CS6
...
63	7		

Untagged IP DSCP=46 packet → Packet local-priority = 5 & PCP = 5 & DSCP = 46
 Tagged PCP=3 IP DSCP=46 packet → Packet local-priority = 5 & PCP = 5 & DSCP = 46

When QoS trust mode = None

Tagged PCP=3 IP DSCP=32 packet	→	Packet local-priority = 1 (Cos Map 0) & PCP = 3 & DSCP = 32
	→	qos cos 6 → Packet local-priority = 1 (Cos Map 0) & PCP = 6 & DSCP = 32
	→	qos dscp 48 → Packet local-priority = 1 (Cos Map 0) & PCP = 3 & DSCP = 48

Dynamic QoS trust mode

The device profile feature can dynamically set the QoS trust mode on an interface based on the LLDP information exchanged with a link partner. The device profile's trust mode temporarily overrides the static trust mode configured for an interface. The override remains in place as long as that link partner is connected and its link state is **up**. Use command `show interface IFNAME qos` to view the current QoS trust mode for an interface.

Port rate limiting

Port rate limiting helps control undesirable traffic. Its purpose is to allow enough broadcast, multicast, and ICMP rate-limit traffic for the network to function properly, while preventing flooding and traffic storms.

A certain amount of each type of traffic is required for normal network operation. Broadcast packets may include ARP and DHCP traffic, for instance. Video streams, and certain types of network protocol packets, are multicasts. Configuring rate limits can help provide the balance between necessary and flooded traffic.

Queue profiles

A queue profile defines the queues that are associated with an interface to control the transmission of packets. Each profile supports up to eight queues, numbered 0 to 7. The larger the queue number, the higher its priority during transmission scheduling. Packets are assigned to a queue based on their local priority value (0 to 7). A queue profile must map all eight local priority values to whatever queues are being

used on the switch, and a schedule profile must specify the configuration for those same queues. A queue without a local priority value assigned to it is not used to store packets.

The switch is automatically provisioned with an initial queue profile named `factory-default` which assigns each local priority to the queue of the same number. To see the default queue profile, use the command `show qos queue-profile factory-default`:

```
switch# show qos queue-profile factory-default
queue_num local_priorities name
-----
0          0                  Scavenger_and_backup_data
1          1
2          2
3          3
4          4
5          5
6          6
7          7
```

More than one local priority value can be assigned to the same queue. For example,

Local Priority	Queue
0	0
1	1
2	2
3	3
4	4
5	5
6	5
7	5

Queues must be consecutively defined starting at queue number zero. For example, a four-queue profile with priority values defined for queues 0, 1, 2, 3 is valid, but a four-queue profile which defines priority values for queues 1, 3, 5, and 7 is not.

Table 1: *Valid queue*

Local Priority	Queue
0	0
1	1
2	1
3	1

Local Priority	Queue
4	2
5	2
6	3
7	3

Table 2: *Invalid queue*

Local Priority	Queue
0	0
1	1
2	1
3	3
4	3
5	5
6	5
7	7

Commonly used commands for working with QoS queues are as follows:

- `qos queue-profile`: Creates an empty queue-profile and enters the profile configuration context.
- `name queue`: Assigns a descriptive name to a queue.
- `map queue`: Assigns a local-priority to a queue.
- `apply qos queue-profile`: Applies a queue-profile globally to all interfaces.

Schedule profiles

A schedule profile determines the order in which queues are selected for transmission, and the amount of service available for each queue. A schedule profile must be configured on every interface at all times. A schedule profile can be applied globally to all interfaces, or only to specific interfaces.

A weighted schedule profile assigns relative servicing for each queue. The amount of service per weight is relative to the underlying hardware implementation, and to the weights assigned to the other non-empty queues. Strict scheduling can be used to service queues purely on the basis of highest priority first (at the risk of starving lower-priority queues during high stress periods). A combination of strict and weighted scheduling offers more service to the highest priority queue when needed, while preserving scheduling between the remaining queues, thus decreasing the risk of starvation.

The switch is automatically provisioned with a schedule profile named **factory-default**, which assigns DWRR to all queues with a weight of 1. Use the command `show schedule-profile factory-default` to

view the default schedule profile. (Do not use `show running-configuration`, as it only displays changes from the initial settings.) The 6400 Switch Series only has four queues.

```
switch# show qos schedule-profile factory-default
queue_num algorithm weight
-----
0          dwrr        1
1          dwrr        1
2          dwrr        1
3          dwrr        1
4          dwrr        1
5          dwrr        1
6          dwrr        1
7          dwrr        1
```

Egress queue shaping

Egress queue shaping limits the amount of traffic transmitted per strict output queue. The buffer associated with each egress queue stores excess traffic to absorb bursts and smooths the output rate. For example, an administrator might limit strict-priority queue traffic to prevent low-priority queue starvation in the event that a device inappropriately sends too many higher-priority packets.

Egress queue shaping can be configured on an Ethernet port or on a link aggregation group (LAG). To configure egress queue shaping, define a schedule profile with the strict priority algorithm assigned to each queue.

Egress port shaping

Egress port shaping limits the amount of aggregate traffic transmitted through a port. To be effective, the egress port-shaping rate must be less than the port's line rate. By default, the egress port-shaping rate is the same as the line-rate of the port. Buffers associated with each port store excess traffic. When both egress port-shaping and egress queue-shaping are configured on the same interface, the switch respects the minimum of both configurations.

Virtual output queues

6200 and 6300 switch series

The 6200 and 6300 Switch Series contain 8 MB of packet buffer memory. Approximately 5 MB are used for buffering traffic arriving on network ports. Each port can buffer at most 1 MB of packets divided equally across all the active queues in a queue profile. For example, using the default eight-queue profile, each virtual output queue (VOQ) can buffer a maximum of 128 KB of packets.

Replicated traffic (L2 broadcasts, L2 multicasts, IP multicasts, mirrored packets, and unknown-unicasts) requires an extra step. There are eight additional replication queues initially used for replicated traffic. The switch services the replication queues strictly at the rate of 100 G. For ports on the same switch, pointers to the same packet are placed in the VOQ for each destination port. For ports on remote VSF member switches, only one copy of the packet data is moved across the VSF link to the destination replication VOQ to optimize link bandwidth. From there, the destination switch puts pointers to the same packet on the VOQ for each destination port.

When 3-8 queues are configured on the 6300 switch series the egress buffer limit is 2 MB. The most buffer a single port can use in a given queue is 2MB and that 2 MB is shared between the ports in the queue that need it. If QoS is not needed then up 1-2 queues can be configured with an egress buffer limit of 4 MB.



The 6300 switch series is the only platform where collapsing the number of queues increases the buffer allocation limit.

6400 switch series

The 6400 Switch Series use a VOQ architecture where most packet buffering occurs on the ingress line module. Only shallow buffers are used on the egress line module. This optimizes usage of fabric bandwidth. Traffic destined for one port (unicast) uses different buffering and scheduling than traffic destined for multiple ports (replication).

Each line module contains eight VOQs for every destination port in the chassis, and eight egress physical queues (EPQs) for the ports local to the line module, regardless of the number of queues in the queue profile. The 6400 Switch Series R0X38-R0X45 line modules contain 8 MB of packet buffer memory. Approximately 5 MB is used for buffering traffic arriving on network ports. The remainder is used for buffering traffic arriving across the fabric. On each line module, at most 1 MB of buffering is available for every destination port in the chassis. This 1 MB buffer is equally divided across all the active queues in the queue profile.

For example, using the default eight-queue profile, each VOQ can buffer at most 128 KB of packets. Packets wait in VOQs on the arrival line module until the destination port scheduler selects them to be moved to one of its EPQs. If the destination port is on another line module, the packet is moved across the fabric. If the destination port EPQ is on the same line module, only a pointer to the packet is moved from the VOQ to the EPQ.

The schedule profile determines the order the destination port scheduler uses for VOQ servicing. The scheduler moves only enough packets to EPQs as are needed to keep the destination port transmitting at the configured rate.

Replicated traffic (L2 broadcasts, L2 multicasts, IP multicasts, mirrored packets, and unknown-unicasts) requires an extra step. There are eight additional internal replication VOQs per line module initially used for replicated traffic. The line module services these replication VOQs strictly at a rate of 100 G. For ports on the same line module, pointers to the same packet are placed in the VOQ for each destination port. For ports on remote line modules, only one copy of the packet data is moved across the fabric to the destination replication VOQ, to optimize fabric bandwidth. From there, the destination line module will put pointers to the same packet on the VOQ for each destination port.

Terms

Class

For networking, a set of packets sharing a common characteristic. For example, all IPv4 packets.

Code point

The name of a packet header field, or the value carried within a packet header field:

- Example 1: Priority code point (PCP) is the name of a field in the IEEE 802.1Q VLAN tag.
- Example 2: Differentiated services code point (DSCP) is the name of a field carried within the DS field of an IP packet header.

Color

A metadata label associated with each packet within the switch. It has two values: green (0) or yellow (1). When packets encounter congestion for a resource (queue), the switch uses packet color to distinguish which packets must be dropped, and is mostly used for packets marked with Assured Forwarding (AF) DSCP values.

Not supported in this release.

Class of service (CoS)

A 3-bit value used to mark packets with one of eight classes (levels of priority). It is carried within the priority code point (PCP) field of the IEEE 802.1Q VLAN tag.

Differentiated services code point (DSCP)

A 6-bit value used to mark packets for different per-hop behavior as originally defined by IETF RFC 2474. It is carried within the differentiated services (DS) field of the IPv4 or IPv6 header.

Local priority

A meta-data label associated with a packet within the switch which is used to classify packets for different treatment (such as queue assignment). Eight local priorities are defined on the switch, numbered from 0 to 7. A queue profile must map all eight local priorities to whatever queues are in use on the switch, and a schedule profile must specify the configuration for these same queues.

Metadata

Information labels associated with each packet in the switch, separate from the packet headers and data. These labels are used by the switch in its handling of the packet. For example: arrival port, egress port, VLAN membership, and local priority.

Priority code point (PCP)

The name of a 3-bit field in the IEEE 802.1Q VLAN tag. It carries the CoS value to mark a packet with one of eight classes (priority levels).

Quality of service (QoS)

General term used when describing or measuring performance. For networking, it means how different classes of packets are treated when traversing a network or device.

Traffic class (TC)

General term for a set of packets sharing a common characteristic. It used to be the name of an 8-bit field in the IPv6 header originally defined by IETF RFC 2460. This field name was changed to differentiated services by IETF RFC 2474.

Type of service (ToS)

General term when there are different levels of treatment (fare class). It used to be the name of an 8-bit field in the IPv4 header originally defined by IETF RFC 791. This field name was changed to differentiated services by IETF RFC 2474.

Configuring QoS

Procedure

1. Configure how local priority values are assigned to ingress packets with the commands `qos dscp-map`, and `qos trust`.
2. Optionally, add a rate limit for ingress traffic on one or more interfaces with the command `rate-limit`.
3. If you do not want to use the default QoS queue profile to map local priority to queue, create one or more custom queue profiles with the command `qos queue-profile`. For each queue in a custom queue profile:
 - a. Assign a local priority value with the command `map queue`.
 - b. Optionally, define a descriptive name with the command `name queue`. All local priorities (0 to 7) must be mapped to a queue, and the queues selected for use must be in contiguous order starting at 0.
4. If you do not want to use the default QoS schedule profile to determine the order in which queues are selected to transmit a packet, create one or more custom schedule profiles with the command `qos schedule-profile`. For each queue in a custom schedule queue profile, define scheduling priority with the commands `strict queue` and `dwrr queue`.
5. Optionally for strict queues, configure egress queue shaping to limit egress bandwidth on an interface to a value that is less than its line rate. Use the `max-bandwidth` parameter of the `strict queue` command.
6. Activate QoS settings with the command `apply qos`. This command lets you apply a queue profile and schedule profile globally to all interfaces, or a schedule profile override to individual interfaces.
7. View QoS configuration settings with the provided `show` commands.

Examples

This example creates the following configuration:

- Configures CoS to be used to assign local priority to ingress packets.
- Modifies the default CoS map to assign CoS 1 to local priority **1**.
- Creates a queue profile named `q1` and assigns local priorities as follows:

Queue	Local Priority
0	0
1	1

Queue	Local Priority
1	2
2	3
3	4
4	5
5	6
5	7

- Creates a schedule profile named `s1` and assigns DWRR to all queues in the schedule profile with the following weights:

Queue	Weight
0	5
1	10
2	15
3	20
4	25
5	50

- Applies **Q1** and **S1** to all interfaces that do not have a QoS override applied.

```

switch(config)# qos trust cos
switch(config)# qos cos-map 1 local-priority 1
switch(config)# qos queue-profile Q1
switch(config-queue)# map queue 0 local-priority 0
switch(config-queue)# map queue 1 local-priority 1
switch(config-queue)# map queue 1 local-priority 2
switch(config-queue)# map queue 2 local-priority 3
switch(config-queue)# map queue 3 local-priority 4
switch(config-queue)# map queue 4 local-priority 5
switch(config-queue)# map queue 5 local-priority 6
switch(config-queue)# map queue 5 local-priority 7
switch(config-queue)# qos schedule-profile S1
switch(config-schedule)# dwrr queue 0 weight 5
switch(config-schedule)# dwrr queue 1 weight 10
switch(config-schedule)# dwrr queue 2 weight 15
switch(config-schedule)# dwrr queue 3 weight 20
switch(config-schedule)# dwrr queue 4 weight 25
switch(config-schedule)# dwrr queue 5 weight 50
switch(config)# apply qos queue-profile Q1 schedule-profile S1

```

Configuring expedited forwarding for VoIP traffic

Voice over IP (VoIP) traffic is delay and jitter sensitive. For optimum transmission of VoIP traffic, dwell time in network devices must be kept to a minimum and all network devices in the data path must have identical per-hop behaviors. To configure a dedicated queue on the switch to handle VoIP traffic with priority service before all other queues, follow these steps.

Prerequisites

This scenario assumes that VoIP packets are uniquely identified using DiffServ code point 46, Expedited Forwarding (EF).

Procedure

1. Map DSCP EF packets exclusively to local priority 5. The default DSCP map has eight code points (40 through 47), that are mapped to local priority 5. To reserve local priority 5 for VoIP traffic, the other code points must be reassigned. In this scenario, local priority 6 is used for all reassignments, including for code point 40, Call Signaling protocol (CS5).

```
switch(config)# qos dscp-map 40 local-priority 6 name CS5
switch(config)# qos dscp-map 41 local-priority 6
switch(config)# qos dscp-map 42 local-priority 6
switch(config)# qos dscp-map 43 local-priority 6
switch(config)# qos dscp-map 44 local-priority 6
switch(config)# qos dscp-map 45 local-priority 6
switch(config)# qos dscp-map 47 local-priority 6
```

2. Queue 7 is the highest priority queue, so for best throughput, create a queue profile that maps local priority to queue 7.

```
switch(config)# qos queue-profile ef_priority
switch(config-queue)# name queue 7 Voice_Priority_Queue
switch(config-queue)# map queue 7 local-priority 5
switch(config-queue)# map queue 6 local-priority 7
switch(config-queue)# map queue 5 local-priority 6
switch(config-queue)# map queue 4 local-priority 4
switch(config-queue)# map queue 3 local-priority 3
switch(config-queue)# map queue 2 local-priority 2
switch(config-queue)# map queue 1 local-priority 1
switch(config-queue)# map queue 0 local-priority 0
```

3. Create a schedule profile that services queue 7 using strict priority (SP), and the remaining queues with DWRR. This scenario gives all DWRR queues equal weight.

```
switch(config)# qos schedule-profile voip
switch(config-schedule)# strict queue 7
switch(config-schedule)# dwrr queue 6 weight 1
switch(config-schedule)# dwrr queue 5 weight 1
switch(config-schedule)# dwrr queue 4 weight 1
switch(config-schedule)# dwrr queue 3 weight 1
switch(config-schedule)# dwrr queue 2 weight 1
switch(config-schedule)# dwrr queue 1 weight 1
switch(config-schedule)# dwrr queue 0 weight 1
switch(config-schedule)# exit
switch(config)#
```

4. Apply the profiles to all interfaces.

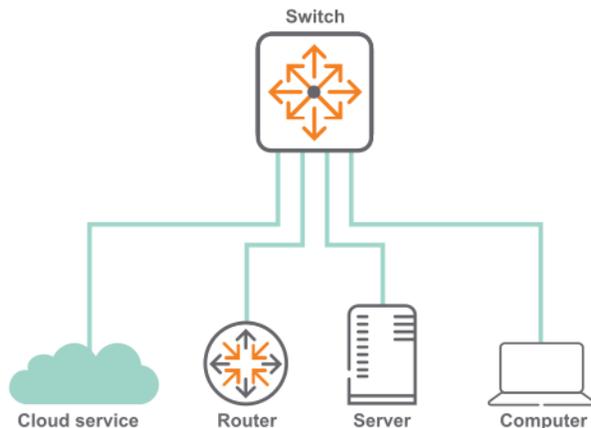
```
switch(config)# apply qos queue-profile ef_priority schedule-profile voip
```

5. Configure DSCP trust mode on all ports

```
switch(config)# qos trust dscp
```

Configuring rate limiting

This scenario illustrates how to use rate limiting to manage the traffic from various devices connected to a switch. The physical topology of the network looks like this:



A certain amount of broadcast traffic is necessary to maintain healthy network operation, particularly from routers and across service boundaries. In this scenario, both the service cloud and the router connections limit this traffic to 1 Gbps. The server has a smaller limit, as it does not require as much network protocol traffic as the service cloud and router.

A multicast server needs to be able to stream multicast traffic to clients, so a multicast rate limit may not be helpful. A computer, however, should not be generating large amounts of multicast traffic (it may be receiving streams, but typically not sending them). In this example, the computer is configured with a multicast rate limit to prevent malicious traffic from taking up network bandwidth.

Finally, while the service cloud and router may need to send traffic for unknown unicast addresses to resolve address forwarding, the server and computer should send very little of this type of traffic. Rate limiting unknown unicast traffic on those two devices enforces that.

On the 6400 Switch Series, interface identification differs.

Procedure

1. Configure broadcast and multicast rate limiting for the service cloud connection.

```
switch# config
switch(config)# interface 1/1/1
switch(config-if)# rate-limit broadcast 500 pps
switch(config-if)# rate-limit multicast 500 pps
switch(config-if)# exit
```

2. Configure broadcast rate limiting for the router connection.

```
switch(config-if)# interface 1/1/2
switch(config-if)# rate-limit broadcast 500 pps
switch(config-if)# exit
```

3. Configure broadcast rate limiting for the server connection.

```
switch(config-if)# interface 1/1/5
switch(config-if)# rate-limit broadcast 100 pps
switch(config-if)# exit
```

4. Configure broadcast, and multicast rate limiting for the computer connection.

```
switch(config-if)# interface 1/1/10
switch(config-if)# rate-limit broadcast 50 pps
switch(config-if)# rate-limit multicast 50 pps
```

Configuring egress queue shaping

This example shows how to apply egress queue shaping to an interface. First, a schedule profile is created that has per-queue bandwidth limits set on all queues with `strict` as the scheduling algorithm. Next, this profile is applied to an interface or LAG.

On the 6400 Switch Series, interface identification differs.

The following example creates a schedule profile named **EQSEExample**, which services all seven queues using `strict` priority. This profile configures queues 1, 4, and 7 with a bandwidth limit of 10 Gbps, 20 Gbps, and 30 Gbps respectively. The profile is then applied to interface 1/1/1.

```
switch(config)# qos schedule-profile EQSEExample
switch(config-schedule)# strict queue 0
switch(config-schedule)# strict queue 1 max-bandwidth 10000000
switch(config-schedule)# strict queue 2
switch(config-schedule)# strict queue 3
switch(config-schedule)# strict queue 4 max-bandwidth 20000000
switch(config-schedule)# strict queue 5
switch(config-schedule)# strict queue 6
switch(config-schedule)# strict queue 7 max-bandwidth 30000000
switch(config-schedule)# exit
switch(config)# interface 1/1/1
switch(config-if)# apply qos schedule-profile EQSEExample
```

Configuring egress port shaping

This example shows how to apply egress port shaping to an interface to limit the rate of egress traffic. Egress port shaping is configured by specifying the desired bandwidth rate in kilobits per second (kbps). To be effective, the egress rate must be less than the line rate of the egress interface. If the configured egress rate exceeds the interface's line rate, then egress shaping has no effect.

The configured egress rate on a specific interface can be found by using the `show interface` and `show interface qos` commands.

On the 6400 Switch Series, interface identification differs.

The following example configures an egress rate of 100 Mbps .

```
switch(config)# interface 1/1/1
switch(config-if)# qos shape 100000
```

In the next example, both egress port shaping and egress queue shaping are configured on the same interface.

The example creates a schedule profile named `EQSEExample` with strict priority for all seven queues. Queue 7 is configured with a bandwidth limit of 300 Mbps. The profile is then applied to interface `1/1/1` with egress port shaping of 400 Mbps . As egress queue shaping and egress port shaping are both configured on port `1/1/1`, egress queue shaping is subject to the lower port or queue shape rate. The effective bandwidth for the traffic egressing on queue 7 will be 300 Mbps .

```
switch(config)# qos schedule-profile EQSEExample
switch(config-schedule)# strict queue 0
switch(config-schedule)# strict queue 1
switch(config-schedule)# strict queue 2
switch(config-schedule)# strict queue 3
switch(config-schedule)# strict queue 4
switch(config-schedule)# strict queue 5
switch(config-schedule)# strict queue 6
switch(config-schedule)# strict queue 7 max-bandwidth 300000
switch(config-schedule)# exit
switch(config)# interface 1/1/1
switch(config-if)# apply qos schedule-profile EQSEExample
switch(config-if)# qos shape 400000
```

Supporting Ethernet 802.1D Class of Service

IEEE 802.1Q is the most current Ethernet standard for Class of Service (CoS). It superseded an earlier standard, 802.1D, in 2005. IEEE 802.1Q slightly changed the ordering of the classes of service from its predecessor IEEE 802.1D for CoS 2 and CoS 0:

CoS 802.1Q	CoS 802.1D
7 Network Control	7 Network Control
6 Internetwork Control	6 Voice (<10ms latency)
5 Voice (<10ms latency)	5 Video (<100ms latency)
4 Video (<100ms latency)	4 Controlled Load
3 Critical Applications	3 Excellent Effort

CoS 802.1Q	CoS 802.1D
2 Excellent Effort	0 Best Effort
0 Best Effort	2 Spare
1 Background	1 Background

In 802.1D, both CoS 2 and CoS 1 are below CoS 0 (Best Effort).

When a switch is installed in a network of devices following 802.1D Class of Service, the QoS Cos map must be re-configured to follow the 802.1D standard by swapping the assignments of CoS 0 and 2:

```
switch# config
switch(config)# qos cos-map 0 local-priority 2 color green name Best_Effort
switch(config)# qos cos-map 2 local-priority 1 color green name Spare
switch(config)# exit
switch# show qos cos-map
code_point local_priority color name
-----
0          2          green Best_Effort
1          0          green Background
2          1          green Spare
3          3          green Critical_Applications
4          4          green Video
5          5          green Voice
6          6          green Internetwork_Control
7          7          green Network_Control
```

Monitoring queue operation

Use the show interface queues command to display the traffic transmitted per queue, and the number of packets dropped due to the queue being full. (Tx Bytes is available on the 6200 and 6300 Switch Series.) For example:

```
switch# show interface 1/1/1 queues
Interface 1/1/1 is (Administratively down)
Admin state is down
State information: admin_down
      Tx Packets      Tx Bytes      Tx Drops
Q0          100          8000          0
Q1       1234567      12345678908          5
Q2           0           0          0
Q3           0           0          0
Q4           0           0          0
Q5           0           0          0
Q6           0           0          0
Q7           0           0          0
```

Tx Bytes: Total bytes transmitted (available on the 6200 and 6300 Switch Series). The byte count may include packet headers and internal metadata that are removed before the packet is transmitted. Packet headers added when the packet is transmitted may not be included.

Tx Packets: Total packets transmitted.

Tx Drops: The number of packets dropped by a queue before it was sent. When traffic cannot be forwarded out an egress interface, it backs up at ingress. The more servicing assigned to a queue by a schedule profile, the less likely traffic destined for that queue will back up and be dropped. Tx Drops shows the sum of packets that were dropped across all line modules (due to insufficient capacity) by the ingress Virtual Output Queues (VOQs) destined for the egress port. Applicable to the 6400 Switch Series: Because the counts are read separately from each line module, the sum is not an instantaneous snapshot.

apply qos

```
apply qos [queue-profile <QUEUE-NAME>] schedule-profile <SCHEDULE-NAME>
no apply qos schedule-profile <SCHEDULE-NAME>
```

Description

Applies a queue profile and schedule profile globally to all Ethernet and LAG interfaces on the switch, or applies a schedule profile to a specific interface. When applied globally, the specified schedule profile is configured only on Ethernet interfaces and LAGs that do not already have their own schedule profile.

The same profile can be applied both globally and locally to an interface. This guarantees that an interface always uses the specified profile, even if the global profile is changed.

The `no` form of this command removes the specified schedule profile from an interface and the interface uses the global schedule profile. This is the only way to remove a schedule profile override from the interface.



Interfaces may shut down briefly during reconfiguration.

Parameter	Description
<code>queue-profile <QUEUE-NAME></code>	Specifies the name of the queue profile to apply. Range: 1 to 64 alphanumeric characters, including period (.), underscore (_), and hyphen (-). This parameter is not supported in the <code>config-if</code> context.
<code>schedule-profile <SCHEDULE-NAME></code>	Specifies the name of the schedule profile to apply. Range: 1 to 64 alphanumeric characters, including period (.), underscore (_), and hyphen (-).

Usage

- The switch must always have a globally-applied queue and schedule profile. To stop using a given profile, apply a different profile.
- For a queue profile to be complete and ready to be applied, all eight local priorities must be mapped to a queue.
- For a schedule profile to be complete and ready to be applied, it must define all queues specified in the queue profile. All queues must use the same algorithm, except for the highest numbered queue, which can be **strict**.
- Both the queue profile and the schedule profile must specify the same number of queues.
- Schedule profiles can be modified while applied, but only in ways where a single command will not result in the profile becoming invalid. For example, queue 7 can have the algorithm changed, and weighted queues can have their weights changed.

- Queues must be consecutively defined starting at queue number zero. For example, a four-queue profile with priority values defined for queues 0, 1, 2, 3 is valid, but a four-queue profile which defines priority values for queues 1, 3, 5, and 7 is not.

If the number of queues was changed from the previous queue profile to the new one, any Ethernet or LAG interfaces with locally applied schedule profiles will program the newly applied global schedule-profile. The *show running-config interface* command will list the existing *apply qos schedule-profile* command with a comment describing the actual profile applied:

```
apply qos schedule-profile Old_Schedule
!actual schedule-profile New_Schedule
```

Examples

The following commands illustrate a valid configuration where every local priority value is assigned to a queue and all assigned queues are defined:

```
switch(config)# qos cos-map 1 local-priority 1
switch(config)# qos queue-profile Q1
switch(config)# map queue 0 local-priority 0
switch(config)# map queue 1 local-priority 1
switch(config)# map queue 2 local-priority 2
switch(config)# map queue 3 local-priority 3
switch(config)# map queue 4 local-priority 4
switch(config)# map queue 5 local-priority 5
switch(config)# map queue 6 local-priority 6
switch(config)# map queue 7 local-priority 7
switch(config)# qos schedule-profile S1
switch(config)# dwrr queue 0 weight 5
switch(config)# dwrr queue 1 weight 10
switch(config)# dwrr queue 2 weight 15
switch(config)# dwrr queue 3 weight 20
switch(config)# dwrr queue 4 weight 25
switch(config)# dwrr queue 5 weight 50
```

The following commands illustrate an invalid configuration because local priority 2 is not assigned to a queue:

```
switch(config)# qos cos-map 1 local-priority 1
switch(config)# qos queue-profile Q1
switch(config)# map queue 0 local-priority 0
switch(config)# map queue 1 local-priority 1
switch(config)# map queue 3 local-priority 3
switch(config)# map queue 4 local-priority 4
switch(config)# map queue 5 local-priority 5
switch(config)# map queue 5 local-priority 6
switch(config)# map queue 5 local-priority 7
switch(config)# qos schedule-profile S1
switch(config)# dwrr queue 0 weight 5
switch(config)# dwrr queue 1 weight 10
switch(config)# dwrr queue 3 weight 15
switch(config)# dwrr queue 4 weight 25
switch(config)# dwrr queue 5 weight 50
```

Applying the QoS profile **Q1** and the schedule profile **S1** to all interfaces that do not have an applied interface-specific schedule profile:

```
switch(config)# apply qos queue-profile Q1 schedule-profile S1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config config-if config-lag-if	Administrators or local user group members with execution rights for this command.

dwrr queue

```
dwrr queue <QUEUE-NUMBER> weight <WEIGHT>  
no dwrr queue <QUEUE-NUMBER>
```

Description

Assigns the deficit weighted round robin (DWRR) algorithm and its weight to a queue in a schedule profile. DWRR allocates available bandwidth among all non-empty queues in relation to the queue weights.

The `no` form of this command removes the DWRR algorithm from a queue in a schedule profile.

Parameter	Description
<QUEUE-NUMBER>	Specifies the queue number. Range: 0 to 7.
weight <WEIGHT>	Specifies the scheduling weight. Range: 1 to 1023.

Examples

Assigning DWRR with a weight of **17** to queue **2** in the schedule profile **myschedule**:

```
switch(config)# qos schedule-profile myschedule  
switch(config-schedule)# dwrr queue 2 weight 17
```

Deleting DWRR for queue **2** from the schedule profile **myschedule**:

```
switch(config)# qos schedule-profile myschedule  
switch(config-schedule)# no dwrr queue 2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6200 6300 6400	config-schedule-<NAME>	Administrators or local user group members with execution rights for this command.

map queue

```
map queue <QUEUE-NUMBER> local-priority <PRIORITY-NUMBER>
no map queue <QUEUE-NUMBER> [local-priority <PRIORITY-NUMBER>]
```

Description

Assigns a local priority to a queue in a queue profile. By default, the larger the queue number the higher its priority. A queue without a local priority value assigned to it is not used to store packets. The same queue can be assigned multiple local priorities.

The `no` form of this command removes the specified local priority from a specific queue. If no local priority number is specified, then all local priorities are removed from the queue.

Parameter	Description
<QUEUE-NUMBER>	Specifies the queue number. Range: 0 to 7.
<PRIORITY-NUMBER>	Specifies the local priority. Range: 0 to 7, where 0 is the lowest priority and 7 is the highest.

Usage

The following commands illustrate a valid configuration, where every local priority value is assigned to a queue:

```
map queue 0 local-priority 0
map queue 1 local-priority 1
map queue 1 local-priority 2
map queue 3 local-priority 3
map queue 4 local-priority 4
map queue 5 local-priority 5
map queue 5 local-priority 6
map queue 5 local-priority 7
```

The following commands illustrate an invalid configuration, because local priority 2 is not assigned to a queue:

```
map queue 0 local-priority 0
map queue 1 local-priority 1
map queue 2 local-priority 3
map queue 3 local-priority 4
map queue 4 local-priority 5
map queue 5 local-priority 6
map queue 5 local-priority 7
```

Examples

Assigning priority 7 to queue 7 in profile **myprofile**:

```
switch(config)# qos queue-profile myprofile
switch(config-queue)# map queue 7 local-priority 7
```

Removing priority 7 from queue 7 in profile **myprofile**:

```
switch(config)# qos queue-profile myprofile
switch(config-queue)# no map queue 7 local-priority 7
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config-queue	Administrators or local user group members with execution rights for this command.

min-bandwidth



The min-bandwidth command only applies to Aruba 6300 Series Switches.

```
min-bandwidth queue <QUEUE-NUMBER> percent <VALUE>
no min-bandwidth queue <QUEUE-NUMBER>
```

Description

Assigns the Guaranteed Minimum Bandwidth (GMB) algorithm and a percentage of bandwidth to a queue. GMB allocates available bandwidth among all non-empty queues in relation to their configured minimum bandwidth. Non-empty queues are serviced first in strict order up to their minimum bandwidth. If there is any remaining bandwidth, the scheduler will strictly service any remaining non-empty queues.

The `no` form of this command only clears the algorithm for a queue if GMB has been assigned.

Parameter	Description
<QUEUE-NUMBER>	Specifies the queue number. Range: 0 to 7.
<VALUE>	Specifies bandwidth percentage used for GMB scheduling. Range: 0 to 100.

Examples

Assigning queue 0 of schedule profile S1 the GMB scheduling algorithm with minimum bandwidth of 5 percent:

```
switch(config)# qos schedule-profile S1
switch(config-schedule)# min-bandwidth queue 0 percent 5
```

Removing GMB from queue 0:

```
switch(config)# qos schedule-profile s1
switch(config-schedule)# no min-bandwidth queue 0
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6200 6300 6400	config-schedule- <i><NAME></i>	Administrators or local user group members with execution rights for this command.

name queue

```
name queue <QUEUE-NUMBER> <DESCRIPTION>
no name queue <QUEUE-NUMBER>
```

Description

Assigns a description to a queue in a queue profile. This is for identification purposes and has no effect on configuration.

The `no` form of this command removes the description associated with a queue.

Parameter	Description
<i><QUEUE-NUMBER></i>	Specifies the queue number. Range: 0 to 7.
<i><DESCRIPTION></i>	Specifies a queue description for identification purposes. Range: 1 to 64 alphanumeric characters, including period (.), underscore (_), and hyphen (-).

Examples

Assigning the description **priority-traffic** to queue 7:

```
switch(config)# qos queue-profile myprofile
switch(config-queue)# name queue 7 priority-traffic
```

Removing the description from queue 7:

```
switch(config)# qos queue-profile myprofile
switch(config-queue)# no name queue 7
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config-queue	Administrators or local user group members with execution rights for this command.

qos cos

```
qos cos <CODE-POINT>
no qos cos
```

Description

Configures a CoS PCP remark for an Ethernet or LAG interface. Packets that ingress on the interface are remarked at egress using the configured CoS PCP value.

The remark only occurs when QoS trust mode on the interface is set to `none`.

If QoS trust mode is not set to `none`, then the remark is ignored, and the following commands will show the CoS remark status as `ignored (incompatible Port Access Trust configuration)` OR `not applied' (incompatible QoS global/port Trust configuration)`:

- `show running-configuration`
- `show interface <PORT-NUM>`
- `show interface <PORT-NUM> qos`

The `no` form of this command removes a CoS remark on an interface.

Parameter	Description
<CODE-POINT>	Specifies an 802.1 VLAN priority CoS value. Range: 0 to 7.

Examples

On the 6400 Switch Series, interface identification differs.

Configuring a CoS remark of **3** on interface **1/1/1**:

```
switch(config)# interface 1/1/1
switch(config-if)# qos trust none
switch(config-if)# qos cos 3
```

Deleting a CoS remark of **3** on interface **1/1/1**:

```
switch(config)# interface 1/1/1
switch(config-if)# no qos cos
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6200 6300 6400	config-if	Administrators or local user group members with execution rights for this command.

qos cos-map

```
qos cos-map <CODE-POINT> local-priority <PRIORITY-NUMBER> [color <COLOR>] [name <DESCRIPTION>]
no qos cos-map <CODE-POINT>
```

Description

Defines the local priority assigned to incoming packets for a specific 802.1 VLAN priority code point (CoS) value. The CoS map values are used to mark incoming packets when QoS trust mode is set to **cos**. In **trust none** mode, CoS map entry 0 is used to set the port default local priority and color.

To see the default CoS map settings, use the following command:

```
switch# show qos cos-map default
code_point local_priority color name
-----
0          1             green Best_Effort
1          0             green Background
2          2             green Excellent_Effort
3          3             green Critical_Applications
4          4             green Video
5          5             green Voice
6          6             green Internetwork_Control
7          7             green Network_Control
```

The **no** form of this command restores the assignments for a CoS map value to the default setting.

Parameter	Description
<CODE-POINT>	Specifies an 802.1 VLAN priority CoS value. Range: 0 to 7. Default 0.
local-priority <PRIORITY-NUMBER>	Specifies a local priority value to associate with the CODE-POINT value. Range: 0 to 7. Default: 0.

Parameter	Description
color <COLOR>	Reserved for future use.
name <DESCRIPTION>	Specifies a description for the CoS setting. The name is for identification only, and has no effect on queue configuration. Range: 1 to 64 alphanumeric characters, including period (.), underscore (_), and hyphen (-).

Usage

Examples

Mapping CoS value **1** to a local priority of **2**:

```
switch(config)# qos cos-map 1 local-priority 2
```

Mapping CoS value **1** to the default local priority value:

```
switch(config)# no qos cos-map 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6200 6300 6400	config	Administrators or local user group members with execution rights for this command.

qos dscp

```
qos dscp <CODE-POINT>
no qos dscp
```

Description

Configures a differentiated services code point (DSCP) remark for an Ethernet or LAG interface. IPV4 and IPV6 packets that ingress on the interface are remarked at egress using the configured DSCP value.

The remark only occurs when QoS trust mode on the interface is set to `none`. If a DSCP remark is configured and then trust mode is subsequently set to `cos` or `dscp`, then the DSCP remark is ignored.

The following commands will show the remark status as *ignored* (incompatible Port Access Trust configuration) or *not applied* (incompatible QoS global or port trust configuration):

- `show running-configuration`
- `show interface <INTERFACE-NAME>`
- `show interface <INTERFACE-NAME> qos`

The `no` form of this command removes a CoS remark on an interface.

Parameter	Description
<code><CODE-POINT></code>	Specifies an IP differentiated services code point value. Range: 0 to 63.

Usage

Order of operation for arriving IPv4 or IPv6 packets:

1. Trust none is applied with initial local-priority and color metadata assigned from the CoS Map entry index 0.
2. The local-priority value and the queue profile are then used to determine the queue for the packet.
3. The remark of the packet's DSCP metadata field is performed. When the packet is transmitted, its IPv4 or IPv6 DS header is remarked with the DSCP metadata.

For arriving non-IP packets:

Trust none is applied with initial local-priority and color metadata assigned from the CoS Map entry index 0. This selects the queue for packet scheduling. The PCP of any tagged non-IP packets is unchanged.

Examples

On the 6400 Switch Series, interface identification differs.

Configuring a DSCP remark of **43** on interface **1/1/1**:

```
switch(config)# interface 1/1/1
switch(config-if)# qos trust none
switch(config-if)# qos dscp 43
```

Deleting a DSCP remark of **43** on interface **1/1/1**:

```
switch(config)# interface 1/1/1
switch(config-if)# no dscp 43
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6200 6300 6400	config-if config-lag-if	Administrators or local user group members with execution rights for this command.

qos dscp-map

```
qos dscp-map <CODE-POINT> local-priority <PRIORITY-NUMBER> [color <COLOR>] [cos <PCP-VALUE>]
[name <DESCRIPTION>]
no qos dscp-map <CODE-POINT>
```

Description

Defines the local priority assigned to incoming packets for a specific IP differentiated services code point (DSCP) value. The DSCP map values are used to prioritize incoming packets when QoS trust mode is set to **dscp**.

The **no** form of this command restores the assignments for a code point to the default setting.

Use `show qos dscp-map` to view the current settings. To see the default DSCP map settings, use the following command:

```
switch# show qos dscp-map default
code_point local_priority cos color name
-----
0          1              green CS0
1          1              green
2          1              green
3          1              green
4          1              green
5          1              green
...
45         5              green
46         5              green EF
47         5              green
48         6              green CS6
...
61         7              green
62         7              green
63         7              green
```

Parameter	Description
<CODE-POINT>	Specifies an IP differentiated services code point. Range: 0 to 63. Default: 0.
local-priority <PRIORITY-NUMBER>	Specifies a local priority value to associate with the CODE-POINT value. Range: 0 to 7. Default: 0.
color <COLOR>	Reserved for future use
cos <PCP-VALUE>	Specifies an optional 802.1p VLAN Priority Code Point remark value. Range: 0 to 7. Default: No remark.
name <DESCRIPTION>	Specifies a description for the DSCP setting. The name is used for identification only, and has no effect on queue configuration. Range: 1 to 64 alphanumeric characters, including period (.), underscore (_), and hyphen (-).

Examples

Setting code point **1** to a local priority of **2** and a CoS of **0**:

```
switch(config)# qos dscp-map 1 local-priority 2 cos 0
```

Setting code point 1 to the default value:

```
switch(config)# no qos dscp-map 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config	Administrators or local user group members with execution rights for this command.

qos queue-profile

```
qos queue-profile <NAME>  
no qos queue-profile <NAME>
```

Description

Creates a new QoS queue profile and switches to the `config-queue` context for the profile. Or, if the specified QoS queue profile exists, this command switches to the `config-queue` context for the profile. A queue profile maps queues to local-priority values. Each profile has one to eight queues numbered 0 to 7. The larger the queue number, the higher its priority during transmission scheduling.

A queue profile named **factory-default** is defined by default and is automatically applied to all interfaces. It cannot be edited or deleted. Use the command `show qos queue-profile factory-default` to view this profile. Do not use `show running-configuration`, as it will only display changes from the initial values. Use `show qos queue-profile [NAME]` to view the status of all queue profiles or the settings of the named profile.

The `no` form of this command removes the specified QoS queue profile. Only profiles that are not currently applied can be removed.

Parameter	Description
<NAME>	Specifies the name of the QoS queue profile to create or configure. Range: 1 to 64 alphanumeric characters, including period (.), underscore (_), and hyphen (-).

Examples

Creating the profile **myprofile**:

```
switch(config)# qos queue-profile myprofile  
switch(config-queue)#
```

Deleting the profile **myprofile**:

```
switch(config)# no qos queue-profile myprofile
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config	Administrators or local user group members with execution rights for this command.

qos schedule-profile

```
qos schedule-profile <NAME>  
no qos schedule-profile <NAME>
```

Description

Creates a QoS schedule profile and switches to the `config-schedule` context for the profile. If the specified schedule profile exists, this command switches to the `config-schedule` context for the profile. The schedule profile determines the order in which queues are selected to transmit a packet, and the amount of service defined for each queue.

Parameter	Description
<NAME>	Specifies the name of the QoS queue profile to create or configure. Range: 1 to 64 alphanumeric characters, including period (.), underscore (_), and hyphen (-).

Usage

Queues in a schedule profile are numbered consecutively starting from zero. Queue zero is the lowest priority queue. The larger the queue number, the higher priority the queue has in scheduling algorithms.

A profile named **factory-default** is defined by default and applied to all interfaces. It cannot be edited or deleted. To see its settings, use the command:

```
switch# show qos schedule-profile factory-default  
queue_num algorithm weight  
-----  
0          dwrr        1  
1          dwrr        1  
2          dwrr        1  
3          dwrr        1  
4          dwrr        1  
5          dwrr        1  
6          dwrr        1  
7          dwrr        1
```

A profile named **strict** is predefined and cannot be edited or deleted. The strict profile services all queues of the queue profile to which it is applied, using the strict priority algorithm.

A schedule profile must be defined on all interfaces at all times.

There are two permitted configurations for a schedule profile:

1. All queues use the same scheduling algorithm (for example, DWRR).
2. The highest queue number uses strict priority, and all remaining (lower) queues use the same algorithm (for example, DWRR). This supports priority scheduling behavior necessary for the IETF RFC 3246 Expedited Forwarding specification (<https://tools.ietf.org/html/rfc3246>).

Only limited changes can be made to an applied schedule profile:

- The weight of a dwrr queue.
- The bandwidth of a strict queue.
- The algorithm of the highest numbered queue can be swapped between dwrr and strict, and vice versa.

Applicable to REST: Any other changes will result in an unusable schedule profile, and the switch will revert to the `factory-default` profile until the profile is corrected.

The `no` form of this command removes the specified QoS schedule profile when it is not applied. Only profiles that are not currently applied to an interface can be removed.

Examples

Creating the schedule profile **myschedule**:

```
switch(config)# qos schedule-profile myschedule
switch(config-schedule)#
```

Deleting the schedule profile **myschedule**:

```
switch(config)# no qos schedule-profile myschedule
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	<code>config</code>	Administrators or local user group members with execution rights for this command.

qos shape

```
qos shape <RATE>
no qos shape
```

Description

Limits the egress bandwidth on an interface to a value that is lower than its line rate.

The `no` form of this command removes shaping from an interface.

Parameter	Description
<RATE>	Specifies the maximum traffic rate in kbps. Range 64 to 100000000.

Usage

When the traffic rate destined for the port exceeds the configured egress bandwidth, the switch will buffer the excess up to the limit of the queues. Rates larger than the interface line rate will have no effect. When set on a LAG, each member Ethernet port independently shapes its egress bandwidth to the specified rate.

Examples

On the 6400 Switch Series, interface identification differs.

Configuring an egress port shaping rate of 400 Mbps on interface **1/1/1**:

```
switch(config)# interface 1/1/1
switch(config-if)# qos shape 400000
```

Deleting egress port shaping on interface **1/1/1**:

```
switch(config)# interface 1/1/1
switch(config-if)# no qos shape
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6200 6300 6400	config-if	Administrators or local user group members with execution rights for this command.

qos trust

```
qos trust {none | cos | dscp}
no qos trust
```

Description

In the `config` context:

- This command sets the trust mode that is globally applied to all interfaces that do not have a trust mode configured.

- The `no` form of this command restores all interfaces that do not currently have a trust mode configured to the default setting.

In the `config-if` context:

- This command sets the trust mode override for a specific interface.
- The `no` form of this command clears a trust mode override. The interface then uses the global setting. This is the only way to remove a trust mode override.

Parameter	Description
<code>none</code>	Ignores all packet headers. Ingress packets are assigned the local priority and color values configured for CoS map entry 0. Default.
<code>cos</code>	For 802.1 VLAN tagged packets, use the priority code point field from the outermost VLAN header as the index into the CoS map to obtain the local priority and color values for the packet. If the packet is untagged, use the local priority and color values configured for CoS map entry 0.
<code>dscp</code>	For IP packets, use the DSCP field from the IP header as the index into the DSCP Map. For non-IP packets with 802.1 VLAN tags, use the priority code point field from the outermost VLAN header as the index into the CoS map to obtain the local priority and color values for the packet. For untagged, non-IP packets, use the local priority and color values configured for CoS map entry 0.

Example

Setting the global trust mode to **dscp**, which is applied to all interfaces that do not already have an individual trust mode configured. An override is then applied to interface **2/2/2**, and LAG 100, setting trust mode to **cos**:

```
switch(config)# qos trust dscp
switch(config)# interface 2/2/2
switch(config-if)# qos trust cos
switch(config-if)# interface lag 100
switch(config-if)# qos trust cos
```



WARNING: QoS port remark configurations are not applied when the QoS trust mode is *mode*. This warning message is seen if a port trust command other than *trust none* is attempted when there is already a remark configuration on the port. To restore the old remark configuration, configure the port trust mode to *none*.



WARNING: QoS port remark configurations are not applied when the global QoS trust mode is *mode*. This warning message is seen if a port *no qos trust* command is attempted when there is already a remark configuration on the port and the global trust mode is not *none*. To re-apply the remark configuration, set the port trust mode to *none*.

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config config-if config-lag-if	Administrators or local user group members with execution rights for this command.

rate-limit

```
rate-limit {broadcast | multicast | unknown-unicast | icmp {ip-all | ipv4 | ipv6}} <RATE>
{kpbs | pps}
no rate-limit {broadcast | multicast | icmp}
```

Description

Sets the amount of traffic of a specific type that can ingress on an Ethernet interface, or on each port of a LAG interface. Rate limits are enforced separately on each individual member of a LAG, not on the LAG as a whole.

The `no` form of this command removes the traffic limit for the specified traffic type.

Parameter	Description
{broadcast multicast icmp {ip-all ipv4 ipv6}}	Specifies the type of ingress traffic to which the rate limit applies: broadcast, multicast, or ICMP. The multicast rate limit affects multicast and broadcast traffic. The broadcast rate limit only affects broadcast traffic. When both types are applied to the same interface, broadcast packets are limited to the lower of the two rate values. Layer 2 BPDU packets, like spanning tree, are also included in the multicast rate limit. The ICMP rate limit can be configured to apply to IPv4, IPv6, or all IP traffic. Only one ICMP rate-limit can be configured at a time. Applying a new ICMP rate-limit replaces any previous ICMP rate-limit.
<RATE> {kpbs pps}	Specifies the rate limit. Range: 64 to 100000000 kbps (in steps of 64 kbps), or 64 to 209090910 kbps (in steps of 64 pps). The actual rate limit varies with steps approximately equal to the minimum value. Verify the actual rate limit using the command <code>show interface <INTERFACE-NAME></code> .

Examples

On the 6400 Switch Series, interface identification differs.

Limiting broadcast traffic to **2000pps** on interface **1/1/3**:

```
switch(config)# interface 1/1/3
switch(config-if)# rate-limit broadcast 500 kbps
```

Limiting all ICMP IPv4 traffic to **10000kbps** on interface **1/1/3**:

```
switch(config)# interface 1/1/3
switch(config-if)# rate-limit icmp ipv4 10000 kbps
```

Viewing the results of the previous configuration settings:

```
switch# show interface 1/1/3 qos
Interface 1/1/3 is up
Admin state is up
Description:
Hardware: Ethernet, MAC Address: 08:97:34:b1:20:00
MTU 1500
Type 1000BT
qos trust none
rate-limit broadcast 2000 pps
(2000 actual)
rate-limit icmp ip-all 10000 kbps
(10000 actual)
Speed 1000 Mb/s
L3 Counters: Rx Disabled, Tx Disabled
Auto-Negotiation is on
Flow-control: off
Rx
    0 input packets          0 bytes
    0 input error           0 dropped
    0 CRC/FCS
L3:
    0 packets, 0 bytes
Tx
    127 output packets      16510 bytes
    0 input error           0 dropped
    0 collision
L3:
    0 packets, 0 bytes
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6200 6300 6400	config-if	Administrators or local user group members with execution rights for this command.

show interface queues

show interface <INTERFACE-NAME> queues

Description

Displays interface-level queue statistics.

Parameter	Description
<INTERFACE-NAME>	Specifies the name of an Ethernet port or LAG on the switch. Format: member/slot/port or lag number.

Usage

Statistics include:

- **Tx Bytes:** Total bytes transmitted. The byte count may include packet headers and internal metadata that are removed before the packet is transmitted. Packet headers added when the packet is transmitted may not be included. The byte count includes any packets subsequently dropped by an egress ACL (6300 Series Switch only).
- **Tx Packets:** Total packets transmitted. The count includes packets subsequently dropped by an egress ACL (6300 Series Switch only).
- **Tx Drops:** Total packets dropped by an egress queue due to insufficient capacity.
- **Tx Drops:** For the 6300 Series Switch total packets dropped by an egress queue due to insufficient capacity. For the 6400 Series Switch sum of packets that were dropped across all line modules by Virtual Output Queues (VOQs) destined for the egress port queue due to insufficient capacity. As the counts are read separately from each line module, the sum is not an instantaneous snapshot.

Examples

On the 6400 Switch Series, interface identification differs.

Showing queue statistics for interface **1/1/5**:

```
switch# show interface 1/1/5 queues
Interface 1/1/5 is down
Admin state is up
      Tx Bytes      Tx Packets      Tx Drops
Q0                0                0                3
Q1             15356                73                0
Q2                0                0                0
Q3                0                0                0
Q4                0                0                0
Q5                0                0                0
Q6                0                0                0
Q7                0                0                0
```

Showing queue statistics for interface **lag 1**:

```
switch# show interface lag 1 queues
Aggregate-name lag1
Aggregated-interfaces :
1/1/6 1/1/7
Speed 20000 Mb/s
      Tx Bytes      Tx Packets      Tx Drops
```

Q0	0	0	0
Q1	0	0	0
Q2	0	0	0
Q3	0	0	0
Q4	0	0	0
Q5	0	0	0
Q6	0	0	0
Q7	3450	25	0

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show interface qos

show interface <INTERFACE-NAME> qos

Description

Shows various QoS settings for a specific interface.

Parameter	Description
<INTERFACE-NAME>	Specifies the name of an interface on the switch. Format: member/slot/port or lag number.

Examples

On the 6400 Switch Series, interface identification differs.

Showing QoS settings for interface **1/1/5**:

```
switch# show interface 1/1/5 qos
Interface 1/1/5 is up
Admin state is up
qos trust none (global)
qos queue-profile factory-default (global)
qos schedule-profile factory-default (global)
qos cos 5
qos dscp 47
rate-limit broadcast 40000 kbps (40000 actual)
rate-limit icmp ip-all 10000 kbps (10000 actual)
Forwarded Pkts          Dropped Pkts          Forwarded Bytes          Dropped Bytes
Broadcast:              944468                  1044                    1044658890
      85662408
```

```
ICMP:                82210                0                2689008
                    0
qos shape 200000 (199999 actual)
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show qos cos-map

```
show qos cos-map [default] [vsx-peer]
```



NOTE: The `vsx-peer` parameter is not supported by the 6300 Series Switch

Description

Shows the global QoS CoS code point settings, or the factory default settings.

Parameter	Description
default	Shows the factory default CoS code point settings.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

Showing the current CoS map:

```
switch# show qos cos-map
code_point local_priority color name
-----
0          2             green Best_Effort
1          0             green Background
2          1             green Spare
3          3             green Excellent_Effort
4          4             green Controlled_Load
5          5             green Video
6          6             green Voice
7          7             green Network_Control
```

Showing the default CoS map:

```

switch# show qos cos-map default
code_point local_priority color name
-----
0          1              green Best_Effort
1          0              green Background
2          2              green Excellent_Effort
3          3              green Critical_Applications
4          4              green Video
5          5              green Voice
6          6              green Intersetwork_Control
7          7              green Network_Control

```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6200 6300 6400	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show qos dscp-map

```
show qos dscp-map [default] [vsx-peer]
```



NOTE: The `vsx-peer` parameter is not supported by the 6300 Series Switch

Description

Displays the current or default global QoS dscp-map.

Parameter	Description
default	Shows the factory default DSCP code point settings.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

Showing the current QoS DSCP map:

```

switch# show qos dscp-map
code_point local_priority cos color name

```

```

-----
0      1      green  CS0
1      1      green
2      1      3    green
3      1      green
4      1      green
5      1      green
6      1      green
7      1      green
8      0      green  CS1
...
45     5      green
46     7      green  EF
47     5      7    green
48     6      green  CS6
...
61     7      green
62     7      green
63     7      green

```

Showing the default QoS DSCP map:

```

switch# show qos dscp-map default
code_point local_priority cos color name
-----
0      1      green  CS0
1      1      green
2      1      green
3      1      green
4      1      green
5      1      green
...
45     5      green
46     5      green  EF
47     5      green
48     6      green  CS6
...
61     7      green
62     7      green
63     7      green

```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show qos queue-profile

show qos queue-profile [<NAME> | factory-default] [vsx-peer]



NOTE: The `vsx-peer` parameter is not supported by the 6300 Series Switch

Description

Shows the status of all queue profiles, or a specific queue profile.

Parameter	Description
<NAME>	Specifies the name of a queue profile. Range 1 to 64 alphanumeric characters, including period (.), underscore (_), and hyphen (-).
[factory-default]	Specifies the factory default queue profile.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Usage

The status of a queue profile can be:

- Applied - The profile is actively being used by the switch.
- Complete - The profile meets the criteria to be applied.
- Incomplete - The profile does not meet the criteria to be applied.

For a queue profile to be complete and ready to be applied:

- All eight local priorities must be mapped to some queue.
- There can be 1 to 8 queues.
- The queues must be consecutively numbered starting at zero.

Examples

Showing the settings of the factory default queue profile:

```
switch# show qos queue-profile factory-default
queue_num local_priorities name
-----
0          0
1          1
2          2
3          3
4          4
5          5
6          6
7          7
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show qos schedule-profile

```
show qos schedule-profile [<NAME> | factory-default | strict] [vsx-peer]
```



NOTE: The `vsx-peer` parameter is not supported by the 6300 Series Switch

Description

Shows the status of all schedule profiles, or a specific schedule profile.

Parameter	Description
<NAME>	Specifies the name of a queue or schedule profile. Range: 1 to 64 alphanumeric characters, including period (.), underscore (_), and hyphen (-).
[factory-default]	Specifies the factory default queue profile.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Usage

The status of a schedule profile can be:

- Applied - The profile is actively being used by one or more ports.
- Complete - The profile meets the criteria to be applied.
- Incomplete - The profile does not meet the criteria to be applied.

For a schedule profile to be complete and ready to be applied it must have:

- An algorithm for each queue defined by the applied queue profile.
- All queues must use the same algorithm except for the highest numbered queue, which may be strict.

Example

Showing the status of all schedule profiles:

```
switch# show qos schedule-profile
profile_status profile_name
-----
applied        myschedule
complete       factory-default
complete       Test
```

Showing the configuration of factory default schedule profile:

```
switch# show qos schedule-profile factory-default
queue_num algorithm weight
-----
0          dwrr        1
1          dwrr        1
2          dwrr        1
3          dwrr        1
4          dwrr        1
5          dwrr        1
6          dwrr        1
7          dwrr        1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show qos trust

```
show qos trust [default] [vsx-peer]
```



NOTE: The `vsx-peer` parameter is not supported by the 6300 Series Switch

Description

Shows the global QoS trust settings, or the factory default settings.

Parameter	Description
default	Shows the factory default QoS trust settings.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

Showing the current QoS trust settings:

```
switch# show qos trust
qos trust cos
```

Showing the default QoS trust settings:

```
switch# show qos trust default
qos trust none
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

strict queue

```
strict queue <QUEUE-NUMBER> [[max-bandwidth <BANDWIDTH>]]
no strict queue <QUEUE-NUMBER>
```

Description

Assigns the strict priority algorithm to a queue. Strict priority services all packets waiting in a queue, before servicing the packets in lower priority queues.

Egress queue shaping can be configured using the `max-bandwidth` option to limit the amount of traffic transmitted per output queue. The buffer associated with each egress queue stores the excess traffic to smooth the output rate. Sustained rates of traffic above the maximum bandwidth will eventually fill the output queue causing tail drops. Use the command `show interface` to determine if any tail drop errors have occurred.

The `no` form of this command removes the queue configuration from the schedule profile. To remove only egress queue shaping, re-enter the `strict queue` command without the `max-bandwidth` parameter.

Parameter	Description
<QUEUE-NUMBER>	Specifies the number of the queue. Range: 0 to 7.
max-bandwidth <BANDWIDTH>	Specifies the maximum bandwidth allowed on the queue in Kbps. Range: 64 to 100000000.

Usage

Either all the queues of the schedule profile can be *strict* or just the highest numbered queue. When applied to a LAG, each member Ethernet port independently schedules its egress transmissions using the strict settings. Only limited changes can be made to a *strict* queue that is part of an applied schedule profile:

- The max-bandwidth settings.
- The highest numbered queue can be swapped between *strict* and *dwrr`* or *min-bandwidth* (only applicable for the Aruba 6300 Series Switch)
- The highest numbered queue can be swapped between *strict* and *dwrr* (only applicable for the Aruba 6400 Series Switch)

Any other changes or removing a queue (no *strict* queue) will result in an unusable schedule profile. If that schedule profile is applied in the interface context, the switch will revert to the schedule profile applied in the global context until the profile is corrected. If that schedule profile is applied in the global context, the switch will revert to using the factory-default profile until the profile is corrected.

Examples

Assigning strict priority to queue **7** in the schedule profile **myschedule**:

```
switch(config)# qos schedule-profile myschedule
switch(config-schedule)# strict queue 7
```

Deleting strict priority from queue **7** in the schedule profile **myschedule**:

```
switch(config)# qos schedule-profile myschedule
switch(config-schedule)# no strict queue 7
```

Assigning strict priority to queue **7** in the schedule profile **myschedule** with a maximum bandwidth of 10000 Kbps:

```
switch(config)# qos schedule-profile myschedule
switch(config-schedule)# strict queue 7 max-bandwidth 10000
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config-schedule- <i><NAME></i>	Administrators or local user group members with execution rights for this command.

Accessing Aruba Support

Aruba Support Services	https://www.arubanetworks.com/support-services/
Aruba Support Portal	https://asp.arubanetworks.com/
North America telephone	1-800-943-4526 (US & Canada Toll-Free Number) +1-408-754-1200 (Primary - Toll Number) +1-650-385-6582 (Backup - Toll Number - Use only when all other numbers are not working)
International telephone	https://www.arubanetworks.com/support-services/contact-support/

Be sure to collect the following information before contacting Support:

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Other useful sites

Other websites that can be used to find information:

Airheads social forums and Knowledge Base	https://community.arubanetworks.com/
AOS-CX Switch Software Documentation Portal	https://www.arubanetworks.com/techdocs/AOS-CX/help_portal/Content/home.htm
Aruba Hardware Documentation and Translations	https://www.arubanetworks.com/techdocs/hardware/DocumentationPortal/Content/home.htm

Portal	
Aruba software	https://asp.arubanetworks.com/downloads
Software licensing	https://lms.arubanetworks.com/
End-of-Life information	https://www.arubanetworks.com/support-services/end-of-life/
Aruba Developer Hub	https://developer.arubanetworks.com/

Accessing Updates

You can access updates from the Aruba Support Portal or the HPE My Networking Website.

Aruba Support Portal

<https://asp.arubanetworks.com/downloads>

If you are unable to find your product in the Aruba Support Portal, you may need to search My Networking, where older networking products can be found:

My Networking

<https://www.hpe.com/networking/support>

To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

<https://support.hpe.com/portal/site/hpsc/aae/home/>

Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HP Passport set up with relevant entitlements.

Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.

To subscribe to eNewsletters and alerts:

<https://asp.arubanetworks.com/notifications/subscriptions> (requires an active Aruba Support Portal (ASP) account to manage subscriptions). Security notices are viewable without an ASP account.

Warranty Information

To view warranty information for your product, go to <https://www.arubanetworks.com/support-services/product-warranties/>.

Regulatory Information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at <https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts>

Additional regulatory information

Aruba is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements, environmental data (company programs, product recycling, energy efficiency), and safety information and compliance data, (RoHS and WEEE). For more information, see <https://www.arubanetworks.com/company/about-us/environmental-citizenship/>.

Documentation Feedback

Aruba is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (docsfeedback-switching@hpe.com). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.