

AOS-CX 10.09 Quality of Service Guide

4100i, 6000, 6100 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 |
| 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 | 14 |
| Egress queue shaping | 14 |
| Terms | 15 |
| QoS configuration | 16 |
| Configuring QoS | 16 |
| Configuring expedited forwarding for VoIP traffic | 17 |
| Configuring rate limiting | 19 |
| Configuring egress queue shaping | 20 |
| Supporting Ethernet 802.1D Class of Service | 20 |
| Monitoring queue operation | 21 |
| QoS commands | 23 |
| apply qos | 23 |
| map queue | 25 |
| min-bandwidth | 26 |
| name queue | 27 |
| qos cos | 28 |
| qos dscp | 29 |
| qos dscp-map | 30 |
| qos queue-profile | 32 |
| qos schedule-profile | 32 |
| qos trust | 34 |
| rate-limit | 35 |
| show interface queues | 38 |
| show interface qos | 39 |
| show qos dscp-map | 40 |
| show qos queue-profile | 41 |
| show qos schedule-profile | 43 |

| | |
|----------------------|----|
| show qos trust | 44 |
| strict queue | 45 |

Support and Other Resources 47

| | |
|-------------------------------|----|
| Accessing Aruba Support | 47 |
| Accessing Updates | 48 |
| Aruba Support Portal | 48 |
| My Networking | 48 |
| Warranty Information | 48 |
| Regulatory Information | 48 |
| Documentation Feedback | 49 |

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 4100i Switch Series (JL817A, JL818A)
- Aruba 6000 Switch Series (R8N85A, R8N86A, R8N87A, R8N88A, R8N89A)
- Aruba 6100 Switch Series (JL675A, JL676A, JL677A, JL678A, JL679A)

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>■ <code><example-text></code>■ <i>example-text</i>■ <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. |
| { } | Braces. Indicates that at least one of the enclosed items is required. |

| Convention | Usage |
|---------------|--|
| [] | 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 4100i Switch Series

- *member*: Always 1. VSF is not supported on this switch.
- *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 the switch.

On the 6000 and 6100 Switch Series

- *member*: Always 1. VSF is not supported on this switch.
- *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 the switch.

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.



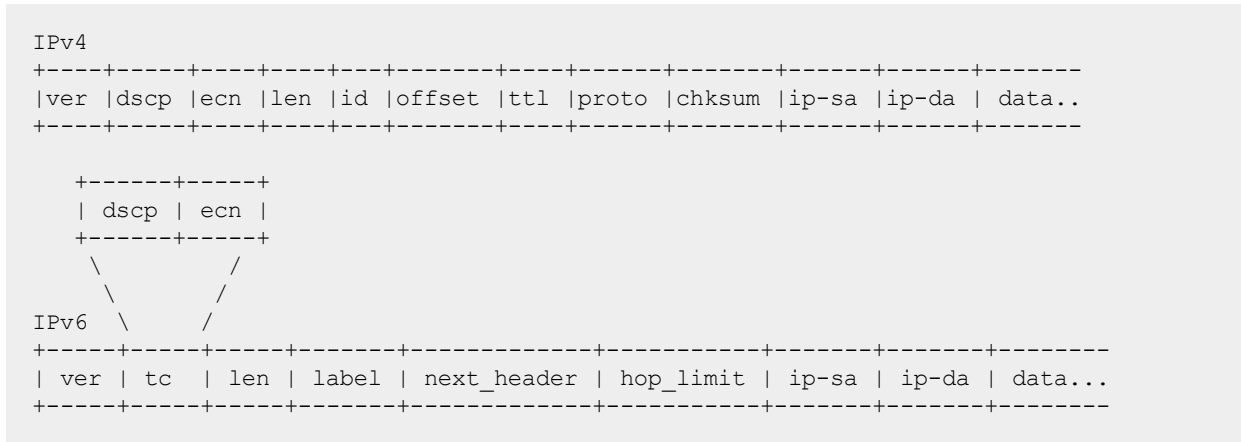
Differentiated services

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 |
|------------|------------------|-------------------------|------|
| 56 | 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.

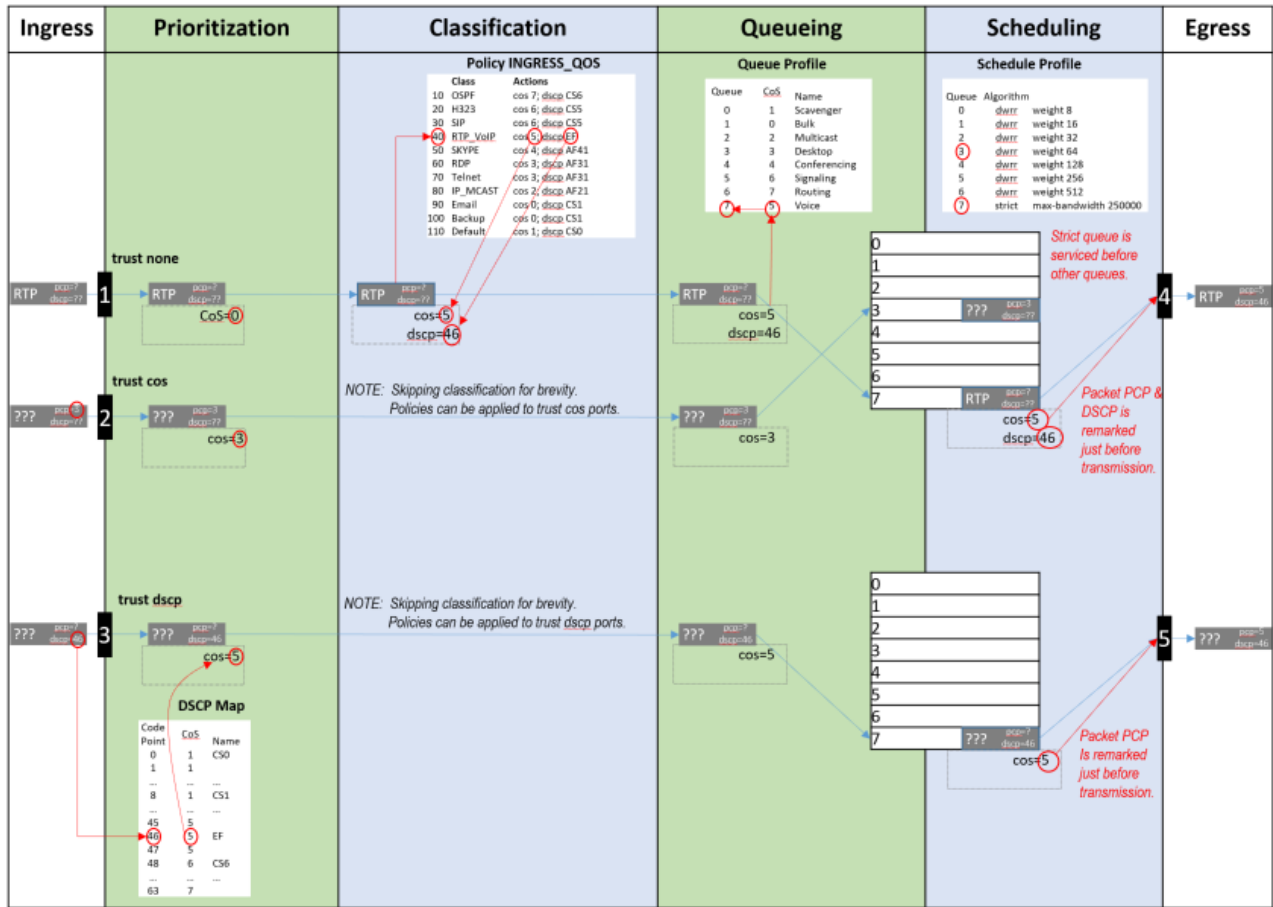


Switches with at least 52 ports will experience negative performance if a flood occurs where at least 42 ports are members of the same VLAN and all 52 ports have QoS rules applied to them.

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

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



QoS trust


Traffic priorities for networks can be carried in VLAN tags using the CoS Priority Code Point (PCP), or in IP packet headers using the Differentiated Services Code Point (DSCP). Whether these priorities affect how traffic is serviced depends on how the QoS trust mode is configured on the switch. The QoS trust mode specifies how the switch assigns values to ingress dscp packets and can be set globally for all interfaces or individually for each interface. By default the trust mode is set to `cos`, meaning the existing QoS information in the packet will be trusted. The priority values of DSCP packets are mapped to corresponding CoS values. If the trust mode is set to `none` the CoS value in the incoming packet will be marked as zero.

For example:

When QoS trust mode = CoS



When QoS trust mode = DSCP



When QoS trust mode = None



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.



Unknown unicast rate limiting is supported.

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. Packets are assigned to the queue based on CoS value. The queue profile defines queue number to CoS mapping. The `factory-default` map cross maps the 0 and 1 values.

```
switch# show qos queue-profile factory-default
queue_num cos          name
-----
0          1          Scavenger_and_backup_data
1          0
2          2
3          3
4          4
5          5
6          6
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`: Maps queue to CoS.
- `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. The switch described in this guide supports the Guaranteed Minimum Bandwidth (GMB), Strict, and Strict EQS scheduling algorithms.

The switch described in this guide supports 8 queues.

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

switch# show qos schedule-profile factory-default
queue_num algorithm      percent max-bandwidth_kbps
-----
0          min-bandwidth 2
1          min-bandwidth 3
2          min-bandwidth 30
3          min-bandwidth 10
4          min-bandwidth 10
5          min-bandwidth 10
6          min-bandwidth 15
7          min-bandwidth 20
```

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.

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.

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.

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 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 for queue mapping, create one or more custom queue profiles with the command `qos queue-profile`. For each queue in a custom queue profile:
 - a. Assign a CoS value mapping with the command `map queue`.
 - b. Optionally, define a descriptive name with the command `name queue`. All queues must be mapped to a CoS value, 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 `min-bandwidth`.
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

The following example creates a queue profile named Q1 and assigns queue to CoS value mapping:

| Queue | CoS Value |
|-------|-----------|
| 0 | 0 |
| 1 | 1 |
| 2 | 2 |
| 3 | 3 |
| 4 | 4 |

| Queue | CoS Value |
|-------|-----------|
| 5 | 5 |
| 6 | 6 |
| 7 | 7 |

Creates a schedule profile named S1 and assigns minimum bandwidth for each queue with a corresponding percentage value:

| Queue | Minimum bandwidth (percentage) |
|-------|--------------------------------|
| 0 | 5 |
| 1 | 5 |
| 2 | 10 |
| 3 | 10 |
| 4 | 20 |
| 5 | 20 |
| 6 | 10 |
| 7 | 20 |

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

```
switch(config)# qos trust cos
switch(config)# qos queue-profile Q1
switch(config-queue)# map queue 0 cos 0
switch(config-queue)# map queue 1 cos 1
switch(config-queue)# map queue 2 cos 2
switch(config-queue)# map queue 3 cos 3
switch(config-queue)# map queue 4 cos 4
switch(config-queue)# map queue 5 cos 5
switch(config-queue)# map queue 6 cos 6
switch(config-queue)# map queue 7 cos 7
switch(config)# qos schedule-profile S1
switch(config-schedule)# min-bandwidth queue 0 percent 5
switch(config-schedule)# min-bandwidth queue 1 percent 5
switch(config-schedule)# min-bandwidth queue 2 percent 10
switch(config-schedule)# min-bandwidth queue 3 percent 10
switch(config-schedule)# min-bandwidth queue 4 percent 20
switch(config-schedule)# min-bandwidth queue 5 percent 20
switch(config-schedule)# min-bandwidth queue 6 percent 10
switch(config-schedule)# min-bandwidth queue 7 percent 20
switch(config-schedule)# 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 CoS value 6. The default DSCP map has eight code points (40 through 47), that are mapped to CoS value 5. To reserve CoS value 6 for VoIP traffic, the other code points must be reassigned. In this scenario, CoS value 6 is used for all reassignments, including for code point 40, Call Signaling protocol (CS5).

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

2. Queue 7 is the highest priority queue, so for best throughput, create a queue profile that maps CoS 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 cos 5
switch(config-queue)# map queue 6 cos 7
switch(config-queue)# map queue 5 cos 6
switch(config-queue)# map queue 4 cos 4
switch(config-queue)# map queue 3 cos 3
switch(config-queue)# map queue 2 cos 2
switch(config-queue)# map queue 1 cos 1
switch(config-queue)# map queue 0 cos 0
```

3. Create a schedule profile that services queue 7 using strict priority (SP), and the remaining queues with minimum bandwidth. This scenario defines queue 7 as strict and gives queues 0-6 an equal minimum bandwidth.

```
switch(config)# qos schedule-profile voip
switch(config-schedule)# strict queue 7
switch(config-schedule)# min-bandwidth queue 6 percent 10
switch(config-schedule)# min-bandwidth queue 5 percent 10
switch(config-schedule)# min-bandwidth queue 4 percent 10
switch(config-schedule)# min-bandwidth queue 3 percent 10
switch(config-schedule)# min-bandwidth queue 2 percent 10
switch(config-schedule)# min-bandwidth queue 1 percent 10
switch(config-schedule)# min-bandwidth queue 0 percent 10
```

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.

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 1000000 kbps
switch(config-if)# rate-limit multicast 2000000 kbps
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 1000000 kbps
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 200000 kbps
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 1000 kbps
switch(config-if)# rate-limit multicast 500 kbps
```

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.

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
```

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 | 5 Video (<100ms latency) |

| CoS 802.1Q | CoS 802.1D |
|--------------------------|--------------------|
| (<10ms latency) | |
| 4 Video (<100ms latency) | 4 Controlled Load |
| 3 Critical Applications | 3 Excellent Effort |
| 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 to queue mapping must be reconfigured to follow the 802.1D standard by swapping the assignments of CoS 0 and 2:

1. Create a new queue profile with the desired queue to CoS mapping changes.
2. Apply the created queue profile, using the `apply qos` command.

```
switch# config
switch(config)# queue-profile newProfile
switch(config-queue)# map queue 0 cos 1
switch(config-queue)# map queue 2 cos 0
switch(config-queue)# map queue 1 cos 2
switch(config-queue)# map queue 3 cos 3
switch(config-queue)# map queue 4 cos 4
switch(config-queue)# map queue 5 cos 5
switch(config-queue)# map queue 6 cos 6
switch(config-queue)# map queue 7 cos 7

switch# show qos queue-profile newProfile
queue_num    cos      name
-----
0            1
1            2
2            0
3            3
4            4
5            5
6            6
7            7
```

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 6000 and 6100 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. 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.

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 CoS values 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.
- There can be only 2, 4, or 8 queues in a queue profile configuration and the queues must be consecutively numbered starting at zero.
- All queues must use the same algorithm except for the highest numbered queue, which may be strict.

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:

Examples

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

```
switch(config)# qos trust cos
switch(config)# qos queue-profile Q1
switch(config)# map queue 0 cos 0
switch(config)# map queue 1 cos 1
switch(config)# map queue 2 cos 2
switch(config)# map queue 3 cos 3
switch(config)# map queue 4 cos 4
switch(config)# map queue 5 cos 5
switch(config)# map queue 6 cos 6
switch(config)# map queue 7 cos 7
switch(config)# qos schedule-profile S1
switch(config)# min-bandwidth queue 0 percent 5
switch(config)# min-bandwidth queue 1 percent 5
switch(config)# min-bandwidth queue 2 percent 10
switch(config)# min-bandwidth queue 3 percent 10
switch(config)# min-bandwidth queue 4 percent 20
switch(config)# min-bandwidth queue 5 percent 20
switch(config)# min-bandwidth queue 6 percent 10
switch(config)# min-bandwidth queue 7 percent 20
```

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

```
switch(config)# qos trust cos
switch(config)# qos queue-profile Q1
switch(config)# map queue 0 cos 0
switch(config)# map queue 1 cos 1
switch(config)# map queue 3 cos 3
switch(config)# map queue 4 cos 4
switch(config)# map queue 5 cos 5
switch(config)# map queue 6 cos 6
switch(config)# map queue 7 cos 7
switch(config)# qos schedule-profile S1
switch(config)# min-bandwidth queue 0 percent 5
switch(config)# min-bandwidth queue 1 percent 5
switch(config)# min-bandwidth queue 3 percent 10
switch(config)# min-bandwidth queue 4 percent 20
switch(config)# min-bandwidth queue 5 percent 20
switch(config)# min-bandwidth queue 6 percent 10
switch(config)# min-bandwidth queue 7 percent 20
```


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. |

map queue

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

Description

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

The `no` form of this command removes the specified cos value from a specific queue. If no CoS value is specified, then all CoS values are removed from the queue.

| Parameter | Description |
|-------------------|--|
| <QUEUE-NUMBER> | Specifies the queue number. Range: 0 to 7. |
| <PRIORITY-NUMBER> | Specifies the CoS value. 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

```
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 |
|-----------------------|------------------------|--|
| 4100i 6000 6100 | config-schedule-<NAME> | 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 |
|----------------|--|
| <QUEUE-NUMBER> | Specifies the queue number. Range: 0 to 7. |
| <DESCRIPTION> | 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

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 |
|-----------------------|-----------------|--|
| 4100i 6000 6100 | config-if | 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-POINT> | Specifies an IP differentiated services code point value. Range: 0 to 63. |

Usage

Order of operation for arriving IPv4 or IPv6 packets:

1. The CoS metadata is assigned from the DSCP map entry indexed by the DSCP remark value.
2. If a CoS remark is also configured along with the DSCP remark, the CoS remark value will be assigned to the packet's CoS metadata.
3. The CoS metadata and queue profile are then used to determine the queue for the packet. If the packet is transmitted with an 802.1Q VLAN tag, the PCP will be remarked to match the CoS metadata.

For arriving non-IP packets:

The CoS metadata is assigned from the DSCP map entry indexed by the DSCP remark value. This CoS value and the queue profile are used to select the queue for packet scheduling. The PCP of tagged non-IP packets will be remarked to this CoS value.

Examples

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 |
|-----------------------|----------------------------|--|
| 4100i 6000 6100 | 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> cos <COS-VALUE> [name <DESCRIPTION>]
no qos dscp-map <CODE-POINT>
```

Description

Defines the CoS value 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
DSCP      code_point  cos   name
-----
000000    0           0     CS0
000001    1           0
000010    2           0
000011    3           0
000100    4           0
000101    5           0
...
101101    45          5
101110    46          5
101111    47          5
110000    48          6     CS6
...
111100    60          7
111101    61          7
111110    62          7
111111    63          7
```

| Parameter | Description |
|---------------------------------------|---|
| <code><CODE-POINT></code> | Specifies an IP differentiated services code point. Range: 0 to 63. Default: 0. |
| <code>cos <COS-VALUE></code> | Specifies an 802.1p VLAN priority CoS remark value. Range: 0 to 7. Default 0. |
| <code>cos <PCP-VALUE></code> | Specifies an optional 802.1p VLAN Priority Code Point remark value. Range: 0 to 7. Default: No remark. |
| <code>name <DESCRIPTION></code> | 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 **41** to a CoS of **6**:

```
switch(config)# qos dscp-map 41 cos 6
```

Setting code point **41** to the default value:

```
switch(config)# no qos dscp-map 41
```

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 CoS values. Each profile has two, four, or eight queues numbered 0 to 7. The larger the queue number, the higher its priority during transmission scheduling.

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      percent max-bandwidth_kbps
-----
0      min-bandwidth 2
1      min-bandwidth 3
2      min-bandwidth 30
3      min-bandwidth 10
4      min-bandwidth 10
5      min-bandwidth 10
6      min-bandwidth 15
7      min-bandwidth 20
```

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, GMB).
2. The highest queue number uses strict priority, and all remaining (lower) queues use the same algorithm (for example, GMB). 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 percentage of a GMB queue.
- The bandwidth of a strict queue.
- The algorithm of the highest numbered queue can be swapped between GMB 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 | config | Administrators or local user group members with execution rights for this command. |

qos trust

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

Description

Configures one of three modes that are applied globally on all Ethernet interfaces and LAGs that have not applied their own trust mode. Trust mode determines whether VLAN or IP headers are used to assign CoS values to ingress packets.

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 |
|-----------|--|
| none | Ignores all packet headers. Ingress packets are assigned CoS |

| Parameter | Description |
|-----------|--|
| | value zero. |
| cos | For 802.1 VLAN-tagged packets, use the priority code point field from the outermost VLAN header to assign the CoS value. For untagged packets, the CoS value is assigned to zero. Default. |
| dscp | For IP packets, use the DSCP as the index into the DSCP Map table to obtain the CoS value for the packet. For non-IP packets, the CoS value is assigned to zero. |

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>
kbps
no rate-limit {broadcast | multicast | unknown-unicast | 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 unknown-unicast icmp {ip-all ipv4 ipv6}} | <p>Specifies the type of ingress traffic to which the rate limit applies: broadcast, multicast, unknown-unicast, or ICMP.</p> <p>The multicast rate limit affects multicast and broadcast traffic. The broadcast rate limit only affects broadcast traffic.</p> <p>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.</p> <p>Unknown-unicast packets may be intended for devices whose addresses have temporarily aged out of network forwarding caches. Configuring rate limits can help provide balance between necessary and flooded traffic.</p> <p>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.</p> |
| <RATE> | <p>Specifies the rate limit. Range: 64 to 100000000 kbps (in</p> |

| Parameter | Description |
|-----------|--|
| | steps of 64 kbps). 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

Limiting broadcast traffic to **1024kbps** on interface **1/1/3**:

```
switch(config)# interface 1/1/3
switch(config-if)# rate-limit broadcast 1024 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 1024 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
qos trust cos (global)
qos queue-profile factory-default (global)
qos schedule-profile factory-default (global)
rate-limit unknown-unicast 1024 kbps (1024 actual)
rate-limit broadcast 1024 kbps (1100 actual)
rate-limit multicast 1024 kbps (1100 actual)
rate-limit icmp ip-all 1024 kbps (1024 actual)

switch# show interface 1/1/3

Interface 1/1/3 is up
Admin state is up
Link state: up for 3 minutes (since Thu Nov 26 17:56:14 UTC 2020)
Link transitions: 1
Description:
Hardware: Ethernet, MAC Address: f8:60:f0:c9:21:bc
MTU 1500
Type 1GbT
Full-duplex
qos trust cos
rate-limit unknown-unicast 1024 kbps (1024 actual)
rate-limit broadcast 1024 kbps (1100 actual)
rate-limit multicast 1024 kbps (1100 actual)
rate-limit icmp ip-all 1024 kbps (1024 actual)
Speed 1000 Mb/s
Auto-negotiation is on
```

```

Energy-Efficient Ethernet is disabled
Flow-control: off
Error-control: off
MDI mode: MDIX
VLAN Mode: access
Access VLAN: 1
Rx
    0 total packets                0 total bytes
    0 unicast packets
    0 multicast packets
    0 broadcast packets
    0 errors                        0 dropped
    0 CRC/FCS                      0 pause
Tx
1057962 total packets            366066962 total bytes
    0 unicast packets
    0 multicast packets
1058039 broadcast packets
    0 errors                        0 dropped
    0 collision                    0 pause

```

Command History

| Release | Modification |
|------------------|--------------|
| 10.07 or earlier | -- |

Command Information

| Platforms | Command context | Authority |
|-----------------------|-----------------|--|
| 4100i 6000 6100 | 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 .

- **Tx Packets:** Total packets transmitted. The count includes packets subsequently dropped by an egress ACL.
- **Tx Drops:** Total packets dropped by an egress queue due to insufficient capacity.

Examples

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

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 cos (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)
```

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 dscp-map

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

Description

Displays the current or default global QoS dscp-map.

| Parameter | Description |
|-----------|---|
| default | Shows the factory default DSCP code point settings. |

Examples

Showing the current QoS DSCP map:


```

switch# show qos dscp-map
DSCP      code_point  cos   name
-----
000000    0           0     CS0
000001    1           0
000010    2           0
000011    3           0
000100    4           0
000101    5           0
...
101101    45          5
101110    46          6     new
101111    47          5
110000    48          6     CS6
...
111101    61          7
111110    62          7
111111    63          7

```

Showing the default QoS DSCP map:

```

switch# show qos dscp-map default
DSCP      code_point  cos   name
-----
000000    0           0     CS0
000001    1           0
000010    2           0
000011    3           0
000100    4           0
000101    5           0
...
101101    45          5
101110    46          5
101111    47          5
110000    48          6     CS6
...
111100    60          7
111101    61          7
111110    62          7
111111    63          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 queue-profile

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

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. |

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 cos values must be mapped to some queue.
- There can be only 2, 4, or 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 cos name
-----
0          1
1          0
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]
```

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. |

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      percent max-bandwidth_kbps
-----
0          min-bandwidth 2
1          min-bandwidth 3
2          min-bandwidth 30
3          min-bandwidth 10
4          min-bandwidth 10
5          min-bandwidth 10
6          min-bandwidth 15
7          min-bandwidth 20
```

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]

Description

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

| Parameter | Description |
|-----------|---|
| default | Shows the factory default QoS trust settings. |

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 cos
```

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 |
|--|--|
| <code><QUEUE-NUMBER></code> | Specifies the number of the queue. Range: 0 to 7. |
| <code>max-bandwidth <BANDWIDTH></code> | Specifies the maximum bandwidth allowed on the queue in Kbps. Range: 468 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 *min-bandwidth*

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.