



Quality of Service

The Quality of Service feature is applied throughout the network to ensure that network traffic is prioritized according to required criteria and that the desired traffic receives preferential treatment. This chapter contains the following sections:

- [General, on page 1](#)
- [QoS Basic Mode, on page 10](#)
- [QoS Advanced Mode, on page 11](#)
- [QoS Statistics, on page 19](#)

General

Quality of Service (QoS) is a feature on the switch which prioritizes traffic resulting in a performance improvement for critical network traffic. QoS varies by switch, as the higher the level switch, the higher the network application layer it works with. The number of queues differ, as well as the kind of information used to prioritize.

QoS Properties

Quality of Service (QoS) prioritizes the traffic flow based on the type of traffic and can be applied to prioritize traffic for latency-sensitive applications (such as voice or video) and to control the impact of latency-insensitive traffic.

To configure QoS properties, follow these steps:

Procedure

Step 1 Click **Quality of Service > General > QoS Properties**.

Step 2 Set the QoS mode. The following options are available:

- Disable—QoS is disabled on the device.
- Basic—QoS is enabled on the device in Basic mode.
- Advanced—QoS is enabled on the device in Advanced mode.

- Step 3** Select **Port/LAG** and click **Go** to display/modify all ports/LAGs on the device and their CoS information.
- The following fields are displayed for all ports/LAGs:
- Interface—Type of interface.
 - Default CoS—Default VPT value for incoming packets that do not have a VLAN Tag. The default CoS is 0.
- Step 4** Click **Apply**. The Running Configuration file is updated.
- To set QoS on an interface, select it, and click **Edit**.
- Step 5** Enter the parameters.
- Interface—Select the port or LAG.
 - Default CoS—Select the default CoS (Class-of-Service) value to be assigned for incoming packets (that do not have a VLAN tag).
- Step 6** Click **Apply**. The interface default CoS value is saved to Running Configuration file.
- To restore the default CoS values, click **Restore CoS Defaults**.
-

Queues

The device supports 8 queues for each interface. Queue number eight is the highest priority queue. Queue number one is the lowest priority queue.

There are two ways of determining how traffic in queues is handled, Strict Priority and Weighted Round Robin (WRR).

- Strict Priority—Egress traffic from the highest-priority queue is transmitted first. Traffic from the lower queues is processed only after the highest queue has been transmitted, thus providing the highest level of priority of traffic to the highest numbered queue.
- Weighted Round Robin (WRR)—In WRR mode the number of packets sent from the queue is proportional to the weight of the queue (the higher the weight the more frames are sent). For example, if there are a maximum of four queues possible and all four queues are WRR and the default weights are used, queue 1 receives 1/15 of the bandwidth (assuming all queues are saturated and there's congestion), queue 2 receives 2/15, queue 3 receives 4/15 and queue 4 receives 8 /15 of the bandwidth. The type of WRR algorithm used in the device isn't the standard Deficit WRR (DWRR), but rather Shaped Deficit WRR (SDWRR).

The queuing modes can be selected in the Queue page. When the queuing mode is by strict priority, the priority sets the order in which queues are serviced, starting with the highest priority queue and going to the next lower queue when each queue is completed.

When the queuing mode is Weighted Round Robin, queues are serviced until their quota has been used up and then another queue is serviced. It's also possible to assign some of the lower queues to WRR, while keeping some of the higher queues in strict priority. In this case traffic for the strict-priority queues is always sent before traffic from the WRR queues. Only after the strict-priority queues have been emptied is traffic from the WRR queues forwarded. (The relative portion from each WRR queue depends on its weight).

To select the priority method and enter WRR data, complete the following steps:

Procedure

Step 1 Click **Quality of Service > General > Queue**.

Step 2 Enter the parameters.

- Queue—Displays the queue number.
- Scheduling Method—Select one of the following options:
 - Strict Priority—Traffic scheduling for the selected queue and all higher queues is based strictly on the queue priority.
 - WRR—Traffic scheduling for the selected queue is based on WRR. The period time is divided between the WRR queues that aren't empty, meaning they have descriptors to egress. This division happens only if the strict-priority queues are empty.
 - WRR Weight—If WRR is selected, enter the WRR weight assigned to the queue.
 - % of WRR Bandwidth—Displays the amount of bandwidth assigned to the queue. These values represent the percent of the WRR weight.

Step 3 Click **Apply**. The queues are configured, and the Running Configuration file is updated.

CoS/802.1p to Queue

The CoS/802.1p to Queue page maps 802.1p priorities to egress queues. The CoS/802.1p to Queue Table determines the egress queues of the incoming packets based on the 802.1p priority in their VLAN Tags. For incoming untagged packets, the 802.1p priority is the default CoS/802.1p priority assigned to the ingress ports.

To map CoS values to egress queues, follow these steps:

Procedure

Step 1 Click **Quality of Service > General > CoS/802.1p to Queue**.

Step 2 Enter the parameters.

- 802.1p—Displays the 802.1p priority tag values to be assigned to an egress queue, where 0 is the lowest and 7 is the highest priority.
- Output Queue—Select the egress queue to which the 802.1p priority is mapped. Either four or eight egress queues are supported, where Queue 4 or Queue 8 is the highest priority egress queue and Queue 1 is the lowest priority.

Step 3 For each 802.1p priority, select the Output Queue to which it is mapped.

Step 4 Click **Apply**, **Cancel** or **Restore Defaults**. 801.1p priority values to queues are mapped, and the Running Configuration file is updated, the changes that entered are canceled, or previously defined values are restored.

DSCP to Queue

The DSCP (IP Differentiated Services Code Point) to Queue page maps DSCP values to egress queues. The DSCP to Queue Table determines the egress queues of the incoming IP packets based on their DSCP values. The original VPT (VLAN Priority Tag) of the packet is unchanged.

By simply changing the DSCP to Queue mapping and the Queue schedule method and bandwidth allocation, it's possible to achieve the desired quality of services in a network.

The DSCP to Queue mapping is applicable to IP packets if:

- The device is in QoS Basic mode and DSCP is the trusted mode.
- The device is in QoS Advanced mode and the packets belongs to flows that are DSCP trusted.

Non-IP packets are always classified to the best-effort queue.

The following tables describe the default DSCP to queue mapping for an 8-queue system where 7 is highest and 8 is used for stack control purposes.

DSCP	63	55	47	39	31	23	15	7
Queue	6	6	7	5	4	3	2	1
DSCP	62	54	46	38	30	22	14	6
Queue	6	6	7	5	4	3	2	1
DSCP	61	53	45	37	29	21	13	5
Queue	6	6	7	5	4	3	2	1
DSCP	60	52	44	36	28	20	12	4
Queue	6	6	7	5	4	3	2	1
DSCP	59	51	43	35	27	19	11	3
Queue	6	6	7	5	4	3	2	1
DSCP	58	50	42	34	26	18	10	2
Queue	6	6	7	5	4	3	2	1
DSCP	57	49	41	33	25	17	9	1
Queue	6	6	7	5	4	3	2	1
DSCP	56	48	40	32	24	16	8	0
Queue	6	6	6	7	6	6	1	1

The following tables describe the default DSCP to queue mapping for an 8-queue system where 8 is highest:

DSCP	63	55	47	39	31	23	15	7
Queue	7	7	8	6	5	4	3	1

DSCP	62	54	46	38	30	22	14	6
Queue	7	7	8	6	5	4	3	1
DSCP	61	53	45	37	29	21	13	5
Queue	7	7	8	6	5	4	3	1
DSCP	60	52	44	36	28	20	12	4
Queue	7	7	8	6	5	4	3	1
DSCP	59	51	43	35	27	19	11	3
Queue	7	7	8	6	5	4	3	1
DSCP	58	50	42	34	26	18	10	2
Queue	7	7	8	6	5	4	3	1
DSCP	57	49	41	33	25	17	9	1
Queue	7	7	8	6	5	4	3	1
DSCP	56	48	40	32	24	16	8	0
Queue	7	7	7	8	7	7	1	2

To map DSCP to queues, follow these steps:

Procedure

Step 1 Click **Quality of Service > General > DSCP to Queue**.

The DSCP to Queue page contains Ingress DSCP. It displays the DSCP value in the incoming packet and its associated class.

Step 2 Select the Output Queue (traffic forwarding queue) to which the DSCP value is mapped.

Step 3 Click **Apply**. The Running Configuration file is updated. Click **Restore Defaults** to restore the default settings.

Bandwidth



Note This setting is only available in the Advanced Setting view.

The Bandwidth page displays bandwidth information for each interface. To view the bandwidth information, complete the following steps:

Procedure

Step 1 Click **Quality of Service > General > Bandwidth**.

The fields in this page are described in the Edit page below, except for the following fields:

- **Ingress Rate Limit:**

- Status—Displays whether Ingress Rate Limit is enabled.
- Rate Limit (kbits/sec)—Displays the ingress rate limit for the port.
- %—Displays the ingress rate limit for the port divided by the total port bandwidth.
- CBS (Bytes)—Maximum burst size of data for the ingress interface in bytes of data

- **Egress Shaping Rates:**

- Status—Displays whether Egress Shaping Rates is enabled.
- CIR (kbits/sec)—Displays the maximum bandwidth for the egress interface.
- CBS (Bytes)—Maximum burst size of data for the egress interface in bytes of data

Step 2 Select an interface, and click **Edit**.

Step 3 Select the Port or LAG interface.

Step 4 Enter the fields for the selected interface:

Option	Description
Ingress Rate Limit	Select to enable the ingress rate limit, which is defined in the field below. (Not relevant for LAGs).
Ingress Rate Limit (kbits per sec)	Enter the maximum amount of bandwidth allowed on the interface. (Not relevant for LAGs).
Ingress Committed Burst Size (CBS)	Enter the maximum burst size of data for the ingress interface in bytes of data. This amount can be sent even if it temporarily increases the bandwidth beyond the allowed limit. This field is only available if the interface is a port. (Not relevant for LAGs).
Egress Shaping Rate	Select to enable egress shaping on the interface.
Committed Information Rate (CIR) (kbits/sec)	Enter the maximum bandwidth for the egress interface.
Egress Committed Burst Size (CBS)	Enter the maximum burst size of data for the egress interface in bytes of data. This amount can be sent even if it temporarily increases the bandwidth beyond the allowed limit.

Step 5 Click **Apply**. The bandwidth settings are written to the Running Configuration file.

Egress Shaping per Queue



Note This setting is only available in the Advanced Setting view.

In addition to limiting the transmission rate per port, which is done in the Bandwidth page, the device can limit the transmission rate of selected egressing frames on a per-queue per-port basis. Egress rate limiting is performed by shaping the output load.

The device limits all frames except for management frames. Any frames that aren't limited are ignored in the rate calculations, meaning that their size isn't included in the limit total.

To configure the egress shaping per queue, complete the following steps:

Procedure

Step 1 Click **Quality of Service > General > Egress Shaping per Queue**.

The Egress Shaping Per Queue page displays the rate limit (CIR) and burst size (CBS) for each queue.

Step 2 Select an interface type (Port or LAG), and click **Go**.

Step 3 Select a Port/LAG, and click **Edit**.

This page enables shaping the egress for up to eight queues on each interface.

Step 4 Select the Interface.

Step 5 For each queue that is required, enter the following fields:

- Enable Shaping—Select to enable egress shaping on this queue.
- Committed Information Rate (CIR)—Enter the maximum rate (CIR) in Kbits per second (Kbps). CIR is the average maximum amount of data that can be sent.
- Committed Burst Size (CBS)—Enter the maximum burst size (CBS) in bytes. CBS is the maximum burst of data allowed to be sent even if a burst exceeds CIR.

Step 6 Click **Apply**. The bandwidth settings are written to the Running Configuration file.

VLAN Ingress Rate Limit



Note This setting is only available in the Advanced Setting view.

Rate limiting per VLAN, performed in the VLAN Ingress Rate Limit page, enables traffic limiting on VLANs. When VLAN ingress rate limiting is configured, it limits aggregate traffic from all the ports on the device.

The following constraints apply to rate limiting per VLAN:

- It has a lower precedence than any other traffic policing defined in the system. For example, if a packet is subject to QoS rate limits but is also subject to VLAN rate limiting, and the rate limits conflict, the QoS rate limits take precedence.
- It's applied at the device level and within the device at the packet processor level. If there's more than one packet processor on the device, the configured VLAN rate limit value is applied to each of the packet processors, independently. Devices with up to 24 ports have a single packet processor, while devices of 48 ports or more have two packet processors.

Rate limiting is calculated separately for each packet processor in a unit.

To define the VLAN ingress rate limit, complete the following steps:

Procedure

Step 1 Click **Quality of Service > General > VLAN Ingress Rate Limit**.

This page displays the VLAN Ingress Rate Limit Table.

Step 2 Click **Add**.

Step 3 Enter the parameters.

- VLAN ID—Select a VLAN.
- Committed Information Rate (CIR)—Enter the average maximum amount of data that can be accepted into the VLAN in Kilobits per second.
- Committed Burst Size (CBS)—Enter the maximum burst size of data for the egress interface in bytes of data. This amount can be sent even if it temporarily increases the bandwidth beyond the allowed limit. This can't be entered for LAGs.

Step 4 Click **Apply**. The VLAN rate limit is added, and the Running Configuration file is updated.

iSCSI

This setting is only available in the Advanced Setting view.

This page enables activating iSCSI optimization, which means setting up a mechanism for giving priority to iSCSI traffic over other types of traffic. If this feature is enabled on a device, iSCSI traffic on any interface will be assigned the defined priority, and iSCSI traffic won't be subject to ACL or Policy rules set on interface.

iSCSI traffic is identified by the TCP port on which iSCSI targets listen to requests and optionally also by the IPv4 address on which iSCSI targets listen to requests. Two iSCSI IPv4 flows with well-known TCP ports 3260 and 860 are defined by default on device. iSCSI flow optimization is bi-directional, which means that it's applied to streams in both directions – from and to targets.

To enable and configure the mechanism for prioritizing and, optionally, marking iSCSI traffic, complete the following steps:

Procedure

Step 1 Click **Quality of Service > General > iSCSI**.

Step 2 Check **Enable** in the iSCSI Status field to enable processing iSCSI traffic on the device.

Step 3 Enter the fields under Quality of Service Settings:

- VPT Assignment—Select either Unchanged to leave the original VLAN Priority Tag (VPT) value in the packet or enter a new value in the Reassigned field.
- DSCP Assignment—Select either Unchanged to leave the original DSCP value in the packet or enter a value in the Reassigned field.
- Queue Assignment—Enter the Queue assignment for iSCSI traffic. By default it's assigned to Queue 7.

Step 4 Click **Apply** to save the settings.

The iSCSI Flow Table displays the various iSCSI flows that have been defined. Two iSCSI flows, with well-known TCP ports 3260 and 860, are displayed. The Flow Type of these flows is Default. If you add a new flow, its Flow Type is Static.

To add a new flow:

Step 5 Select the TCP port and click **Add** and enter the following fields:

- TCP Port—This is the TCP port number on which the iSCSI target listens to requests. You can configure up to 8 target TCP ports on the switch.
- Target IP Address—Specifies the IP address of the iSCSI target (where data is stored). This is also the source of the iSCSI traffic. You can select **Any** to define a flow according to the TCP port parameter, or enter an IP address in **User-Defined** field to define a specific target address.

Step 6 Click **Apply** to save the settings.

Click **Restore Default Flows** to restore the default flows.

TCP Congestion Avoidance



Note This setting is only available in the Advanced Setting view.

The TCP Congestion Avoidance page enables activating a TCP congestion avoidance algorithm. The algorithm breaks up or avoids TCP global synchronization in a congested node, where the congestion is due to various sources sending packets with the same byte count.

To configure TCP congestion avoidance, complete the following steps:

Procedure

-
- Step 1** Click **Quality of Service > General > TCP Congestion Avoidance**.
- Step 2** Click **Enable** to enable TCP congestion avoidance, and click **Apply**.
-

QoS Basic Mode

In QoS Basic mode, a specific domain in the network can be defined as trusted. Within that domain, packets are marked with 802.1p priority and/or DSCP to signal the type of service they require. Nodes within the domain use these fields to assign the packet to a specific output queue. The initial packet classification and marking of these fields is done in the ingress of the trusted domain.

Global Settings



Note If the QoS Mode is set to Advanced, the Global Settings menu displays "The device is currently not in CoS/QoS Basic mode." and no settings are available.

The Global Settings page contains information for enabling Trust on the device (see the Trust Mode field below). This configuration is active when the QoS mode is Basic mode. Packets entering a QoS domain are classified at the edge of the QoS domain.

To define the Trust configuration, complete the following steps:

Procedure

-
- Step 1** Click **Quality of Service > QoS Basic Mode > Global Settings**.
- Step 2** Select the Trust Mode while the device is either in Basic or Advanced mode. If a packet CoS level and DSCP tag are mapped to separate queues, the Trust mode determines the queue to which the packet is assigned:
- CoS/802.1p—Traffic is mapped to queues based on the VPT field in the VLAN tag, or based on the per-port default CoS/802.1p value (if there's no VLAN tag on the incoming packet), the actual mapping of the VPT to queue can be configured in the mapping CoS/802.1p to Queue page.
 - DSCP—All IP traffic is mapped to queues based on the DSCP field in the IP header. The actual mapping of the DSCP to queue can be configured in the DSCP to Queue page. If traffic isn't IP traffic, it's mapped to the best effort queue.
 - CoS/802.1p-DSCP—Either CoS/802.1p or DSCP whichever has been set.
- Step 3** Select **Override Ingress DSCP** to enable and override the original DSCP values in the incoming packets with the new values entered in the DSCP Override table. When Override Ingress DSCP is enabled, the device uses the new DSCP values for egress queuing. It also replaces the original DSCP values in the packets with the new DSCP values.

Note

The frame is mapped to an egress queue using the new, rewritten value, and not by the original DSCP value.

- Step 4** Click **DSCP Override Table** to reconfigure DSCP. (See DSCP Override Table).
- Step 5** DSCP In displays the DSCP value of the incoming packet that needs to be re-marked to an alternative value. Select the DSCP Out value to indicate the outgoing value is mapped.
- Step 6** Click **Apply**. The Running Configuration file is updated with the new DSCP values. Click **Restore Defaults** to go back to the default settings.

Interface Settings

The Interface Settings page enables configuring QoS on each port of the device, as follows:

- QoS State Disabled on an Interface—All inbound traffic on the port is mapped to the best effort queue and no classification/prioritization takes place.
- QoS State of the Port is Enabled—Port prioritize traffic on ingress is based on the system-wide configured trusted mode, which is either CoS/802.1p trusted mode or DSCP trusted mode.

To enter QoS settings per interface, complete the following steps:

Procedure

- Step 1** Click **Quality of Service > QoS Basic Mode > Interface Settings**.
- Step 2** Use the filter to select the Interface Type (Port or Lag) and click **Go** to display the current settings. QoS State displays whether QoS is enabled on the interface
- Step 3** Select an interface, and click **Edit**.
- Step 4** Select the Port or LAG interface.
- Step 5** Click to enable or disable QoS State for this interface.
- Step 6** Click **Apply**. The Running Configuration file is updated.

QoS Advanced Mode



Note This setting is only available in the Advanced Setting view.

Frames that match an ACL and permitted entrance are implicitly labeled with the name of the ACL that permitted their entrance. Advanced mode QoS actions can then be applied to these flows.

In QoS advanced mode, the device uses policies to support per flow QoS. A policy and its components have the following characteristics and relationships:

- A policy contains one or more class maps.

- A class map defines a flow with one or more associating ACLs. Packets that match only ACL rules (ACE) in a class map with Permit (forward) action are considered belonging to the same flow, and are subjected to the same quality of services. Thus, a policy contains one or more flows, each with a user-defined QoS.
- The QoS of a class map (flow) is enforced by the associating policer. There are two types of policers, single policer and aggregate policer. Each policer is configured with a QoS specification. A single policer applies the QoS to a single class map, and thus to a single flow, based on the policer QoS specification. An aggregate policer applies the QoS to one or more class maps, and thus one or more flows. An aggregate policer can support class maps from different policies.

The 2 Rate 3 Color (2R3C) feature is supported on the device. In this feature, every policer has two thresholds. If the first threshold is reached, a user-configured Exceed action is performed. If the second threshold is reached, a user-configured Violate action is performed.

- Per flow QoS is applied to flows by binding the policies to the desired ports. A policy and its class maps can be bound to one or more ports, but each port is bound with at most one policy.

Global Settings

The Global Settings page contains information for enabling Trust on the device. Packets entering a QoS domain are classified at the edge of the QoS domain.

To define the Trust configuration:

Procedure

-
- Step 1** Click **Quality of Service > QoS Advanced Mode > Global Settings**.
- Step 2** Select the Trust Mode while the device is in Advanced mode. If a packet CoS level and DSCP tag are mapped to separate queues, the Trust mode determines the queue to which the packet is assigned:
- CoS/802.1p—Traffic is mapped to queues based on the VPT field in the VLAN tag, or based on the per-port default CoS/802.1p value (if there's no VLAN tag on the incoming packet), the actual mapping of the VPT to queue can be configured in the mapping CoS/802.1p to Queue page.
 - DSCP—All IP traffic is mapped to queues based on the DSCP field in the IP header. The actual mapping of the DSCP to queue can be configured in the DSCP to Queue page. If traffic isn't IP traffic, it's mapped to the best effort queue.
 - CoS/802.1p-DSCP—Select to use Trust CoS mode for non-IP traffic and Trust DSCP for IP traffic.
- Step 3** Select the default Advanced mode QoS trust mode (either trusted or untrusted) for interfaces in the Default Mode Status field. This provides basic QoS functionality on Advanced QoS, so that you can trust CoS/DSCP on Advanced QoS by default (without having to create a policy).
- Step 4** In QoS Advanced Mode, when the Default Mode Status is set to Not Trusted, the default CoS values configured on the interface is ignored and all the traffic goes to queue 1. See the Quality of Service > QoS Advanced Mode > Global Settings page for details.
- Step 5** If you have a policy on an interface then the Default Mode is irrelevant, the action is according to the policy configuration and unmatched traffic is dropped.

- Step 6** Select **Override Ingress DSCP** to override the original DSCP values in the incoming packets with the new values according to the DSCP Override Table. When Override Ingress DSCP is enabled, the device uses the new DSCP values for egress queuing. It also replaces the original DSCP values in the packets with the new DSCP values.

Note

The frame is mapped to an egress queue using the new, rewritten value, and not by the original DSCP value.

- Step 7** If Override Ingress DSCP was enabled, click **DSCP Override Table** to reconfigure DSCP.

a) In The DSCP Override Table, enter the following fields:

- DSCP In—Displays the DSCP value of the incoming packet that needs to be remarked to an alternative value.
- DSCP Out—Select the DSCP Out value to indicate the outgoing value is mapped.

b) Click **Apply**. To go back to the default settings, click **Restore Defaults**.

Out-of-Profile DSCP Mapping

When a policer is assigned to a class maps (flows), you can specify the action to take when the amount of traffic in one or more flows exceeds the QoS-specified limits. The portion of the traffic that causes the flow to exceed its QoS limit is referred to as out-of-profile packets. If the exceed/violate action is Out of Profile DSCP, the device remaps the original DSCP value of the out-of-profile IP packets with a new value based on the Out of Profile DSCP Remarking Table. The device uses the new values to assign resources and the egress queues to these packets. The device also physically replaces the original DSCP value in the out of profile packets with the new DSCP value.

To use the out-of-profile DSCP exceed action, remap the DSCP value in the Out Of Profile DSCP Remarking Table. Otherwise the action is null, because the DSCP value in the table remaps the packets to itself by factory default. This feature changes the DSCP tags for incoming traffic switched between trusted QoS domains. Changing the DSCP values used in one domain, sets the priority of that type of traffic to the DSCP value used in the other domain to identify the same type of traffic. These settings are active when the system is in the QoS Advance mode, and once activated they are active globally. This can be configured in the [QoS Properties, on page 1](#).

To map DSCP values, follow these steps:

Procedure

-
- Step 1** Click **Quality of Service > QoS Advanced Mode > Out of Profile DSCP Mapping**. This page enables setting the DSCP-value of traffic entering or leaving the device.

DSCP In displays the DSCP value of the incoming packet that needs to be re-marked to an alternative value.

You may filter according to **Action Type** to display all **Exceed** or **Violate**. This enables you to configure remarking when the traffic exceeds wither the Exceed or Violate threshold of a policer.

- Step 2** Select the DSCP Out value to where the incoming value is mapped.

- Step 3** Click **Apply**. The Running Configuration file is updated with the new DSCP Remarking table.

Step 4 Click **Restore Defaults** to restore the factory CoS default setting for this interface.

Class Mapping

A Class Map defines a traffic flow with ACLs (Access Control Lists) defined on it. A MAC ACL, IP ACL, and IPv6 ACL can be combined into a class map. Class maps are configured to match packet criteria on a match-all or match-any basis. They are matched to packets on a first-fit basis, meaning that the action associated with the first-matched class map is the action performed by the system. Packets that match the same class map are considered to belong to the same flow.



Note Defining class maps doesn't have any effect on QoS; it's an interim step, enabling the class maps to be used later.

If more complex sets of rules are needed, several class maps can be grouped into a supergroup called a policy.

In the same class map, a MAC ACL can't be used with an IPv6 ACE that has a Destination IPv6 address as a filtering condition.

The Class Mapping page shows the list of defined class maps and the ACLs comprising each, and enables you to add/delete class maps.

To define a Class Map, complete the following steps:

Procedure

Step 1 Click **Quality of Service > QoS Advanced Mode > Class Mapping**.

For each class map, the ACLs defined on it are displayed along with the relationship between them. Up to three ACLs can be displayed along with their Match, which can be either And or Or. This indicates the relationship between the ACLs. The Class Map is then the result of the three ACLs combined with either And or Or.

Step 2 Click **Add**.

A new class map is added by selecting one or two ACLs and giving the class map a name. If a class map has two ACLs, you can specify that a frame must match both ACLs, or that it must match either one or both of the ACLs selected.

Step 3 Enter the parameters.

- Class Map Name—Enter the name of a new class map.
- Match ACL Type—The criteria that a packet must match in order to be considered to belong to the flow defined in the class map. The options are:
 - IP—A packet must match either of the IP-based ACLs in the class map.
 - MAC—A packet must match the MAC-based ACL in the class map.
 - IP and MAC—A packet must match the IP-based ACL and the MAC-based ACL in the class map.
 - IP or MAC—A packet must match either the IP-based ACL or the MAC-based ACL in the class map.

- IP—Select the IPv4 based ACL or the IPv6 based ACL for the class map.
- MAC—Select the MAC-based ACL for the class map.
- Preferred ACL—Select whether packets are first matched to an IP or MAC.

Step 4 Click **Apply**. The Running Configuration file is updated.

Aggregate Policer

You can measure the rate of traffic that matches a predefined set of rules. To enforce limits, use ACLs in one or more class maps to match the desired traffic, and use a policer to apply the QoS on the matching traffic.

A policer is configured with a QoS specification. There are two kinds of policers:

- **Single (Regular) Policer**—A single policer applies the QoS to a single class map, and to a single flow based on the policer's QoS specification. When a class map using single policer is bound to multiple ports, each port has its own instance of single policer. Thus, each applying the QoS on the class map (flow) at ports that are otherwise independent of each other. A single policer is created in the Policy Table page.
- **Aggregate Policer**—An aggregate policer applies the QoS to one or more class maps, and one or more flows. An aggregation policer can support class maps from different policies. An aggregate policer applies QoS to all its flows in aggregation regardless of policies and ports. An aggregate policer is created in the Aggregate Policer page.

An aggregate policer is defined if the policer is to be shared with more than one class. Policers on a port can't be shared with other policers in another device.

Each policer is defined with its own QoS specification with a combination of the following parameters:

- **Peak Enforcement**—Select to enable action if peak burst size is exceeded.
- **Peak Information Rate (PIR)**—Enter the peak traffic rate (PIR) in kbits per second (kbps).
- **Peak Burst Size (PBS)**—Enter the peak burst size (PIR) in bytes.
- **Violate Action**—Select one of the following actions if peak size is exceeded:
 - **Drop**—Drop the frames violating the peak size.
 - **Out-of-Profile DSCP**—Mark frames violating the peak size with the DSCP value with previously set DSCP value
- A maximum allowed rate, called a Committed Information Rate (CIR), measured in Kbps.
- An amount of traffic, measured in bytes, called a Committed Burst Size (CBS). This is traffic that is allowed to pass as a temporary burst even if it's above the defined maximum rate.
- An action to be applied to frames that are over the limits (called out-of-profile traffic), where such frames can be passed as is, dropped, or passed, but remapped to a new DSCP value that marks them as lower-priority frames for all subsequent handling within the device.
- Configures traffic policing on the basis of the specified rates and optional actions Enter the CIR and these optional values and actions

Assigning a policer to a class map is done when a class map is added to a policy. If the policer is an aggregate policer, you must create it using the Aggregate Policar page.

To define an aggregate policer, complete the following steps:

Procedure

Step 1 Click **Quality of Service > QoS Advanced Mode > Aggregate Policar**.

This page displays the existing aggregate policers.

Step 2 Click **Add**.

Step 3 Enter the parameters.

- Aggregate Policar Name—Enter the name of the Aggregate Policar.
- Ingress Committed Information Rate (CIR)—Enter the maximum bandwidth allowed in bits per second. See the description of this in the [Bandwidth, on page 5](#).
- Ingress Committed Burst Size (CBS)—Enter the maximum burst size (even if it goes beyond the CIR) in bytes. See the description of this in the [Bandwidth, on page 5](#).
- Exceed Action—Select the action to be performed on incoming packets that exceed the CIR. Possible values are:
 - Drop—Packets exceeding the defined CIR value are dropped.
 - Out of Profile DSCP—The DSCP values of packets exceeding the defined CIR value are remapped to a value based on the Out Of Profile DSCP Remarking Table.

Step 4 Click **Apply**. The Running Configuration file is updated.

Policy Table

The Policy Table Map page displays the list of advanced QoS policies defined in the system. The page also allows you to create and delete policies. Only those policies that are bound to an interface are active (see [Policy Binding, on page 18](#)).

Each policy consists of:

- One or more class maps of ACLs which define the traffic flows in the policy.
- One or more aggregates that applies the QoS to the traffic flows in the policy.

After a policy has been added, class maps can be added by using the Policy Table page. To add a QoS policy, complete the following steps:

Procedure

Step 1 Click **Quality of Service > QoS Advanced Mode > Policy Table**.

This page displays the list of defined policies.

- Step 2** Click **Policy Class Map Table** to display the Policy Class Maps page or click **Add** to open the Add Policy Table page.
- Step 3** Enter the name of the new policy in the New Policy Name field.
- Step 4** Click **Apply**. The QoS policy profile is added, and the Running Configuration file is updated.

Policy Class Maps

One or more class maps can be added to a policy. A class map defines the type of packets that are considered to belong to the same traffic flow.

To add a class map to a policy:

Procedure

- Step 1** Click **Quality of Service > QoS Advanced Mode > Policy Class Maps**.
- Step 2** Select a policy in the Filter, and click **Go**. All class maps in that policy are displayed.
- Step 3** To add a new class map, click **Add**.
- Step 4** Enter the following parameters.

Policy Name	Displays the policy to which the class map is being added.
Class Map Name	Select an existing class map to be associated with the policy. Class maps are created in the Class Mapping page.
Action Type	Select the action regarding the ingress CoS/802.1p and/or DSCP value of all the matching packets. • Use default trust mode—If this option is selected, use the default mode status in Global Trust mode. If the default mode status is “Not Trusted”, ignore the ingress CoS/802.1p and/or DSCP value and the matching packets are sent as best effort. • Always Trust—If this option is selected, the device trusts the matching packet based on the Global Trust mode (selected in the Global Settings page). It ignores the Default Mode status (selected in the Global Settings page). • Set—If this option is selected, use the value entered in the New Value box to determine the egress queue of the matching packets as follows: If the new value (0..7) is a CoS/802.1p priority, use the priority value and the CoS/802.1p to Queue Table to determine the egress queue of all the matching packets. If the new value (0..63) is a DSCP, use the new DSCP and the DSCP to Queue Table to determine the egress queue of the matching IP packets. Otherwise, use the new value (1..8) as the egress queue number for all the matching packets.

Police Type	Select the policer type for the policy. The options are: • None—No policy is used. • Single—The policer for the policy is a single policer. • Aggregate—The policer for the policy is an aggregate policer.
-------------	---

Step 5 If Police Type is Aggregate, select the Aggregate Policer, and enter the following QoS parameters:

Ingress Committed Information Rate (CIR)	Enter the CIR in Kbps. See a description of this in the Bandwidth page.
Ingress Committed Burst Size (CBS)	Enter the CBS in bytes. See a description of this in the Bandwidth page.
Exceed Action	Select the action assigned to incoming packets exceeding the CIR. The options are: • Drop—Packets exceeding the defined CIR value are dropped. • Out of Profile DSCP—IP packets exceeding the defined CIR are forwarding with a new DSCP derived from the Out Of Profile DSCP Remarking Table.

Step 6 Click **Apply**.

Policy Binding

The Policy Binding page shows which policy profile is bound and to which port. A policy can be bound to an interface as an ingress (input) policy or as an egress (output) policy. When a policy profile is bound to a specific port, it's active on that port. Only one policy profile can be configured per port and per direction. However, a single policy can be bound to more than one port.

When a policy is bound to a port, it filters and applies QoS to traffic that belongs to the flows defined in the policy.

To edit a policy, it must first be removed (unbound) from all those ports to which it's bound.



Note It's possible to either bind a port to a policy or to an ACL but both can't be bound.

To define policy binding, complete the following steps:

Procedure

- Step 1** Click **Quality of Service > QoS Advanced Mode > Policy Binding**.
- Step 2** Select an Interface Type if required.
- Step 3** Click **Go**. The policies for that interface are displayed.
- Step 4** Click **Edit**.

Step 5 Select the interface (Port or Lag) and configure the following:

- Input Policy Binding—Select to bind the input policy to the interface.
- Policy Name—Select the input policy being bound.
- Default Action—Select action if packet matches policy:
 - Deny Any—Select to forward packets on the interface if they match any policy.
 - Permit Any—Select to forward packets on the interface if they don't match any policy.

Note

Permit Any can be defined only if IP Source Guard isn't activated on the interface.

Step 6 Select Enable to enable the Output Policy Binding and configure the following:

- Policy Name—Select the output policy being bound.
- Default Action—Select action if packet matches policy:
 - Deny Any—Select to forward packets on the interface if they match any policy.
 - Permit Any—Select to forward packets on the interface if they don't match any policy.

Note

Permit Any can be defined only if IP Source Guard isn't activated on the interface.

Step 7 Click **Apply**. The QoS policy binding is defined, and the Running Configuration file is updated.

QoS Statistics

QoS statistics feature allows you to gather statistics for the rate at which packets are forwarded out of a queue and for the rate at which committed, conformed, or exceeded packets are dropped on the device.

Single Policer Statistics

The Single Policer Statistics page indicates the number of in-profile and out-of-profile packets that are received from an interface that meet the conditions defined in the class map of a policy.



Note This page isn't displayed when the device is in Layer 3 mode and is visible only in the Advanced mode view.

To view policer statistics:

Procedure

Step 1 Click **Quality of Service > QoS Statistics > Single Policer Statistics**.

This page displays the following fields:

- Interface—Statistics for this interface.
- Policy—Statistics for this policy.
- Class Map—Statistics for this class map.
- In-Profile Bytes—Number of in-profile bytes received.
- Out-of-Profile Bytes—Number of out-of-profile bytes received.

Step 2 Click **Add**.

Step 3 Enter the parameters.

- Interface—Select the interface for which statistics are accumulated.
- Policy Name—Select the policy name.
- Class Map Name—Select the class name.

Step 4 Click **Apply**. An additional request for statistics is created and the Running Configuration file is updated.

Step 5 Click **Delete** to delete the data or click **Clear Counters** to clear the data from the Single Policer Statistic Table.

Aggregate Policer Statistics



Note This setting is only available in the Advanced Setting view.

To view aggregated policer statistics:

Procedure

Step 1 Click **Quality of Service > QoS Statistics > Aggregate Policer Statistics**.

This page displays the following fields:

- Aggregate Policer Name—Policer on which statistics are based.
- In-Profile Bytes—Number of in-profile packets that were received.
- Out-of-Profile Bytes—Number of out-of-profile packets that received.

Step 2 Click **Add**.

- Step 3** Select an Aggregate Policer Name, one of the previously created Aggregate Policers for which statistics are displayed.
- Step 4** Click **Apply**. An additional request for statistics is created, and the Running Configuration file is updated.
- Step 5** Click **Delete** to remove a specific statistic.
- Step 6** Click **Clear Counters** to clear the counters of the selected policer.
-

Queue Statistics

The Queues Statistics page displays queue statistics, including statistics of forwarded and dropped packets, based on interface, queue, and drop precedence.

To view Queues Statistics and define what statistics to display (Counter Set):

Procedure

- Step 1** Click **Quality of Service > QoS Statistics > Queue Statistics**.

This page displays the following fields:

- Refresh Rate—Select the time period that passes before the interface Ethernet statistics are refreshed. The available options are:
 - No Refresh—Statistics aren't refreshed.
 - 15 Sec—Statistics are refreshed every 15 seconds.
 - 30 Sec—Statistics are refreshed every 30 seconds.
 - 60 Sec—Statistics are refreshed every 60 seconds.

To view a specific unit and interface, select the unit/interface in the filter and click **Go**.

To view a specific interface, select the interface in the filter and click **Go**.

The Queue Statistics Table displays the following fields for each queue:

- Queue—Packets forwarded or tail dropped from this queue.
- Transmitted Packets—Number of packets that were transmitted.
- Tail Dropped Packets—Number of packets that were tail dropped.
- Transmitted Bytes—Number of bytes that were transmitted.
- Tail Dropped Bytes—Number of bytes that were tail dropped.

- Step 2** Click **Clear Interface Counters** to clear the statistic counters for the selected interface.
- Step 3** Click **Clear All Interface Counters** to clear the statistic counters for all interfaces.
-

