

---

# Configuration Commands

---

## Generic Commands

### description

**Syntax**    **description** *description-string*  
              **no description**

**Context**    config>qos>sap-egress  
              config>qos>sap-egress>ip-criteria>entry  
              config>qos>sap-ingress  
              config>qos>sap-ingress>ip-criteria>entry  
              config>qos>sap-ingress>ipv6-criteria>entry  
              config>qos>sap-ingress>mac-criteria>entry

**Description**    This command creates a text description stored in the configuration file for a configuration context.  
                  The **no** form of this command removes any description string from the context.

**Default**        No description is associated with the configuration context.

**Parameters**    *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

---

# Operational Commands

## copy

**Syntax** **copy sap-egress** *src-pol dst-pol* [**overwrite**]  
**copy sap-ingress** *src-pol dst-pol* [**overwrite**]  
**hsmda-pool-policy** *src-name dst-name* [**overwrite**]  
**hsmda-scheduler-policy** *src-name dst-name* [**overwrite**]  
**hsmda-slope-policy** *src-name dst-name* [**overwrite**]  
**named-pool-policy** *src-name dst-name* [**overwrite**]

**Context** config>qos

**Description** This command copies existing QoS policy entries for a QoS policy-id to another QoS policy-id.

The **copy** command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

**Parameters** **sap-egress** *src-pol dst-pol* — Indicates that the source policy ID and the destination policy ID are sap-egress policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

**Values** 1 — 65535

**sap-ingress** *src-pol dst-pol* — Indicates that the source policy ID and the destination policy ID are SAP ingress policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

**Values** 1 — 65535

**hsmda-pool-policy** *src-name dst-name* — Indicates that the source HSMDA pool policy ID and the destination policy ID are HSMDA pool policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

**hsmda-scheduler-policy** *src-name dst-name* — Indicates that the source HSMDA scheduler policy ID and the destination policy ID are HSMDA scheduler policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

**hsmda-slope-policy** *src-name dst-name* — Indicates that the source HSMDA slope policy ID and the destination policy ID are HSMDA slope policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

**overwrite** — Specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, an error will occur if the destination policy ID exists.

```
SR>config>qos# copy sap-egress 1 1010
MINOR: CLI Destination "1010" exists use {overwrite}.
SR>config>qos# copy sap-egress 1 1010 overwrite
```

## renum

**Syntax** `renum old-entry-number new-entry-number`

**Context**  
config>qos>sap-ingress>ip-criteria  
config>qos>sap-ingress>ipv6-criteria  
config>qos>sap-egress>ipv6-criteria  
config>qos>sap-ingress>mac-criteria

**Description** This command renumbers existing QoS policy criteria entries to properly sequence policy entries. This can be required in some cases since the router exits when the first match is found and executes the actions in accordance with the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.

**Parameters** *old-entry-number* — Enter the entry number of an existing entry.

**Default** none

**Values** 1 — 65535

*new-entry-number* — Enter the new entry-number to be assigned to the old entry.

**Default** none

**Values** 1 — 65535

## type

**Syntax** `type filter-type`

**Context** config>qos>sap-ingress>mac-criteria

**Description** This command sets the mac-criteria type.

**Default** normal

**Parameters** *filter-type* — Specifies which type of entries this MAC filter can contain.

**Values** **normal** — Regular match criteria are allowed; ISID match not allowed.

**vid** — Configures the VID filter type used to match on ethernet\_II frame types. This allows matching VLAN tags for explicit filtering.

---

# Service Ingress QoS Policy Commands

## sap-ingress

**Syntax** [no] **sap-ingress** *policy-id* | *policy-name*

**Context** config>qos

**Description** This command is used to create or edit the ingress policy. The ingress policy defines the Service Level Agreement (SLA) enforcement service packets receive as they ingress a SAP. SLA enforcement is accomplished through the definition of queues that have Forwarding Class (FC), Committed Information Rate (CIR), Peak Information Rate (PIR), Committed Burst Size (CBS), and Maximum Burst Size (MBS) characteristics. The simplest policy defines a single queue that all ingress traffic flows through. Complex policies have multiple queues combined with specific IP or MAC match criteria that indicate which queue a packet will flow through.

Policies in effect are templates that can be applied to multiple services as long as the **scope** of the policy is template. Queues defined in the policy are not instantiated until a policy is applied to a service SAP.

A SAP ingress policy is considered incomplete if it does not include definition of at least one queue and does not specify the default action. 7750 SR OS software does not allow incomplete SAP ingress policies to be applied to services.

SAP ingress policies can be defined with either IP headers as the match criteria or MAC headers as the match criteria. The IP and MAC criteria are mutually exclusive and cannot be part of the same SAP ingress policy.

It is possible that a SAP ingress policy will include the **dscp** map command, the **dot1p** map command and an IP or MAC match criteria. When multiple matches occur for the traffic, the order of precedence will be used to arrive at the final action. The order of precedence is as follows:

1. 802.1p bits
2. DSCP
3. IP Quintuple or MAC headers

The SAP ingress policy with *policy-id* 1 is a system-defined policy applied to services when no other policy is explicitly specified. The system SAP ingress policy can be modified but not deleted. The **no sap-ingress** command restores the factory default settings when used on *policy-id* 1. The default SAP ingress policy defines one queue associated with the best effort (be) forwarding class, with CIR of zero and PIR of line rate.

Any changes made to the existing policy, using any of the sub-commands are applied immediately to all services where this policy is applied. For this reason, when many changes are required on a policy, it is recommended that the policy be copied to a work area policy ID. That work-in-progress policy can be modified until complete and then written over the original policy-id. Use the **config qos copy** command to maintain policies in this manner.

The **no sap-ingress** *policy-id* command deletes the SAP ingress policy. A policy cannot be deleted until it is removed from all services where it is applied. The system default SAP ingress policy is a special case; the **no** command restores the factory defaults to policy-id 1.

**Parameters** *policy-id* — The *policy-id* uniquely identifies the policy.

**Values** 1 — 65535

*policy-name* — The *policy-name* uniquely identifies the policy.

**Values** Valid names consist of any string up to 64 characters long. Policies must first be created with a *policy-id*, after which a *policy-name* can be assigned and used as an alias to reference the policy during configuration changes. Policy names may not begin with a number (0-9) or the underscore “\_” character (e.g. *\_myPolicy*). “Default” or “default” can not be used as policy names. Saved configurations and display output from the "info" and most "show" commands will show the *policy-id* (not the *policy-name*) where the policies are referenced.

## policy-name

**Syntax** **policy-name** *policy-name*  
**no policy-name**

**Context** cconfig>qos>sap-ingress  
config>qos>sap-egress

**Description** Policies must first be created with a *policy-id*, after which a *policy-name* can be assigned and used as an alias to reference the policy during configuration changes. Saved configurations and display output from the **info** and most **show** commands will show the *policy-id* (not the *policy-name*) where the policies are referenced.

**Default** no *policy-name*

**Parameters** *policy-name* — Policy names may not begin with a number (0-9) or the underscore “\_” character (e.g. *\_myPolicy*). “Default” or “default” cannot be used as policy names. Specify a character string 64 characters or less.

## scope

**Syntax** **scope** {**exclusive** | **template**}  
**no scope**

**Context** config>qos>sap-ingress *policy-id*

**Description** This command configures the Service Ingress QoS policy scope as exclusive or template. The policy’s scope cannot be changed if the policy is applied to a service. The **no** form of this command sets the scope of the policy to the default of **template**.

**Default** template

**Parameters** **exclusive** — When the scope of a policy is defined as exclusive, the policy can only be applied to one SAP. If a policy with an exclusive scope is assigned to a second SAP an error message is generated. If the

policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP.

The system default policies cannot be put into the exclusive scope. An error will be generated if scope exclusive is executed in any policies with a policy-id equal to 1.

**template** — When the scope of a policy is defined as template, the policy can be applied to multiple SAPs on the router.

Default QoS policies are configured with template scopes. An error is generated when the template scope parameter to exclusive scope on default policies is modified.

## default-fc

**Syntax** `default-fc fc-name`

**Context** `config>qos>sap-ingress`

**Description** This command configures the default forwarding class for the policy. In the event that an ingress packet does not match a higher priority (more explicit) classification command, the default forwarding class or sub-class will be associated with the packet. Unless overridden by an explicit forwarding class classification rule, all packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class. Optionally, the default ingress enqueueing priority for the traffic can be overridden as well.

The default forwarding class is best effort (be). The **default-fc** settings are displayed in the **show configuration** and **save** output regardless of inclusion of the **detail** keyword.

**Context** be

**Parameters** *fc-name* — Specify the forwarding class name for the queue. The value given for *fc-name* must be one of the predefined forwarding classes in the system.

The sub-class-name parameter is optional and used with the fc-name parameter to define a preexisting sub-class. The fc-name and sub-class-name parameters must be separated by a period (dot). If sub-class-name does not exist in the context of fc -name, an error will occur. If sub-class-name is removed using the no fc fc-name.sub-class-name force command, the default-fc command will automatically drop the sub-class-name and only use fc-name (the parent forwarding class for the sub-class) as the forwarding class.

**Values** fc: class[.sub-class]  
class: be, l2, af, l1, h2, ef, h1, nc  
sub-class: 29 characters max

**Default** None (Each sub-class-name must be explicitly defined)

## default-priority

<b>Syntax</b>	<b>default-priority</b> { <b>high</b>   <b>low</b> }
<b>Context</b>	config>qos>sap-ingress
<b>Description</b>	This command configures the default enqueueing priority for all packets received on an ingress SAP using this policy. To change the default priority for the policy, the <b>fc-name</b> must be defined whether it is being changed or not.
<b>Default</b>	low
<b>Parameters</b>	<b>high</b> — Setting the enqueueing parameter to <b>high</b> for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing, once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost. <b>low</b> — Setting the enqueueing parameter to <b>low</b> for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing, once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

## fc

<b>Syntax</b>	[ <b>no</b> ] <b>fc</b> <i>fc-name</i>
<b>Context</b>	config>qos>sap-ingress
<b>Description</b>	<p>The <b>fc</b> command creates a class or sub-class instance of the forwarding class <i>fc-name</i>. Once the <i>fc-name</i> is created, classification actions can be applied and the sub-class can be used in match classification criteria. Attempting to use an undefined sub-class in a classification command will result in an execution error and the command will fail.</p> <p>The <b>no</b> form of the command removes all the explicit queue mappings for <i>fc-name</i> forwarding types. The queue mappings revert to the default queues for <i>fc-name</i>. To successfully remove a sub-class, all associations with the sub-class in the classification commands within the policy must first be removed or diverted to another forwarding class or sub-class.</p>
<b>Parameters</b>	<p><i>fc-name</i> — The parameter sub-class-name is optional and must be defined using a dot separated notation with a preceding valid system-wide forwarding class name. Creating a sub-class follows normal naming conventions. Up to sixteen ASCII characters may be used. If the same sub-name is used with two or more forwarding class names, each is considered a different instance of sub-class. A sub-class must always be specified with its preceding forwarding class name. When a forwarding class is created or specified without the optional sub-class, the parent forwarding class is assumed.</p> <p>Within the SAP ingress QoS policy, up to 56 sub classes may be created. Each of the 56 sub-classes may be created within any of the eight parental forwarding classes. Once the limit of 56 is reached, any further sub-class creations will fail and the sub-class will not exist.</p> <p>Successfully creating a sub-class places the CLI within the context of the sub-class for further sub-class parameter definitions. Within the sub-class context, commands may be executed that define sub-class</p>





SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

- Parameters**
- policer-id* — When the forwarding class **policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-ingress** QoS policy.
    - Values** 1—32
    - Default** None
  - fp-redirect-group** — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

## broadcast-policer

**Syntax** **broadcast-policer** *policer-id* [**fp-redirect-group**]  
**no broadcast-policer**

**Context** config>qos>sap-ingress>fc

**Description** Within a **sap-ingress** QoS policy forwarding class context, the **broadcast-policer** command is used to map packets that match the forwarding class and are considered broadcast in nature to the specified *policer-id*. The specified *policer-id* must already exist within the **sap-ingress** QoS policy. While the system is determining the forwarding class of a packet, it is also looking up its forwarding destination based on the ingress service type and the service instance forwarding records. If the service type is VPLS and the destination MAC address is the broadcast address (ff:ff:ff:ff:ff:ff), the packet is classified into the broadcast forwarding type.

Broadcast forwarding type packets are mapped to either an ingress multipoint queue (using the **broadcast queue-id** or **broadcast queue-id group ingress-queue-group** commands) or an ingress policer (**broadcast-policer** *policer-id*). The **broadcast** and **broadcast-policer** commands within the forwarding class context are mutually exclusive. By default, the broadcast forwarding type is mapped to the SAP ingress default multipoint queue. If the **broadcast-policer** *policer-id* command is executed, any previous policer mapping or queue mapping for the broadcast forwarding type within the forwarding class is overridden if the policer mapping is successful.

A policer defined within the **sap-ingress** policy is not actually created on an ingress SAP or a subscriber using an **sla-profile** where the policy is applied until at least one forwarding type (unicast, broadcast, unknown or multicast) from one of the forwarding classes is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or ingress policing is not supported on the port associated with the SAP or subscriber, the initial forwarding class forwarding type mapping will fail.

The **broadcast-policer** command is ignored for instances of the policer applied to SAPs or subscribers where broadcast packets are not supported.

When the broadcast forwarding type within a forwarding class is mapped to a policer, the broadcast packets classified to the sub-classes within the forwarding class are also mapped to the policer.

The **no** form of this command is used to restore the mapping of the broadcast forwarding type within the forwarding class to the default multipoint queue. If all forwarding class forwarding types had been removed from the default multipoint queue, the queue will not exist on the SAPs or subscribers associated with the

QoS policy and the **no broadcast-policer** command will cause the system to attempt to create the default multipoint queue on each object. If the system cannot create the queue on each instance, the **no broadcast-policer** command will fail and the broadcast forwarding type within the forwarding class will continue its mapping to the existing policer-id. If the **no broadcast-policer** command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

- Parameters** *policer-id* — When the forwarding class **broadcast-policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the sap-ingress QoS policy.
- Values** 1—32
- Default** None
- fp-redirect-group** — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

## multicast-policer

**Syntax** **multicast-policer** *policer-id* [**fp-redirect-group**]  
**no multicast-policer**

**Context** config>qos>sap-ingress>fc

**Description** Within a **sap-ingress** QoS policy forwarding class context, the **multicast-policer** command is used to map packets that match the forwarding class and are considered multicast in nature to the specified policer-id. The specified policer-id must already exist within the **sap-ingress** QoS policy. While the system is determining the forwarding class of a packet, it is also looking up its forwarding destination based on the ingress service type and the service instance forwarding records. Two basic types of services support multicast packets; routed services (IES and VPRN) and L2 multipoint services (VPLS, I-VPLS and B-VPLS). For the routed service types, a multicast packet is destined to an IPv4 or IPv6 multicast address. For the L2 multipoint services, a multicast packet is a packet destined to a multicast MAC address (multicast bit set in the destination MAC address but not the ff:ff:ff:ff:ff:ff broadcast address). The VPLS services also support two other multipoint forwarding types (broadcast and unknown) which are considered separate from the multicast forwarding type.

If ingress forwarding logic has resolved a packet to the multicast forwarding type within the forwarding class, it will be mapped to either an ingress multipoint queue (using the **multicast queue-id** or **multicast queue-id group ingress-queue-group** commands) or an ingress policer (**multicast-policer policer-id**). The **multicast** and **multicast-policer** commands within the forwarding class context are mutually exclusive. By default, the multicast forwarding type is mapped to the SAP ingress default multipoint queue. If the **multicast-policer policer-id** command is executed, any previous policer mapping or queue mapping for the multicast forwarding type within the forwarding class is overridden if the policer mapping is successful.

A policer defined within the **sap-ingress** policy is not actually created on an ingress SAP or a subscriber using an **sla-profile** where the policy is applied until at least one forwarding type (unicast, broadcast, unknown or multicast) from one of the forwarding classes is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or ingress policing is not supported on the port associated with the SAP or subscriber, the initial forwarding class forwarding type mapping will fail.

The multicast-policer command is ignored for instances of the policer applied to SAPs subscribers where broadcast packets are not supported.

When the multicast forwarding type within a forwarding class is mapped to a policer, the multicast packets classified to the sub-classes within the forwarding class are also mapped to the policer.

The **no** form of this command is used to restore the mapping of the multicast forwarding type within the forwarding class to the default multipoint queue. If all forwarding class forwarding types had been removed from the default multipoint queue, the queue will not exist on the SAPs subscribers associated with the QoS policy and the no multicast-policer command will cause the system to attempt to create the default multipoint queue on each object. If the system cannot create the queue on each instance, the no multicast-policer command will fail and the multicast forwarding type within the forwarding class will continue its mapping to the existing policer-id. If the no multicast-policer command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

<b>Parameters</b>	<i>policer-id</i> — When the forwarding class <b>multicast-policer</b> command is executed, a valid <i>policer-id</i> must be specified. The parameter <i>policer-id</i> references a <i>policer-id</i> that has already been created within the <b>sap-ingress</b> QoS policy.
<b>Values</b>	1—32
<b>Default</b>	None
	<b>fp-redirect-group</b> — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

## unknown-policer

<b>Syntax</b>	<b>unknown-policer</b> <i>policer-id</i> [ <b>fp-redirect-group</b> ] <b>no unknown-policer</b>
<b>Context</b>	config>qos>sap-ingress>fc
<b>Description</b>	<p>Within a <b>sap-ingress</b> QoS policy forwarding class context, the <b>unknown-policer</b> command is used to map packets that match the forwarding class and are considered unknown in nature to the specified policer-id. The specified policer-id must already exist within the <b>sap-ingress</b> QoS policy. While the system is determining the forwarding class of a packet, it is also looking up its forwarding destination based on the ingress service type and the service instance forwarding records. If the service type is VPLS and the destination MAC address is unicast but the MAC has not been learned and populated within the VPLS services FDB, the packet is classified into the unknown forwarding type.</p> <p>Unknown forwarding type packets are mapped to either an ingress multipoint queue (using the <b>unknown queue-id</b> or <b>unknown queue-id group ingress-queue-group</b> commands) or an ingress policer (<b>unknown-policer policer-id</b>). The <b>unknown</b> and <b>unknown-policer</b> commands within the forwarding class context are mutually exclusive. By default, the unknown forwarding type is mapped to the SAP ingress default multipoint queue. If the <b>unknown-policer policer-id</b> command is executed, any previous policer mapping or queue mapping for the unknown forwarding type within the forwarding class is overridden if the policer mapping is successful.</p> <p>A policer defined within the <b>sap-ingress</b> policy is not actually created on an ingress SAP or a subscriber using an <b>sla-profile</b> where the policy is applied until at least one forwarding type (unicast, broadcast,</p>

unknown or multicast) from one of the forwarding classes is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or ingress policing is not supported on the port associated with the SAP or subscriber, the initial forwarding class forwarding type mapping will fail.

The **unknown-policer** command is ignored for instances of the policer applied to SAPs or subscribers where unknown packets are not supported.

When the unknown forwarding type within a forwarding class is mapped to a policer, the unknown packets classified to the sub-classes within the forwarding class are also mapped to the policer.

The **no** form of this command is used to restore the mapping of the unknown forwarding type within the forwarding class to the default multipoint queue. If all forwarding class forwarding types had been removed from the default multipoint queue, the queue will not exist on the SAPs or subscriber associated with the QoS policy and the no broadcast-policer command will cause the system to attempt to create the default multipoint queue on each object. If the system cannot create the queue on each instance, the no unknown-policer command will fail and the unknown forwarding type within the forwarding class will continue its mapping to the existing policer-id. If the no unknown-policer command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscriber will be lost.

**Parameters** *policer-id* — When the forwarding class **unknown-policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-ingress** QoS policy.

**Values** 1—32

**Default** None

**fp-redirect-group** — Redirects a forwarding class to a forwarding plane queue-group as specified in a SAP QoS policy.

## dot1p

**Syntax** **dot1p** *dot1p-value* [**fc** *fc-name*] [**profile** {**in** | **out** | **use-de**}]  
**no dot1p** *dot1p-value*

**Context** config>qos>sap-ingress

**Description** This command explicitly sets the forwarding class or sub-class or enqueueing priority when a packet is marked with a *dot1p-priority* specified. Adding a dot1p rule on the policy forces packets that match the *dot1p-priority* specified to override the forwarding class and enqueueing priority based on the parameters included in the dot1p rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing priority derived from earlier matches in the classification hierarchy.

The *dot1p-priority* is derived from the most significant three bits in the IEEE 802.1Q or IEEE 802.1P header. The three dot1p bits define 8 Class-of-Service (CoS) values commonly used to map packets to per-hop Quality-of-Service (QoS) behavior.

The **no** form of this command removes the explicit dot1p classification rule from the SAP ingress policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs using the policy.

**Parameters** *dot1p-value* — This value is a required parameter that specifies the unique IEEE 802.1P value that will match the dot1p rule. If the command is executed multiple times with the same *dot1p-value*, the previous forwarding class and enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.

A maximum of eight dot1p rules are allowed on a single policy.

**Values** 0 — 7

**fc *fc-name*** — The value given for the *fc-name* parameter must be one of the predefined forwarding classes in the system. Specifying the *fc-name* is optional. When a packet matches the rule, the forwarding class is only overridden when the *fc fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

**Default** None

**priority {*in*|*out*|*use-de*}** — All frames that are assigned to this forwarding class will be considered in or out of profile based on this command or to use the default. In case of congestion, the in- profile frames are preferentially queued over the out-of-profile frames

**Values** **in** — All frames are treated as in-profile.

**out** — All frames are treated as out of profile.

**use-de** — The profile of all frames is set according to the DEI bit.

## dscp

**Syntax** **dscp** *dscp-name* **fc** *fc-name* [**profile** {*in* | *out*}]  
**no dscp** *dscp-name*

**Context** config>qos>sap-ingress

**Description** This command explicitly sets the forwarding class or subclass or enqueueing priority when a packet is marked with the DiffServ Code Point (DSCP) value contained in the *dscp-name*. Adding a DSCP rule on the policy forces packets that match the DSCP value specified to override the forwarding class and enqueueing priority based on the parameters included in the DSCP rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing priority derived from earlier matches in the classification hierarchy.

The DSCP value (referred to by *dscp-name*) is derived from the most significant six bits in the IP header ToS byte field (DSCP bits). The six DSCP bits define 64 DSCP values used to map packets to per-hop Quality-of-Service (QoS) behavior. The most significant three bits in the IP header ToS byte field are also commonly used in a more traditional manner to specify an IP precedence value, causing an overlap between the precedence space and the DSCP space. Both IP precedence and DSCP classification rules are supported. DSCP rules have a higher match priority than IP precedence rules and where a *dscp-name* DSCP value overlaps an *ip-prec-value*, the DSCP rule takes precedence.

Multiple *dscp-name* entries in the DSCP master table can point to the same DSCP bits value. When two DSCP rules specify different *dscp-name*, but the names contain the same DSCP bits value, the previous

DSCP name is removed and replaced by the latest entry.

The **no** form of the command removes the explicit DSCP classification rule from the SAP ingress policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs using the policy.

### Parameters

*dscp-name* — The *dscp-name* is a required parameter that specifies the unique IP header ToS byte DSCP bits value that will match the DSCP rule. If the command is executed multiple times with the same *dscp-name* or a *dscp-name* that contains the same DSCP bit value, the previous forwarding class and enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.

A maximum of 64 DSCP rules are allowed on a single policy.

The specified name must exist as a *dscp-name*. SR OS software provides names for the well-known code points. The system-defined names available are as follows. The system-defined names are referenced lower case letters or exactly as shown in the first column in the following tables.

User-defined name to code point value associations can be added using the **dscp-name** *dscp-name* *dscp-value* command.

<b>dscp-name</b>	<b>dscp-value Decimal</b>	<b>dscp-value Hexadecimal</b>	<b>dscp-value Binary</b>
nc1	48	0x30	0b110000
nc2	56	0x38	0b111000
ef	46	0x2e	0b101110
af41	34	0x22	0b100010
af42	36	0x24	0b100100
af43	38	0x26	0b100110
af31	26	0x1a	0b011010
af32	28	0x1c	0b011100
af33	30	0x1d	0b011110
af21	18	0x12	0b010010
af22	20	0x14	0b010100
af23	22	0x16	0b010110
af11	10	0x0a	0b001010
af12	12	0x0c	0b001100
af13	14	0x0e	0b001110
default	0	0x00	0b000000



**high** — This parameter is used in conjunction with the **priority** parameter. Setting the enqueueing parameter to **high** for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. Once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

**Default** low priority

**low** — This parameter is used in conjunction with the **priority** parameter. Setting the enqueueing parameter to **low** for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing, once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

**Default** Inherit (When **priority** is not defined, the rule preserves the previous enqueueing priority of the packet.)

## dscp

**Syntax** **dscp** *dscp-name* [**fc** *fc-name*] [**hsmda-counter-override** *counter-id*] [**profile** {**in** | **out**}]  
**no dscp** *dscp-name*

**Context** config>qos>sap-egress

**Description** This command defines a specific IP Differentiated Services Code Point (DSCP) value that must be matched to perform the associated reclassification actions. If an egress packet on the SAP matches the specified IP DSCP value, the forwarding class, profile or HSMMDA egress queue accounting behavior may be overridden. By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions. The default behavior for HSMMDA queue accounting is to use the counters associated with the queue to which the packet is mapped. Matching a DHCP based reclassification rule will override all IP precedence based reclassification rule actions.

The IP DSCP bits used to match against dscp reclassification rules come from the Type of Service (ToS) field within the IPv4 header or the Traffic Class field from the IPv6 header. If the packet does not have an IP header, dscp based matching is not performed.

The reclassification actions from a dscp reclassification rule may be overridden by a an IP flow match event.

The **fc** keyword is optional. When specified, the egress classification rule will overwrite the forwarding class derived from ingress. The new forwarding class is used for egress remarking and queue mapping decisions. If an ip-criteria match occurs after the prec match, the new forwarding class may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new fc, the fc from the dscp match will be used.

The **profile** keyword is optional. When specified, the egress classification rule will overwrite the profile of the packet derived from ingress. The new profile value is used for egress remarking and queue congestion behavior. If an ip-criteria match occurs after the dscp match, the new profile may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new profile, the profile from the dscp match will be used.

The hsmmda-counter-override keyword is optional. When specified and the egress SAP is created on an HSMMDA, the egress classification rule will override the default queue accounting function for the packet. By default, the HSMMDA uses each queues default queue counters for packets mapped to the queue. The hsmmda-counter-override keyword is used to map the packet to an explicit exception counter. One of eight counters may be used. When the packet is mapped to an exception counter, the packet will not increment the



queues discard or forwarding counters, instead the exception discard and forwarding counters will be used. The dscp based counter override decision may be overwritten by an ip-criteria reclassification rule match if the higher priority classification rule has an hsmda-counter-override action defined.

The **no** form of the command removes the reclassification rule from the SAP egress QoS policy.

## Parameters

*dscp-value*: — The dscp-value parameter is required when defining a dscp reclassification rule. The value must be specified as an integer from 0 to 63.

**Values** 0 — 63

*fc fc-name*: — The **fc** reclassification action is optional. When specified, packets matching the IP DSCP value will be explicitly reclassified to the forwarding class specified as fc-name regardless of the ingress classification decision. The explicit forwarding class reclassification may be overwritten by an ip-criteria reclassification match. The fc-name defined must be one of the eight forwarding classes supported by the system. To remove the forwarding class reclassification action for the specified dscp-value, the dscp command must be re-executed without the fc reclassification action defined.

**Values** be, l1, af, l2, h1, ef, h2 or nc

**profile {in | out}** — The profile reclassification action is optional. When specified, packets matching the IP DSCP value will be explicitly reclassified to either in-profile or out-of-profile regardless of the ingress profiling decision. The explicit profile reclassification may be overwritten by an ip-criteria reclassification match. To remove the profile reclassification action for the specified dscp-value, the dscp command must be re-executed without the profile reclassification action defined.

**in** — The **in** parameter is mutually exclusive to the out parameter following the profile reclassification action keyword. Either in or out must be specified when the profile keyword is present. When in is specified, any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane. This value may be overwritten by an explicit profile action in an ip-criteria reclassification match.

**out**: — The **out** parameter is mutually exclusive to the in parameter following the profile reclassification action keyword. Either in or out must be specified when the profile keyword is present. When out is specified, any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane. This value may be overwritten by an explicit profile action in an ip-criteria reclassification match.

**hsmda-counter-override counter-id** — The hsmda-counter-override reclassification action is optional and only has significance on SAPs which are created on an HSMDA. When specified, packets matching the IP precedence value will be mapped to the defined HSMDA exception counter-id for the packets queue group. The default behavior is to use the default counter on the queue group for the queue to which the packet is mapped. The hsmda-counter-override action may be overwritten by an ip-criteria reclassification rule match. The specified counter-id must be specified as an integer between 1 and 8. To remove the ESMDA exception counter reclassification action for the specified dscp-value, the dscp command must be re-executed without the hsmda-counter-override reclassification action defined.

**Values** 1 — 8

## ethernet-ctag

**Syntax** **ethernet-ctag**  
**no ethernet-ctag**

**Context** config>qos>sap-egress

**Description** This command specifies that the top customer tag should be used for egress reclassification based on dot1p criteria. This command applies to all dot1p criteria configured in a given SAP egress QoS policy.

The no form of this command means that a service delimiting tag will be used for egress reclassification based on dot1p criteria.

**Default** noethernet-ctag

## ip-criteria

**Syntax** **[no] ip-criteria**

**Context** config>qos>sap-ingress

**Description** IP criteria-based SAP ingress policies are used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.

This command is used to enter the context to create or edit policy entries that specify IP criteria such as IP quintuple lookup or DiffServ code point.

7750 SR OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. Once IP criteria entries are removed from a SAP ingress policy, the IP criteria is removed from all services where that policy is applied.

## ip-criteria

**Syntax** **[no] ip-criteria**

**Context** config>qos>sap-egress

**Description** IP criteria-based SAP egress policies are used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.

This command is used to enter the context to create or edit policy entries that specify IP criteria such as IP quintuple lookup or DiffServ code point.

7750 SR OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. Once IP criteria entries are removed from a SAP ingress policy, the IP criteria is removed from all services where that policy is applied.

## ipv6-criteria

<b>Syntax</b>	<b>[no] ipv6-criteria</b>
<b>Context</b>	config>qos>sap-ingress
<b>Description</b>	<p>IPv6 criteria-based SAP ingress policies are used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.</p> <p>This command is used to enter the node to create or edit policy entries that specify IPv6 criteria such as IP quintuple lookup or DiffServ code point.</p> <p>7750 SR OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.</p> <p>The <b>no</b> form of this command deletes all the entries specified under this node. Once ipv6-criteria entries are removed from a SAP ingress policy, the ipv6-criteria is removed from all services where that policy is applied.</p>
<b>Parameters</b>	<i>policy-id</i> — The policy-id that uniquely identifies the policy.
<b>Values</b>	1 — 65535

## lsp-exp

<b>Syntax</b>	<b>lsp-exp</b> <i>lsp-exp-value</i> [ <b>fc</b> <i>fc-name</i> ] [ <b>priority</b> { <b>low</b>   <b>high</b> }] [ <b>hsmdda-counter-override</b> <i>counter-id</i> ] <b>no</b> <b>lsp-exp</b> <i>lsp-exp-value</i>
<b>Context</b>	config>qos>sap-ingress
<b>Description</b>	<p>This command explicitly sets the forwarding class or sub-class enqueueing priority when a packet is marked with a MPLS EXP bits specified. Adding a lsp-exp rule on the policy forces packets that match the MPLS LSP EXP specified to override the forwarding class and enqueueing priority based on the parameters included in the lsp-exp rule. When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing priority derived from earlier matches in the classification hierarchy..</p> <p>The <i>lsp-exp-value</i> is derived from the MPLS LSP EXP bits of the top label.</p> <p>Multiple commands can be entered to define the association of some or all eight LSP EX bit values to the forwarding class.</p> <p>The <b>no</b> form of this command removes the explicit lsp-exp classification rule from the SAP ingress policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs using the policy.</p> <p>This command applies to Ethernet Layer 2 SAPs only.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>lsp-exp-value</i> — This value is a required parameter that specifies the unique MPLS LSP EXP value that will match the lsp-exp rule. If the command is executed multiple times with the same lsp-exp-value, the

previous forwarding class and enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.

A maximum of eight lsp-exp rules are allowed on a single policy.

**Values** 0 — 7

**fc** *fc-name* — The value given for the fc-name parameter must be one of the predefined forwarding classes in the system. Specifying the fc-name is optional. When a packet matches the rule the forwarding class is only overridden when the fc fc-name parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

The sub-class-name parameter is optional and used with the fc-name parameter to define a preexisting sub-class. The fc-name and sub-class-name parameters must be separated by a period (dot). If sub-class-name does not exist in the context of fc -name, an error will occur.

**Values** fc: class[.sub-class]  
class: be, l2, af, l1, h2, ef, h1, nc  
sub-class: 29 characters max

**Default** None (Each sub-class-name must be explicitly defined)

**priority** — The priority parameter is used to override the default enqueueing priority for all packets received on an ingress SAP using this policy that match this rule. Specifying the priority is optional. When a packet matches the rule, the enqueueing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueueing priority is inherited based on previous rule matches.

**high** — The high parameter is used in conjunction with the priority parameter. Setting the enqueueing parameter to high for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing, once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

**low** — The low parameter is used in conjunction with the priority parameter. Setting the enqueueing parameter to low for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing, once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

**Default** No override.

**hsmda-counter-override** *counter-id* — The hsmda-counter-override reclassification action is optional and only has significance on SAPs which are created on an HSMDA. When specified, packets matching the MPLS EXP value will be mapped to the defined HSMDA exception counter-id for the packets queue group. The default behavior is to use the default counter on the queue group for the queue to which the packet is mapped. The specified counter-id must be specified as an integer between 1 and 8. To remove the HSMDA exception counter reclassification action for the specified lsp-exp-value, the lsp-exp command must be re-executed without the hsmda-counter-override reclassification action defined.

**Values** 1 — 8

## mac-criteria

**Syntax** [no] **mac-criteria**

**Context** config>qos>sap-ingress

**Description** The **mac-criteria** based SAP ingress policies are used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.

This command is used to enter the node to create or edit policy entries that specify MAC criteria.

Router implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. Once mac-criteria entries are removed from a SAP ingress policy, the mac-criteria is removed from all services where that policy is applied.

## policer

**Syntax** **policer** *policer-id* [**create**]  
**no policer** *policer-id*

**Context** config>qos>sap-ingress

**Description** This command is used in the sap-ingress and sap-egress QoS policies to create, modify or delete a policer. Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. The sap-ingress policy may have up to 32 policers (numbered 1 through 32) may be defined while the sap-egress QoS policy supports a maximum of 8 (numbered 1 through 8). While a policer may be defined within a QoS policy, it is not actually created on SAPs or subscribers associated with the policy until a forwarding class is mapped to the policer's ID.

All policers must be created within the QoS policies. A default policer is not created when a sap-ingress or sap-egress QoS policy is created.

Once a policer is created, the policer's metering rate and profiling rates may be defined as well as the policer's maximum and committed burst sizes (MBS and CBS respectively). Unlike queues which have dedicated counters, policers allow various stat-mode settings that define the counters that will be associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet based on a defined number of bytes.

Once a policer is created, it cannot be deleted from the QoS policy unless any forwarding classes that are mapped to the policer are first moved to other policers or queues.

The system will allow a policer to be created on a SAP QoS policy regardless of the ability to support policers on objects where the policy is currently applied. The system only scans the current objects for policer support and sufficient resources to create the policer when a forwarding class is first mapped to the

policer ID. If the policer cannot be created due to one or more instances of the policy not supporting policing or having insufficient resources to create the policer, the forwarding class mapping will fail.

The **no** form of this command is used to delete a policer from a sap-ingress or sap-egress QoS policy. The specified policer cannot currently have any forwarding class mappings for the removal of the policer to succeed. It is not necessary to actually delete the policer ID for the policer instances to be removed from SAPs or subscriber associated with the QoS policy once all forwarding classes have been moved away from the policer. It is automatically deleted from each policing instance although it still appears in the QoS policy.

**Parameters** *policer-id* — The *policer-id* must be specified when executing the policer command. If the specified ID already exists, the system enters that policer's context to allow the policer's parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the policer ID will be created (depending on the system's current create keyword requirements which may require the create keyword to actually add the new policer ID to the QoS policy) and the system will enter that new policer's context for possible parameter modification.

**Values** 1—32

## description

**Syntax** **description** *description string*  
**no description**

**Context** config>qos>sap-ingress>policer

**Description** The **description** command is used to define an informational ASCII string associated with the policer control policy. The string value can be defined or changed at any time once the policy exists. The **no** form of this command is used to remove an explicit description string from the policer.

**Default** **no description**

**Parameters** *description string* — The *description-string* parameter defines the ASCII description string for the policer control policy. The description string can be up to 80 characters. If the string contains spaces, it must be placed within beginning and ending double quotation marks. Beginning and ending quotation marks are not considered part of the description string. Only printable ASCII characters are allowed in the string. The string does not need to be unique and may be repeated in the descriptions for other policer control policies or other objects. If the command is executed without the description-string present, any existing description string will be unaffected.

**Default** None

## adv-config-policy

<b>Syntax</b>	<b>[no] adv-config-policy</b> <i>policy-name</i>
<b>Context</b>	config>qos>sap-ingress>policer config>qos>sap-egress>policer
<b>Description</b>	<p>This command specifies the advanced QoS policy. The advanced QoS policy contains only queue and policer child control parameters within a child-control node.</p> <p>Once a policy is created, it may be applied to a queue or policer defined within a <b>sap-egress</b> or <b>sap-ingress</b> QoS policy. It may also be applied to a queue or policer defined within an ingress or <b>egress queue-group</b> template. When a policy is currently associated with a QoS policy or template, the policy may be modified but not deleted (even in the event that the QoS policy or template is not in use).</p> <p>The <b>no</b> form of this command removes the specified advanced policy.</p>
<b>Default</b>	None
<b>Parameters</b>	<i>policy-name</i> — The name of the advanced QoS policy.
<b>Values</b>	Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## adaptation-rule

<b>Syntax</b>	<b>adaptation-rule</b> [ <b>pir</b> { <b>max</b>   <b>min</b>   <b>closest</b> }] [ <b>cir</b> { <b>max</b>   <b>min</b>   <b>closest</b> }] <b>no adaptation-rule</b>
<b>Context</b>	config>qos>sap-ingress>policer
<b>Description</b>	<p>This command is used to define how the policer's configuration parameters are translated into the underlying hardware capabilities used to implement each policer instance. For instance, the configured rates for the policer need to be mapped to the timers and decrement granularity used by the hardware's leaky bucket functions that actually perform the traffic metering. If a rate is defined that cannot be exactly matched by the hardware, the adaptation-rule setting provides guidance for which hardware rate should be used.</p> <p>The <b>max</b> keyword tells the system that the defined rate is the maximum rate that should be given to the policer. If the hardware cannot exactly match the given rate, the next lowest hardware supported rate is used.</p> <p>The <b>min</b> keyword tells the system that the defined rate is the minimum rate that should be given to the policer. If the hardware cannot exactly match the given rate, the next highest hardware supported rate is used.</p> <p>The <b>closest</b> keyword tells the system that the defined rate is the target rate for the policer. If the hardware cannot exactly match the given rate, the system will use the closest hardware supported rate compared to the target rate.</p> <p>The hardware also needs to adapt the given mbs and cbs values into the PIR bucket violate threshold (discard) and the CIR bucket exceed threshold (out-of-profile). The hardware may not have an exact</p>

threshold match which it can use. In R8.0, the system treats the mbs and cbs values as minimum threshold values.

The **no** form of this command is used to return the policer's metering and profiling hardware adaptation rules to closest.

## Parameters

**pir** {**max** | **min** | **closest**} — When the optional **pir** parameter is specified, the **max**, **min** or **closest** keyword qualifier must follow.

**max** — The **max** keyword is used to inform the system that the metering rate defined for the policer is the maximum allowed rate. The system will choose a hardware supported rate that is closest but not exceeding the specified rate.

**min** — The **min** keyword is used to inform the system that the metering rate defined for the policer is the minimum allowed rate. The system will choose a hardware supported rate that is closest but not lower than the specified rate.

**closest** — The **closest** keyword is used to inform the system that the metering rate defined for the policer is the target rate. The system will choose a hardware supported rate that is closest to the specified rate.

**Default**      closest

**cir** {**max** | **min** | **closest**} — When the optional **cir** parameter is specified, the **max**, **min** or **closest** keyword qualifier must follow.

**max** — The **max** keyword is used to inform the system that the profiling rate defined for the policer is the maximum allowed rate. The system will choose a hardware supported rate that is closest but not exceeding the specified rate.

**min** — The **min** keyword is used to inform the system that the profiling rate defined for the policer is the minimum allowed rate. The system will choose a hardware supported rate that is closest but not lower than the specified rate.

**closest** — The **closest** keyword is used to inform the system that the profiling rate defined for the policer is the target rate. The system will choose a hardware supported rate that is closest to the specified rate.

**Default**      closest

## Values

## cbs

**Syntax**      **cbs** {*size* [**bytes** | **kilobytes**] | **default**}  
**no cbs**

**Context**      config>qos>sap-ingress>policer

**Description**      This command is used to configure the policer's CIR leaky bucket's exceed threshold. The CIR bucket's exceed threshold represents the committed burst tolerance allowed by the policer. If the policer's forwarding rate is equal to or less than the policer's defined CIR, the CIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the forwarding rate increases beyond the profiling rate, the amount of data allowed to be in-profile above the rate is capped by the threshold.

The policer's **cbs** size defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.



The no form of this command returns the policer to its default CBS size.

**Default** 64 kilobytes when CIR = **max**, otherwise 10ms volume of traffic for a configured non zero/non max CIR.

**Parameters** *size* [**bytes** | **kilobytes**] — The *size* parameter is required when specifying **cbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

**byte** — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

**kilobyte** — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

**Values** 1 — 4194304

**Default** **kilobyte**

## high-prio-only

**Syntax** **high-prio-only** *percent-of-mbs*  
**no high-prio-only**

**Context** config>qos>sap-ingress>policer

**Description** This command is used to configure the percentage of the policer's PIR leaky bucket's MBS (maximum burst size) that is reserved for high priority traffic. While the **mbs** value defines the policer's high priority violate threshold, the percentage value defined is applied to the **mbs** value to derive the bucket's low priority violate threshold. See the **mbs** command details for information on which types of traffic is associated with each violate threshold.

**Default** **high-prio-only 10**

**Parameters** *percent-of-mbs* — The *percent-of-mbs* parameter is required when specifying **high-prio-only** and is expressed as a percentage with granularity of 1,000th of a percent.

**Values** 0—100

**Default** 10

## mbs

**Syntax** **mbs** {*size* [**bytes** | **kilobytes**] | **default**}  
**no mbs**

**Context** config>qos>sap-ingress>policer

**Description** This command is used to configure the policer's PIR leaky bucket's high priority violate threshold. The **high-prio-only** command is applied to the MBS value to derive the bucket's low priority violate threshold. For ingress, trusted in-profile packets and un-trusted high priority packets use the policer's high priority violate threshold while trusted out-of-profile and un-trusted low priority packets use the policer's low

priority violate threshold. At egress, in-profile packets use the policer's high priority violate threshold and out-of-profile packets use the policer's low priority violate threshold.

The PIR bucket's violate threshold represent the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by **high-prio-only** is available for the higher priority traffic.

The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied.

The no form of this command returns the policer to its default MBS size.

**Default** 64 kilobytes when PIR = **max**, otherwise 10ms volume of traffic for a configured non zero/non max PIR.

**Parameters** *size* [**bytes** | **kilobytes**] — The *size* parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

**byte** — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

**kilobyte** — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

**Values** 1—4194304

**Default** kilobyte

## packet-byte-offset

**Syntax** **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*}  
**no packet-byte-offset**

**Context** config>qos>sap-ingress>policer

**Description** This command is used to modify the size of each packet handled by the policer by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed. The **packet-byte-offset** command is meant to be an arbitrary mechanism that can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the policing metering and profiling throughput is affected by the offset as well as the stats associated with the policer.

When child policers are adding to or subtracting from the size of each packet, the parent policer's **min-thresh-separation** value should also need to be modified by the same amount.

The policer's **packet-byte-offset** defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** version of this command is used to remove per packet size modifications from the policer.

**Parameters** **add** *bytes* — The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined the corresponding bytes parameter specifies the number of bytes that

is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

**Values** 0 — 31

**Default** None

**subtract bytes** — The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When *b* is defined the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

**Values** 1 — 32

**Default** None

## parent

**Syntax** **parent** {*root* | *arbiter-name*} [**level** *level*] [**weight** *weight-within-level*]  
**no parent**

**Context** config>qos>sap-ingress>policer

**Description** This command is used to create a child to parent mapping between each instance of the policer and either the **root** arbiter or a specific tiered **arbiter** on the object where the policy is applied. Defining a parent association for the policer causes the policer to compete for the parent policer's available bandwidth with other child policers mapped to the policer control hierarchy.

Policer control hierarchies may be created on SAPs or on a subscriber context. To create a policer control hierarchy on an ingress or egress SAP context, a **policer-control-policy** must be applied to the SAP. Once applied, the system will create a parent policer that is bandwidth limited by the policy's **max-rate** value under the root arbiter. The root arbiter in the policy also provides the information used to determine the various priority level discard-unfair and discard-all thresholds. Besides the root arbiter, the policy may also contain user defined tiered arbiters that provide arbitrary bandwidth control for subsets of child policers that are either directly or indirectly parented by the arbiter.

When the QoS policy containing the policer with a **parent** mapping to an arbiter name exists on the SAP, the system will scan the available arbiters on the SAP. If an arbiter exists with the appropriate name, the policer to arbiter association is created. If the specified arbiter does not exist either because a **policer-control-policy** is not currently applied to the SAP or the arbiter name does not exist within the applied policy, the policer is placed in an orphan state. Orphan policers operate as if they are not parented and are not subject to any bandwidth constraints other than their own PIR. When a policer enters the orphan state, it is flagged as operationally degraded due to the fact that it is not operating as intended and a trap is generated. Whenever a **policer-control-policy** is added to the SAP or the existing policy is modified, the SAP's policer's parenting configurations must be reevaluated. If an orphan policer becomes parented, the degraded flag should be cleared and a resulting trap should be generated.

For subscribers, the policer control hierarchy is created through the **policer-control-policy** applied to the **sub-profile** used by the subscriber. A unique policer control hierarchy is created for each subscriber associated with the **sub-profile**. The QoS policy containing the policer with the parenting command comes

into play through the subscriber **sla-profile** which references the QoS policy. The combining of the **sub-profile** and the **sla-profile** at the subscriber level provides the system with the proper information to create the policer control hierarchy instance for the subscriber.

Executing the **parent** command will fail if:

- The policer's stat-mode in the QoS policy is set to no-stats
- A stat-mode no-stats override exists on an instance of the policer on a SAP or subscriber context

A policer with a **parent** command applied cannot be configured with **stat-mode no-stats** in either the QoS policy or as an override on an instance of the policer

Once a policer is successfully parented to an arbiter, the **parent** commands **level** and **weight** parameters are used to determine at what priority level and at which weight in the priority level that the child policer competes with other children (policers or other arbiters) for bandwidth.

The **no** form of this command is used to remove the parent association from all instances of the policer.

#### Parameters

{**root** | *arbiter-name*} — When the **parent** command is executed, either the keyword **root** or an *arbiter-name* must be specified.

**root** — The **root** keyword specifies that the policer is intended to become a child to the **root** arbiter where an instance of the policer is created. If the **root** arbiter does not exist, the policer will be placed in the orphan state.

**Default**     **root**

*arbiter-name* — The *arbiter-name* parameter specifies that the policer is intended to become a child to one of the tiered arbiters with the given arbiter-name where an instance of the policer is created. If the specified arbiter-name does not exist, the policer will be placed in the orphan state.

**Default**     None

**weight** *weight-within-level* — The **weight** *weight-within-level* keyword and parameter are optional when executing the **parent** command. When **weight** is not specified, a default level of 1 is used in the parent arbiters priority level. When **weight** is specified, the *weight-within-level* parameter must be specified as an integer value from 1 through 100.

**Default**     1

## percent-rate

**Syntax**     **percent-rate** *pir-percent* [**cir** *cir-percent*]  
**no percent-rate**

**Context**     config>qos>sap-egress>policer  
              config>qos>sap-ingress>policer

**Description**     The **percent-rate** command within the egress queue group template and on the port queue group queue overrides enables supports a queue's shaping rate and CIR rate as a percentage of the egress port's line rate. When the rates are expressed as a percentage within the template, the actual rate used per instance of the queue group queue-id will vary based on the port speed. For example, when the same template is used to create a queue group on a 1-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit port due to the difference in port speeds. This enables the same template to be used

on multiple ports without needing to use port based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's shaping and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a queue is defined using the percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

An egress port queue group queue rate override may be expressed as either a percentage or an explicit rate independent on how the queue's template rate is expressed.

The **no** form of this command returns the queue to its default shaping rate and cir rate. When **no percent-rate** is defined within a port egress queue group queue override, the queue reverts to the defined shaping and CIR rates within the egress queue group template associated with the queue.

**Parameters** *percent-of-line-rate* — The *percent-of-line-rate* parameter is used to express the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

**Values** Percentage ranging from 0.01 to 100.00. The default is 100.00.

**cir** *percent-of-line-rate* — The **cir** keyword is optional and when defined the required *percent-of-line-rate* CIR parameter expresses the queue's committed scheduling rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

## profile-capped

**Syntax** [no] profile-capped

**Context** config>qos>sap-ingress>policer  
config>qos>sap-egress>policer  
config>qos>queue-group-templates>ingress>queue-group  
config>qos>queue-group-templates>ingress>queue-group

**Description** Profile capped mode enforces an overall in-profile burst limit to the CIR bucket for ingress undefined, ingress explicit in-profile, egress soft-in-profile, and egress explicit in-profile packets. The default behavior when profile-capped mode is not enabled is to ignore the CIR output state when an explicit in-profile packet is handled by an ingress or egress policer.

The profile-capped mode makes two changes:

- At egress, soft-in-profile packets (packets received from ingress as in-profile) are treated the same as explicit in-profile (unless explicitly reclassified as out-of-profile) and have an initial policer state of in-profile.

- At both ingress and egress, any packet output from the policer with a non-conforming CIR state are treated as out-of-profile (out-of-profile state is ignored for initial in-profile packets when profile capped mode is not enabled)

**Default** no profile-capped

## rate

**Syntax** `rate {max | kilobits-per-second} [cir {max | kilobits-per-second}]`  
**no rate**

**Context** config>qos>sap-ingress>policer

**Description** This command is used to configure the policer's metering and optional profiling rates. The metering rate is used by the system to configure the policer's PIR leaky bucket's decrement rate while the profiling rate configures the policer's CIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on the each packets size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches it's exceed (CIR) or violate (PIR) threshold. The **cbs**, **mbs**, and **high-prio-only** commands are used to configure the policer's PIR and CIR thresholds.

If a packet arrives at the policer while the bucket's depth is less than the threshold associated with the packet, the packet is considered to be conforming to the bucket's rate. If the bucket depth is equal to or greater than the threshold, the packet is considered to be in the exception state. For the CIR bucket, the exception state is exceeding the CIR rate while the PIR bucket's exception state is violating the PIR bucket rate. If the packet is violating the PIR, the packet is marked red and will be discarded. If the packet is not red, it may be green or yellow based on the conforming or exceeding state from the CIR bucket.

When a packet is red neither the PIR or CIR bucket depths are incremented by the packets size. When the packet is yellow the PIR bucket is incremented by the packet size, but the CIR bucket is not. When the packet is green, both the PIR and CIR buckets are incremented by the packet size. This ensures that conforming packets impact the bucket depth while exceeding or violating packets do not.

The policer's **adaptation-rule** command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.

By default, the policer's metering rate is **max** and the profiling rate is 0 Kbps (all packets out-of-profile).

The **rate** settings defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command is used to restore the default metering and profiling rate to a policer.

**Parameters** {**max** | *kilobits-per-second*} — Specifying the keyword **max** or an explicit *kilobits-per-second* parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the policer is first created, the metering rate defaults to max. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second.

**Values** **max** or 0—20,000,000

**cir** {**max** | *kilobits-per-second*} — The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit *kilobits-per-second* parameter directly following the

cir keyword is required and identifies the policer's profiling rate for the CIR leaky bucket. When the policer is first created, the profiling rate defaults to 0 Kbps. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second.

**Values** max or 0—20,000,000

## stat-mode

**Syntax** **stat-mode** {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-priority-no-cir** | **offered-limited-profile-cir** | **offered-profile-cir** | **offered-priority-cir** | **offered-total-cir** | **offered-profile-capped-cir** | **offered-limited-capped-cir** }  
**no stat mode**

**Context** config>qos>sap-ingress>policer  
config>qos>queue-group-templates>ingress>queue-group

**Description** This command is used to configure the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. An ingress policer has multiple types of offered packets (explicit in-profile, explicit out-of-profile, high priority or low priority) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the large number of policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly profiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported which prevents any packet accounting, the use of the policer's **parent** command requires at the policer's **stat-mode** to be set at least to the **minimal** setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. Once a policer has been made a child to a parent policer, the **stat-mode** cannot be changed to **no-stats** unless the policer parenting is first removed.

Each time the policer's **stat-mode** is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. You can view the the total/allocated/free stats by using the **tools dump system-resources** command. If insufficient counters exist to implement a mode on any policer instance, the **stat-mode** change will fail and the previous mode will continue unaffected for all instances of the policer.

The default **stat-mode** when a policer is created within the policy is **minimal**.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the **stat-mode** override command will fail. The previous **stat-mode** setting active for the policer will continue to be used by the policer.

The **no** form of this command attempts to return the policer's stat-mode setting to minimal. The command will fail if insufficient policer counter resources exist to implement minimal where the QoS policer is currently applied and has a forwarding class mapping.

**no-stats** — Counter resource allocation:0

The policer does not have any forwarding plane counters allocated and cannot provide offered, discard and forward statistics. A policer using no-stats cannot be a child to a parent policer and the policer's parent command will fail.

When **collect-stats** is enabled, the lack of counters causes the system to generate the following statistics:

- a. offered-in = 0
- b. offered-out = 0
- c. discard-in = 0
- d. discard-out = 0
- e. forward-in = 0
- f. forward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

**minimal** — Counter resource allocation: 1

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (profile or priority) and do not count green or yellow output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.

This counter mode is useful when only the most basic accounting information is required.

The counters are used in the following manner:

- 1. offered = profile in/out, priority high/low
- 2. discarded = Same as 1
- 3. forwarded = Derived from 1 - 2

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 0
- c. discard-in = 2
- d. discard-out = 0
- e. forward-in = 3
- f. forward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

With **minimal** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-in = 1
- ii. offered-out = 0



- iii. offered-undefined = 0
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

**offered-profile-no-cir** — Counter resource allocation:2

The **offered-profile-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-profile-no-cir** mode is most useful when the policer is receiving only in-profile and out-of-profile pre-marked (and trusted) packets. It is expected that in this instance a CIR rate will not be defined since all packets are already pre-marked. This mode does not prevent the policer from receiving un-trusted (color undefined) nor does it prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

- 1. offered-in = profile in
- 2. offered-out = profile out, priority high/low
- 3. dropped-in = Same as 1
- 4. dropped-out = Same as 2
- 5. forwarded-in = Derived from 1 - 3
- 6. forwarded-out = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 2
- c. discard-in = 3
- d. discard-out = 4
- e. forward-in = 5
- f. forward-out = 6

With **offered-profile-no-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-in = 1
- ii. offered-out = 2
- iii. offered-undefined = 0
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

**offered-priority-no-cir** — Counter resource allocation:2

The **offered-priority-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-priority-no-cir** mode is most useful when the policer is receiving only un-trusted packets and the ingress priority high and priority low classification options are being used without a CIR profiling rate defined. This mode does not prevent the policer from receiving trusted packets that are pre-marked in-profile or out-of-profile nor does it prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

1. offered-high = profile in, priority high
2. offered-low = profile out, priority low
3. dropped-high = Same as 1
4. dropped-low = Same as 2
5. forwarded-high = Derived from 1 - 3
6. forwarded-low = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-high = 1
- b. offered-low = 2
- c. discard-high = 3
- d. discard-low = 4
- e. forward-high = 5
- f. forward-low = 6

With **offered-priority-no-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-high = 1
- ii. offered-low = 2
- iii. offered-undefined = 0
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

#### **offered-limited-profile-cir** — Counter resource allocation:3

The **offered-limited-profile-cir** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-limited-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile (profile out but no profile in) traffic and un-trusted packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile packets.

The counters are used in the following manner:

1. offered-undefined-that-turned-green = profile in, priority high/low
2. offered-undefined-that-turned-yellow-or-red = priority high/low
3. offered-out-that-stayed-yellow-or-turned-red = profile out

- 4. dropped-undefined-that-turned-green = Same as 1
- 5. dropped-undefined-that-turned-yellow-or-red = Same as 2
- 6. dropped-out-that-turned-yellow-or-red = Same as 3
- 7. forwarded-undefined-that-turned-green = Derived from 1 - 4
- 8. forwarded-undefined-that-turned-yellow = Derived from 2 - 5
- 9. forwarded-out-that-turned-yellow = Derived from 3 - 6

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 0
- b. offered-out = 1 + 2 + 3
- c. discard-in = 0
- d. discard-out = 4 + 5 + 6
- e. forward-in = 7
- f. forward-out = 8 + 9

With **offered-limited-profile-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-in = 0
- ii. offered-out = 3
- iii. offered-undefined = 1 + 2
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

**offered-profile-cir** — Counter resource allocation:4

The **offered-profile-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile and in-profile traffic and is also receiving un-trusted packets that are being applied to a defined CIR profiling rate. This mode differs from **offered-limited-profile-cir** mode in that it expects both trusted in-profile and out-of-profile packets while still performing CIR profiling on packets with un-trusted markings. It is expected that in most cases where both trusted and un-trusted packets are received, the predominate case will not include trusted in-profile packets making the offered-limited-profile-cir accounting mode acceptable.

The counters are used in the following manner:

- 1. offered-in-that-stayed-green-or-turned-red = profile in
- 2. offered-undefined-that-turned-green = priority high/low
- 3. offered-undefined-that-turned-yellow-or-red = priority high/low
- 4. offered-out-that-stayed-yellow-or-turned-red = profile out
- 5. dropped-in-that-stayed-green-or-turned-red = Same as 1

- 6. dropped-undefined-that-turned-green = Same as 2
- 7. dropped-undefined-that-turned-yellow-or-red = Same as 3
- 8. dropped-out-that-turned-yellow-or-red = Same as 4
- 9. forwarded-in-that-stayed-green = Derived from 1 - 5
- 10. forwarded-undefined-that-turned-green = Derived from 2 - 6
- 11. forwarded-undefined-that-turned-yellow = Derived from 3 - 7
- 12. forwarded-out-that-turned-yellow = Derived from 4 - 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 2 + 3 + 4
- c. discard-in = 5 + 6
- d. discard-out = 7 + 8
- e. forward-in = 9 + 10
- f. forward-out = 11 + 12

With **offered-profile-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-high = 1
- ii. offered-low = 4
- iii. offered-undefined = 2 + 3
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

**offered-priority-cir** — Counter resource allocation:4

The **offered-priority-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-priority-cir** mode is most useful when the policer is receiving only un-trusted packets that are being classified as high priority or low priority and are being applied to a defined CIR profiling rate. This mode differs from **offered-profile-cir** mode in that it does not expect trusted in-profile and out-of-profile packets but does not exclude the ability of the policer to receive them.

The counters are used in the following manner:

- 1. offered-high-that-turned-green = profile in, priority high
- 2. offered-high-that-turned-yellow-or-red = profile in, priority high
- 3. offered-low-that-turned-green = profile out, priority low
- 4. offered-low-that-turned-yellow-or-red = profile out, priority low
- 5. dropped-high-that-turned-green = Same as 1
- 6. dropped-high-that-turned-yellow-or-red = Same as 2

- 7. dropped-low-that-turned-green = Same as 3
- 8. dropped-low-that-turned-yellow-or-red = Same as 4
- 9. forwarded-high-that-turned-green = Derived from 1 - 5
- 10. forwarded-high-that-turned-yellow = Derived from 2 - 6
- 11. forwarded-low-that-turned-green = Derived from 3 - 7
- 12. forwarded-low-that-turned-yellow = Derived from 4 - 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-high = 1 + 2
- b. offered-low = 3 + 4
- c. discard-in = 5 + 7
- d. discard-out = 6 + 8
- e. forward-in = 9 + 11
- f. forward-out = 10 + 12

With **offered-priority-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-high = 1 + 2
- ii. offered-low = 3 + 4
- iii. offered-undefined = 0
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

**offered-total-cir** — Counter resource allocation:2

The **offered-total-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-total-cir** mode is most useful when the policer is not receiving trusted in-profile or out-of-profile traffic and both high and low priority classifications are not being used on the un-trusted packets and the offered packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile or out-of-profile packets and does not prevent the use of priority high or low classifications on the un-trusted packets.

The counters are used in the following manner:

- 1. offered-that-turned-green = profile in/out, priority high/low
- 2. offered- that-turned-yellow-or-red = profile in/out, priority high/low
- 3. dropped-offered-that-turned-green = Same as 1
- 4. dropped-offered-that-turned-yellow-or-red = Same as 2
- 5. forwarded-offered-that-turned-green = Derived from 1 - 3
- 6. forwarded-offered-that-turned-yellow = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1 + 2
- b. offered-out = 0
- c. discard-in = 3
- d. discard-out = 4
- e. forward-in = 5
- f. forward-out = 6

Counter 0 indicates that the accounting statistic returns a value of zero.

With **offered-total-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

- i. offered-high = 1 + 2
- ii. offered-low = 0
- iii. offered-undefined = 0
- iv. offered-managed = 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

**offered-profile-capped-cir** — Counter resource allocation:2

When **offered-profile-capped-cir** is defined, the system creates five offered-output counters in the forwarding plane and five discard counters in the traffic manager.

The **offered-profile-capped-cir** mode is similar to the offered-profile-cir mode except that it includes support for profile in and **soft-in-profile** that may be output as ‘out-of-profile’ due to enabling **profile-capped** mode on the ingress policer.

The impact of using **offered-profile-capped-cir** stat-mode while **profile-capped** mode is disabled are that one of the counting resources in the forwarding plane and traffic manager will not be used and soft-in-profile will be treated as ‘offered-in’ instead of ‘offered-undefined’.

The counters are used in the following manner:

- 1. ‘offered-in-that-stayed-green’ = profile in, soft-in-profile
- 2. ‘offered-in-that-turned-yellow-or-red’ = profile in, soft-in-profile
- 3. ‘offered-soft-out-that-turned-green’ = soft-out-of-profile
- 4. ‘offered-soft-out- that-turned-yellow-or-red’= soft-out-of-profile
- 5. ‘offered-out-that-turned-yellow-or-red’ = profile out
- 6. ‘dropped-in-that-stayed-green’ = Same as 1
- 7. ‘dropped-in-that-turned-yellow-or-red’ = Same as 2
- 8. ‘dropped-soft-out-that-turned-green’ = Same as 3
- 9. ‘dropped-soft-out-that-turned-yellow-or-red’ = Same as 4
- 10. ‘dropped-out-that-turned-yellow-or-red’ = Same as 5

- |     |   |                       |
|-----|---|-----------------------|
| 11. | ‘forwarded-in-that-stayed-green’        | = Derived from 1 - 6  |
| 12. | ‘forwarded-in-that-turned-yellow’       | = Derived from 2 - 7  |
| 13. | ‘forwarded-soft-out-that-turned-green’  | = Derived from 3 - 8  |
| 14. | ‘forwarded-soft-out-that-turned-yellow’ | = Derived from 4 - 9  |
| 15. | ‘forwarded-out-that-turned-yellow’      | = Derived from 5 - 10 |

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

- |    |                     |                |
|----|---------------------|----------------|
| a. | ‘offered-undefined’ | = 3 + 4        |
| b. | ‘offered-in’        | = 1 + 2        |
| c. | ‘offered-out’       | = 5            |
| d. | ‘discard-in’        | = 6 + 8        |
| e. | ‘discard-out’       | = 7 + 9 + 10   |
| f. | ‘forward-in’        | = 11 + 13      |
| g. | ‘forward-out’       | = 12 + 14 + 15 |

**offered-limited-capped-cir** — Counter resource allocation:2

**offered-limited-capped-cir** is defined, the system creates four forwarding plane offered-output counters in the network processor and four discard counters in the traffic manager.

The **offered-limited-capped-cir** mode is similar to the **offered-profile-capped-cir** mode except that it combines **profile out** and **soft-out-of-profile** and eliminates the ‘offered-undefined’ statistic.

The impact of using **offered-limited-capped-cir** stat-mode while **profile-capped** mode is disabled are that one of the counting resources in the forwarding plane and traffic manager will not be used and **soft-in-profile** will be treated as ‘offered-in’ instead of ‘offered-undefined’.

The counters are used in the following manner:

- |     |  |                                    |
|-----|--|------------------------------------|
| 1.  | ‘offered-in-that-stayed-green’           | = profile in, soft-in-profile      |
| 2.  | ‘offered-in-that-turned-yellow-or-red’   | = profile in, soft-in-profile      |
| 3.  | ‘offered-out-that-turned-green’          | = soft-out-of-profile              |
| 4.  | ‘offered-out- that-turned-yellow-or-red’ | = profile out, soft-out-of-profile |
| 5.  | ‘dropped-in-that-stayed-green’           | = Same as 1                        |
| 6.  | ‘dropped-in-that-turned-yellow-or-red’   | = Same as 2                        |
| 7.  | ‘dropped-out-that-turned-green’          | = Same as 3                        |
| 8.  | ‘dropped-out-that-turned-yellow-or-red’  | = Same as 4                        |
| 9.  | ‘forwarded-in-that-stayed-green’         | = Derived from 1 - 5               |
| 10. | ‘forwarded-in-that-turned-yellow’        | = Derived from 2 - 6               |
| 11. | ‘forwarded-out-that-turned-green’        | = Derived from 3 - 7               |

12. 'forwarded-out-that-turned-yellow' = Derived from 4 - 8

When collect-stats is enabled, the counters are used by the system to generate the following statistics:

a.	'offered-in'	= 1 + 2
b.	'offered-out'	= 3 + 4
c.	'discard-in'	= 5 + 7
d.	'discard-out'	= 6 + 8
e.	'forward-in'	= 9 + 11
f.	'forward-out'	= 10 + 12

## prec

**Syntax** **prec** *ip-prec-value* **fc** *fc-name* [**priority** {**high** | **low**}]  
**no prec** *ip-prec-value*

**Context** config>qos>sap-ingress

**Description** This command explicitly sets the forwarding class or enqueueing priority when a packet is marked with an IP precedence value (*ip-prec-value*). Adding an IP precedence rule on the policy forces packets that match the specified *ip-prec-value* to override the forwarding class and enqueueing priority based on the parameters included in the IP precedence rule.

When the forwarding class is not specified in the rule, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy.

When the enqueueing priority is not specified in the rule, a matching packet preserves (or inherits) the existing enqueueing priority derived from earlier matches in the classification hierarchy.

The *ip-prec-value* is derived from the most significant three bits in the IP header ToS byte field (precedence bits). The three precedence bits define 8 Class-of-Service (CoS) values commonly used to map packets to per-hop Quality-of-Service (QoS) behavior. The precedence bits are also part of the newer DiffServ Code Point (DSCP) method of mapping packets to QoS behavior. The DSCP uses the most significant six bits in the IP header ToS byte and so overlaps with the precedence bits. Both IP precedence and DSCP classification rules are supported. DSCP rules have a higher match priority than IP precedence rules and where a *dscp-name* DSCP value overlaps an *ip-prec-value*, the DSCP rule takes precedence.

The HSMDA queue offered stats represent all packets sent to a specific ingress or egress queue regardless of the HSMDA counter override. This results in accurate queue offered stats, while the discard and forwarding stats per queue only represent packets that have not been associated with an exception counter. If the queue discard and forwarding stats are subtracted from the queue offered stats, an approximation of the number of packets handled by the queue that have been associated with an exception counter may be calculated. This is an approximation due to the possible presence of packets currently in the queue that are not represented by the discard or forwarding stats at the time the stats are collected but had been included in the queue offered stats. This discrepancy is minimized when the stats are collected over time and disappears completely once the queue drains.

The **no** form of the command removes the explicit IP precedence classification rule from the SAP ingress



policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs using the policy.

## Parameters

***ip-prec-value*** — The *ip-prec-value* is a required parameter that specifies the unique IP header ToS byte precedence bits value that will match the IP precedence rule. If the command is executed more than once with the same *ip-prec-value*, the previous forwarding class and enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.

A maximum of eight IP precedence rules are allowed on a single policy.

The precedence is evaluated from the lowest to highest value.

**Values** 0 — 7

***fc fc-name*** — The value given for the *fc-name* parameter must be one of the predefined forwarding classes in the system. Specifying the *fc-name* is optional. When a packet matches the rule the forwarding class is only overridden when the ***fc fc-name*** parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

The sub-class-name parameter is optional and used with the *fc-name* parameter to define a preexisting sub-class. The *fc-name* and sub-class-name parameters must be separated by a period (dot). If sub-class-name does not exist in the context of *fc -name*, an error will occur. If sub-class-name is removed using the `no fc fc-name.sub-class-name force` command, the `default-fc` command will automatically drop the sub-class-name and only use *fc-name* (the parent forwarding class for the sub-class) as the forwarding class.

**Values** **fc:** class[.sub-class]  
class: be, l2, af, l1, h2, ef, h1, nc  
sub-class: 29 characters max

**Default** Inherit (When ***fc*** is not defined, the rule preserves the previous forwarding class of the packet.)

***priority*** — The priority parameter overrides the default enqueueing priority for all packets received on an ingress SAP using this policy that match this rule. Specifying the priority is optional. When a packet matches the rule the enqueueing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueueing priority is inherited based on previous rule matches.

***high*** — This parameter is used in conjunction with the ***priority*** parameter. Setting the enqueueing parameter to ***high*** for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

***low*** — This parameter is used in conjunction with the ***priority*** parameter. Setting the enqueueing parameter to ***low*** for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

**Default** Inherit (When priority is not defined, the rule preserves the previous enqueueing priority of the packet.)

**Values** high, low

## prec

**Syntax** `prec ip-prec-value fc fc-name [hsmda-counter-override counter-id] [profile {in | out}]`  
`no prec ip-prec-value`

**Context** config>qos>sap-egress

**Description** This command defines a specific IP precedence value that must be matched to perform the associated reclassification actions. If an egress packet on the SAP matches the specified IP precedence value, the forwarding class, profile or HSMMDA egress queue accounting behavior may be overridden. By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions. The default behavior for HSMMDA queue accounting is to use the counters associated with the queue to which the packet is mapped.

The IP precedent bits used to match against prec reclassification rules come from the Type of Service (ToS) field within the IPv4 header. If the packet does not have an IPv4 header, prec based matching is not performed.

The reclassification actions from a prec reclassification rule may be overridden by a DHCP or IP flow matching events.

The **fc** keyword is optional. When specified, the egress classification rule will overwrite the forwarding class derived from ingress. The new forwarding class is used for egress remarking and queue mapping decisions. If a dhcp or ip-criteria match occurs after the prec match, the new forwarding class may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new fc, the fc from the prec match will be used.

The **profile** keyword is optional. When specified, the egress classification rule will overwrite the profile of the packet derived from ingress. The new profile value is used for egress remarking and queue congestion behavior. If a dhcp or ip-criteria match occurs after the prec match, the new profile may be overridden by the higher priority match actions. If the higher priority match actions do not specify a new profile, the profile from the prec match will be used.

The **hsmda-counter-override** keyword is optional. When specified and the egress SAP is created on an HSMMDA, the egress classification rule will override the default queue accounting function for the packet. By default, the HSMMDA uses each queues default queue counters for packets mapped to the queue. The **hsmda-counter-override** keyword is used to map the packet to an explicit exception counter. One of eight counters may be used. When the packet is mapped to an exception counter, the packet will not increment the queues discard or forwarding counters, instead the exception discard and forwarding counters will be used. The dscp based counter override decision may be overwritten by an ip-criteria reclassification rule match if the higher priority classification rule has an hsmda-counter-override action defined.

The **no** form of the command removes the reclassification rule from the SAP egress QoS policy.

**Parameters** **fc fc-name** — This keyword is optional. When specified, packets matching the IP precedence value will be explicitly reclassified to the forwarding class specified as fc-name regardless of the ingress classification decision. The explicit forwarding class reclassification may be overwritten by a higher priority dscp or ip-criteria reclassification match. The fc-name defined must be one of the eight forwarding classes supported by the system. To remove the forwarding class reclassification action for

the specified prec-value, the prec command must be re-executed without the fc reclassification action defined.

**Values** be, l1, af, l2, h1, ef, h2 or nc

**Default** None

**profile {in | out}** — This keyword is optional. When specified, packets matching the IP precedence value will be explicitly reclassified to either in-profile or out-of-profile regardless of the ingress profiling decision. The explicit profile reclassification may be overwritten by a higher priority dscp or ip-criteria reclassification match. To remove the profile reclassification action for the specified prec-value, the prec command must be re-executed without the profile reclassification action defined.

**in** — The in parameter is mutually exclusive to the out parameter following the profile reclassification action keyword. Either in or out must be specified when the profile keyword is present. When in is specified, any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane. This value may be overwritten by an explicit profile action in a higher priority dscp or ip-criteria reclassification match.

**out** — The out parameter is mutually exclusive to the in parameter following the profile reclassification action keyword. Either in or out must be specified when the profile keyword is present. When out is specified, any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane. This value may be overwritten by an explicit profile action in a higher priority dscp or ip-criteria reclassification match.

**hsmda-counter-override counter-id** — The hsmda-counter-override reclassification action is optional and only has significance on SAPs which are created on an HSMDA. When specified, packets matching the IP precedence value will be mapped to the defined HSMDA exception counter-id for the packets queue group. The default behavior is to use the default counter on the queue group for the queue to which the packet is mapped. The hsmda-counter-override action may be overwritten by an ip-criteria reclassification rule match. The specified counter-id must be specified as an integer between 1 and 8. To remove the ESMDA exception counter reclassification action for the specified dscp-value, the dscp command must be re-executed without the hsmda-counter-override reclassification action defined.

**Values** 1 —8

## queue

**Syntax** **queue** *queue-id* [**multipoint**] [*queue-type*] [*queue-mode*] **pool** *pool-name*  
**queue** *queue-id* [**multipoint**] [*queue-type*] **pool** *pool-name*  
**no queue** *queue-id*

**Context** config>qos>sap-ingress  
config>qos>sap-egress

**Description** This command creates the context to configure an ingress service access point (SAP) QoS policy queue. Explicit definition of an ingress queue's hardware scheduler status is supported. A single ingress queue allows support for multiple forwarding classes. The default behavior automatically chooses the expedited or non-expedited nature of the queue based on the forwarding classes mapped to it. As long as all forwarding classes mapped to the queue are expedited (nc, ef, h1 or h2), the queue is treated as an expedited queue by the hardware schedulers. When any non-expedited forwarding classes are mapped to the queue (be, af, l1

or 12), the queue is treated as best effort (be) by the hardware schedulers. The expedited hardware schedulers are used to enforce expedited access to internal switch fabric destinations. The hardware status of the queue must be defined at the time of queue creation within the policy.

The **queue** command allows the creation of multipoint queues. Only multipoint queues can receive ingress packets that need flooding to multiple destinations. By separating the unicast for multipoint traffic at service ingress and handling the traffic on separate multipoint queues special handling of the multipoint traffic is possible. Each queue acts as an accounting and (optionally) shaping device offering precise control over potentially expensive multicast, broadcast and unknown unicast traffic. Only the back-end support of multipoint traffic (between the forwarding class and the queue based on forwarding type) needs to be defined. The individual classification rules used to place traffic into forwarding classes are not affected. Queues must be defined as multipoint at the time of creation within the policy.

The multipoint queues are for multipoint-destined service traffic. Within non-multipoint services, such as Epipe services, all traffic is considered unicast due to the nature of the service type. Multicast and broadcast-destined traffic in an Epipe service will not be mapped to a multipoint service queue.

When an ingress SAP QoS policy with multipoint queues is applied to an Epipe SAP, the multipoint queues are not created. When an ingress SAP QoS policy with multipoint queues is applied to an IES SAP, a multipoint queue will be created when PIM is enabled on the IES interface.

Any billing or statistical queries about a multipoint queue on a non-multipoint service returns zero values. Any queue parameter information requested about a multipoint queue on a non-multipoint service returns the queue parameters in the policy. Buffers will not be allocated for multipoint queues on non-multipoint services. Buffer pool queries return zero values for actual buffers allocated and current buffer utilization.

The **no** form of this command removes the *queue-id* from the SAP ingress QoS policy and from any existing SAPs using the policy. If any forwarding class forwarding types are mapped to the queue, they revert to their default queues. When a queue is removed, any pending accounting information for each SAP queue created due to the definition of the queue in the policy is discarded.

The **pool** keyword is a create time parameter that allows the queue to receive buffers from an explicit buffer pool instead of the default buffer pool. The specified pool-name must have been explicitly created in a named-pool-policy and the policy must have been applied to the MDA or port on which the queue resides.

If the specified pool-name does not exist on the MDA, the queue will be treated as 'pool orphaned' and will be mapped to the appropriate default pool. Once the pool comes into existence on the MDA or port, the queue will be mapped to the new pool.

Once the queue is created within the policy, the **pool** command may be used to either remove the queue from the pool, or specify a new pool name association for the queue. The **pool** command does not appear in **save** or **show** command output. Instead, the current pool name for the queue will appear (or not appear) on the queue command output using the **pool** keyword.

## Parameters

*queue-id* — The *queue-id* for the queue, expressed as an integer. The *queue-id* uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.

**Values**      1 — 32

*queue-type* — The **expedite**, **best-effort** and **auto-expedite** queue types are mutually exclusive to each other. Each defines the method that the system uses to service the queue from a hardware perspective. While parental virtual schedulers can be defined for the queue, they only enforce how the queue interacts for bandwidth with other queues associated with the same scheduler hierarchy. An internal mechanism that provides access rules when the queue is vying for bandwidth with queues in other virtual schedulers is also needed. A keyword must be specified at the time the queue is created in the SAP ingress policy. If an attempt

to change the keyword after the queue is initially defined, an error is generated.

**expedite** — This keyword ensures that the queue is treated in an expedited manner independent of the forwarding classes mapped to the queue.

**best-effort** — This keyword ensures that the queue is treated in a non-expedited manner independent of the forwarding classes mapped to the queue.

**auto-expedite** — This keyword allows the system to auto-define the way the queue is serviced by the hardware. When **auto-expedite** is defined on the queue, the queue is treated in an expedited manner when all forwarding classes mapped to the queue are configured as expedited types *nc*, *ef*, *h1* or *h2*. When a single non-expedited forwarding class is mapped to the queue (*be*, *af*, *l1* and *l2*) the queue automatically falls back to non-expedited status.

**Values** expedite, best-effort, auto-expedite

**Default** auto-expedite

**multipoint** — This keyword specifies that this *queue-id* is for multipoint forwarded traffic only. This *queue-id* can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. If you attempt to map forwarding class unicast traffic to a multipoint queue, an error is generated and no changes are made to the current unicast traffic queue mapping.

A queue must be created as multipoint. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated and the command will not execute.

The **multipoint** keyword can be entered in the command line on a pre-existing multipoint queue to edit *queue-id* parameters.

**Values** multipoint or not present

**Default** Present (the queue is created as non-multipoint)

*queue-mode* — Specifies the mode in which the queue is operating. This attribute is associated with the queue at the time of creation and cannot be modified thereafter.

**Values** **profile-mode**: When the queue is operating in the profile mode (or, the color aware mode), the queue tries to provide the appropriate bandwidth to the packets with different profiles. The profiles are assigned according to the configuration of the forwarding class or the sub-forwarding class.

**priority-mode**: The queue is capable of handling traffic differently with two distinct priorities. These priorities are assigned by the stages preceding the queueing framework in the system. In priority mode, the queue does not have the functionality to support the profiled traffic and in such cases the queue will have a degraded performance. However, the converse is not valid and a queue in profile mode should be capable of supporting the different priorities of traffic.

**Default** **priority-mode**

*pool-name* — The specified pool-name identifies a named pool where the policy will be applied. Each queue created within the system is tied to a physical port. When the policy is applied and the queue is created, the system will scan the named pools associated with the port to find the specified pool name. If the pool is not found on the port, the system will then look at named pools defined at the ports MDA level. If the pool name is not found on either the port or MDA, the queue will be marked as ‘pool-orphaned’ and will be mapped to the appropriate default pool. If the pool comes into existence, the queue will be

moved from the default pool to the new named pool and the 'pool-orphaned' state will be cleared. The specified name must be an ASCII name string up to 16 characters long.

**Values** Any valid ASCII name string

**Default** None

The queue's pool association may only be removed by either re-executing the queue command without the pool keyword or by executing the no pool command within the queue's CLI context. When the pool name is removed, the queue will be placed on the appropriate default pool.

---

# Service Ingress QoS Policy Forwarding Class Commands

## broadcast-queue

**Syntax** `broadcast-queue queue-id [group queue-group-name]`

**Context** `config>qos>sap-ingress>fc fc-name`

**Description** This command overrides the default broadcast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all broadcast traffic on a SAP using this policy will be forwarded using the *queue-id*.

The broadcast forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of the command sets the broadcast forwarding type *queue-id* back to the default of tracking the multicast forwarding type queue mapping.

**Parameters** *queue-id* — The *queue-id* parameter must be an existing, multipoint queue defined in the `config>qos>sap-ingress` context.

**Values** Any valid multipoint queue ID in the policy including 2 through 32.

**Default** 11

**group queue-group-name** — This optional parameter is used to redirect the forwarding type within the forwarding class to the specified *queue-id* within the *queue-group-name*. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the specified queue group. The *queue-group-name* are configured in the `config>qos>queue-group-templates egress` and `ingress` contexts.

## de-1-out-profile

**Syntax** `[no] de-1-out-profile`

**Context** `config>qos>sap-ingress>fc`

**Description** This command, when enabled on a parent forwarding class, applies a color profile mode to the packets stored in the queue associated with this forwarding class. The queue associated with the parent forwarding class **MUST** be of type **profile-mode**.

When this QoS policy is applied to the ingress of a Frame Relay VLL SAP, the system will treat the received FR frames with DE bit set as out-of-profile regardless of their previous marking as the result of the default classification or on a match with an IP filter. It also adjusts the CIR of the ingress SAP queue to take into account out-of-profile frames which were sent while the SAP queue was in the “< CIR” state of the bucket. This makes sure that the CIR of the SAP is achieved in the long run.

All received DE=0 frames which are classified into this parent forwarding class or any of its sub-classes have their profile unchanged by enabling this option. That is the DE=0 frame profile could be undetermined

(default), in-profile, or out-of-profile as per previous classification. The DE=0 frames which have a profile of undetermined will be evaluated by the system CIR marking algorithm and will be marked appropriately.

The **priority** option if used has no effect. All FR VLL DE=1 frames have automatically their priority set to low while DE=0 frames have their priority set to high. Furthermore, DE=1 frames have drop-preference bit set in the internal header. The internal settings of the priority bit and of the drop-preference bit of the frame is independent of the use or not of the profile mode.

All other capabilities of the Fpipe service are maintained. This includes remarking of the DE bit on egress SAP, and FR PW control word on egress network port for the packets which were classified into “out-of-profile” at ingress SAP.

This **de-1-out-profile** keyword has an effect when applied to the ingress of a SAP which is part of an fpipe service. It can also be used on the ingress of an epipe or vpls SAP.

The **no** form of the command disables the color profile mode of operation on all SAPs this ingress QoS policy is applied.

**Default** no de-1-out-profile

## egress-fc

**Syntax** **egress-fc** *fc-name*  
**no egress-fc**

**Context** config>qos>sap-ingress>fc

**Description** This command configures the forwarding class to be used by the egress QoS processing. It overrides the forwarding class determined by ingress classification but not the QoS Policy Propagation via BGP.

The forwarding class and/or forwarding sub-class can be overridden.

The new egress forwarding class is applicable to both SAP egress and network egress.

**Default** no egress-fc

**Parameters** *fc-name* — Specifies the forwarding class name to be used by the egress QoS processing.

**Default** None. The fc name must be specified.

**Values** be, l2, af, l1, h2, ef, h1, nc

## in-remark

**Syntax** **in-remark dscp** *dscp-name*  
**in-remark prec** *ip-prec-value*  
**no in-remark**

**Context** config>qos>sap-ingress>fc *fc-name*

**Description** This command is used in a SAP ingress QoS policy to define an explicit in-profile remark action for a forwarding class or sub-class. While the SAP ingress QoS policy may be applied to any SAP, the remarking functions are only enforced when the SAP is associated with an IP interface (IES or VPRN SAPs). When the



policy is applied to a Layer 2 SAP (i.e., Epipe or VPLS), the remarking definitions are silently ignored.

In the case where the policy is applied to a Layer 3 SAP, the in-profile remarking definition will be applied to packets that have been classified to the forwarding class or sub-class. It is possible for a packet to match a classification command that maps the packet to a particular forwarding class or sub-class only to have a more explicit (higher priority match) override the association. Only the highest priority match forwarding class or sub-class association will drive the in-profile marking.

The in-remark command is only applicable to ingress IP routed packets that are considered in-profile. The profile of a SAP ingress packet is affected by either the explicit in-profile/out-of-profile definitions or the ingress policing function applied to the packet. The following table shows the effect of the in-remark command on received SAP ingress packets. Within the in-profile IP packet's ToS field, either the six DSCP bits or the three precedence bits are remarked.

SAP Ingress Packet State	'in-remark' Command Effect
Non-Routed, Policed In-Profile	No Effect (non-routed packet)
Non-Routed, Policed Out-of-Profile	No Effect (non-routed packet)
Non-Routed, Explicit In-Profile	No Effect (non-routed packet)
Non-Routed, Explicit Out-of-Profile	No Effect (non-routed packet)
IP Routed, Policed In-Profile	in-remark value applied to IP header ToS field
IP Routed, Policed Out-of-Profile	No Effect (out-of-profile packet)
IP Routed, Explicit In-Profile	in-remark value applied to IP header ToS field
IP Routed, Explicit Out-of-Profile	No Effect (out-of-profile packet)

The **no** form of the command disables ingress remarking of in-profile packets classified to the forwarding class or sub-class.

**Parameters** **dscp** *dscp-name* — This parameter is one of two mutually exclusive settings that are applicable to the in-remark command. The in-remark command can be configured to either remark the DiffServ Code Point (DSCP) six bit value or the three Precedence bits. The dscp parameter specifies that the matching packets DSCP bits should be overridden with the value represented by dscp-name.

32 characters, maximum, The name specified by dscp-name is used to refer to the six bit value represented by dscp-name. It must be one of the predefined DSCP names defined on the system.

**Values** be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

**Default** None (an explicit valid DSCP name must be specified)

**prec** *ip-prec-value* — This parameter is one of two mutually exclusive settings that are applicable to the in-remark command. The in-remark command can be configured to either remark the DiffServ Code Point (DSCP) six bit value or the three Precedence bits. The prec parameter specifies that the matching packets Precedence bits should be overridden with the value represented by prec-value.

**Values** 0 — 7

**Default** None (an explicit precedence value must be specified)

## multicast-queue

**Syntax** **multicast-queue** *queue-id* [**group** *queue-group-name*]

**Context** config>qos>sap-ingress>fc *fc-name*

**Description** This command overrides the default multicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all multicast traffic on a SAP using this policy is forwarded using the *queue-id*.

The multicast forwarding type includes the **unknown** unicast forwarding type and the **broadcast** forwarding type unless each is explicitly defined to a different multipoint queue. When the unknown and broadcast forwarding types are left as default, they will track the defined queue for the multicast forwarding type.

The **no** form of the command sets the multicast forwarding type *queue-id* back to the default queue for the forwarding class. If the **broadcast** and **unknown** forwarding types were not explicitly defined to a multipoint queue, they will also be set back to the default multipoint queue (queue 11).

**Parameters** *queue-id* — The *queue-id* parameter specified must be an existing, multipoint queue defined in the config>qos>sap-ingress context.

**Values** Any valid multipoint queue-ID in the policy including 2 through 32.

**Default** 11

**group** *queue-group-name* — This optional parameter is used to redirect the forwarding type within the forwarding class to the specified *queue-id* within the *queue-group-name*. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the specified queue group. The *queue-group-name* are configured in the **config>qos>queue-group-templates** egress and ingress contexts.

## out-remark

**Syntax** **out-remark dscp** *dscp-name*  
**out-remark prec** *ip-prec-value*  
**no out-remark**

**Context** config>qos>sap-ingress>fc *fc-name*

**Description** This command is used in a SAP ingress QoS policy to define an explicit out-of-profile remark action for a forwarding class or sub-class. While the SAP ingress QoS policy may be applied to any SAP, the remarking

functions are only enforced when the SAP is associated with an IP interface (IES or VPRN SAPs). When the policy is applied to a Layer 2 SAP (for example, Epipe or VPLS), the remarking definitions are silently ignored.

In the case where the policy is applied to a Layer 3 SAP, the out-of-profile remarking definition will be applied to packets that have been classified to the forwarding class or sub-class. It is possible for a packet to match a classification command that maps the packet to a particular forwarding class or sub-class only to have a more explicit (higher priority match) override the association. Only the highest priority match forwarding class or sub-class association will drive the out-of-profile marking.

The out-remark command is only applicable to ingress IP routed packets that are considered out-of-profile. The profile of a SAP ingress packet is affected by either the explicit in-profile/out-of-profile definitions or the ingress policing function applied to the packet. The following table shows the effect of the out-remark command on received SAP ingress packets. Within the out-of-profile IP packet's ToS field, either the six DSCP bits or the three precedence bits are remarked.

**Table 32: Out-remark command effect**

SAP Ingress Packet State	'out-remark' Command Effect
Non-Routed, Policed In-Profile	No Effect (non-routed packet)
Non-Routed, Policed Out-of-Profile	No Effect (non-routed packet)
Non-Routed, Explicit In-Profile	No Effect (non-routed packet)
Non-Routed, Explicit Out-of-Profile	No Effect (non-routed packet)
IP Routed, Policed In-Profile	No Effect (in-profile packet)
IP Routed, Policed Out-of-Profile	out-remark value applied to IP header ToS field
IP Routed, Explicit In-Profile	No Effect (in-of-profile packet)
IP Routed, Explicit Out-of-Profile	out-remark value applied to IP header ToS field

A packet that is explicitly remarked at ingress will not be affected by any egress remarking decision. Explicit ingress remarking has highest priority.

The **no** form of the command disables ingress remarking of out-of-profile packets classified to the forwarding class or sub-class.

**Default** none

**Parameters** **dscp** *dscp-name* — This parameter is one of two mutually exclusive settings that are applicable to the out-remark command. The out-remark command can be configured to either remark the DiffServ Code Point (DSCP) six bit value or the three Precedence bits. The dscp parameter specifies that the matching packets DSCP bits should be overridden with the value represented by dscp-name.

32 characters, maximumThe name specified by dscp-name is used to refer to the six bit value represented by dscp-name. It must be one of the predefined DSCP names defined on the system.

**Values** be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, e f, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57 , cp58, cp59, cp60, cp61, cp62, cp63

**Default** None (an explicit valid DSCP name must be specified)

**prec** *ip-prec-value* — This parameter is one of two mutually exclusive settings that are applicable to the out-remark command. The out-remark command can be configured to either remark the DiffServ Code Point (DSCP) six bit value or the three Precedence bits. The prec parameter specifies that the matching packets Precedence bits should be overridden with the value represented by prec-value.

The value specified by prec-value is used to overwrite the Precedence bits within a matching routed packets IP header ToS field.

**Values** 0 — 7

**Default** None (an explicit Precedence value must be specified)

An explicit dscp name or prec value must be specified for out-of-profile remarking to be applied.

## policer

**Syntax** **policer** *policer-id* [**fp-redirect-group**]  
**no policer** *policer-id*

## multicast-policer

**Syntax** **multicast-policer** *policer-id* [**fp-redirect-group**]  
**no multicast-policer** *policer-id*

## broadcast-policer

**Syntax** **broadscast-policer** *policer-id* [**fp-redirect-group**]  
**no broadcast-policer** *policer-id*

## unknown-policer

**Syntax** **unknown-policer** *policer-id* [**fp-redirect-group**]  
**no unknown-policer** *policer-id*

## profile

**Syntax** **profile** {in | out}  
**no profile**

**Context** config>qos>sap-igress>fc

**Description** This command places a forwarding class or sub-class into a color aware profile mode. Normally, packets associated with a class are considered in-profile or out-of-profile solely based on the dynamic rate of the ingress queue relative to its CIR. Explicitly defining a class as in-profile or out-of-profile overrides this function by handling each packet with the defined profile state.

The profile command may only be executed when the forwarding class or the parent forwarding class (for a sub-class) is mapped to a queue that has been enabled to support color aware profile packets. The queue may

only be configured for profile-mode at the time the queue is created in the SAP ingress QoS policy.

A queue operating in profile-mode may support in-profile, out-of-profile and non-profiled packets simultaneously. However, the high and low priority classification actions are ignored when the queue is in profile-mode.

The **no** form of the command removes an explicit in-profile or out-of-profile configuration on a forwarding class or sub-class.

**Default** **no profile** — The default profile state of a forwarding class or sub-class is not to treat ingress packets as color aware. An explicit definition for in-profile or out-of-profile must be specified on the forwarding class or sub-class.

**Parameters** **in** — The **in** keyword is mutually exclusive to the **out** keyword. When the profile in command is executed, all packets associated with the class will be handled as in-profile. Packets explicitly handled as in-profile or out-of-profile still flow through the ingress service queue associated with the class to preserve order within flows. In-profile packets will count against the CIR of the queue, diminishing the amount of CIR available to other classes using the queue that are not configured with an explicit profile.

**out** — The **out** keyword is mutually exclusive to the **in** keyword. When the profile out command is executed, all packets associated with the class will be handled as out-of-profile. Packets explicitly handled as in-profile or out-of-profile still flow through the ingress service queue associated with the class to preserve order within flows. Out-of-profile packets will not count against the CIR of the queue, allowing other classes using the queue that are not configured with an explicit profile to be measured against the full CIR.

## unknown-queue

**Syntax** **unknown-queue** *queue-id* [**group** *queue-group-name*]  
**no unknown-queue**

**Context** config>qos>sap-ingress>fc *fc-name*

**Description** This command overrides the default unknown unicast forwarding type queue mapping for **fc** *fc-name*. The specified *queue-id* must exist within the policy as a multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all unknown traffic on a SAP using this policy is forwarded using the *queue-id*.

The unknown forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior.

The **no** form of this command sets the unknown forwarding type *queue-id* back to the default of tracking the multicast forwarding type queue mapping.

**Parameters** *queue-id* — Specifies an existing multipoint queue defined in the **config>qos>sap-ingress** context.

**Values** Any valid multipoint *queue-id* in the policy including 2 through 32.

**Default** 11

**group** *queue-group-name* — This optional parameter is used to redirect the forwarding type within the forwarding class to the specified *queue-id* within the *queue-group-name*. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the

specified queue group. The *queue-group-name* are configured in the `config>qos>queue-group-templates` egress and ingress contexts.

## queue

**Syntax** `queue queue-id [{group queue-group-name [instance instance-id]} | port-redirect-group-queue]`  
`no queue`

**Context** `config>qos>sap-ingress>fc`  
`config>qos>sap-egress>fc`

**Description** This command overrides the default queue mapping for `fc` *fc-name*. The specified *queue-id* must exist within the policy before the mapping can be made. Once the forwarding class mapping is executed, all traffic classified to *fc-name* on a SAP using this policy.

The `no` form of this command sets the *queue-id* back to the default queue for the forwarding class (queue 1).

**Default** `no queue`

**Parameters** *queue-id* — Specifies the SAP egress *queue-id* to be associated with the forwarding class. The *queue-id* must be an existing queue defined in *sap-egress* policy-*id*.

**Values** 1 — 8.

**Default** 1

**group** *queue-group-name* — This optional parameter is used to redirect the forwarding type within the forwarding class to the specified *queue-id* within the *queue-group-name*. When the policy is applied, all packets matching the forwarding class and forwarding type will be redirected to the queue within the specified queue group. The *queue-group-name* are configured in the `config>qos>queue-group-templates` egress and ingress contexts. This parameter is used when policy based queue group redirection is desired. That is, the specific queue group to redirect to is named in the QoS policy.

**instance** *instance-id* — This parameter is used to specify the specific instance of a queue group with template *queue-group-name* to which this queue should be redirected. This parameter is only valid for queue groups on egress ports where policy based redirection is required.

**Values** 1 — 40960

**Default** 1

**port-redirect-group-queue** — This keyword is used to mark a given forwarding class queue for redirection to an egress queue group *queue*. This is only used when the specific queue group instance is assigned at the time the QoS policy is applied to the SAP. This redirection model is known as SAP based redirection.

## hsmda-queues

**Syntax** **hsmda-queues**

**Context** config>qos>sap-egress

**Description** This command enables the context to configure queue definitions for use on SAPs or subscribers on HSMDAs. A single QoS policy simultaneously defines queues for both standard MDA and for HSMDA subscribers and SAPs. This allows the policy association decision to be ignorant of the type of hardware the SAP or subscriber is traversing.

## queue

**Syntax** **queue** *queue-id* [**port-redirect-group-queue**]  
**no queue** *queue-id*

**Context** config>qos>sap-egress>hsmda-queues

**Description** This command, within the QoS policy HSMDA-queues context, is a container for the configuration parameters controlling the behavior of an HSMDA queue. Unlike the standard QoS policy queue command, this command is not used to actually create or dynamically assign the queue to the object which the policy is applied. The queue identified by *queue-id* always exists on the SAP or subscriber context whether the command is executed or not. In the case of HSMDA SAPs and subscribers, all eight queues exist at the moment the system allocates an HSMDA queue group to the object (both ingress and egress).

### **Best-Effort, Expedited and Auto-Expedite Queue Behavior Based on Queue-ID**

With standard service queues, the scheduling behavior relative to other queues is based on two items, the queues Best-Effort or Expedited nature and the dynamic rate of the queue relative to the defined CIR. HSMDA queues are handled differently. The create time auto-expedite and explicit expedite and best-effort qualifiers have been eliminated and instead the scheduling behavior is based solely on the queues identifier. Queues with a *queue-id* equal to 1 are placed in scheduling class 1. Queues with *queue-id* 2 are placed in scheduling class 2. And so on up to scheduling class 8. Each scheduling class is either mapped directly to a strict scheduling priority level based on the class ID, or the class may be placed into a weighted scheduling class group providing byte fair weighted round robin scheduling between the members of the group. Two weighted groups are supported and each may contain up to three consecutive scheduling classes. The weighed group assumes its highest member class's inherent strict scheduling level for scheduling purposes. Strict priority level 8 has the highest priority while strict level 1 has the lowest. When grouping of scheduling classes is defined, some of the strict levels will not be in use.

### **Single Type of HSMDA Queues**

Another difference between HSMDA queues and standard service queues is the lack of Multipoint queues. At ingress, an HSMDA SAP or subscriber does not require multi-point queues since all forwarding types (broadcast, multicast, unicast and unknown) forward to a single destination, the ingress forwarding plane on the IOM. Instead of a possible eight queues per forwarding type (for a total of up to 32) within the SAP ingress QoS policy, the *hsmda-queues* node supports a maximum of eight queues.



### Every HSMDA Queue Supports Profile Mode Implicitly

Unlike standard service queues, the HSMDA queues do not need to be placed into the special mode profile at create time in order to support ingress color aware policing. Each queue may handle in-profile, out-of-profile and profile undefined packets simultaneously. As with standard queues, the explicit profile of a packet is dependant on ingress sub-forwarding class to which the packet is mapped.

### Queue sharing and redirection

Redirection to an egress port queue group specified for the HSMDA is possible using the port-redirect-group parameter. If this is specified, then packets are redirected to the queue-id in the HSMDA queue instance named at the the time the egress QoS policy is applied to the SAP.

The **no** form of the command restores the defined queue-id to its default parameters. All HSMDA queues having the queue-id and associated with the QoS policy are re-initialized to default parameters.

### Parameters

*queue-id* — Defines the context of which of the eight ingress or egress queues will be entered for editing purposes.

**port-redirect-group** — This parameter is used to mark a given forwarding class queue for redirection to an egress port queue group. This is only used when the specific queue group instance is assigned at the time the qos policy is applied to the SAP. This redirection model is known as SAP based redirection.

## packet-byte-offset

**Syntax** **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}  
**no packet-byte-offset**

**Context** config>qos>sap-egress>hsmda-queues

**Description** This command adds or subtracts the specified number of bytes to the accounting function for each packet handled by the HSMDA queue. Normally, the accounting and leaky bucket functions are based on the Ethernet DLC header, payload and the 4 byte CRC (everything except the preamble and inter-frame gap). As an example, the packet-byte-offset command can be used to add the frame encapsulation overhead (20 bytes) to the queues accounting functions.

The accounting functions affected include:

- Offered High Priority / In-Profile Octet Counter
- Offered Low Priority / Out-of-Profile Octet Counter
- Discarded High Priority / In-Profile Octet Counter
- Discarded Low Priority / Out-of-Profile Octet Counter
- Forwarded In-Profile Octet Counter
- Forwarded Out-of-Profile Octet Counter
- Peak Information Rate (PIR) Leaky Bucket Updates
- Committed Information Rate (CIR) Leaky Bucket Updates
- Queue Group Aggregate Rate Limit Leaky Bucket Updates

The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based

and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.

The packet-byte-offset command accepts either **add** or **subtract** as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. Up to 20 bytes may be added to the packet and up to 43 bytes may be removed from the packet. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation, the command `packet-byte-offset subtract 14` would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.

As inferred above, the variable accounting size offered by the packet-byte-offset command is targeted at the queue and queue group level. When the queue group represents the last-mile bandwidth constraints for a subscriber, the offset allows the HSMDA queue group to provide an accurate accounting to prevent overrun and underrun conditions for the subscriber. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscribers packets on an Ethernet aggregation network.

The packet-byte-offset value may be overridden for the HSMDA queue at the SAP or subscriber profile level.

The **no** form of the command removes any accounting size changes to packets handled by the queue. The command does not affect overrides that may exist on SAPs or subscriber profiles associated with the queue.

**Parameters** **add** *add-bytes* — Indicates that the byte value should be added to the packet for queue and queue group level accounting functions. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command. The **add** keyword is mutually exclusive with the **subtract** keyword.

**Values** 1 — 31

**subtract** *sub-bytes* — Indicates that the byte value should be subtracted from the packet for queue and queue group level accounting functions. The **subtract** keyword is mutually exclusive with the **add** keyword. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command.

**Values** 1 — 32

## wrr-policy

**Syntax** **wrr-policy** *hsmda-wrr-policy-name*  
**no wrr-policy**

**Context** config>qos>sap-egress>hsmda-queues

**Description** This command associates an existing HSMDA weighted-round-robin (WRR) scheduling loop policy to the HSMDA queue.

**Parameters** *hsmda-wrr-policy-name* — Specifies the existing HSMDA WRR policy name to associate to the queue.

## low-burst-max-class

**Syntax** **low-burst-max-class** *class-id*  
**no low-burst-max-class**

**Context** config>qos>sap-egress>hsmda-queues>queue

**Description** This command assigns the low burst maximum class to associate with the HSMDA queue.  
The **no** form of the command returns the class id for the queue to the default value.

**Parameters** *class-id* — Specifies the class identifier of the low burst max class for the HSMDA queue.

**Values** 1— 32

## wrr-weight

**Syntax** **wrr-weight** *value*  
**no wrr-weight**

**Context** config>qos>sap-egress>hsmda-queues>queue

**Description** This command assigns the weight value to the HSMDA queue.

The **no** form of the command returns the weight value for the queue to the default value.

**Parameters** *percentage* — Specifies the weight for the HSMDA queue.

**Values** 1— 32

## slope-policy

**Syntax** **slope-policy** *hsmda-slope-policy-name*  
**no slope-policy**

**Context** config>qos>sap-egress>hsmda-queues>queue

**Description** This command associates an existing HSMDA slope policy to the QoS policy HSMDA queue. The specified *hsmda-slope-policy-name* must exist for the command to succeed. If the policy name does not exist, the command has no effect on the existing slope policy association. Once a slope policy is associated with a QoS policy queue, subscriber profile override or SAP override, the slope policy cannot be removed from the system. Any edits to an associated slope policy are immediately applied to the queues using the slope policy.

Within the ingress and egress QoS policies, packets are classified as high priority or low-priority. For color aware policies, packets are also potentially classified as in-profile, out-of-profile or profile-undefined. Based on these classifications, packets are mapped to the RED slopes in the following manner:

Ingress Slope Mapping

- In-Profile — High Slope (priority ignored)

- Profile-Undefined, High Priority — High Slope
- Out-of-Profile Low Slope (priority ignored)
- Profile-Undefined, Low Priority — Low Slope

#### Egress Slope Mapping

- In-Profile from ingress — High Slope
- Out-of-Profile from ingress — Low Slope

The specified policy contains a value that defines the queue's MBS value (queue-mbs). This is the maximum depth of the queue specified in bytes where all packets start to discard. The high and low priority RED slopes provide congestion control mechanisms that react to the current depth of the queue and start a random discard that increases in probability as the queue depth increases. The start point and end point for each discard probability slope is defined as follows:

- Start-Utilization — This is defined as percentage of MBS and specifies where the discard probability for the slope begins to rise above 0%. (A corresponding Start-Probability parameter is not needed as the start probability is always 0%.)
- Maximum-Utilization — This is also defined as a percentage of MBS and specifies where (based on MBS utilized) the discard probability rises to 100%. This is the first portion of the knee coordinates and is meaningless without the Maximum-Probability parameter.
- Maximum-Probability — This is defined as a percentage of discard probability and in conjunction with maximum-utilization completes the knee coordinate where the discard probability deviates from the slope and rises to 100%.

Up to 1024 HSMDA slope policies may be configured on a system.

The system maintains a slope policy named **hsmdda-default** which acts as a default policy when an explicit slope policy has not been defined for an HSMDA queue. The default policy may be edited, but it cannot be deleted. If a no slope-policy hsmdda-default command is executed, the default slope policy returns to the factory default settings. The factory default settings are as follows:

#### High Slope:

- Start-Utilization 100%
- Max-Utilization 100%
- Max-Probability 100%
- Shutdown

#### Low Slope:

- Start-Utilization 90%
- Max-Utilization 90%
- Max-Probability 1
- No Shutdown

Time-Average-Factor: 0

The **no** form of the command restores the association between the queue and the HSMDA default slope policy. The command has no immediate effect for queues that have a local override defined for the slope policy.

**Parameters** *hsmda-slope-policy-name* — Specifies an existing slope policy within the system. If a slope policy with the specified name does not exist, the slope-policy command will fail without modifying the slope behavior on the queue. Once a slope policy is associated with an HSMDA queue, the policy cannot be deleted.

**Default** hsmda-default

---

## Service Ingress QoS Policy Entry Commands

### action

**Syntax** **action** [**fc** *fc-name*] [**priority** {**high** | **low**}]  
**no action**

**Context** config>qos>sap-ingress>ip-criteria>entry  
config>qos>sap-ingress>ipv6-criteria>entry  
config>qos>sap-ingress>mac-criteria>entry

**Description** This mandatory command associates the forwarding class or enqueueing priority with specific IP, IPv6 or MAC criteria entry ID. The action command supports setting the forwarding class parameter to a sub-class. Packets that meet all match criteria within the entry have their forwarding class and enqueueing priority overridden based on the parameters included in the **action** parameters. When the forwarding class is not specified in the **action** command syntax, a matching packet preserves (or inherits) the existing forwarding class derived from earlier matches in the classification hierarchy. When the enqueueing priority is not specified in the action, a matching packet preserves (or inherits) the existing enqueueing priority derived from earlier matches in the classification hierarchy.

The **action** command must be executed for the match criteria to be added to the active list of entries. If the entry is designed to prevent more explicit (higher entry ID) entries from matching certain packets, the **fc** *fc-name* and **match** *protocol* fields should not be defined when executing action. This allows packets matching the entry to preserve the forwarding class and enqueueing priority derived from previous classification rules.

Each time action is executed on a specific entry ID, the previous entered values for **fc** *fc-name* and **priority** are overridden with the newly defined parameters or inherits previous matches when a parameter is omitted.

The **no** form of the command removes the entry from the active entry list. Removing an entry on a policy immediately removes the entry from all SAPs using the policy. All previous parameters for the action is lost.

**Default** Action specified by the **default-fc**.

**Parameters** **fc** *fc-name* — The value given for **fc** *fc-name* must be one of the predefined forwarding classes in the system. Specifying the **fc** *fc-name* is required. When a packet matches the rule, the forwarding class is only overridden when the **fc** *fc-name* parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

The sub-class-name parameter is optional and used with the *fc-name* parameter to define a preexisting sub-class. The *fc-name* and sub-class-name parameters must be separated by a period (dot). If sub-class-name does not exist in the context of *fc -name*, an error will occur. If sub-class-name is removed using the **no** *fc* *fc-name.sub-class-name* **force** command, the **default-fc** command will automatically drop the sub-class-name and only use *fc-name* (the parent forwarding class for the sub-class) as the forwarding class.

**Values**      `fc: class[.sub-class]`  
                   class: be, l2, af, l1, h2, ef, h1, nc  
                   sub-class: 29 characters max

**Default**     Inherit (When `fc fc-name` is not defined, the rule preserves the previous forwarding class of the packet.)

**priority** — The **priority** parameter overrides the default enqueueing priority for all packets received on a SAP using this policy that match this rule. Specifying the priority (**high** or **low**) is optional. When a packet matches the rule, the enqueueing priority is only overridden when the priority parameter is defined on the rule. If the packet matches and priority is not explicitly defined in the rule, the enqueueing priority is inherited based on previous rule matches.

**Default**      Inherit (When the **priority (high or low)** is not defined, the rule preserves the previous enqueueing priority of the packet)

**high** — The **high** parameter is used in conjunction with the **priority** parameter. Setting the **priority** enqueueing parameter to **high** for a packet increases the likelihood to enqueue the packet when the queue is congested. The enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the queue, the significance of the enqueueing priority is lost.

**low** — The **low** parameter is used in conjunction with the **priority** parameter. Setting the **priority** enqueueing parameter to **low** for a packet decreases the likelihood to enqueue the packet when the queue is congested. The enqueueing priority only affects ingress SAP queuing. When the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

**Default**      Inherit

## action

**Syntax**      `action [fc fc-name] [hsmda-counter-override counter-id] [profile {in | out}]`  
**no action**

**Context**      `config>qos>sap-egress>ip-criteria>entry`  
                   `config>qos>sap-egress>ipv6-criteria>entry`

**Description**    This command defines the reclassification actions that should be performed on any packet matching the defined IP flow criteria within the entries match node. When defined under the ip-criteria context, the reclassification only applies to IPv4 packets. When defined under the ipv6-criteria context, the reclassification only applies to IPv6 packets.

If an egress packet on the SAP matches the specified IP flow entry, the forwarding class, profile or HSMDA egress queue accounting behavior may be overridden. By default, the forwarding class and profile of the packet is derived from ingress classification and profiling functions. The default behavior for HSMDA queue accounting is to use the counters associated with the queue to which the packet is mapped. Matching an IP flow reclassification entry will override all IP precedence or DSCP based reclassification rule actions when an explicit reclassification action is defined for the entry.

When an IP flow entry is first created, the entry will have no explicit behavior defined as the reclassification actions to be performed. In show and info commands, the entry will display no action as the specified reclassification action for the entry. When the entry is defined with no action, the entry will not be populated in the IP flow reclassification list used to evaluate packets egressing a SAP with the SAP egress policy

defined. An IP flow reclassification entry is only added to the evaluation list when the action command for the entry is executed either with explicit reclassification entries or without any actions defined. Specifying action without any trailing reclassification actions allows packets matching the entry to exist the evaluation list without matching entries lower in the list. Executing no action on an entry removes the entry from the evaluation list and also removes any explicitly defined reclassification actions associated with the entry.

The **fc** keyword is optional. When specified, the egress classification rule will overwrite the forwarding class derived from ingress. The new forwarding class is used for egress remarking and queue mapping decisions.

The **profile** keyword is optional. When specified, the egress classification rule will overwrite the profile of the packet derived from ingress. The new profile value is used for egress remarking and queue congestion behavior.

The **hsmda-counter-override** keyword is optional. When specified and the egress SAP is created on an HSM DA, the egress classification rule will override the default queue accounting function for the packet. By default, the HSM DA uses each queues default queue counters for packets mapped to the queue. The **hsmda-counter-override** keyword is used to map the packet to an explicit exception counter. One of eight counters may be used. When the packet is mapped to an exception counter, the packet will not increment the queues discard or forwarding counters, instead the exception discard and forwarding counters will be used.

The **no** form of this command removes all reclassification actions from the IP flow reclassification entry and also removes the entry from the evaluation list. An entry removed from the evaluation list will not be matched to any packets egress a SAP associated with the SAP egress QoS policy.

**Default** Action specified by the **default-fc**.

**Parameters** **fc** *fc-name* — The **fc** reclassification action is optional. When specified, packets matching the IP flow reclassification entry will be explicitly reclassified to the forwarding class specified as *fc-name* regardless of the ingress classification decision. The *fc-name* defined must be one of the eight forwarding classes supported by the system. To remove the forwarding class reclassification action for the IP flow entry, the action command must be re-executed without the **fc** reclassification action defined.

**Values**      **fc:** class[.sub-class]  
                  class: be, l2, af, l1, h2, ef, h1, nc  
                  sub-class: 29 characters max

**Default**      none

**profile** {**in** | **out**} — The profile reclassification action is optional. When specified, packets matching the IP flow reclassification entry will be explicitly reclassified to either in-profile or out-of-profile regardless of the ingress profiling decision. To remove the profile reclassification action for the IP flow reclassification entry, the action command must be re-executed without the profile reclassification action defined.

**in** — The **in** parameter is mutually exclusive to the **out** parameter following the profile reclassification action keyword. Either **in** or **out** must be specified when the profile keyword is present. When **in** is specified, any packets matching the reclassification rule will be treated as in-profile by the egress forwarding plane.

**out** — The **out** parameter is mutually exclusive to the **in** parameter following the profile reclassification action keyword. Either **in** or **out** must be specified when the profile keyword is present. When **out** is specified, any packets matching the reclassification rule will be treated as out-of-profile by the egress forwarding plane.



**hsmda-counter-override** *counter-id* — The hsmda-counter-override reclassification action is optional and only has significance on SAPs which are created on an HSMDA. When specified, packets matching the IP precedence value will be mapped to the defined HSMDA exception counter-id for the packets queue group. The default behavior is to use the default counter on the queue group for the queue to which the packet is mapped. The hsmda-counter-override action may be overwritten by an ip-criteria reclassification rule match. The specified counter-id must be specified as an integer between 1 and 8. To remove the ESMDA exception counter reclassification action for the specified dscp-value, the dscp command must be re-executed without the hsmda-counter-override reclassification action defined.

**Values** 1—8

**Default** None

## entry

**Syntax** [no] entry *entry-id* [create]

**Context** config>qos>sap-ingress>ip-criteria  
config>qos>sap-egress>ip-criteria  
config>qos>sap-ingress>ipv6-criteria  
config>qos>sap-ingress>mac-criteria

**Description** This command is used to create or edit an IP or IPv6 criteria entry for the policy. Multiple entries can be created using unique *entry-id* numbers.

The list of flow criteria is evaluated in a top down fashion with the lowest entry ID at the top and the highest entry ID at the bottom. If the defined match criteria for an entry within the list matches the information in the egress packet, the system stops matching the packet against the list and performs the matching entries reclassification actions. If none of the entries match the packet, the IP flow reclassification list has no effect on the packet.

An entry is not populated in the list unless the action command is executed for the entry. An entry that is not populated in the list has no effect on egress packets. If the action command is executed without any explicit reclassification actions specified, the entry is populated in the list allowing packets matching the entry to exit the list, preventing them from matching entries lower in the list. Since this is the only flow reclassification entry that the packet matched and this entry explicitly states that no reclassification action is to be performed, the matching packet will not be reclassified.

The **no** form of this command removes the specified entry from the policy. Entries removed from the policy are immediately removed from all services where that policy is applied.

**Default** none

**Parameters** *entry-id* — The *entry-id*, expressed as an integer, uniquely identifies a match criterion and the corresponding action. It is recommended that multiple entries be given *entry-ids* in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

An entry cannot have any match criteria defined (in which case, everything matches) but must have at least the keyword **action fc** *fc-name* for it to be considered complete. Entries without the action keyword will be considered incomplete and hence will be rendered inactive.

**Default** none

**Values** 1— 65535

**create** — Required parameter when creating a flow entry when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the flow entry already exists.

## match

**Syntax** [no] match [protocol *protocol-id*]

**Context** config>qos>sap-egress>ip-criteria>entry  
config>qos>sap-ingress>ip-criteria>entry

**Description** This command creates a context to configure match criteria for SAP QoS policy match criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, then all criteria must be satisfied (AND function) before the action associated with the match is executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

It is possible that a SAP policy includes the **dscp** map command, the **dot1p** map command, and an IP match criteria. When multiple matches occur for the traffic, the order of precedence is used to arrive at the final action. The order of precedence is as follows:

1. 802.1p bits
2. DSCP
3. IP Quintuple or MAC headers

The **no** form of this command removes the match criteria for the *entry-id*.

**Parameters** **protocol** *protocol-id* — Specifies an IP protocol to be used as a SAP QoS policy match criterion.

The protocol type such as TCP / UDP / OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17).

**Values** protocol-id: 0 — 255 protocol numbers accepted in DHB  
keywords: none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-frag, ipv6-icmp, ipv6-no-nxt, ipv6-opts, ipv6-route, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp\* — udp/tcp wildcard

**Table 33: IP Protocol Names**

Protocol	Protocol ID	Description
icmp	1	Internet Control Message
igmp	2	Internet Group Management
ip	4	IP in IP (encapsulation)
tcp	6	Transmission Control
egp	8	Exterior Gateway Protocol
igp	9	any private interior gateway (used by Cisco for their IGRP)
udp	17	User Datagram
rdp	27	Reliable Data Protocol
ipv6	41	IPv6
ipv6-route	43	Routing Header for IPv6
ipv6-frag	44	Fragment Header for IPv6
idrp	45	Inter-Domain Routing Protocol
rsvp	46	Reservation Protocol
gre	47	General Routing Encapsulation
ipv6-icmp	58	ICMP for IPv6
ipv6-no-nxt	59	No Next Header for IPv6
ipv6-opts	60	Destination Options for IPv6
iso-ip	80	ISO Internet Protocol
eigrp	88	EIGRP
ospf-igp	89	OSPF/IGP
ether-ip	97	Ethernet-within-IP Encapsulation
encap	98	Encapsulation Header
pnni	102	PNNI over IP
pim	103	Protocol Independent Multicast
vrrp	112	Virtual Router Redundancy Protocol
l2tp	115	Layer Two Tunneling Protocol
stp	118	Schedule Transfer Protocol

**Table 33: IP Protocol Names (Continued)**

Protocol	Protocol ID	Description
ptp	123	Performance Transparency Protocol
isis	124	ISIS over IPv4
crtp	126	Combat Radio Transport Protocol
crudp	127	Combat Radio User Datagram

## match

**Syntax** **match** [**next-header** *next-header*]  
**no match**

**Context** config>qos>sap-ingress>ipv6-criteria>entry  
config>qos>sap-egress>ipv6-criteria>entry

**Description** This command creates a context to configure match criteria for ingress SAP QoS policy match IPv6 criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, then all criteria must be satisfied (AND function) before the action associated with the match is executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

It is possible that a SAP ingress policy includes the **dscp** map command, the **dot1p** map command, and an IPv6 match criteria. When multiple matches occur for the traffic, the order of precedence is used to arrive at the final action. The order of precedence is as follows:

1. 802.1p bits
2. DSCP
3. IP Quintuple or MAC headers

The **no** form of this command removes the match criteria for the *entry-id*.

**Parameters** **next-header** *next-header* — Specifies the next meader to match.

The protocol type such as TCP / UDP / OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17).

**Values** protocol numbers accepted in DHB: 0 — 42, 45 — 49, 52 — 59, 61 — 255

**keywords:** none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-icmp, ipv6-no-nxt, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp  
\* — udp/tcp wildcard

## match

<b>Syntax</b>	<b>match</b> [ <b>frame-type</b> { <b>802dot3</b>   <b>802dot2-llc</b>   <b>802dot2-snap</b>   <b>ethernet-II</b>   <b>atm</b> }] <b>no match</b>				
<b>Context</b>	config>qos>sap-ingress>mac-criteria>entry				
<b>Description</b>	<p>This command creates a context for entering/editing match MAC criteria for ingress SAP QoS policy match criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.</p> <p>If more than one match criteria (within one match statement) are configured then all criteria must be satisfied (AND function) before the action associated with the match will be executed.</p> <p>A <b>match</b> context can consist of multiple match criteria, but multiple <b>match</b> statements cannot be entered per entry.</p> <p>The <b>no</b> form of the command removes the match criteria for the <i>entry-id</i>.</p>				
<b>Parameters</b>	<p><b>frame-type</b> <i>keyword</i> — The <b>frame-type</b> keyword configures an Ethernet frame type or an ATM frame type to be used for the MAC filter match criteria.</p> <table><tr><td><b>Default</b></td><td>802dot3</td></tr><tr><td><b>Values</b></td><td>802dot3, 802dot2-llc, 802dot2-snap, ethernet_II, atm</td></tr></table> <p><b>802dot3</b> — Specifies the frame type is Ethernet IEEE 802.3.</p> <p><b>802dot2-llc</b> — Specifies the frame type is Ethernet IEEE 802.2 LLC.</p> <p><b>802dot2-snap</b> — Specifies the frame type is Ethernet IEEE 802.2 SNAP.</p> <p><b>ethernet_II</b> — Specifies the frame type is Ethernet Type II.</p> <p><b>atm</b> — Specifies the frame type as ATM cell. Note that the user is not allowed to configure entries with frame type of atm and a frame type of other supported values in the same QoS policy.</p>	<b>Default</b>	802dot3	<b>Values</b>	802dot3, 802dot2-llc, 802dot2-snap, ethernet_II, atm
<b>Default</b>	802dot3				
<b>Values</b>	802dot3, 802dot2-llc, 802dot2-snap, ethernet_II, atm				

## atm-vci

<b>Syntax</b>	<b>atm-vci</b> <i>vci-value</i> <b>no atm-vci</b>
<b>Context</b>	config>qos>sap-ingress>mac-criteria>entry>match
<b>Description</b>	<p>This command configures a VCI based filter entry in the SAP ingress QoS policy.</p> <p>This new criterion has only take affect when applied to a VPI SAP of an apipe VLL service of type atm-vpc. The application of this criterion to the ATM SAP of any other ATM VLL service, any other VLL service, VPLS service, or IES/VRPN service has no effect.</p> <p>The user is not allowed to configure a MAC matching criterion other than atm-vci once a MAC criteria filter entry that includes the frame type of atm has been configured.</p> <p>When the policy is applied to the ingress ATM VPI SAP of an atm-vpc VLL service and a received packet matches the VCI value configured in the atm-vci parameter, it is assigned the FC in the fc option of the action part of the filter. This determines which forwarding class queue this packet will be stored. Note that if</p>

the user entered a priority value in the priority option, it is ignored as the priority and profile of ATM VLL service packets is solely determined based on the ATM conformance definition configured in the ATM QoS traffic descriptor profile applied to this ATM SAP.

On egress ATM SAP, the Q-chip will queue the packet on the egress SAP queue corresponding to the packet's FC and forwards the packet to the ATM MDA (CMA). The ATM MDA (CMA) stores the individual cells in the VP queue corresponding to the SAP.

It is strongly recommended that the user does not enable cell-concatenation on the spoke-SDP when a VCI QoS filter is applied to the SAP. The filter will match against the VCI in the header of the first cell in the concatenated packet. Cell concatenation is disabled by default on a spoke-sdp of all ATM VLL service types.

The **no** form of this command removes the VCI value as the match criterion.

**Parameters**

*vci-value* — The value of the VCI field in the received ATM cell header.

**Values** 1, 2, 5 — 65535

---

## IP QoS Policy Match Commands

### dscp

<b>Syntax</b>	<b>dscp</b> <b>no dscp</b>
<b>Context</b>	config>qos>sap-ingress>ip-criteria>entry>match config>qos>sap-egress>ip-criteria>entry>match config>qos>sap-egress>fc config>qos>sap-ingress>ipv6-criteria>entry>match config>qos>sap-egress>ipv6-criteria>entry>match
<b>Description</b>	This command configures a DiffServ Code Point (DSCP) code point to be used for remarking of packets from the specified FC. If the optional in/out-profile is specified, the command will remark different DSCP code points depending on whether the packet was classified to be in or out-of-profile ingress to the node.  The <b>no</b> form of this command removes the DSCP match criterion.
<b>Default</b>	none
<b>Parameters</b>	<i>dscp-name</i> — Specifies a dscp name that has been previously mapped to a value using the <b>dscp-name</b> command. The DiffServ code point can only be specified by its name.  <b>Values</b> be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

### hsmda

<b>Syntax</b>	<b>hsmda</b>
<b>Context</b>	config>qos>sap-egress>fc
<b>Description</b>	This command defines how packets matching the forwarding class will be mapped to an HSMDA queue ID. The SAP QoS policies simultaneously support both standard service queue mappings and ESDMA queue mappings for the same forwarding class and the hsmda node is used to separate the HSMDA mappings from the standard mappings This allows the same QoS policy to be used on a standard MDA attached SAP and an HSMDA attached SAP.

## queue

**Syntax** `queue [1..8]`  
`no queue`

**Context** `config>qos>sap-egress>fc>hsmda`

**Description** This command specifies the HSM DA queue mapping for all packets in point-to-point services and unicast destined packets in multipoint services. Point-to-point services include `epipe` and other VLL type services. Multipoint services include IES, VPLS and VPRN services. The `queue` command does not apply to multicast, broadcast or unknown unicast packets within multipoint services (the `multicast`, `broadcast` and `unknown` commands must be used to define the queue mapping for non-unicast packets within a forwarding class). For `Epipes`, the `queue queue-id` mapping applies to all packets, regardless of the packets destination MAC address.

Each forwarding class has a default queue ID based on the intrinsic hierarchy between the forwarding classes as represented in [Table 34](#). Executing the `queue` command within the HSM DA context of a forwarding class with a different queue ID than the default overrides the default mapping. Multiple forwarding classes may be mapped to the same HSM DA queue ID.

**Table 34: Default FC HSM DA Queue ID Mappings**

Forwarding Class	Default HSM DA Queue ID
NC	queue 8
H1	queue 7
EF	queue 6
H2	queue 5
L1	queue 4
AF	queue 3
L2	queue 2
BE	queue 1

[Table 35](#) presents the way that packets are mapped to queues based on the type of service and the various forwarding types.

**Table 35: Ingress HSM DA Queue Mapping Behavior Based on Forwarding Type**

Service Type	Queue	Queue Mappings For Each Forwarding Type		
		Broadcast	Multicast	Unknown
Epipes	All packets matching the FC	None	None	None



**Table 35: Ingress HSMDA Queue Mapping Behavior Based on Forwarding Type**

Queue Mappings For Each Forwarding Type				
IES	All packets matching the FC	Packets with Broadcast DA	IP Multicast Packets	None
VPLS	All packets matching the FC	Packets with Broadcast DA	Packets with Multicast DA	Packets with Unicast DA but Unknown in FIB
VPRN	All packets matching the FC	Packets with Broadcast DA	IP Multicast Packets	None

The forwarding class queue mappings may be modified at anytime. The sub-forwarding classes inherit the parent forwarding classes queue mappings.

The no form of the command returns the HSMDA queue mapping for queue to the default mapping for the forwarding class.

**Parameters** *queue-id* — Configures a specific HSMDA queue.

<b>Values</b>	1 — 8
	BE Default: 1
	L2 Default: 2
	AF Default: 3
	L1 Default: 4
	H2 Default: 5
	EF Default: 6
	H1 Default: 7
	NC Default: 8

## dst-ip

**Syntax** **dst-ip** *{ip-address/mask | ip-address netmask}*  
**no dst-ip**

**Context** config>qos>sap-ingress>ip-criteria>entry>match  
 config>qos>sap-egress>ip-criteria>entry>match  
 config>qos>sap-ingress>ipv6-criteria>entry>match  
 config>qos>sap-egress>ipv6-criteria>entry>match

**Description** This command configures a destination address range to be used as a SAP QoS policy match criterion. To match on the destination address, specify the address and its associated mask, e.g., 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used.

The **no** form of this command removes the destination IP address match criterion.

**Default** none

**Parameters** *ip-address* — The IP address of the destination IP or IPv6 interface. This address must be unique within the subnet and specified in dotted decimal notation.

**Values**

ip-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
x:	[0 — FFFF]H
d:	[0 — 255]D

prefix-length: 1 — 128

## dst-port

**Syntax** **dst-port** {**lt** | **gt** | **eq**} *dst-port-number*  
**dst-port range** *start end*  
**no dst-port**

**Context** config>qos>sap-ingress  
config>qos>sap-ingress>ip-criteria>entry>match  
config>qos>sap-egress>ip-criteria>entry>match  
config>qos>sap-ingress>ipv6-criteria>entry>match  
config>qos>sap-egress>ipv6-criteria>entry>match

**Description** This command configures a destination TCP or UDP port number or port range for a SAP QoS policy match criterion.

The **no** form of this command removes the destination port match criterion.

**Default** none

**Parameters** **lt** | **gt** | **eq** *dst-port-number* — The TCP or UDP port numbers to match specified as less than (**lt**), greater than (**gt**) or equal to (**eq**) to the destination port value specified as a decimal integer.

**Values** 1 — 65535 (decimal)

**range** *start end* — The range of TCP or UDP port values to match specified as between the *start* and *end* destination port values inclusive.

**Values** 1 — 65535 (decimal)

## fragment

**Syntax** **fragment** {**true** | **false**}  
**no fragment**

**Context** config>qos>sap-ingress>ip-criteria>entry>match  
config>qos>sap-egress>ip-criteria>entry>match

**Description** This command configures fragmented or non-fragmented IP packets as a SAP QoS policy match criterion.

The **no** form of this command removes the match criterion.

**Default** fragment false

**Parameters** **true** — Configures a match on all fragmented IP packets. A match will occur for all packets that have either the MF (more fragment) bit set OR have the Fragment Offset field of the IP header set to a non-zero value.

**false** — Configures a match on all non-fragmented IP packets. Non-fragmented IP packets are packets that have the MF bit set to zero and have the Fragment Offset field also set to zero.

## src-ip

**Syntax** **src-ip** {*ip-address/mask* | *ip-address netmask*}  
**no src-ip**

**Context** config>qos>sap-ingress>ip-criteria>entry>match  
config>qos>sap-egress>ip-criteria>entry>match  
config>qos>sap-ingress>ipv6-criteria>entry>match  
config>qos>sap-egress>ipv6-criteria>entry>match

**Description** This command configures a source IP or IPv6 address range to be used as an SAP QoS policy match criterion.

To match on the source IP or IPv6 address, specify the address and its associated mask, e.g. 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used.

The **no** form of the command removes the source IP or IPv6 address match criterion.

**Default** No source IP match criterion.

**Parameters** *ip-address* | *ipv6-address* — The IP or IPv6 address of the source IP interface. This address must be unique within the subnet and specified in dotted decimal notation.

**Values** ip-address: a.b.c.d  
ipv6-address: x:x:x:x:x:x:x (eight 16-bit pieces)  
x:x:x:x:x.d.d.d  
x: [0 — FFFF]H  
d: [0 — 255]D  
prefix-length: 1 — 128

*mask* — The subnet mask length, expressed as an integer or in dotted decimal notation.

**Values** 0 — 32

*netmask* — Specify the subnet mask in dotted decimal notation.

**Values** a.b.c.d (dotted quad equivalent of mask length)

## src-port

**Syntax** **src-port** {**lt** | **gt** | **eq**} *src-port-number*  
**src-port range** *start end*  
**no src-port**

**Context** config>qos>sap-ingress>ip-criteria>entry>match  
config>qos>sap-egress>ip-criteria>entry>match  
config>qos>sap-ingress>ipv6-criteria>entry>match

**Description** This command configures a source TCP or UDP port number or port range for a SAP QoS policy match criterion.

The **no** form of this command removes the source port match criterion.

**Default** No src-port match criterion.

**Parameters** **lt** | **gt** | **eq** *src-port-number* — The TCP or UDP port numbers to match specified as less than (**lt**), greater than (**gt**) or equal to (**eq**) to the source port value specified as a decimal integer.

**Values** 1 — 65535 (decimal)

**range** *start end* — The range of TCP or UDP port values to match specified as between the *start* and *end* source port values inclusive.

**Values** 1 — 65535 (decimal)

---

## Service Ingress MAC QoS Policy Match Commands

### dot1p

**Syntax** **dot1p** *dot1p-value* [*dot1p-mask*]  
**no dot1p**

**Context** config>qos>sap-ingress>mac-criteria>entry

**Description** The IEEE 802.1p value to be used as the match criterion.  
Use the **no** form of this command to remove the dot1p value as the match criterion.

**Default** None

**Parameters** *dot1p-value* — Enter the IEEE 802.1p value in decimal.

**Values** 0 — 7

*dot1p-mask* — This 3-bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	D	4
Hexadecimal	0xH	0x4
Binary	0bBBB	0b100

To select a range from 4 up to 7 specify *p-value* of 4 and a *mask* of 0b100 for value and mask.

**Default** 7 (decimal) (exact match)

**Values** 1 — 7 (decimal)

### dsap

**Syntax** **dsap** *dsap-value* [*dsap-mask*]  
**no dsap**

**Context** config>qos>sap-ingress>mac-criteria>entry

**Description** Configures an Ethernet 802.2 LLC DSAP value or range for an ingress SAP QoS policy match criterion.  
This is a one-byte field that is part of the 802.2 LLC header of the IEEE 802.3 Ethernet Frame.  
The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and cannot be part of the same match criteria.  
Use the no form of this command to remove the dsap value as the match criterion.

**Default** None

**Parameters** *dsap-value* — The 8-bit dsap match criteria value in hexadecimal.

**Values** 0x00 — 0xFF (hex)

*dsap-mask* — This is optional and can be used when specifying a range of dsap values to use as the match criteria.

This 8 bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	DDD	240
Hexadecimal	0xHH	0xF0
Binary	0BBBBBBBB	0b11110000
<b>Default</b>	FF (hex) (exact match)	
<b>Values</b>	0x00 — 0xFF (hex)	

## dst-mac

**Syntax** **dst-mac** *ieee-address* [*ieee-address-mask*]  
**no dst-mac**

**Context** config>qos>sap-ingress>mac-criteria>entry

**Description** Configures a destination MAC address or range to be used as a Service Ingress QoS policy match criterion. The no form of this command removes the destination mac address as the match criterion.

**Default** none

**Parameters** *ieee-address* — The MAC address to be used as a match criterion.

**Values** HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

*ieee-address-mask* — A 48-bit mask to match a range of MAC address values.

This 48-bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHH	0xFFFFFFFF00000
Binary	0BBBBBBB...B	0b11110000...B

All packets with a source MAC OUI value of 00-03-FA subject to a match condition should be specified as: 0003FA000000 0x0FFFFFFF000000

**Default** 0xFFFFFFFFFFFF (hex) (exact match)

**Values** 0x0000000000000000 — 0xFFFFFFFFFFFF (hex)

## etype

**Syntax** **etype** *etype-value*  
**no etype**

**Context** config>qos>sap-ingress>mac-criteria>entry

**Description** Configures an Ethernet type II value to be used as a service ingress QoS policy match criterion.

The Ethernet type field is a two-byte field used to identify the protocol carried by the Ethernet frame. For e.g. 0800 is used to identify the IP v4 packets.

The Ethernet type field is used by the Ethernet version-II frames. IEEE 802.3 Ethernet frames do not use the type field. For IEEE 802.3 frames use the dsap, ssap or snap-pid fields as match criteria.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and cannot be part of the same match criteria.

The no form of this command removes the previously entered etype field as the match criteria.

**Default** None

**Parameters** *etype-value* — The Ethernet type II frame Ethertype value to be used as a match criterion expressed in hexadecimal.

**Values** 0x0600 — 0xFFFF

## inner-tag

**Syntax** **inner-tag** *value* [*vid-mask*]  
**no inner-tag**

**Context** config>qos>sap-ingress>mac-criteria>entry

**Description** This command configures the matching of the second tag that is carried transparently through the service. The inner-tag on ingress is the second tag on the frame if there are no service delimiting tags. Inner tag is the second tag before any service delimiting tags on egress but is dependent in the ingress configuration and may be set to 0 even in cases where additional tags are on the frame. This allows matching VLAN tags for explicit filtering or QoS setting when using default or null encapsulations.

The inner-tag is not applicable in ingress on dot1Q SAPs. The inner-tag may be populated on egress depending on the ingress SAP type.

On QinQ SAPs of null and default that do not strip tags inner-tag will contain the second tag (which is still the second tag carried transparently through the service.) On ingress SAPs that strip any tags, inner-tag will contain 0 even if there are more than 2 tags on the frame.

The optional `vid_mask` is defaulted to 4095 (exact match) but may be specified to allow pattern matching. The masking operation is  $((\text{value and vid-mask}) == (\text{tag and vid-mask}))$ . A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6.

Note for QoS the VID type cannot be specified on the default QoS policy.

The default `vid-mask` is set to 4095 for exact match.

## outer-tag

**Syntax** `outer-tag value [vid-mask]`  
`no outer-tag`

**Context** `config>qos>sap-ingress>mac-criteria>entry`

**Description** This command configures the matching of the first tag that is carried transparently through the service. Service delimiting tags are stripped from the frame and outer tag on ingress is the first tag after any service delimiting tags. Outer tag is the first tag before any service delimiting tags on egress. This allows matching VLAN tags for explicit filtering or QoS setting when using default or null encapsulations.

On dot1Q SAPs `outer-tag` is the only tag that can be matched. On dot1Q SAPs with exact match (`sap 2/1/1:50`) the `outer-tag` will be populated with the next tag that is carried transparently through the service or 0 if there is no additional VLAN tags on the frame.

On QinQ SAPs that strip a single service delimiting tag `outer-tag` will contain the next tag (which is still the first tag carried transparently through the service.) On SAPs with two service delimiting tags (two tags stripped) `outer-tag` will contain 0 even if there are more than 2 tags on the frame.

The optional `vid_mask` is defaulted to 4095 (exact match) but may be specified to allow pattern matching. The masking operation is  $((\text{value \& vid-mask}) == (\text{tag \& vid-mask}))$ . A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6.

Note for QoS the VID type cannot be specified on the default QoS policy.

The default `vid-mask` is set to 4095 for exact match.

## snap-oui

**Syntax** `snap-oui {zero | non-zero}`  
`no snap-oui`

**Context** `config>qos>sap-ingress>mac-criteria>entry`

**Description** Configures an IEEE 802.3 LLC SNAP Ethernet frame OUI zero or non-zero value to be used as a service ingress QoS policy match criterion.

The `no` form of this command removes the criterion from the match criteria.

**Default** none

**Parameters** `zero` — Specifies to match packets with the three-byte OUI field in the SNAP-ID set to zero.

`non-zero` — Specifies to match packets with the three-byte OUI field in the SNAP-ID not set to zero.



## snap-pid

**Syntax** **snap-pid** *snap-pid*  
**no snap-pid**

**Context** config>qos>sap-ingress>mac-criteria>entry

**Description** Configures an IEEE 802.3 LLC SNAP Ethernet frame PID value to be used as a service ingress QoS policy match criterion.

This is a two-byte protocol id that is part of the IEEE 802.3 LLC SNAP Ethernet Frame that follows the three-byte OUI field.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and cannot be part of the same match criteria.

Note: **snap-pid** match criteria is independent of the OUI field within the SNAP header. Two packets with different three-byte OUI fields but the same PID field will both match the same policy entry based on a snap-pid match criteria.

The **no** form of this command removes the snap-pid value as the match criteria.

**Default** none

**Parameters** *snap-pid* — The two-byte snap-pid value to be used as a match criterion in hexadecimal.

**Values** 0x0000 — 0xFFFF

## src-mac

**Syntax** **src-mac** *ieee-address* [*ieee-address-mask*]  
**no src-mac**

**Context** config>qos>sap-ingress>mac-criteria>entry

**Description** This command configures a source MAC address or range to be used as a service ingress QoS policy match criterion.

The **no** form of this command removes the source mac as the match criteria.

**Default** none

**Parameters** *ieee-address* — Enter the 48-bit IEEE mac address to be used as a match criterion.

**Values** HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

*ieee-address-mask* — This 48-bit mask can be configured using:

This 48 bit mask can be configured using the following formats

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDDD	281474959933440

Format Style	Format Syntax	Example
Hexadecimal	0xHHHHHHHHHHHH	0x0FFFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

To configure all packets with a source MAC OUI value of 00-03-FA are subject to a match condition, then the entry should be specified as: 003FA000000 0xFFFFFFFF000000

**Default** 0xFFFFFFFFFFFF (hex) (exact match)  
**Values** 0x00000000000000 — 0xFFFFFFFFFFFF (hex)

## ssap

**Syntax** **ssap** *ssap-value* [*ssap-mask*]  
**no ssap**

**Context** config>qos>sap-ingress>mac-criteria>entry

**Description** This command configures an Ethernet 802.2 LLC SSAP value or range for an ingress SAP QoS policy match criterion.

This is a one-byte field that is part of the 802.2 LLC header of the IEEE 802.3 Ethernet Frame.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and cannot be part of the same match criteria.

The **no** form of this command removes the ssap match criterion.

**Default** none

**Parameters** *ssap-value* — The 8-bit ssap match criteria value in hex.

**Values** 0x00 — 0xFF (hex)

*ssap-mask* — This is optional and can be used when specifying a range of ssap values to use as the match criteria.

This 8 bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	DDD	240
Hexadecimal	0xHH	0xF0
Binary	0bBBBBBBBB	0b11110000

**Default** none

**Values** 0x00 — 0xFF

---

## Service Egress QoS Policy Forwarding Class Commands

### fc

**Syntax**    *fc fc-name*  
              *no fc fc-name*

**Context**    config>qos>sap-egress

**Description**    The **fc** *fc-name* node within the SAP egress QoS policy is used to contain the explicitly defined queue mapping and dot1p marking commands for *fc-name*. When the mapping for *fc-name* points to the default queue and the dot1p marking is not defined, the node for *fc-name* is not displayed in the **show configuration** or **save configuration** output unless the detail option is specified.

The **no** form of the command removes the explicit queue mapping and dot1p marking commands for *fc-name*. The queue mapping reverts to the default queue for *fc-name* and the dot1p marking (if appropriate) uses the default of 0.

**Default**    none

**Parameters**    *fc-name* — This parameter specifies the forwarding class queue mapping or dot1p marking is to be edited. The value given for *fc-name* must be one of the predefined forwarding classes in the system.

**Values**        be, 12, af, 11, h2, ef, h1, nc

### policer

**Syntax**        **policer** *policer-id* [ **{**port-redirect-group-queue **}** [queue *queue-id*] | **{**group *queue-group-name* **}** **[**instance *instance-id* **]** [queue *group-queue-id*] **]**  
**no policer**

**Context**        config>qos>sap-egress>fc

**Description**    Within a sap-egress QoS policy forwarding class context, the policer command is used to map packets that match the forwarding class to the specified policer-id. The specified policer-id must already exist within the sap-egress QoS policy. The forwarding class of the packet is first discovered at ingress based on the ingress classification rules. When the packet arrives at egress, the sap-egress QoS policy may match a forwarding class reclassification rule which overrides the ingress derived forwarding class. The forwarding class context within the sap-egress QoS policy is then used to map the packet to an egress queue (using the queue *queue-id*, or port-redirect-group queue *queue-id*, or group *queue-group-name* instance *instance-id* queue *queue-id* commands) or an egress policer (*policer policer-id*). The queue and policer commands within the forwarding class context are mutually exclusive. By default, the forwarding class is mapped to the SAP egress default queue (queue 1). If the policer *policer-id* command is executed, any previous policer mapping or queue mapping for the forwarding class is overridden if the policer mapping is successful.

A policer defined within the sap-egress policy is not actually created on an egress SAP or a subscriber using an sla-profile where the policy is applied until at least one forwarding class is mapped to the policer. If insufficient policer resources exist to create the policer for a SAP or subscriber or egress policing is not

supported on the port associated with the SAP or subscriber, the initial forwarding class mapping will fail.

Packets that are mapped to an egress policer that are not discarded by the policer must be placed into a default queue on the packets destination port. The system uses egress port queue groups for this purpose. An egress queue group named `policer-output-queues` is automatically created on each port that support egress policers. By default, the system uses the forwarding class mappings within this queue group to decide which queue within the group will receive each packet output from the policer. This default policer output queuing behavior may be overridden for non-subscriber packets by redirection to a queue group. The name and instance of the queue group to redirect to is either specified in the QoS policy itself, or the fact that a forwarding class must be redirected is simply identified in the QoS policy and the specific queue group instance is only identified at the time the QoS policy is applied:

- If the **policer** *policer-id* command is successfully executed, the default egress queuing is performed for the forwarding class using the `policer-output-queues` queue group and the *queue-id* within the group based on the forwarding class map from the group template
- If the **policer** *policer-id* **queue** *queue-id* command is successfully executed, the specified SAP *queue-id* within egress QoS policy is used instead of the default policer output queues.
- If the **policer** *policer-id* **port-redirect-group-queue** keyword is successfully executed, the system will map the forwarding class to the queue within the egress queue group instance specified at the time the QoS policy is applied to the SAP, using the forwarding class map from the queue group template.
- If the **policer** *policer-id* **port-redirect-group queue** *queue-id* command is successfully executed, the system will map the forwarding class to the configured *queue-id* within the egress queue group instance that is specified at the time the QoS policy is applied to the SAP (ignoring using the forwarding class map from the queue group template).
- If the **policer** *policer-id* **group** *queue-group-name* **instance** *instance-id* command is successfully executed, the system will map the forwarding class to the queue within the specified egress queue group instance using the forwarding class map from the group template.
- If the **policer** *policer-id* **group** *queue-group-name* **instance** *instance-id* **queue** *queue-id* command is successfully executed, the system will map the forwarding class to the specified *queue-id* within the specified egress queue group instance (ignoring the forwarding class map in the group template).

If the specified **group** *queue-group-name* is not defined as an egress queue-group-template, the **policer** command will fail. Further, if the specified group does not exist on the port for the SAPs or subscribers associated with the **sap-egress** QoS policy, the policer command will fail. While a group *queue-group-name* is specified in a **sap-egress** QoS policy, the groups corresponding egress template cannot be deleted. While a port egress queue group is associated with a policer instance, the port queue group cannot be deleted.

If the specified queue group-*queue-id* is not defined in the egress queue-group-template *queue-group-name*, the policer command will fail. While a *queue-id* within an egress queue group template is referenced by a **sap-egress** QoS policy forwarding class policer command, the queue cannot be deleted from the queue group template.

If an egress policed packet is discarded by the egress port queue group queue, the source policer discard stats are incremented. This means that the discard counters for the policer represent both the policer discard events and the destination queue drop tail events associated with the policer.

The **no** form of this command is used to restore the mapping of the forwarding class to the default queue. If all forwarding classes have been removed from the default queue, the queue will not exist on the SAPs or subscribers associated with the QoS policy and the **no** policer command will cause the system to attempt to

create the default queue on each object. If the system cannot create the default queue in each instance, the no policer command will fail and the forwarding class will continue its mapping to the existing *policer-id*. If the no policer command results in a policer without any current mappings, the policer will be removed from the SAPs and subscribers associated with the QoS policy. All statistics associated with the policer on each SAP and subscribers will be lost.

**Default** none

**Parameters** *policer-id* — When the forwarding class **policer** command is executed, a valid *policer-id* must be specified. The parameter *policer-id* references a *policer-id* that has already been created within the **sap-ingress** QoS policy.

**Values** 1—8

**Default** none

**port-redirect-group-queue** — Used to override the forwarding class default egress queue destination to an egress port queue group. The specific egress queue group instance to use is specified at the time that the QoS policy is applied to the SAP. Therefore, this parameter is only valid if SAP based redirection is required.

**queue queue-id** — This parameter overrides the forwarding class default egress queue destination to a specified *queue-id*. If port-redirect-group is not configured, then this will be a local SAP queue of that *queue-id*. A queue of ID *queue-id* must exist within the egress QoS policy. If **port-redirect-group-queue** is configured then the **queue queue-id** in the egress port queue group instance is used.

**Values** 1—8

**Default** Derived from forwarding class assignment in queue-group definition.

**group queue-group-name** — The **group queue-group-name** is optional and is used to override the forwarding class's default egress queue destination. If the queue group-queue-id parameter is not specified, the forwarding class map within the specified group's template is used to derive which queue within the group will receive the forwarding class's packets. An egress queue group template must exist for the specified queue-group-name or the policer command will fail. The specified queue-group-name must also exist as an egress queue group on the ports where SAPs and subscribers associated with the sap-egress policy is applied or the policer command will fail.

**Values** Any qualifying egress queue group name

**Default** **policer-output-queues**

**queue group-queue-id** — The **queue group-queue-id** is optional when the group queue-group-name parameter is specified and is used to override the forwarding class mapping within the group's egress queue group template. The specified group-queue-id must exist within the group's egress queue group template or the policer command will fail.

**Values** 1—8

**Default** Derived from forwarding class assignment in queue-group definition

**instance** *instance-id* — This parameter is used to specify the specific instance of a queue group with template queue-group-name to which this queue should be redirected. This parameter is only valid for queue groups on egress ports where policy based redirection is required.

**Values** 1 — 40960

**Default** 1

## description

**Syntax** **description** *description string*  
**no description**

**Context** config>qos>sap-egress>policer

**Description** The **description** command is used to define an informational ASCII string associated with the policer control policy. The string value can be defined or changed at any time once the policy exists.

The **no** form of this command is used to remove an explicit description string from the policer.

**Default** no description

**Parameters** *description string* — The *description-string* parameter defines the ASCII description string for the policer control policy. The description string can be up to 80 characters. If the string contains spaces, it must be placed within beginning and ending double quotation marks. Beginning and ending quotation marks are not considered part of the description string. Only printable ASCII characters are allowed in the string. The sting does not need to be unique and may be repeated in the descriptions for other policer control policies or other objects. If the command is executed without the description-sting present, any existing description string will be unaffected.

**Default** None

## adaptation-rule

**Syntax** **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]  
**no adaptation-rule**

**Context** config>qos>sap-egress>policer

**Description** This command is used to define how the policer's configuration parameters are translated into the underlying hardware capabilities used to implement each policer instance. For instance, the configured rates for the policer need to be mapped to the timers and decrement granularity used by the hardware's leaky bucket functions that actually perform the traffic metering. If a rate is defined that cannot be exactly matched by the hardware, the adaptation-rule setting provides guidance for which hardware rate should be used.

The **max** keyword tells the system that the defined rate is the maximum rate that should be given to the policer. If the hardware cannot exactly match the given rate, the next lowest hardware supported rate is used.

The **min** keyword tells the system that the defined rate is the minimum rate that should be given to the policer. If the hardware cannot exactly match the given rate, the next highest hardware supported rate is

used.

The **closest** keyword tells the system that the defined rate is the target rate for the policer. If the hardware cannot exactly match the given rate, the system will use the closest hardware supported rate compared to the target rate.

The hardware also needs to adapt the given mbs and cbs values into the PIR bucket violate threshold (discard) and the CIR bucket exceed threshold (out-of-profile). The hardware may not have an exact threshold match which it can use. In R8.0, the system treats the mbs and cbs values as minimum threshold values.

The **no** form of this command is used to return the policer's metering and profiling hardware adaptation rules to closest.

<b>Parameters</b>	<p><b>pir</b> {<b>max</b>   <b>min</b>   <b>closest</b>} — When the optional <b>pir</b> parameter is specified, the <b>max</b>, <b>min</b> or <b>closest</b> keyword qualifier must follow.</p> <p><b>max</b> — The <b>max</b> keyword is used to inform the system that the metering rate defined for the policer is the maximum allowed rate. The system will choose a hardware supported rate that is closest but not exceeding the specified rate.</p> <p><b>min</b> — The <b>min</b> keyword is used to inform the system that the metering rate defined for the policer is the minimum allowed rate. The system will choose a hardware supported rate that is closest but not lower than the specified rate.</p> <p><b>closest</b> — The <b>closest</b> keyword is used to inform the system that the metering rate defined for the policer is the target rate. The system will choose a hardware supported rate that is closest to the specified rate.</p> <p><b>Default</b>      closest</p> <p><b>cir</b> {<b>max</b>   <b>min</b>   <b>closest</b>} — When the optional <b>cir</b> parameter is specified, the <b>max</b>, <b>min</b> or <b>closest</b> keyword qualifier must follow.</p> <p><b>max</b> — The <b>max</b> keyword is used to inform the system that the profiling rate defined for the policer is the maximum allowed rate. The system will choose a hardware supported rate that is closest but not exceeding the specified rate.</p> <p><b>min</b> — The <b>min</b> keyword is used to inform the system that the profiling rate defined for the policer is the minimum allowed rate. The system will choose a hardware supported rate that is closest but not lower than the specified rate.</p> <p><b>closest</b> — The <b>closest</b> keyword is used to inform the system that the profiling rate defined for the policer is the target rate. The system will choose a hardware supported rate that is closest to the specified rate.</p> <p><b>Default</b>      closest</p>
-------------------	--

## cbs

<b>Syntax</b>	<b>cbs</b> { <i>size</i> [bytes   kilobytes]   <b>default</b> } <b>no cbs</b>
<b>Context</b>	config>qos>sap-egress>policer
<b>Description</b>	This command is used to configure the policer's CIR leaky bucket's exceed threshold. The CIR bucket's exceed threshold represents the committed burst tolerance allowed by the policer. If the policer's forwarding

rate is equal to or less than the policer's defined CIR, the CIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the forwarding rate increases beyond the profiling rate, the amount of data allowed to be in-profile above the rate is capped by the threshold.

The policer's **cbs** size defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The no form of this command returns the policer to its default CBS size.

**Default**    **cbs 16 kilobytes**

**Parameters**    *size* [**bytes** | **kilobytes**] — The *size* parameter is required when specifying **cbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

**byte** — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

**kilobyte** — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

**Values**        1 — 4194304

**Default**       **kilobyte**

## high-prio-only

**Syntax**        **high-prio-only** *percent-of-mbs*  
**no high-prio-only**

**Context**        config>qos>sap-egress>policer

**Description**    This command is used to configure the percentage of the policer's PIR leaky bucket's MBS (maximum burst size) that is reserved for high priority traffic. While the **mbs** value defines the policer's high priority violate threshold, the percentage value defined is applied to the **mbs** value to derive the bucket's low priority violate threshold. See the **mbs** command details for information on which types of traffic is associated with each violate threshold.

**Default**        **high-prio-only 10**

**Parameters**    *percent-of-mbs* — The *percent-of-mbs* parameter is required when specifying **high-prio-only** and is expressed as a percentage with granularity of 1,000th of a percent.

**Values**        0—100

**Default**        10



## mbs

**Syntax** **mbs** {*size* [bytes | kilobytes] | default}  
**no mbs**

**Context** config>qos>sap-egress>policer  
config>qos>sap-egress>hsm-da-queues>queues

**Description** This command is used to configure the policer's PIR leaky bucket's high priority violate threshold. The **high-prio-only** command is applied to the MBS value to derive the bucket's low priority violate threshold. For egress, trusted in-profile packets and un-trusted high priority packets use the policer's high priority violate threshold while trusted out-of-profile and un-trusted low priority packets use the policer's low priority violate threshold. At egress, in-profile packets use the policer's high priority violate threshold and out-of-profile packets use the policer's low priority violate threshold.

The PIR bucket's violate threshold represent the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by **high-prio-only** is available for the higher priority traffic.

The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied.

The no form of this command returns the policer to its default MBS size.

**Default** None

**Parameters** *size* [bytes | kilobytes] — The *size* parameter is required when specifying **mbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

**byte** — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

**kilobyte** — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

**Values** 1—3932160

**Default** kilobyte

## packet-byte-offset

**Syntax** **packet-byte-offset** {add bytes | subtract bytes}  
**no packet-byte-offset**

**Context** config>qos>sap-egress>policer

**Description** This command is used to modify the size of each packet handled by the policer by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth

impact is changed. The **packet-byte-offset** command is meant to be an arbitrary mechanism that can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the policing metering and profiling throughput is affected by the offset as well as the stats associated with the policer.

When child policers are adding to or subtracting from the size of each packet, the parent policer's **min-thresh-separation** value should also need to be modified by the same amount.

The policer's **packet-byte-offset** defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** version of this command is used to remove per packet size modifications from the policer.

#### Parameters

**add bytes** — The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

**Values** 0—31

**Default** None

**subtract bytes** — The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When **b** is defined the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

**Values** 0—31

**Default** None

## parent

**Syntax** **parent** {**root** | *arbiter-name*} [**level** *level*] [**weight** *weight-within-level*]  
**no parent**

**Context** config>qos>sap-egress>policer

#### Description

This command is used to create a child to parent mapping between each instance of the policer and either the **root** arbiter or a specific tiered **arbiter** on the object where the policy is applied. Defining a parent association for the policer causes the policer to compete for the parent policer's available bandwidth with other child policers mapped to the policer control hierarchy.

Policer control hierarchies may be created on SAPs or on a subscriber context. To create a policer control hierarchy on an ingress or egress SAP context, a **policer-control-policy** must be applied to the SAP. Once applied, the system will create a parent policer that is bandwidth limited by the policy's **max-rate** value under the root arbiter. The root arbiter in the policy also provides the information used to determine the various priority level discard-unfair and discard-all thresholds. Besides the root arbiter, the policy may also contain user defined tiered arbiters that provide arbitrary bandwidth control for subsets of child policers that are either directly or indirectly parented by the arbiter.

When the QoS policy containing the policer with a **parent** mapping to an arbiter name exists on the SAP, the system will scan the available arbiters on the SAP. If an arbiter exists with the appropriate name, the policer

to arbiter association is created. If the specified arbiter does not exist either because a **policer-control-policy** is not currently applied to the SAP or the arbiter name does not exist within the applied policy, the policer is placed in an orphan state. Orphan policers operate as if they are not parented and are not subject to any bandwidth constraints other than their own PIR. When a policer enters the orphan state, it is flagged as operationally degraded due to the fact that it is not operating as intended and a trap is generated. Whenever a **policer-control-policy** is added to the SAP or the existing policy is modified, the SAP's policer's parenting configurations must be reevaluated. If an orphan policer becomes parented, the degraded flag should be cleared and a resulting trap should be generated.

For subscribers, the policer control hierarchy is created through the **policer-control-policy** applied to the **sub-profile** used by the subscriber. A unique policer control hierarchy is created for each subscriber associated with the **sub-profile**. The QoS policy containing the policer with the parenting command comes into play through the subscriber **sla-profile** which references the QoS policy. The combining of the **sub-profile** and the **sla-profile** at the subscriber level provides the system with the proper information to create the policer control hierarchy instance for the subscriber.

Executing the **parent** command will fail if:

- The policer's stat-mode in the QoS policy is set to no-stats
- A stat-mode no-stats override exists on an instance of the policer on a SAP or subscribers context

A policer with a **parent** command applied cannot be configured with **stat-mode no-stats** in either the QoS policy or as an override on an instance of the policer

Once a policer is successfully parented to an arbiter, the **parent** commands **level** and **weight** parameters are used to determine at what priority level and at which weight in the priority level that the child policer competes with other children (policers or other arbiters) for bandwidth.

The **no** form of this command is used to remove the parent association from all instances of the policer.

## Parameters

**{root | arbiter-name}** — When the **parent** command is executed, either the keyword **root** or an *arbiter-name* must be specified.

**root** — The **root** keyword specifies that the policer is intended to become a child to the **root** arbiter where an instance of the policer is created. If the **root** arbiter does not exist, the policer will be placed in the orphan state.

*arbiter-name* — The *arbiter-name* parameter specifies that the policer is intended to become a child to one of the tiered arbiters with the given arbiter-name where an instance of the policer is created. If the specified arbiter-name does not exist, the policer will be placed in the orphan state.

**Default**     None

**weight** *weight-within-level* — The **weight** *weight-within-level* keyword and parameter are optional when executing the **parent** command. When **weight** is not specified, a default level of 1 is used in the parent arbiters priority level. When **weight** is specified, the *weight-within-level* parameter must be specified as an integer value from 1 through 100.

**Default**     1

## rate

**Syntax** `rate {max | kilobits-per-second} [cir {max | kilobits-per-second}]  
no rate`

**Context** `config>qos>sap-egress>policer`

**Description** This command is used to configure the policer's metering and optional profiling rates. The metering rate is used by the system to configure the policer's PIR leaky bucket's decrement rate while the profiling rate configures the policer's CIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on the each packets size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches it's exceed (CIR) or violate (PIR) threshold. The **cbs**, **mbs** and **high-prio-only** commands are used to configure the policer's PIR and CIR thresholds.

If a packet arrives at the policer while the bucket's depth is less than the threshold associated with the packet, the packet is considered to be conforming to the bucket's rate. If the bucket depth is equal to or greater than the threshold, the packet is considered to be in the exception state. For the CIR bucket, the exception state is exceeding the CIR rate while the PIR bucket's exception state is violating the PIR bucket rate. If the packet is violating the PIR, the packet is marked red and will be discarded. If the packet is not red, it may be green or yellow based on the conforming or exceeding state from the CIR bucket.

When a packet is red neither the PIR or CIR bucket depths are incremented by the packets size. When the packet is yellow the PIR bucket is incremented by the packet size, but the CIR bucket is not. When the packet is green, both the PIR and CIR buckets are incremented by the packet size. This ensures that conforming packets impact the bucket depth while exceeding or violating packets do not.

The policer's **adaptation-rule** command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.

By default, the policer's metering rate is **max** and the profiling rate is 0 Kbps (all packets out-of-profile).

The **rate** settings defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command is used to restore the default metering and profiling rate to a policer.

**Parameters** `{max | kilobits-per-second}` — Specifying the keyword **max** or an explicit *kilobits-per-second* parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the policer is first created, the metering rate defaults to max. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second.

**Values** **max** or 0—20,000,000

**cir** `{max | kilobits-per-second}` — The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit *kilobits-per-second* parameter directly following the cir keyword is required and identifies the policer's profiling rate for the CIR leaky bucket. When the policer is first created, the profiling rate defaults to 0 Kbps. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second.

**Values** **max** or 0—20,000,000

## stat-mode

**Syntax** **stat-mode** {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-profile-cir** | **offered-total-cir** | **offered-profile-capped-cir** | **offered-limited-capped-cir**}  
**no stat mode**

**Context** config>qos>sap-egress>policer  
config>qos>queue-group-templates>egress>queue-group

**Description** The **sap-egress** QoS policy's policer **stat-mode** command is used to configure the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. An egress policer has multiple types of offered packets (soft in-profile and out-of-profile from ingress and hard in-profile and out-of-profile due to egress profile overrides) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the potential large number of egress policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly re-profiled offered packets. The **stat-mode** command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters.

While a **no-stats** mode is supported which prevents any packet accounting, the use of the policer's **parent** command requires at the policer's **stat-mode** to be set at least to the **minimal** setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. Once a policer has been made a child to a parent policer, the **stat-mode** cannot be changed to **no-stats** unless the policer parenting is first removed.

Each time the policer's **stat-mode** is changed, any previous counter values are lost and any new counters are set to zero.

Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. You can view the the total/allocated/free stats by using the **tools dump system-resources** command. If insufficient counters exist to implement a mode on any policer instance, the **stat-mode** change will fail and the previous mode will continue unaffected for all instances of the policer.

The default **stat-mode** when a policer is created within the policy is **minimal**.

The **stat-mode** setting defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the **stat-mode** override command will fail. The previous **stat-mode** setting active for the policer will continue to be used by the policer.

The **no** form of this command attempts to return the policer's **stat-mode** setting to **minimal**. The command will fail if insufficient policer counter resources exist to implement **minimal** where the QoS policer is currently applied and has a forwarding class mapping.

**Parameters** **no-stats** — Counter resource allocation:0

The policer does not have any forwarding plane counters allocated and cannot provide offered, discard and forward statistics. A policer using **no-stats** cannot be a child to a parent policer and the policer's **parent** command will fail.

When **collect-stats** is enabled, the lack of counters causes the system to generate the following statistics:

a. offered-in = 0

- b. offered-out = 0
- c. discard-in = 0
- d. discard-out = 0
- e. forward-in = 0
- f. forward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

**minimal** — Counter resource allocation:1

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (profile or priority) and do not count green or yellow output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.

This counter mode is useful when only the most basic accounting information is required.

The counters are used in the following manner:

- 1. offered = soft-in-profile-out-of-profile, profile in/out
- 2. discarded = Same as 1
- 3. forwarded = Derived from 1 - 2

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 0
- c. discard-in = 2
- d. discard-out = 0
- e. forward-in = 3
- f. forward-out = 0

Counter 0 indicates that the accounting statistic returns a value of zero.

**offered-profile-no-cir** — Counter resource allocation:2

The **offered-profile-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-profile-no-cir** mode is most useful when profile based offered, discard and forwarding stats are required from the egress policer, but a CIR is not being used to recolor the soft in-profile and out-of-profile packets. This mode does not prevent the policer from being configured with a CIR rate.

The counters are used in the following manne:

- 1. offered-in = soft-in-profile, profile in
- 2. offered-out = soft-out-of-profile, profile out
- 3. dropped-in = Same as 1

- 4. dropped-out = Same as 2
- 5. forwarded-in = Derived from 1 - 3
- 6. forwarded-out = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 2
- c. discard-in = 3
- d. discard-out = 4
- e. forward-in = 5
- f. forward-out = 6

**offered-profile-cir** — Counter resource allocation: 3

The **offered-profile-cir** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-profile-cir** mode is most useful when profile based offered, discard and forwarding stats are required from the egress policer and a CIR rate is being used to recolor the soft in-profile and out-of-profile packets.

The counters are used in the following manner:

- 1. offered-in-that-stayed-green-or-turned-red = profile in
- 2. offered-soft-that-turned-green = soft-in-profile-out-of-profile
- 3. offered-soft-or-out-that-turned-yellow-or-red = soft-in-profile-out-of-profile, profile out
- 4. dropped-in-that-stayed-green-or-turned-red = Same as 1
- 5. dropped-soft-that-turned-green = Same as 2
- 6. dropped-soft-or-out-that-turned-yellow-or-red = Same as 3
- 7. forwarded-in-that-stayed-green = Derived from 1 - 4
- 8. forwarded-soft-that-turned-green = Derived from 2 - 5
- 9. forwarded-soft-or-out-that-turned-yellow = Derived from 3 - 6

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1
- b. offered-out = 2 + 3
- c. discard-in = 4
- d. discard-out = 5 + 6
- e. forward-in = 7 + 8
- f. forward-out = 9

**offered-total-cir** — Counter resource allocation:2

The **offered-total-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-total-cir** mode is most useful when the policer is not receiving trusted in-profile or out-of-profile traffic and both high and low priority classifications are not being used on the un-trusted packets and the offered packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile or out-of-profile packets and does not prevent the use of priority high or low classifications on the un-trusted packets.

The counters are used in the following manner:

1. offered-that-turned-green = soft-in-profile-out-of-profile, profile in/out
2. offered- that-turned-yellow-or-red = soft-in-profile-out-of-profile, profile in/out
3. dropped-offered-that-turned-green = Same as 1
4. dropped-offered-that-turned-yellow-or-red = Same as 2
5. forwarded-offered-that-turned-green = Derived from 1 - 3
6. forwarded-offered-that-turned-yellow = Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. offered-in = 1 + 2
- b. offered-out = 0
- c. discard-in = 3
- d. discard-out = 4
- e. forward-in = 5
- f. forward-out = 6

Counter 0 indicates that the accounting statistic returns a value of zero.

**offered-profile-capped-cir** — Counter resource allocation:2

When **offered-profile-capped-cir** is defined, the system creates five offered-output counters in the forwarding plane and five discard counters in the traffic manager.

The **offered-profile-capped-cir** mode is similar to the **offered-profile-cir** mode except that it includes support for **profile in** and **soft-in-profile** that may be output as ‘out-of-profile’ due to enabling profile-capped mode on the ingress policer.

The impact of using **offered-profile-capped-cir** stat-mode while **profile-capped** mode is disabled are that one of the counting resources in the forwarding plane and traffic manager will not be used and soft-in-profile will be treated as ‘offered-in’ instead of ‘offered-undefined’.

The counters are used in the following manner:

1. ‘offered-in-that-stayed-green’ = profile in, soft-in-profile
2. ‘offered-in-that-turned-yellow-or-red’ = profile in, soft-in-profile
3. ‘offered-soft-out-that-turned-green’ = soft-out-of-profile
4. ‘offered-soft-out- that-turned-yellow-or-red’ = soft-out-of-profile



5. 'offered-out-that-turned-yellow-or-red' = profile out
6. 'dropped-in-that-stayed-green' = Same as 1
7. 'dropped-in-that-turned-yellow-or-red' = Same as 2
8. 'dropped-soft-out-that-turned-green' = Same as 3
9. 'dropped-soft-out-that-turned-yellow-or-red' = Same as 4
10. 'dropped-out-that-turned-yellow-or-red' = Same as 5
11. 'forwarded-in-that-stayed-green' = Derived from 1 - 6
12. 'forwarded-in-that-turned-yellow' = Derived from 2 - 7
13. 'forwarded-soft-out-that-turned-green' = Derived from 3 - 8
14. 'forwarded-soft-out-that-turned-yellow' = Derived from 4 - 9
15. 'forwarded-out-that-turned-yellow' = Derived from 5 - 10

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

- a. 'offered-undefined' = 3 + 4
- b. 'offered-in' = 1 + 2
- c. 'offered-out' = 5
- d. 'discard-in' = 6 + 8
- e. 'discard-out' = 7 + 9 + 10
- f. 'forward-in' = 11 + 13
- g. 'forward-out' = 12 + 14 + 15

#### **offered-limited-capped-cir** — Counter resource allocation:2

**offered-limited-capped-cir** is defined, the system creates four forwarding plane offered-output counters in the network processor and three discard counters in the traffic manager.

The **offered-limited-capped-cir** mode is similar to the **offered-profile-capped-cir** mode except that it combines **profile out** and **soft-out-of-profile** and eliminates the 'offered-undefined' statistic.

The impact of using **offered-limited-capped-cir** stat-mode while profile-capped mode is disabled are that one of the counting resources in the forwarding plane and traffic manager will not be used and **soft-in-profile** will be treated as 'offered-in' instead of 'offered-undefined'.

The counters are used in the following manner:

1. 'offered-in-that-stayed-green' = profile in, soft-in-profile
2. 'offered-in-that-turned-yellow-or-red' = profile in, soft-in-profile
3. 'offered-out-that-turned-green' = soft-out-of-profile
4. 'offered-out- that-turned-yellow-or-red' = profile out, soft-out-of-profile

5.	'dropped-in-that-stayed-green'	= Same as 1
6.	'dropped-in-that-turned-yellow-or-red'	= Same as 2
7.	'dropped-out-that-turned-green'	= Same as 3
8.	'dropped-out-that-turned-yellow-or-red'	= Same as 4
9.	'forwarded-in-that-stayed-green'	= Derived from 1 - 5
10.	'forwarded-in-that-turned-yellow'	= Derived from 2 - 6
11.	'forwarded-out-that-turned-green'	= Derived from 3 - 7
12.	'forwarded-out-that-turned-yellow'	= Derived from 4 - 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a.	'offered-in'	= 1 + 2
b.	'offered-out'	= 3 + 4
c.	'discard-in'	= 5 + 7
d.	'discard-out'	= 6 + 8
e.	'forward-in'	= 9 + 11
f.	'forward-out'	= 10 + 12

## dscp

**Syntax** **dscp** {*dscp-name* | **in-profile** *dscp-name* **out-profile** *dscp-name*}  
**no dscp**

**Context** config>qos>sap-egress>fc

**Description** This command configures a DiffServ Code Point (DSCP) code point to be used for remarking packets from the specified FC. If the optional in/out-profile is specified, the command will remark different DSCP code points depending on whether the packet was classified to be in or out-of-profile ingress to the node.

**Default** not enabled

**Parameters** *dscp-name* — Specifies a dscp name that has been previously mapped to a value using the **dscp-name** command. The DiffServ code point can only be specified by its name.

**Values** be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

## prec

**Syntax** `prec ip-prec-value [fc fc-name] [hsmda-counter-override counter-id] [profile {in | out}]`  
`no prec ip-prec-value`

**Context** config>qos>sap-egress>fc

**Description** This command defines a value to be used for remarking packets for the specified FC. If the optional in/out-profile is specified, the command will remark different PREC values depending on whether the packet was classified to be in or out-of-profile ingress to the node.

The hsmda-counter-override keyword is optional. When specified and the egress SAP is created on an HSMDA, the egress classification rule will override the default queue accounting function for the packet. By default, the HSMDA uses each queues default queue counters for packets mapped to the queue. The hsmda-counter-override keyword is used to map the packet to an explicit exception counter. One of eight counters may be used. When the packet is mapped to an exception counter, the packet will not increment the queues discard or forwarding counters, instead the exception discard and forwarding counters will be used. The dscp based counter override decision may be overwritten by an ip-criteria reclassification rule match if the higher priority classification rule has an hsmda-counter-override action defined.

not enabled

**Parameters** *ip-prec-value* — The *ip-prec-value* is a required parameter that specifies the unique IP header ToS byte precedence bits value that will match the IP precedence rule. If the command is executed more than once with the same *ip-prec-value*, the previous forwarding class and enqueueing priority is completely overridden by the new parameters or defined to be inherited when a forwarding class or enqueueing priority parameter is missing.

A maximum of eight IP precedence rules are allowed on a single policy.

The precedence is evaluated from the lowest to highest value.

**Values** 0 — 7

**hsmda-counter-override counter-id** — The hsmda-counter-override reclassification action is optional and only has significance on SAPs which are created on an HSMDA. When specified, packets matching the IP precedence value will be mapped to the defined HSMDA exception counter-id for the packets queue group. The default behavior is to use the default counter on the queue group for the queue to which the packet is mapped. The hsmda-counter-override action may be overwritten by an ip-criteria

reclassification rule match. The specified counter-id must be specified as an integer between 1 and 8. To remove the HSMDA exception counter reclassification action for the specified dscp-value, the dscp command must be re-executed without the hsmda-counter-override reclassification action defined.

**Values** 1 — 8

## scope

<b>Syntax</b>	<b>scope</b> { <b>exclusive</b>   <b>template</b> } <b>no scope</b>
<b>Context</b>	config>qos>sap-egress <i>policy-id</i>
<b>Description</b>	Enter the scope of this policy. The scope of the policy cannot be changed if the policy is applied to one or more services.  The no form of this command sets the scope of the policy to the default of template.
<b>Default</b>	template
<b>Parameters</b>	<b>exclusive</b> — When the scope of a policy is defined as exclusive, the policy can only be applied to a single SAP. Attempting to assign the policy to a second SAP will result in an error message. If the policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP.  The system default policies cannot be put into the exclusive scope. An error will be generated if scope exclusive is executed in any policies with a policy-id equal to 1.  <b>template</b> — When the scope of a policy is defined as template, the policy can be applied to multiple SAPs on the router.

## sap-egress

<b>Syntax</b>	[no] <b>sap-egress</b> <i>policy-id</i>   <i>policy-name</i>
<b>Context</b>	config>qos
<b>Description</b>	This command is used to create or edit a Service Egress QoS policy. The egress policy defines the Service Level Agreement (SLA) for service packets as they egress on the SAP.  Policies in effect are templates that can be applied to multiple services as long as the scope of the policy is template. The queues defined in the policy are not instantiated until a policy is applied to a service.  A sap-egress policy differs from sap-ingress policies in the complexity of the QoS parameters that can be defined. At ingress, policies determine queue mappings based on ingress DSCP, Dot1P and IP or MAC match criteria. Multiple queues can be created per forwarding class and each queue can have different CIR or PIR parameters.  At egress, the policies are much simpler, as the forwarding class and in or out of profile determination happened way back at the original service ingress SAP. Egress SAP QoS policies allow the definition of queues and the mapping of forwarding classes to those queues. Each queue needs to have a relative CIR for determining its allocation of QoS resources during periods of congestion. A PIR can also be defined that forces a hard limit on the packets transmitted through the queue. When the forwarding class is mapped to the queue, a Dot1p value can optionally be specified. If specified and the SAP has a Dot1q encapsulation type, the Dot1p value will be used for all packets that egress on that forwarding class. If the Dot1p value is not specified, a Dot1p value of zero will be used. If the SAP is null encapsulated, or on a SONET/SDH interface, the Dot1p value has no meaning.  A <b>default-action</b> parameter is required to specify the default queue used by all forwarding classes not specifically mapped within the queue parameters. A sap-egress policy will be considered incomplete, if it

does not include definition of at least one queue and does not specify the default action. Incomplete sap-egress policies cannot be applied to services.

The sap-egress policy with policy-id 1 is the default sap-egress QoS policy and is applied to service egress SAPs when an explicit policy is not specified or removed. The system sap-egress policy can be modified but not deleted. Using the **no sap-egress** command on **policy-id 1** causes it to revert to its factory default parameters.

The factory default settings for sap-egress policy-id 1 define a single queue with PIR set to the maximum value and a CIR set to 25. The single queue is the default queue and all forwarding classes will map to it. Packets being tagged according to the SAP encapsulation defined will have the Dot1p bits set to zero.

Any changes made to an existing policy, using any of the sub-commands, will be applied immediately to all egress SAPs where this policy is applied. For this reason, when many changes are required on a policy, it is highly recommended that the policy be copied to a work area policy-id. That work-in-progress policy can be modified until complete and then written over the original policy-id. Use the **config qos copy** command to maintain policies in this manner.

The no form of this command to deletes the sap-egress policy. A policy cannot be deleted until it is removed from all service SAPs where it is applied. When a sap-egress policy is removed from a SAP, the SAP will revert to the default sap-egress policy-id 1.

The system default sap-egress policy is a special case. The **no** command restores the factory defaults to policy-id 1.

#### Parameters

*policy-id* — The policy-id uniquely identifies the policy on the router.

**Default** none

**Values** 1 — 65535

*policy-name* — The *policy-name* uniquely identifies the policy.

**Values** Valid names consist of any string up to 64 characters long. Policies must first be created with a policy-id, after which a policy-name can be assigned and used as an alias to reference the policy during configuration changes. Policy names may not begin with a number (0-9) or the underscore “\_” character (e.g. *\_myPolicy*). “Default” or “default” can not be used as policy names. Saved configurations and display output from the “info” and most “show” commands will show the policy-id (not the policy-name) where the policies are referenced.

## de-mark

**Syntax** **[no] de-mark [force de-value]**

**Context** config>qos>sap-egress>fc

#### Description

This command is used to explicitly define the marking of the DE bit for **fc** *fc-name* according to the in and out of profile status of the packet (fc-name may be used to identify the dot1p-value).

If no de-value is present, the default values are used for the marking of the DE bit: for example, 0 for in-profile packets, 1 for out-of-profile ones— see IEEE 802.1ad-2005 standard.

In the PBB case, for a Backbone SAP (B-SAP) and for packets originated from a local I-VPLS, the

command dictates the marking of the DE bit for both the BVID and ITAG.

If this command is not used, the DE bit should be preserved if an ingress TAG exist or set to zero otherwise.

If the de-value is specifically mentioned in the command line it means this value is to be used for all the packets of this forwarding class regardless of their in/out of profile status.

**Values** 0 or 1

## dot1p

**Syntax** [no] dot1p {dot1p-value | in-profile dot1p-value out-profile dot1p-value}

**Context** config>qos>sap-egress>fc fc-name

**Description** This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for *fc-name*. When the marking is set, all packets of *fc-name* that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined *dot1p-value*. If the egress packets for *fc-name* are not IEEE 802.1Q or IEEE 802.1P encapsulated, the dot1p command has no effect.

The optional **in-profile** | *dot1p-value* **out-profile** *dot1p-value* structure added to the existing dot1p command will add the capability to mark on an egress SAP the in and out of profile status via a certain dot1p combination, similarly with the DE options.

The command with the additional structure may be used on the SAP when the internal in and out of profile status needs to be communicated to an access network/customer device that does not support the DE bit. Once the in-profile keyword is added, then the rest of the newly added structure must be specified.

When these commands are used the DE Bit or the equivalent field is left unchanged by the egress processing if a tag exists. If a new tag is added, the related DE bit is set to 0.

When the previous command (dot1p dot1p-value) is used without the new structure, it means that the dot1p-value is used for the entire forwarding class, same as before. The two versions of the command are exclusive.

Independently the in or out profile status may be indicated via the setting of the DE bit setting if the de-mark command is used.

The **no** form of the command sets the IEEE 802.1P or IEEE 802.1Q priority bits to 0.

**Default** 0

**in-profile** *dot1p-value* — Specifies the 802.1p value to set for in-profile frames in this forwarding class.

**Values** 0 — 7

**out-profile** *dot1p-value* — specifies the 802.1p value to set for out-profile frames in this forwarding class.

**Values** 0 — 7

---

## Service Queue QoS Policy Commands

### adaptation-rule

**Syntax**    **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]  
**no adaptation-rule**

**Context**    config>qos>sap-ingress>queue  
              config>qos>sap-egress>queue

This command defines the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

**Default**    adaptation-rule pir closest cir closest

**Parameters**    *adaptation-rule* — Specifies the adaptation rule to be used while computing the operational CIR or PIR value.

**Values**

**pir** — Defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id rate* command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

**cir** — Defines the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id rate* command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

**max** — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

**min** — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

**closest** — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

## adv-config-policy

<b>Syntax</b>	<b>[no] adv-config-policy</b> <i>policy-name</i>
<b>Context</b>	config>qos>sap-ingress>queue config>qos>sap-egress>queue
<b>Description</b>	<p>This command specifies the advanced QoS policy. The advanced QoS policy contains only queue and policer child control parameters within a child-control node.</p> <p>Once a policy is created, it may be applied to a queue or policer defined within a <b>sap-egress</b> or <b>sap-ingress</b> QoS policy. It may also be applied to a queue or policer defined within an ingress or <b>egress queue-group</b> template. When a policy is currently associated with a QoS policy or template, the policy may be modified but not deleted (even in the event that the QoS policy or template is not in use).</p> <p>The <b>no</b> form of this command removes the specified advanced policy.</p>
<b>Default</b>	None
<b>Parameters</b>	<i>policy-name</i> — The name of the advanced QoS policy.
<b>Values</b>	Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## adaptation-rule

<b>Syntax</b>	<b>adaptation-rule</b> [ <b>pir</b> <i>adaptation-rule</i> ] <b>no adaptation-rule</b>
<b>Context</b>	config>qos>sap-egress>hsmda-queues>queue
<b>Description</b>	<p>This command defines the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.</p> <p>The <b>no</b> form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific <b>adaptation-rule</b> is removed, the default constraints for <b>rate</b> and <b>cir</b> apply.</p>
<b>Default</b>	adaptation-rule pir closest cir closest
<b>Parameters</b>	<i>adaptation-rule</i> — Specifies the adaptation rule to be used while computing the operational CIR or PIR value.
<b>Values</b>	<b>pir</b> — Defines the constraints enforced when adapting the PIR rate defined within the <b>queue</b> <i>queue-id</i> <b>rate</b> command. The <b>pir</b> parameter requires a qualifier that defines the



constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

**max** — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

**min** — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

**closest** — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

## avg-frame-overhead

**Syntax** **avg-frame-overhead** *percent*  
**no avg-frame-overhead**

**Context** config>qos>sap-egress>queue

**Description** This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-Load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet-based offered-load.
- Frame-encapsulation overhead — Using the avg-frame-overhead parameter, the frame-encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10,000 octets and the avg-frame-overhead equals 10%, the frame-encapsulation overhead would be 10,000 x 0.1 or 1,000 octets.

For egress Ethernet queues, the frame-encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame-encapsulation overhead would be 50 x 20 or 1,000 octets.

- Frame-based offered-load — The frame-based offered-load is calculated by adding the offered-load to the frame-encapsulation overhead. If the offered-load is 10,000 octets and the encapsulation overhead was 1,000 octets, the frame-based offered-load would equal 11,000 octets.
- Packet to frame factor — The packet -to-frame factor is calculated by dividing the frame-encapsulation overhead by the queue's offered-load (packet based). If the frame-encapsulation

overhead is 1,000 octets and the offered-load is 10,000 octets then the packet to frame factor would be 1,000 / 10,000 or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.

- **Frame-based CIR** — The frame-based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame-based CIR would be 500 x 1.1 or 550 octets.
- **Frame-based within-cir offered-load** — The frame-based within-cir offered-load is the portion of the frame-based offered-load considered to be within the frame-based CIR. The frame-based within-cir offered-load is the lesser of the frame-based offered-load and the frame-based CIR. If the frame-based offered-load equaled 11000 octets and the frame-based CIR equaled 550 octets, the frame-based within-cir offered-load would be limited to 550 octets. If the frame-based offered-load equaled 450 octets and the frame-based CIR equaled 550 octets, the frame-based within-cir offered-load would equal 450 octets (or the entire frame-based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame-based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.

- **Frame-based PIR** — The frame-based PIR is calculated by multiplying the packet to frame-factor with the queue's-configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame-factor equals 0.1, the frame-based PIR would be 7,500 x 1.1 or 8,250 octets.
- **Frame-based within-pir offered-load** — The frame-based within-pir offered-load is the portion of the frame-based offered-load considered to be within the frame-based PIR. The frame-based within-pir offered-load is the lesser of the frame-based offered-load and the frame-based PIR. If the frame-based offered-load equaled 11,000 octets and the frame-based PIR equaled 8250 octets, the frame-based within-pir offered-load would be limited to 8,250 octets. If the frame-based offered-load equaled 7,000 octets and the frame-based PIR equaled 8,250 octets, the frame-based within-pir offered load would equal 7,000 octets.

**Port Scheduler Operation Using Frame Transformed Rates** — The port scheduler uses the frame based rates to figure the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

**SAP and Subscriber SLA-Profile Average Frame Overhead Override** — The average frame overhead parameter on a sap-egress may be overridden on an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers. An avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

**Default** 0

**Parameters** *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

**Values** 0.00 — 100.00

## burst-limit

**Syntax** **burst-limit** {**default** | *size* [**byte** | **kilobyte**]}  
**no burst-limit**

**Context** config>qos>sap-ingress>queue  
config>qos>sap-egress>queue

**Description** The `queue burst-limit` command is used to define an explicit shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate.

The `burst-limit` command is supported under the `sap-ingress` and `sap-egress` QoS policy queues. The command is also supported under the `ingress` and `egress queue-group-templates` queues.

The **no** form of this command is used to restore the default burst limit to the specified queue. This is equivalent to specifying `burst-limit default` within the QoS policies or queue group templates. When specified within a `queue-override queue` context, any current burst limit override for the queue will be removed and the queue's burst limit will be controlled by its defining policy or template.

**Parameters** **default** — The default parameter is mutually exclusive to specifying an explicit size value. When `burst-limit default` is executed, the queue is returned to the system default value.

*size* — When a numeric value is specified (*size*), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and by default is interpreted as the burst limit in Kilobytes. If the value is intended to be interpreted in bytes, the `byte` qualifier must be added following *size*.

**Values** 1 to 14,000 (14,000 or 14,000,000 depending on bytes or kilobytes)

**Default** No default for *size*, use the default keyword to specify default burst limit

**byte** — The **bytes** qualifier is used to specify that the value given for *size* must be interpreted as the burst limit in bytes. The `byte` qualifier is optional and mutually exclusive with the `kilobytes` qualifier.

**kilobyte** — The **kilobyte** qualifier is used to specify that the value given for *size* must be interpreted as the burst limit in Kilobytes. The `kilobyte` qualifier is optional and mutually exclusive with the `bytes` qualifier. If neither `bytes` nor `kilobytes` is specified, the default qualifier is `kilobytes`.

## cbs

**Syntax** **cbs** *size-in-kbytes*  
**no cbs**

**Context** config>qos>sap-egress>queue  
config>qos>sap-ingress>queue

**Description** This command provides a mechanism to override the default reserved buffers for the queue.

It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potentially large number of service queues and the economy of statistical multiplexing the individual queue's CBS settings into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly dropping packets.

If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.

The **no** form of this command returns the CBS size to the default value.

**Default** default

**Parameters** *size-in-kbytes* — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

**Values** 0 — 131072 or default

## high-prio-only

**Syntax** **high-prio-only** *percent*  
**no high-prio-only**

**Context** config>qos>sap-ingress>queue  
config>qos>sap-egress>queue

**Description** The **high-prio-only** command configures the percentage of buffer space for the queue, used exclusively by high priority packets. The specified value overrides the default value for the context.

The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The **high-prio-only** parameter is used to override the default value derived from the **network-queue** command.

The **no** form of this command restores the default high priority reserved size.

**Parameters** *percent* — The percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10.

**Values** 0 — 100, default

## mbs

<b>Syntax</b>	<b>mbs size [bytes   kilobytes]</b> <b>no mbs</b>
<b>Context</b>	config>qos>sap-ingress>queue
<b>Description</b>	<p>The Maximum Burst Size (MBS) command provides the explicit definition of the maximum amount of buffers allowed for a specific queue. The value is given in kilobytes and overrides the default value for the context.</p> <p>The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueueing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.</p> <p>The sap-ingress context for mbs provides a mechanism for overriding the default maximum size for the queue.</p> <p>The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.</p> <p>If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.</p> <p>The <b>no</b> form of this command returns the MBS size assigned to the queue to the value.</p>
<b>Default</b>	default
<b>Parameters</b>	<i>size</i> [ <b>bytes</b>   <b>kilobytes</b> ] — The size parameter is an integer expression of the maximum number of bytes or kilobytes of buffering allowed for the queue. For a value of 100 kbps enter the value 100 and specify the <b>kilobytes</b> parameter. A value of 0 causes the queue to discard all packets.
<b>Values</b>	0 — 131072 or default

## packet-byte-offset

<b>Syntax</b>	<b>packet-byte-offset {add bytes   subtract bytes}</b> <b>no packet-byte-offset</b>
<b>Context</b>	config>qos>sap-egress>queue
<b>Description</b>	<p>This command is used to modify the size of each packet handled by the queue by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed.</p> <p>The packet-byte-offset command is meant to be an arbitrary mechanism the can be used to either add downstream frame encapsulation or remove portions of packet headers.</p> <p>When a packet-byte-offset value is applied to a queue instance, it adjusts the immediate packet size. This means that the queue rates, i.e., operational PIR and CIR, and queue bucket updates use the adjusted packet</p>

size. In addition, the queue statistics will also reflect the adjusted packet size. Scheduler policy rates, which are data rates, will use the adjusted packet size.

The port scheduler max-rate and the priority level rates and weights, if a Weighted Scheduler Group is used, are always on-the-wire rates and thus use the actual frame size. The same goes for the agg-rate-limit on a SAP, a subscriber, or a Multi-Service Site (MSS) when the queue is port-parented.

When the user enables frame-based-accounting in a scheduler policy or queue-frame-based-accounting with agg-rate-limit in a port scheduler policy, the queue rate will be capped to a user configured on-the-wire rate but the packet-byte-offset value is still in effect as explained above.

The **no** form of this command is used to remove per packet size modifications from the queue.

**Parameters** **add bytes** — The **add** keyword is mutually exclusive to the **subtract** keyword. Either parameter must be specified. When **add** is defined, the corresponding bytes parameter specifies the number of bytes that is added to the size of each packet associated with the queue for scheduling and accounting purposes.

**Values** 0—32

**Default** None

**subtract bytes** — The **subtract** keyword is mutually exclusive to the **add** keyword. Either parameter must be specified. When **subtract** is defined, the corresponding bytes parameter specifies the number of bytes that is subtracted to the size of each packet associated with the queue for scheduling and accounting purposes.

**Values** 0—64

**Default** None

## parent

**Syntax** **parent scheduler-name [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level]**  
**no parent**

**Context** config>qos>sap-ingress>queue  
config>qos>sap-egress>queue

**Description** This command defines an optional parent scheduler that further governs the available bandwidth given the queue aside from the queue's PIR setting. When multiple schedulers and/or queues share a child status with the parent scheduler, the **weight** or **level** parameters define how this queue contends with the other children for the parent's bandwidth.

Checks are not performed to see if a *scheduler-name* exists when the parent command is defined on the queue. Scheduler names are configured in the **config>qos>scheduler-policy>tier level** context. Multiple schedulers can exist with the *scheduler-name* and the association pertains to a scheduler that should exist on the egress SAP as the policy is applied and the queue created. When the queue is created on the egress SAP, the existence of the *scheduler-name* is dependent on a scheduler policy containing the *scheduler-name* being directly or indirectly applied (through a multi-service customer site) to the egress SAP. If the *scheduler-name* does not exist, the queue is placed in the orphaned operational state. The queue will accept packets but will not be bandwidth limited by a virtual scheduler or the scheduler hierarchy applied to the SAP. The orphaned state must generate a log entry and a trap message. The SAP which the queue belongs to must also depict an orphan queue status. The orphaned state of the queue is automatically cleared when the *scheduler-*

*name* becomes available on the egress SAP.

The parent scheduler can be made unavailable due to the removal of a scheduler policy or scheduler. When an existing parent scheduler is removed or inoperative, the queue enters the orphaned state mentioned above and automatically return to normal operation when the parent scheduler is available again.

When a parent scheduler is defined without specifying weight or strict parameters, the default bandwidth access method is weight with a value of 1.

The **no** form of the command removes a child association with a parent scheduler. If a parent association does not currently exist, the command has no effect and returns without an error. Once a parent association has been removed, the former child queue attempts to operate based on its configured rate parameter. Removing the parent association on the queue within the policy takes effect immediately on all queues using the SAP egress QoS policy.

## Parameters

*scheduler-name* — The defined *scheduler-name* conforms to the same input criteria as the schedulers defined within a scheduler policy. Scheduler names are configured in the `config>qos>scheduler-policy>tier level` context. There are no checks performed at the time of definition to ensure that the *scheduler-name* exists within an existing scheduler policy. For the queue to use the defined *scheduler-name*, the scheduler exists on each egress SAP the queue is eventually created on. For the duration where *scheduler-name* does not exist on the egress SAP, the queue operates in an orphaned state.

**Values** Any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

**Default** None. Each parental association must be explicitly defined.

**weight** *weight* — *weight* defines the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent *scheduler-name*. Any queues or schedulers defined as weighted receive no parental bandwidth until all strict queues and schedulers on the parent have reached their maximum bandwidth or are idle. In this manner, weighted children are considered to be the lowest priority.

All **weight** values from all weighted active queues and schedulers with a common parent scheduler are added together. Then, each individual active weight is divided by the total, deriving the percentage of remaining bandwidth provided to the queue or scheduler after the strict children are serviced. A weight is considered to be active when the pertaining queue or scheduler has not reached its maximum rate and still has packets to transmit. All child queues and schedulers with a weight of 0 are considered to have the lowest priority level and are not serviced until all strict and non-zero weighted queues and schedulers are operating at the maximum bandwidth or are idle.

**Values** 0 — 100

**Default** 1

**level** *level* — The optional **level** parameter defines the level of hierarchy when compared to other schedulers and queues when vying for bandwidth on the parent *scheduler-name*. Any queues or schedulers defined as **strict** receive no parental bandwidth until all strict queues and schedulers with a higher (numerically larger) priority on the parent have reached their maximum bandwidth or are idle.

Children of the parent scheduler with a lower strict priority or that are weighted will not receive bandwidth until all children with a higher strict priority have either reached their maximum bandwidth or are idle. Children with the same strict level are serviced in a round robin fashion.

**Values** 1 — 100

**Default** 1

**cir-weight *cir-weight*** — Defines the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

**Values** 0 — 100

**cir-level *cir-level*** — Defines the port priority the queue or scheduler will use to receive bandwidth for its within-cir offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.

**Values** 0 — 8 (8 is the highest priority)

**Default** 0

## percent-rate

**Syntax** **percent-rate** *pir-percent* [**cir** *cir-percent*] [**port-limit**|**local-limit**]  
**percent-rate** *pir-percent* **police** [**port-limit**|**local-limit**]  
**no percent-rate**

**Context** config>qos>sap-egress>queue  
config>qos>sap-ingress>queue

**Description** The percent-rate command within the SAP ingress and egress QOS policy enables supports for a queue's PIR and CIR rate to be configured as a percentage of the egress port's line rate or of its parent scheduler's rate. When the rates are expressed as a port-limit, the actual rates used per instance of the queue will vary based on the port speed. For example, when the same QOS policy is used on a 1-Gigabit and a 10-Gigabit Ethernet port, the queue's rates will be 10 times greater on the 10 Gigabit port due to the difference in port speeds. This enables the same QOS policy to be used on SAPs on different ports without needing to use SAP based queue overrides to modify a queue's rate to get the same relative performance from the queue.

If the port's speed changes after the queue is created, the queue's PIR and CIR rates will be recalculated based on the defined percentage value.

When the rates are expressed as a local-limit, the actual rates used per instance of the queue are relative to the queue's parent scheduler rate. This enables the same QOS policy to be used on SAPs with different parent scheduler rates without needing to use SAP based queue overrides to modify a queue's rate to get the same relative performance from the queue. If the parent scheduler rate changes after the queue is created, the queue's PIR and CIR rates will be recalculated based on the defined percentage value.

The rate and percent-rate commands override one another. If the current rate for a queue is defined using the



percent-rate command and the rate command is executed, the percent-rate values are deleted. In a similar fashion, the percent-rate command causes any rate command values to be deleted. A queue's rate may dynamically be changed back and forth from a percentage to an explicit rate at anytime.

Queue rate overrides can only be specified in the form as configured in the QoS policy (a SAP override can only be specified as a percent-rate if the associated QoS policy was also defined as percent-rate). Likewise, a SAP override can only be specified as a rate (kbps) if the associated QoS policy was also defined as a rate. Queue-overrides are relative to the limit type specified in the QoS policy.

The **no** form of this command returns the queue to its default shaping rate and cir rate. When no percent-rate is defined within a SAP ingress or egress queue-override, the queue reverts to the defined shaping and CIR rates within the SAP ingress and egress QoS policy associated with the queue.

## Parameters

*percent-of-line-rate* — The percent-of-line-rate parameter is used to express the queue's shaping rate as a percentage of line rate. The line rate associated with the queue's port may dynamically change due to configuration or auto-negotiation. The line rate may also be affected by an egress port scheduler defined max-rate.

*pir-percent* — The pir-percent parameter is used to express the queue's PIR as a percentage dependant on the use of the port-limit or local-limit.

**Values** Percentage ranging from 0.01 to 100.00. The default is 100.00.

*pir-percent* — The pir-percent parameter is used to express the queue's PIR as a percentage dependant on the use of the port-limit or local-limit.

*cir-percent* — The **cir** keyword is optional and when defined the required cir-percent CIR parameter expresses the queue's CIR as a percentage dependant on the use of the port-limit or local-limit.

**Values** Percentage ranging from 0.00 to 100.00. The default is 100.00

**port-limit** — The **port-limit** keyword specifies that the configure PIR and CIR percentages are relative to the rate of the port to which this queue connects.

**local-limit** — The **local-limit** keyword specifies that the configure PIR and CIR percentages are relative to the rate of the queue's parent scheduler rate.

## pool

**Syntax** **pool** *pool-name*  
**no pool** *pool-name*

**Context** config>qos>sap-ingress>queue  
config>qos>sap-egress>queue

**Description** This command is utilized once the queue is created within the policy. The pool command can be used to either remove the queue from the pool, or specify a new pool name association for the queue. The pool command does not appear in save or show command output. Instead, the current pool name for the queue will appear (or not appear) on the queue command output using the pool keyword.

The **no pool** command is used to remove a named pool association for the queue. When the pool name is removed, the queue will be placed on the appropriate default pool.

**Default** None

**Parameters** *pool-name* — The specified *pool-name* identifies a named pool where the policy will be applied. Each queue created within the system is tied to a physical port. When the policy is applied and the queue is created, the system will scan the named pools associated with the port to find the specified pool name. If the pool is not found on the port, the system will then look at named pools defined at the ports MDA level. If the pool name is not found on either the port or MDA, the queue will be marked as ‘pool-orphaned’ and will be mapped to the appropriate default pool. If the pool comes into existence, the queue will be moved from the default pool to the new named pool and the ‘pool-orphaned’ state will be cleared. The specified name must be an ASCII name string up to 32 characters long.

## port-parent

**Syntax** **port-parent** [**weight** *weight*] [**level** *level*] [**cir-weight** *cir-weight*] [**cir-level** *cir-level*]  
**no port-parent**

**Context** config>qos>sap-egress>queue

**Description** This command specifies whether this queue feeds off a port-level scheduler. When configured, this SAP egress queue is parented by a port-level scheduler. This object is mutually exclusive with SAP egress queue parent. Only one kind of parent is allowed.

The **port-parent** command defines a child/parent association between an egress queue and a port based scheduler or between an intermediate service scheduler and a port based scheduler. The command may be issued in three distinct contexts; **sap-egress queue** *queue-id*, **network-queue queue** *queue-id* and **scheduler-policy scheduler** *scheduler-name*. The **port-parent** command allows for a set of within-cir and above-CIR parameters that define the port priority levels and weights for the queue or scheduler. If the **port-parent** command is executed without any parameters, the default parameters are assumed.

In this context, the **port-parent** command is mutually exclusive to the **parent** command (used to create a parent/child association between a queue and an intermediate scheduler). Executing a **port-parent** command when a parent definition is in place causes the current intermediate scheduler association to be removed and replaced by the defined port-parent association. Executing a **parent** command when a port-parent definition exists causes the port scheduler association to be removed and replaced by the defined intermediate scheduler name.

Changing the parent context on a SAP egress policy queue may cause a SAP or subscriber context of the queue (policy associated with a SAP or subscriber profile) to enter an orphaned state. If an instance of a queue is created on a port or channel that does not have a port scheduler enabled and the sap-egress policy creating the queue has a port-parent association, the queue will be allowed to run according to its own rate parameters and will not be controlled by a virtual scheduling context. If an instance of a queue is on a port or channel that has a port scheduler configured and the sap-egress policy defines the queue as having a non-existent intermediate scheduler parent, the queue will be treated as an orphan and will be handled according to the current orphan behavior on the port scheduler.

The **no** form of this command removes a port scheduler parent association for the queue or scheduler. If a port scheduler is defined on the port which the queue or scheduler instance exists, the queue or scheduler will become orphaned if an port scheduler is configured on the egress port of the queue or scheduler.

**Default** **no port-parent**

<b>Parameters</b>	<b>weight</b> <i>weight</i> — Defines the weight the queue or scheduler will use at the above-cir port priority level (defined by the level parameter).
	<b>Values</b> 0 — 100
	<b>Default</b> 1
	<b>level</b> <i>level</i> — Defines the port priority the queue or scheduler will use to receive bandwidth for its above-cir offered-load.
	<b>Values</b> 1 — 8 (8 is the highest priority)
	<b>Default</b> 1
	<b>cir-weight</b> <i>cir-weight</i> — Defines the weight the queue or scheduler will use at the within-cir port priority level (defined by the cir-level parameter). The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.
	<b>Values</b> 0 — 100
	<b>cir-level</b> <i>cir-level</i> — Defines the port priority the queue or scheduler will use to receive bandwidth for its within-cir offered-load. If the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.
	<b>Values</b> 0 — 8 (8 is the highest priority)
	<b>Default</b> 0

## rate

<b>Syntax</b>	<b>rate</b> <i>pir-rate</i> [ <b>cir</b> <i>cir-rate</i>   <b>police</b> ] <b>no rate</b>
<b>Context</b>	config>qos>sap-ingress>queue
<b>Description</b>	<p>This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the switch fabric (for SAP ingress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.</p> <p>The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. For SAP ingress, the CIR also defines the rate that packets are considered in-profile by the system. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.</p> <p>The CIR can be used by the queue's parent commands <i>cir-level</i> and <i>cir-weight</i> parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.</p>

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP ingress or SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

**Default** **rate max cir 0** — The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

**Parameters** *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

**Values** 1 — 100000000, **max**

**Default** max

*cir-rate* — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and must be given as a positive integer.

**Values** 0 — 100000000, **max**

**Default** 0

**police** — Specifies that the out of profile traffic feeding into the physical queue instance should be dropped. Using this keyword will override the bandwidth specified by the SAP ingress queue's administrative CIR.

## rate

**Syntax** **rate** *pir-rate* [**cir** *cir-rate*]  
**no rate**

**Context** config>qos>sap-egress>queue

**Description** This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the

parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

**Default** **rate max cir 0** — The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

**Parameters** *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed. Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

**Values** 1 — 100000000, **max**

**Default** max

*cir-rate* — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer.

**Values** 0 — 100000000, **max**

**Default** 0

## rate

**Syntax** **rate** *pir-rate* {**max** | *kilobits-per-second*}  
**no rate**

**Context** config>qos>sap-egress>hsmda-queues>queue

**Description** This command configures a rate limit (PIR) for scheduling packets out of the queue and an optional CIR rate used to determine the profile of packets scheduled from the queue. Configuring a rate limit for a queue on an HSMDA is optional; the default rate is set to maximum (max) causing the shaper to have no effect. The **cir** keyword is used to configure the rate threshold between the in-profile and out-of-profile state of the queue during scheduling from the queue.

Since the CIR leaky bucket is updated during scheduling events and not enqueueing events. The profiling function is not based on packet arrival. Instead, the queue absorbs bursts that exceed the queues forwarding rate. In this case, burst tolerance is more heavily affected by the maximum queue depth (mbs) and the PIR shaping behavior then the CIR leaky bucket behavior.

### Ingress Queue Policing

Ingress HSMDA queues support an explicit policing mode configured using the **rate** command. When the queue is configured to police traffic, the defined rate is used to determine whether the scheduled packet removed from the queue is in-profile or out-of-profile. Packets that are scheduled while the

queue is in-profile are forwarded to the ingress forwarding plane. Packet that are scheduled while the queue is out-of-profile are discarded without updating any PIR leaky buckets on the HSMDA associated with the packet (queue PIR, queue group PIR, secondary shaper PIR or scheduling priority PIR).

The advantage to using the policing mode instead of relying on PIR based queue shaping is that the policing mode does not stop scheduling for the queue when the defined rate is reached (as would happen with the queue PIR). Since scheduling is not stopped, the queue does not experience congestion due to the policing rate and this minimizes jitter associated with forwarding packets from the queue.

For best results, the queue should be at a relatively high scheduling priority for proper operation. Since jitter sensitive traffic must be prioritized over other traffic in the system, this should not be a problem. The scheduling priority is based on the queue ID. Queue ID 8 has the highest relative priority while queue ID 1 the lowest. For a complete overview on HSMDA scheduling, refer to the `hsmda-scheduler-policy` command.

### Ingress Color Aware Profiling

At ingress, it is possible to classify packets handled by the queue as explicitly in-profile or out-of-profile (out-of-profile is by far the most used case). In-profile is commonly referred to as green and out-of-profile as yellow colored packets based on two color marking decisions upstream. The ability to identify certain packets as green or yellow while treating other packets as undefined is called color aware profiling. Typically, only yellow (out-of-profile) packets will be treated as color aware, while green and other markings will be treated as undefined. During scheduling from the queue, the undefined packets will be processed by the CIR leaky bucket while the pre-colored packets will not. In this way, the CIR will mark the undefined packets as in-profile or out-of-profile while preventing the out-of-profile yellow packets from consuming in-profile bandwidth for the queue.

Packets may be classified as in-profile or out-of-profile in two ways. The first is to create sub-forwarding classes (such as `af.outi` and `af.ini`), defined them as explicitly in-profile or out-of-profile using the `profile` command and then map the green packets to the in-profile sub-class and the yellow packet to the out-of-profile sub-class.

The second way to enable color aware policing is to configure recognition of the DEI bit within the Dot1Q header. When supported by the network, DEI bit will be set to 0 for undefined packets and set to 1 for out-of-profile yellow packets that are discard eligible. When DEI recognition is enabled, the DEI bit automatically defines the packet as undefined or out-of-profile without the need to configure the sub-class behavior.

### Egress Profiling Based Dot1P Remarking

HSMDA egress queues are capable of remarking Dot1P and DEI bits based on the current state of the queues CIR. Egress Dot1P remarking is enabled at the forwarding class level. Using egress profiling based Dot1P remarking, either two distinct Dot1P values may be used to distinguish in-profile and out-of-profile, or just the DEI bit may be toggled.

### SAP and Subscriber Queue Rate Overrides

The shaping rate and CIR values may be overridden on each SAP or subscriber to which the QoS policy is associated.

The `no` form of the command returns all queues created with the `queue-id` by association with the QoS policy to the default PIR and CIR parameters (`max, 0`).

**Default** *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

**Values** 1 — 100000000, **max**

**Default** max

---

## Show Commands

### sap-ingress

**Syntax** `sap-ingress [policy-id] [detail]`

**Context** show>qos

**Description** This command displays SAP ingress QoS policy information.

**Parameters** *policy-id* — Displays information about the specific policy ID.

**Default** all SAP ingress policies

**Values** 1 — 65535

**detail** — Displays detailed policy information including policy associations.

#### Sample Output

**Show SAP Ingress Output** — The following table describes SAP ingress show command output.

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Scope	<b>Exclusive</b> — Implies that this policy can only be applied to a single SAP. <b>Template</b> — Implies that this policy can be applied to multiple SAPs on the router.
Description	A text string that helps identify the policy's context in the configuration file.
Default FC	Specifies the default forwarding class for the policy.
Priority	Specifies the enqueueing priority when a packet is marked with a <i>dot1p-value</i> specified.
Criteria-type	<b>IP</b> — Specifies that an IP criteria-based SAP ingress policy is used to select the appropriate ingress queue and corresponding forwarding class for matched traffic. <b>MAC</b> — Specifies that a MAC criteria-based SAP is used to select the appropriate ingress queue and corresponding forwarding class for matched traffic.
Mode	Specifies the configured mode of the meter (trTcm or srTcm).



Label	Description (Continued)
CIR Admin	Specifies the administrative Committed Information Rate (CIR) parameters for the queue. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth.
CIR Oper	The operational value derived by computing the CIR value from the administrative CIR and PIR values and their corresponding adaptation rules.
CIR Rule	<p><code>min</code> – The operational CIR for the queue will be equal to or greater than the administrative rate specified using the rate command.</p> <p><code>max</code> – The operational CIR for the queue will be equal to or less than the administrative rate specified using the rate command.</p> <p><code>closest</code> – The operational PIR for the queue will be the rate closest to the rate specified using the rate command without exceeding the operational PIR.</p>
PIR Admin	Specifies the administrative Peak Information Rate (PIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the switch fabric (for SAP ingress queues) or out an egress interface (for SAP egress queues).
PIR Oper	The administrative PIR specified by the user.
PIR Rule	<p><code>min</code> – The operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command.</p> <p><code>max</code> – The operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command.</p> <p><code>closest</code> – The operational PIR for the queue will be the rate closest to the rate specified using the rate command.</p>
CBS	<p><code>def</code> – Specifies the default CBS value for the queue.</p> <p><code>value</code> – Specifies the value to override the default reserved buffers for the queue.</p>
MBS	<code>def</code> – Specifies the default MBS value.

Label	Description (Continued)
	<code>value</code> – Specifies the value to override the default MBS for the queue.
HiPrio	Specifies the percentage of buffer space for the queue, used exclusively by high priority packets.
PIR Lvl/Wt	Specifies the priority level of the scheduler when compared to other child schedulers and queue vying for bandwidth on the parent schedulers during the ‘above CIR’ distribution phase of bandwidth allocation. Weight defines the relative weight of this scheduler in comparison to other child schedulers and queue at the same level.
CIR Lvl/Wt	Specifies the level of hierarchy when compared to other schedulers and queue when vying for bandwidth on the parent scheduler. Weight defines the relative weight of this queue in comparison to other child schedulers and queue while vying for bandwidth on the parent scheduler.
Parent	Specifies the parent scheduler that governs the available bandwidth given the queue aside from the queue’s PIR setting.
Dot1p	Specifies the forwarding class or enqueueing priority when a packet is marked with a <i>dot1p-value</i> specified.
FC	Specifies the forwarding class overrides.
Priority	The optional priority setting overrides the default enqueueing priority for the packets received on an ingress SAP which uses the policy that matches this rule.  <code>High</code> – Specifies that the <code>high</code> enqueueing parameter for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested.  <code>Low</code> – Specifies that the <code>low</code> enqueueing parameter for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested.
DSCP	Specifies the forwarding class or enqueueing priority when a packet is marked with the DiffServ Code Point (DSCP) value.
FC	Specifies one of the predefined forwarding classes in the system. When a packet matches the rule the forwarding class is only overridden when the <code>fc fc-name</code> parameter is defined on the rule.

Label	Description (Continued)
Priority	<p>This parameter specifies the default enqueueing priority overrides for all packets received on an ingress SAP using this policy that match this rule.</p> <p>High – Specifies that the high enqueueing parameter for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested.</p> <p>Low – Specifies that the low enqueueing parameter for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested.</p>
Prec	Specifies the forwarding class or enqueueing priority when a packet is marked with an IP precedence value ( <i>ip-prec-value</i> ).
UCastQ	Specifies the default unicast forwarding type queue mapping.
MCastQ	Specifies the overrides for the default multicast forwarding type queue mapping.
BCastQ	Specifies the default broadcast forwarding type queue mapping.
UnknownQ	Specifies the default unknown unicast forwarding type queue mapping.
Match Criteria	Specifies an IP or MAC criteria entry for the policy.
Entry	
Source IP	Specifies a source IP address range used for an ingress SAP QoS policy match.
Source Port	Specifies a source TCP or UDP port number or port range used for an ingress SAP QoS policy match.
Protocol	Specifies the IP protocol number to be used for an ingress SAP QoS policy match.
DSCP	Specifies a DiffServ Code Point (DSCP) name used for an ingress SAP QoS policy match.
Fragment	<p>True – Configures a match on all fragmented IP packets.</p> <p>False – Configures a match on all non-fragmented IP packets.</p>
FC	Specifies the entry's forwarding class.
Priority	Specifies the default enqueueing priority overrides for all packets received on an ingress SAP using this policy.

Label	Description (Continued)		
Src MAC	Specifies a source MAC address or range to be used as a Service Ingress QoS policy match.		
Dst MAC	Specifies a destination MAC address or range to be used as a Service Ingress QoS policy match.		
Dot1p	Specifies a IEEE 802.1p value to be used as the match.		
Snap-pid	Specifies an IEEE 802.3 LLC SNAP Ethernet Frame PID value to be used as a Service Ingress QoS policy match.		
Ethernet-type	Specifies an Ethernet type II Ethertype value to be used as a Service Ingress QoS policy match.		
ESnap-oui-zero	Specifies an IEEE 802.3 LLC SNAP Ethernet Frame OUI zero or non-zero value to be used as a Service Ingress QoS policy match.		
DSAP	Specifies an Ethernet 802.2 LLC DSAP value or range for an ingress SAP QoS policy match.		
SSAP	Specifies an Ethernet 802.2 LLC DSAP value or range for an ingress SAP QoS policy match.		
FC	Specifies the entry's forwarding class.		
Priority	Specifies the default enqueueing priority overrides for all packets received on an ingress SAP using this policy.		
<b>Service Association</b>			
Service-Id	The unique service ID number which identifies the service in the service domain.		
Customer-Id	Specifies the customer ID which identifies the customer to the service.		
SAP	Specifies the a Service Access Point (SAP) within the service where the SAP ingress policy is applied.		
QoS Sap Ingress			
=====			
Sap Ingress Policy (100)			
-----			
Policy-id	: 100	Scope	: Template
Default FC	: be	Priority	: Low
Criteria-type	: IP		
Description	: Used on VPN sap		
-----			
Queue Mode	CIR Admin PIR Admin CBS	HiPrio	PIR Lvl/Wt Parent
	CIR Rule PIR Rule MBS		CIR Lvl/Wt

1	Prio	0	max	def	def	1/1	None
		closest	closest	def		0/1	
2	Prio	0	max	def	def	1/1	None
		closest	closest	def		0/1	
10	Prio	0	11000	def	def	1/1	VPN_be
		closest	closest	def		0/1	
11	Prio	0	max	def	def	1/1	None
		closest	closest	def		0/1	
12	Prio	0	11000	def	def	1/1	VPN_prio*
		closest	closest	def		0/1	
13	Prio	0	1	def	def	1/1	VPN_rese*
		closest	closest	def		0/1	
15	Prio	1500	1500	def	def	1/1	VPN_video
		closest	closest	def		0/1	
16	Prio	2500	2500	def	def	1/1	VPN_voice
		closest	closest	def		0/1	
17	Prio	36	100	def	def	1/1	VPN_nc
		closest	closest	def		0/1	
20	Prio	0	11000	def	def	1/1	VPN_be
		closest	closest	def		0/1	
22	Prio	0	11000	def	def	1/1	VPN_prio*
		closest	closest	def		0/1	
23	Prio	0	1	def	def	1/1	VPN_rese*
		closest	closest	def		0/1	
25	Prio	1500	1500	def	def	1/1	VPN_video
		closest	closest	def		0/1	
26	Prio	2500	2500	def	def	1/1	VPN_voice
		closest	closest	def		0/1	
27	Prio	36	100	def	def	1/1	VPN_nc
		closest	closest	def		0/1	

FC	UCastQ	MCastQ	BCastQ	UnknownQ
be	10	20	20	20
af	12	22	22	22
h2	16	26	26	26
ef	13	23	23	23
h1	15	25	25	25
nc	17	27	27	27

SubFC	Profile	In-Remark	Out-Remark
af	None	None	None
be	None	None	None
ef	None	None	None
h1	None	None	None
h2	None	None	None
nc	None	None	None

Dot1p	FC	Priority
0	af	High
1	ef	High
7	be	Low

DSCP	FC	Priority
af41	af	High

```

-----
Prec Value      FC                               Priority
-----
0               be                               Default
2               af                               Default
3               ef                               Default
5               h1                               Default
6               h2                               Default
7               nc                               Default
-----

```

Match Criteria

IP Match Criteria

```

-----
Entry          : 10
Description    : Entry 10-FC-AF
Source IP     : 10.10.10.104/24      Source Port   : None
Dest. IP     : Undefined           Dest. Port    : None
Protocol     : 6                   DSCP         : None
Fragment     : Off
FC           : af                   Priority      : High

Entry          : 20
Description    : Entry 20-FC-BE
Source IP     : Undefined           Source Port   : None
Dest. IP     : Undefined           Dest. Port    : eq 255
Protocol     : 17                  DSCP         : None
Fragment     : Off
FC           : Default              Priority      : Default
-----

```

IPv6 Match Criteria

No Match Criteria Entries found.

Associations

```

-----
Service-Id    : 700 (VPLS)           Customer-Id   : 7
- SAP : 1/1/9:0                   override
-----

```

\*A:ALA-48>config>qos#

config>qos# show qos sap-ingress 2 detail

QoS Sap Ingress

Sap Ingress Policy (2)

```

-----
Policy-id     : 2                   Scope         : Template
Default FC   : be                   Priority      : Low
Criteria-type : None
-----

```

Queue	Mode	CIR	Admin Rule	PIR	Admin Rule	CBS	MBS	HiPrio	PIR Lvl/Wt	CIR Lvl/Wt	Parent
1	Prio	0	closest	max	closest	def		def	1/1	0/1	None
11	Prio	0	closest	max	closest	def		def	1/1	0/1	None

```

-----
FC                UCastQ          MCastQ          BCastQ          UnknownQ
-----
af                def            def            def            def
ef                def            def            def            def
-----
SubFC            DE-1-out-profile  Profile          In-Remark        Out-Remark
-----
af                No              None            None            None
ef                Yes             None            None            None
-----
Dot1p            FC              Priority
-----
No Dot1p-Map Entries Found.
-----
DSCP            FC              Priority
-----
No DSCP-Map Entries Found.
-----
Prec Value      FC              Priority
-----
No Prec-Map Entries Found.
-----
Match Criteria
-----
No Matching Criteria.
-----
Associations
-----
No Associations Found.
config>qos#

```

## sap-egress

**Syntax** `sap-egress [policy-id] [detail]`

**Context** `show>qos`

**Description** This command displays SAP egress QoS policy information.

**Parameters** *policy-id* — Displays information about the specific policy ID.

**Values** 1 — 65535

**detail** — Displays detailed policy information including policy associations.

**SAP Egress Output** — The following table describes SAP egress show command output.

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Scope	<code>Exclusive</code> — Implies that this policy can only be applied to a single SAP.

Label	Description (Continued)
	<code>Template</code> – Implies that this policy can be applied to multiple SAPs on the router.
Description	A text string that helps identify the policy's context in the configuration file.
Queue:	
CIR Admin	Specifies the administrative Committed Information Rate (CIR) parameters for the queue. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth.
CIR Oper	The operational value derived by computing the CIR value from the administrative CIR and PIR values and their corresponding adaptation rules.
CIR Rule	<p><code>min</code> – The operational CIR for the queue will be equal to or greater than the administrative rate specified using the rate command except where the derived operational CIR is greater than the operational PIR. If the derived operational CIR is greater than the derived operational PIR, the operational CIR will be made equal to the operational PIR.</p> <p><code>max</code> – The operational CIR for the queue will be equal to or less than the administrative rate specified using the rate command.</p> <p><code>closest</code> – The operational PIR for the queue will be the rate closest to the rate specified using the rate command without exceeding the operational PIR.</p>
PIR Admin	Specifies the administrative Peak Information Rate (PIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the switch fabric (for SAP ingress queues) or out an egress interface (for SAP egress queues).
PIR Oper	The administrative PIR specified by the user.
PIR Rule	<p><code>min</code> – The operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command.</p> <p><code>max</code> – The operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command.</p>



Label	Description (Continued)
	<code>closest</code> – The operational PIR for the queue will be the rate closest to the rate specified using the rate command.
CBS	<code>def</code> – Specifies that the CBS value reserved for the queue. <code>value</code> – Specifies the value to override the default reserved buffers for the queue.
MBS	<code>def</code> – Specifies that the MBS value is set by the <code>def-mbs</code> function. <code>value</code> – Specifies the value to override the default maximum size for the queue.
HiPrio	Specifies the percentage of buffer space for the queue, used exclusively by high priority packets.
PIR Lvl/Wt	Specifies the priority level of the scheduler when compared to other child schedulers and queues vying for bandwidth on the parent schedulers during the ‘above CIR’ distribution phase of bandwidth allocation. Weight defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same level.
CIR Lvl/Wt	Specifies the level of hierarchy when compared to other schedulers and queues when vying for bandwidth on the parent scheduler. Weight defines the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent scheduler.
Parent	Specifies the parent scheduler that governs the available bandwidth given the queue aside from the queue’s PIR setting.
FC Name	Specifies the forwarding class queue mapping or dot1p marking is to be edited.
Queue-id	Specifies the <i>queue-id</i> that uniquely identifies the queue within the policy.
Explicit/Default	<code>Explicit</code> – Specifies the egress IEEE 802.1P (dot1p) bits marking for <i>fc-name</i> . <code>Default</code> – Specifies that the default dot1p value (0) is used.
Service Association	
Service-Id	The unique service ID number which identifies the service in the service domain.
Customer-Id	Specifies the customer ID which identifies the customer to the service.

Label	Description (Continued)
SAP	Specifies the a Service Access Point (SAP) within the service where the policy is applied.
Mirror SAPs:	
Mirror Dest	Specifies the mirror service ID which identifies the service in the service domain.
SAP	Specifies the a Service Access Point (SAP) within the service where the SAP egress policy is applied.

### Sample Output

```
A:ALA-49# show qos sap-egress
=====
Sap Egress Policies
=====
Policy-Id      Scope      Description
-----
1              Template  Default SAP egress QoS policy.
1010           Template
1020           Template
=====
A:ALA-49#
```

```
A:ALA-49# show qos sap-egress 1010
=====
QoS Sap Egress
=====
Sap Scheduler Policy (1010)
-----
Policy-id      : 1010                Scope      : Template
=====
A:ALA-49#
```

```
A:ALA-49# show qos sap-egress 1010 detail
=====
QoS Sap Egress
=====
Sap Scheduler Policy (1010)
-----
Policy-id      : 1010                Scope      : Template
-----
Queue          CIR Admin PIR Admin CBS      HiPrio  PIR Lvl/Wt  Parent
              CIR Rule  PIR Rule  MBS
-----
1              0         max      def      def      1/1      None
              closest  closest  def
8              0         max      def      def      1/1      None
              closest  closest  def
              0/1
```

```

-----
FC Name           Queue-id      Explicit/Default
-----
be                8             Explicit (7)
-----
Associations
-----
Service-Id       : 1 (VPRN)                Customer-Id : 1
- SAP : 1/1/10:1

SLA Profiles :
- test                override
-----
Mirror SAPs
-----
No Mirror SAPs Found.
=====
A:ALA-49#

config>qos# show qos sap-egress 2 detail
=====
QoS Sap Egress
-----
Sap Scheduler Policy (2)
-----
Policy-id        : 2                Scope      : Template
-----
Queue  CIR Admin PIR Admin CBS      HiPrio PIR Lvl/Wt   Parent      AvgOvrhd
      CIR Rule  PIR Rule  MBS      CIR Lvl/Wt
-----
1      0          max      def      def      1/1       None        0.00
      closest  closest  def      0/1
-----
FC Name           Queue-id      Explicit/Default      DE-Mark
-----
af                def           Explicit (4)          Profile
l1                def           Explicit (In:5 Out:6) Force 0
ef                def           Default                None
-----
Associations
-----
No Associations Found.
-----
Mirror SAPs
-----
No Mirror SAPs Found.
=====
config>qos#

```

## queue

**Syntax** `queue from {sap sap-id | queue-group port-id queue-group-name | subscriber subscriber-id | network {mda-id | port-id} | system {card slot-number | mda mda-id port port-id}} {ingress | egress} [id queue-id]`

**Context** show>qos

**Description** The show qos queue command outputs the Burst Control Group (BCG) name and slowest accurate visitation time for the specified queues.

For each queue specified, the system may find multiple hardware queues. This may be true for ingress queues on multipoint services (VPLS, IES, VPRN) or for queues created on an Ethernet Link Aggregation Group (LAG). When this is true, the show command may display the calculated slowest accurate visitation time for the logical queue (all hardware queues will have the same calculated value) but must display the BCG name for each individual hardware queue.

The BCG name associated with a queue may be specified in the show bcg command to display the historical and current visitation time for the BCG managing the burst tolerance of the queue. If the output visitation time is greater (longer time) than the queue returned slowest accurate visitation time, the queue's shaping rate may be negatively impacted.

**Parameters** **from** — The from keyword specifies that the following parameters are match criteria for finding a single or set of ingress or egress queues within the system. The system will accept sap, queue-group, subscriber, network-queues or shared-queues as the match criteria.

**sap** *sap-id* — The sap keyword is used to specify that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues within the specified sap-id. The sap keyword is mutually exclusive with the other from match criteria. If the specified sap-id is not found, the system should return 'The specified SAP ID does not exist'.

**queue-group** *port-id queue-group-name* — The queue-group keyword is used to specify that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues within the specified queue-group-name on the specified port-id. The following ingress or egress keyword further specifies that the targeted queue group is an ingress port or egress port queue group. The queue-group keyword is mutually exclusive with the other from match criteria. If the specified port-id is not provisioned on the system or the specified queue-group-name is not found on the ports specified direction, the system should return 'The specified queue group does not exist'.

**subscriber** *subscriber-id* — The subscriber keyword is used to specify that the system should find and display the BCG and calculate the slowest accurate visitation time for the queues associated with the specified subscriber-id. The queue-group keyword is mutually exclusive with the other from match criteria. If the specified subscriber-id does not exist, the system should return 'The specified subscriber does not exist'.

**network** {*mda-id* | *port-id*} — The network keyword is used to specify that the system should find and display the queue information for the queues associated with the specified mda-id or port-id. If the ingress direction qualifier is specified, an mda-id is required. If the egress direction qualifier is specified, a port-id is required. The network keyword is mutually exclusive with the other from match criteria. If the specified mda-id does not exist, the system should return 'The specified MDA is not provisioned'. If the specified port-id does not exist, the system should return 'The specified port is not provisioned'.

**system** {*card slot-number* |*mda-id* | *port-id*} — The system keyword is used to specify that the system should find and display the queue information for all the system queues associated with the specified card slot-id, mda mda-id or port port-id. If the ingress direction qualifier is specified, the ingress system queues are displayed. If the egress direction qualifier is specified, only the egress system queues are displayed. The system keyword is mutually exclusive with the other from match criteria. If the specified slot-id does not exist, the system should return ‘The specified slot number is not provisioned’. If the specified mda-id does not exist, the system should return ‘The specified MDA is not provisioned’. If the specified port-id does not exist, the system should return ‘The specified port is not provisioned’. The id parameter is not supported when matching system queues.

{**ingress** | **egress**} — The ingress and egress direction qualifiers are mutually exclusive. Either ingress or egress must be specified.

**id queue-id** — The id keyword is used to limit the return queues to a single queue-id. The keyword is not accepted when the system match criteria is used.

## bcg

**Syntax** **bcg** *burst-control-group-name* [**member-queues** [**at-risk-only**]] [**exp-util-bw** *megabits-per-second*]

**Context** show>qos

**Description** The show qos bcg command outputs the current and historical visitation time associated with the specified BCG name.

A Burst Control Group (BCG) represents a list of queues that share the same non-scheduling PIR and CIR bucket target update interval. When a queue’s scheduled rate bursts above its PIR bucket depth, the queue is removed from its scheduling context. The system uses a BCG in order to visit the queues PIR bucket to periodically drain an appropriate amount from the bucket. When the bucket has been drained below the PIR bucket threshold, the queue is allowed back onto its scheduling context. The amount decremented from the bucket is a function of the amount of time that has elapsed since the last bucket update and the queue’s shaping rate (PIR). If the queue’s shaping rate is configured as 1Mbps and 1ms has elapsed since the last bucket update, the system will decrement the PIR bucket by 1,000 bytes. One caveat is that the bucket cannot be decremented past a depth of 0. This fact drives how the system chooses which BCG is used to manage the queue bucket update interval.

If a queue’s shaping rate is 1Mbps and the threshold (burst limit) is set to 10Kbytes, the maximum amount of time that can expire before the queue is updated without resulting in a negative bucket depth is 81.92ms. This can be calculated by taking the number of bits represented by the bucket depth (10Kbytes = 10 \* 1,024 \* 8 = 81,920 bits) and dividing it by the rate (81,920 bits / 1,000,000 bits per second = 81.92ms). The queue will not be removed from the scheduler until the PIR bucket depth has equaled or exceeded the configured burst threshold, so the bucket will be at least 10Kbytes deep. If the system visits the queue PIR bucket within 81.92ms, the resulting decrement operation will leave the bucket. If the system takes longer than 81.92ms, the decrement result will be greater than 10Kbytes and part of the decrement result will be lost. The net result is from less than timely updates is that the queue will not be returned to the scheduler context fast enough and some shaping bandwidth for the queue will be lost (underrun the shaping rate).

Each Q2 based forwarding plane maintains 7 Burst Control Groups, each targeting a certain queue bucket visitation time. A 40ms, 20ms, 10ms, 5ms, 1ms, 500us and 100us BCG is supported. By default, queues are placed on a BCG based on shaping rate and the queue's burst limit (PIR threshold depth) is set based on the BCG visitation time and the queue's specified shaping rate. When all shaping queues on a Q2 are left in a default burst tolerance management state, the system has sufficient BCG visitation resources to ensure that all queues do not experience inaccurate bucket decrement conditions.

When explicit burst-limit threshold values are defined for a shaping queue, the system picks an appropriate BCG based on the queue's configured shaping rate and the explicit threshold to find a BCG with the best target visitation time that results in worst case decrement values that are less than the configured threshold. However, when a queue is placed on a 'faster' BCG, more visitation resources are consumed and it is possible that the system will not meet a queue's decrement constraints.

The **show qos bcg** command allows visibility into a BCG's historic and current visitation time. The system samples the amount of time it takes each list to visit each of its associated queues once each second and stores the last 10 samples. It also keeps the longest visitation time seen since the last time the BCG statistics were cleared, the longest visitation time for the current queue-to-BCG lists associations, calculated longest visitation time based on maximum scheduling bandwidth and lastly the longest visitation time for an optionally defined scheduling rate.

With each sample, the system indirectly calculates the amount of scheduling bandwidth based on how much Q2 resources were diverted from BNG visitation processing. This calculated scheduling bandwidth is useful since it can be used to evaluate the worst case longest visitation times for each BCG. The calculated scheduling bandwidth value is stored with the longest seen visitation time and the longest seen visitation time with the current queue-to-BCG mappings.

#### Parameters

*burst-control-group-name* — The burst-control-group-name is required and specifies which globally unique Burst Control Group will be displayed. If the specified Burst Control Group does not exist, the show command will fail and the system will return 'The specified BCG does not exist'.

**member-queues [at-risk-only]** — The member-queues optional keyword is used to include a list of all queues attached to the specified burst-control-group-name. The optional at-risk-only keyword may be added to limit the displayed queues to only include queues that are considered 'at-risk' for inaccurate shaping based on either the 100% worst case scheduling bandwidth for the current queue mappings. The 100% scheduling bandwidth used in the 'at-risk' determination may be overridden with a specified scheduling bandwidth by using the exp-util-bw parameter.

**exp-util-bw** — An optional keyword used to display a calculated worst case visitation rate for the specified burst-control-group-name based on the specified value for megabits-per-second.

*megabits-per-second* — A value also modifies the member-queues 'at-risk' state output.

## hsmda-pool-policy

**Syntax** **hsmda-pool-policy** [*hsmda-pool-policy-name*] [**associations**] [**detail**]

**Context** show>qos

**Description** This command displays HSMDB pool policy information.

**Parameters** *hsmda-pool-policy-name* — Displays information about the specified HSMDB pool policy up to 32 characters in length.

**associations** — Displays the entities associated with the specified HSMDA pool policy.

**detail** — Displays detailed output for the specified HSMDA pool policy.

### Sample Output

```
*A:ALA-49>config>qos# show qos hsmda-pool-policy
=====
Qos HSMDA Pool Policy
=====
Policy Name          Description
-----
test                 test
default              Default hsmda Pool policy.
=====

*A:ALA-98>config>qos#

*A:ALA-49>config# show qos hsmda-pool-policy test detail
=====
Qos HSMDA Pool Policy
=====
Policy Name   : test
=====
Description   : test
Sys. Reserve  : 10.00
=====
Class Tier
=====
Class Pool    Root Parent    Alloc. Percent
-----
1             1              50.00
2             1              35.00
3             1              30.00
4             1              25.00
5             1              20.00
6             2              50.00
7             2              40.00
8             2              30.00
=====
Root Tier
=====
Root Pool     Root Weight
-----
1             75
2             25
3             0
4             0
5             0
6             0
7             0
8             0
=====
Associations
-----
- MDA Egress: 9/2
=====

*A:ALA-49>config#
```

```

*A:ALA-49>config# show qos hsmda-pool-policy association
=====
Qos HSMDA Pool Policy
=====
Policy Name   : test
=====
Description   : test
-----

Associations
-----
- MDA Egress: 9/2
=====
Policy Name   : default
=====
Description   : Default hsmda Pool policy.
-----

Associations
-----
- MDA Ingress: 9/2
=====
*A:ALA-49>config#

```

## hsmda-pools

**Syntax** **hsmda-pools mda** *mda-id* **{ingress | egress}** **[detail]**

**Context** show>qos

**Description** This command displays information about HSMDA pools.

**Parameters**

- mda-id* — Specifies the chassis slot and MDA slot numbers.
- ingress** — Displays information about ingress MDA HSMDA pools.
- egress** — Displays information about egress MDA HSMDA pools.
- detail** — Displays detailed HSMDA output for the specified MDA.

## hsmda-scheduler-hierarchy

**Syntax** **hsmda-scheduler-hierarchy port** *port-id* **[{shapers | shaper** *shaper-name***}]**  
**hsmda-scheduler-hierarchy mda** *mda-id*  
**hsmda-scheduler-hierarchy sap** *sap-id* **[ingress | egress]**  
**hsmda-scheduler-hierarchy subscriber** *sub-id* **[ingress | egress]**

**Context** show>qos

**Description** This command displays information about HSMDA scheduler hierarchy.



**Parameters** **port** *port-id* — Displays information about the specified port.

**Values** slot[/mda[/port]] or slot/mda/port[.channel]  
aps-id aps-group-id[.channel]  
aps keyword  
group-id 1 — 64  
ccag-id slot/mda/path-id[cc-type]  
path-id a, b  
cc-type .sap-net, .net-sap

**shapers** — Displays all shaper information.

**shaper** *shaper-name* — Displays information for the specified shaper-name.

**sap** *sap-id* — Displays information about the specified SAP ID.

**Values** null *port-id* | *lag-id*  
dot1q *port-id* | *lag-id*:\* | qtag1  
qinq *port-id* | *lag-id*:qtag1.qtag2  
port-id *slot/mda/port*[.channel]  
lag-id *lag-id*  
lag keyword  
*id* 1 — 200  
qtag1 0 — 4094  
qtag2 \*, 0 — 4094

**ingress** | **egress** — Displays information about the ingress or egress SAP ID or the ingress or egress subscriber

**subscriber** *sub-id* — Displays information about the ingress or egress subscriber ID or the ingress or egress subscriber ID.

## hsmda-scheduler-policy

**Syntax** **hsmda-scheduler-policy** [*hsmda-scheduler-policy-name*] [**associations**] [**detail**]

**Context** show>qos

**Description** This command displays HSM DA scheduler policy information.

**Parameters** *hsmda-scheduler-policy-name* — Displays information about the specified HSM DA scheduler policy.

**associations** — Displays the entities associated with the specified HSM DA scheduler policy.

## hsmda-slope-policy

**Syntax** **hsmda-slope-policy** [*hsmda-slope-policy-name*] [**associations**] [**detail**]

**Context** show>qos

**Description** This command displays HSMDA slope policy information.

**Parameters** *hsmda-scheduler-policy-name* — Displays information about the specified HSMDA slope policy.

**associations** — Displays the entities associated with the specified HSMDA slope policy.