Service Egress and Ingress QoS Policies

In This Section

This section provides information to configure SAP ingress and egress QoS policies using the command line interface.

Topics in this section include:

Overview

There is one default service ingress policy and one default service egress policy. Each policy can have up to 32 ingress queues and 8 egress queues per service.

Each policy can have up to 32 ingress queues and 8 egress queues per service. The default policies can be copied and modified but they cannot be deleted. The default policies are identified as policy ID 1.

The default policies are applied to the appropriate interface, by default. For example, the default SAP ingress policy is applied to access ingress SAPs.The default SAP egress policy is applied to access egress SAPs. You must explicitly associate other QoS policies.

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your router, refer to the CLI Usage chapter in the Basic System Configuration Guide.

Egress SAP Forwarding Class and Forwarding Profile Overrides

An access egress packet's forwarding class can be changed to redirect the packet to an alternate queue than the ingress forwarding class determination would have used. An access egress packet's profile (in or out) can also be changed to modifying the congestion behavior within the egress queue. In both cases, egress marking decisions will be based on the new forwarding class and profile as opposed to the egress forwarding class or profile. The exception is when ingress remarking is configured. An ingress remark decision will not be affected by egress forwarding class or egress profile overrides.

SAP Egress QoS Policy Modifications

The SAP egress QoS policy allows reclassification rules that are used to override the ingress forwarding class and profile of packets that egress a SAP where the QoS policy is applied. Only IP-based reclassification rules are supported.

IP precedence, DSCP and IP quintuple entries can be defined, each with an explicit forwarding class or profile override parameters. The reclassification logic for each entry follows the same basic hierarchical behavior as the classification rules within the SAP ingress QoS policy. IP precedence and DSCP have the lowest match priority while the IP criteria (quintuple) entries have the highest. When an optional parameter (such as **profile**) for IP precedence or DSCP entries is not specified, the value from the lower priority IP quintuple match for that parameter is preserved. If the IP precedence values overlap with DSCP values in that they will match the same IP header TOS field, the DSCP entry parameters will override or remove the IP precedence parameters. When none of the matched entries override a parameter, the ingress classification is preserved.

Hardware Support

The egress SAP forwarding class and forwarding profile override is only supported on SAPs configured on IOM2 and IOM3 modules. If a SAP egress QoS policy with forwarding class and forwarding profile overrides are applied to a SAP on an IOM other than the IOM2 and IOM3 (such as an IOM1), no error message is generated, but the forwarding class and forwarding profile override portion of the SAP egress QoS Policy is ignored and has no effect.

DEI Egress Remarking

It is often desirable to meter traffic from different users to ensure fairness or to meet bandwidth guarantees. Dropping all traffic in excess of a committed rate is likely to result in severe under-utilization of the networks, since most traffic sources are bursty in nature. It is burdensome to meter traffic at all points in the network where bandwidth contention occurs. One solution is to mark those frames in excess of the committed rate as drop eligible on admission to the network.

Previously, the discard eligibility was marked / determined using existing QoS fields: for example, the three MPLS EXP and Ethernet dot1p bits. Using certain combination(s) of these bits to indicate both forwarding class (emission priority) and discard eligibility meant decreasing the number of Forwarding Classes that can be differentiated in the network.

IEEE 802.1ad-2005 and IEEE 802.1ah standards allow drop eligibility to be conveyed separately from priority, preserving all the eight forwarding classes (emission priorities) that could be indicated using the 3 802.1p bits. Now all the previously introduced traffic types will be marked as drop eligible. Customers can continue to use the dot1p markings with the enhancement of changing the dot1p value used, in access, based on the in/out profile information.

The following commands can be used to remark the DE values at a SAP egress:

CLI Syntax:
```
sap-egress <policy-id> create
      fc <fc-name> create
            de-mark [force <de-value>]
            de-mark-inner [force <de-value>]
            de-mark-outer [force <de-value>]
      exit
   exit
```

The precedence of the above commands is summarized as, from highest to lowest precedence:

- de-mark-outer used for outer tag markings
- de-mark-inner used for inner tag markings
- existing de-mark used for marking both tags
- markings taken from packet received at ingress

The configuration of qinq-mark-top-only under the SAP egress takes precedence over the use of the de-mark-inner in the policy, i.e. the inner VLAN tag is not remarked when qinq-mark-top-only is configured (the marking used for the inner VLAN tag is based on the current default which is governed by the marking of the packet received at the ingress to the system). If qinq-mark-top-only is omitted, both the inner and outer VLAN tags are remarked.

Note that the egress remarking occurs after any egress classification.

DEI in IEEE 802.1ad

IEEE 802.1ad-2005 standard allows drop eligibility to be conveyed separately from priority in service VLAN TAGs (STAGs) so that all of the previously introduced traffic types can be marked as drop eligible. The service VLAN TAG has a new format where the priority and discard eligibility parameters are conveyed in the three bit priority code point (PCP) field and respectively in the DE bit (Figure 10).

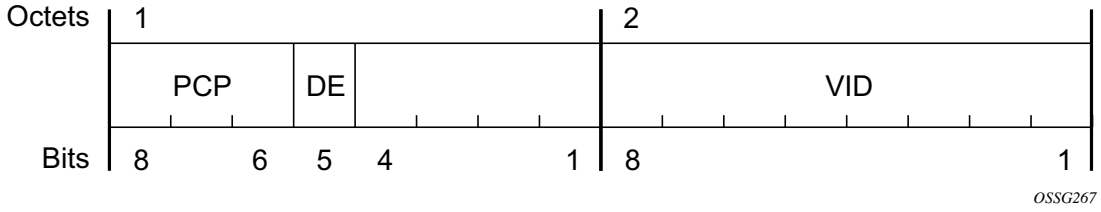

Figure 10: DE Bit in the 802.1ad S-TAG

The introduction of the DE bit allows the S-TAG to convey eight forwarding classes/distinct emission priorities, each with a drop eligible indication.

When DE bit is set to 0 (DE=FALSE) the related packet is not discard eligible. This is the case for the packets that are within the CIR limits and must be given priority in case of congestion. If the DEI is not used or backwards compliance is required the DE bit should be set to zero on transmission and ignored on reception.

When the DE bit is set to 1 (DE=TRUE) the related packet is discard eligible. This is the case for the packets that are sent above the CIR limit (but below the PIR). In case of congestion these packets will be the first ones to be dropped.

DEI in IEEE 802.1ah

IEEE 802.1ah (PBB) standard provides a dedicate bit for DE indication in both the BVID and the ITAG.

The backbone VLAN ID (BVID) is a regular 802.1ad STAG. Its DE bit may be used to convey the related tunnel QoS throughout an Ethernet backbone.

The ITAG header offers also an I-DEI bit that may be used to indicate the service drop eligibility associated with this frame.

These bits must follow the same rules as described in DEI in IEEE 802.1ad on page 207.

IEEE 802.1ad Use Case

Figure 11 illustrates an example of a topology where the new DE feature may be used: a DE aware, 802.1ad access network connected via a regular SAP to a router PE.

In this example, PE1 can ensure coherent processing of the DE indication between the 802.1ad and the MPLS networks: for example, for packets ingressing the SAP connected to 802.1ad access, read the DE indication and perform classification, color-aware metering/policing, marking of the related backbone QoS fields and selective discarding of the frames throughout the queueing system based on their discard eligibility. In addition, packets egressing the SAP towards the 802.1ad access provide proper DE indication by marking the new DE bit in the STAG.

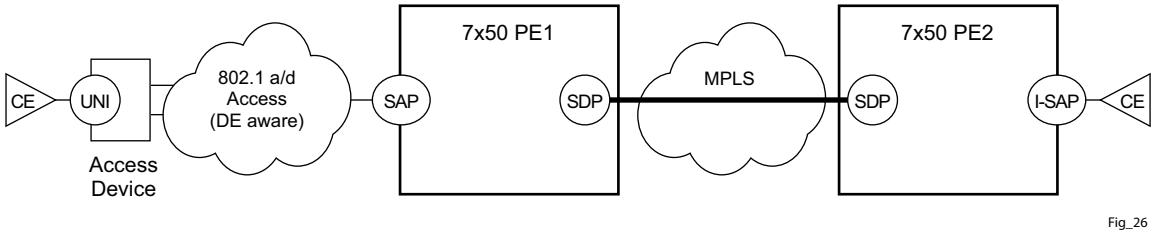

Figure 11: DE Aware 802.1ad Access Network

IEEE 802.1ah Use Case

Figure 12 illustrates an example of a PBB topology where the DE feature can be used. The processing needs highlighted in the 802.1ad use case apply to the 802.1ah BVID, format and etype being identical with the 802.1ad STAG. In addition the DE bit from the 802.1ah ITAG header may need to be processed following the same rules as for the related field in the BVID/STAG: for example, the DE bit from the BVID header represents the QoS associated with the "Ethernet Tunnel" while the DE bit from the ITAG represent the service QoS.

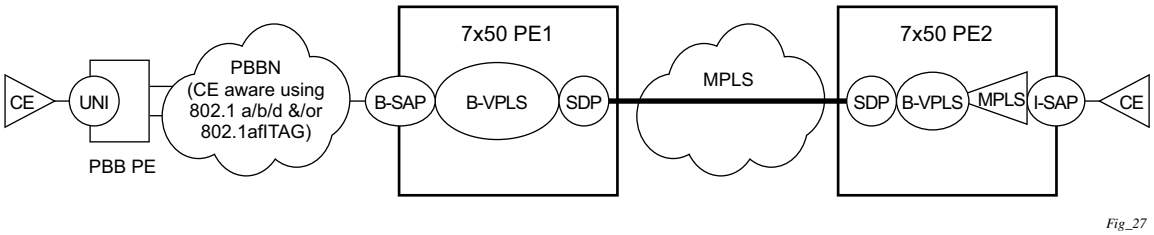

Figure 12: DE Aware PBB Topology

In this example, the BVID is not used for a part of the network leaving only I-DEI bit from the ITAG as the only option for a dedicated DE field. If both are included, then the QoS information from the BVID is to be used.

Egress FC-Based Remarking

FC-based forwarding can be used in a network using core markings of dot1p and may not support DE in all devices. The expectation is that devices beyond the network edge will continue to adhere to the end-to-end QoS policies using dot1p in the packet. Dot1p marking is performed on egress for all services and with respect to in-profile or out-of-profile context.

The following commands can be used to remark the dot1p values at a SAP egress:

CLI Syntax:
```
sap-egress <policy-id> create
        fc <fc-name> create
              dot1p {<dot1p-value>|in-profile <dot1p-value> out-
          profile <dot1p-value>}
               dot1p-inner <dot1p-value>
               dot1p-inner in-profile <dot1p-value> out-profile
          <dot1p-value>
               dot1p-outer <dot1p-value>
               dot1p-outer in-profile <dot1p-value> out-profile
```

```
            <dot1p-value>
        exit
    exit
```

The precedence of the above commands is summarized as, from highest to lowest precedence:

- dot1p-outer used for outer tag markings
- dot1p-inner used for inner tag markings
- existing dot1p used for marking both tags
- markings taken from packet received at ingress

The configuration of qinq-mark-top-only under the SAP egress takes precedence over the use of the dot1p-inner in the policy, i.e. the inner VLAN tag is not remarked when qinq-mark-top-only is configured (the marking used for the inner VLAN tag is based on the current default which is governed by the marking of the packet received at the ingress to the system). If qinq-mark-top-only is omitted, both the inner and outer VLAN tags are remarked.

Note that the egress remarking occurs after any egress classification.

Implementation Requirements

In the 7750 SR series product line, the classification to and (re-)marking from PHB (for example, forwarding class, in/out of profile status) may be described in Table 34.

Table 32: Classification to and (Re-)Marking from PHB

To/From	Classify Ingress Based on	PHB	Mark Egress To
Customer / Access Network (SAP)	dot1p [DE]	FC {in|out}	dot1p [DE]
	DSCP	FC {in|out}	DSCP
	ToS	FC {in|out}	ToS
	IP criteria	FC {in|out}	IP criteria
	MAC criteria	FC {in|out}	MAC criteria
Backbone Network (SDP / B-SAP	dot1p [DE]	FC {in|out}	dot1p [DE]
	DSCP	FC {in|out}	DSCP
	ToS	FC {in|out}	ToS
	EXP	FC {in|out}	EXP

Figure 13 displays a simple example of the DEI processing steps for the IEEE 802.1ad Use Case for both ingress and egress directions (from a PE1 SAP perspective).

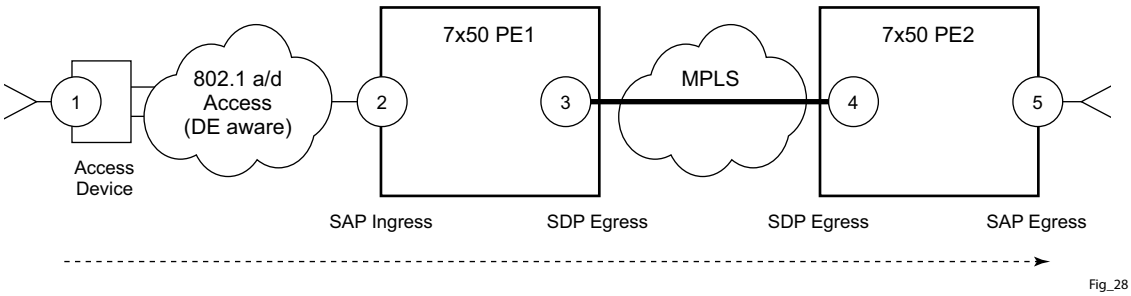

Figure 13: DEI Processing Ingress into the PE1 SAP

The following steps related to DEI are involved in the QoS processing as the packet moves from left to right:

4. The QinQ access device sets the DE bit from the STAG based on the QoS classification or on the results of the metering/policing for the corresponding customer UNI.

4. The SAP on PE1 may use the DE bit from the customer STAG to classify the frames as in/out of profile. Color aware policing/metering can generate addition out of profile packets as the result of packet flow surpassing the CIR.
5. When the packet leaves PE1 via SDP, the DE indication must be copied onto the appropriate tunnel QoS fields (outer VLAN ID and or EXP bits) using the internal PHB (per hop behavior) of the packet (for example, the FC and Profile).
6. As the packet arrives at PE2, ingress into the related SDP, the DE indication is used to classify the packets into an internal PHB.
7. Egress from the PE2 SAP, the internal PHB may be used to perform marking of the DE bit.

A combination of two access networks can be possible. If PBB encapsulation is used, the configuration used for DE in SAP and SDP policies applies to both BVID and ITAG DE bits. When both fields are used the BVID takes precedence.

Default Service Egress and Egress Policy Values

The default service egress and ingress policies are identified as policy-id **1**. The default policies cannot be edited or deleted. The following displays default policy parameters:

SAP Egress Policy

```
A:ALA-7>config>qos>sap-egress$ info detail
----------------------------------------------
            no description
            scope template
            queue 1 auto-expedite create
                no parent
                adaptation-rule pir closest cir closest
                rate max cir 0
                cbs default
                mbs default
                high-prio-only default
            exit
----------------------------------------------
A:ALA-7>config>qos>sap-egress$
```

Table 33: SAP Egress Policy Defaults

Field	Default
description	"Default SAP egress QoS policy."
scope	template
queue 1	1 auto-expedite
parent	no parent
adaptation-rule	adaptation-rule pir closest cir closest
rate	max cir 0
cbs	default
mbs	default
high-prio-only	default

Default SAP Ingress Policy

```
A:ALA-7>config>qos>sap-ingress$ info detail
----------------------------------------------
            description "Default SAP ingress QoS policy"
            scope template
            queue 1 auto-expedite create
                no parent
                adaptation-rule pir closest cir closest
                rate max cir 0
                mbs default
                cbs default
                high-prio-only default
            exit
            queue 2 multipoint auto-expedite create
                no parent
                adaptation-rule pir closest cir closest
                rate max cir 0
                mbs default
                cbs default
                high-prio-only default
            exit
            default-fc be
            default-priority low
----------------------------------------------
A:ALA-7>config>qos>sap-ingress$
```

Table 34: SAP Ingress Policy Defaults

Field	Default
description	"Default SAP ingress QoS policy."
scope	template
queue 1	1 priority-mode auto-expedite
parent	no parent
adaptation-rule	adaptation-rule pir closest cir closest
rate	max cir 0
cbs	default
mbs	default
high-prio-only	default
queue 2	multipoint priority-mode auto-expedite
parent	no parent
adaptation-rule	adaptation-rule pir closest cir closest
rate	max cir 0

Table 34: SAP Ingress Policy Defaults (Continued)

Field	Default
cbs	default
mbs	default
high-prio-only	default
default-fc	be
default-priority	low

Basic Configurations

A basic service egress QoS policy must conform to the following:

- Have a unique service egress QoS policy ID.
- Have a QoS policy scope of template or exclusive.
- Have at least one defined default queue.

A basic service ingress QoS policy must conform to the following:

- Have a unique service ingress QoS policy ID.
- Have a QoS policy scope of template or exclusive.
- Have at least one default unicast forwarding class queue.
- Have at least one multipoint forwarding class queue.

Create Service Egress and Ingress QoS Policies

Configuring and applying QoS policies is optional. If no QoS policy is explicitly applied to a SAP or IP interface, a default QoS policy is applied.

- Percent-Rate Support on page 216
- Service Egress QoS Policy on page 218
- Service Ingress QoS Policy on page 222

Percent-Rate Support

The **percent-rate** command is supported for pir and cir parameters for both queues and policers. Also supported is the capability of specifying the rate as a percentage value of the line rate for sap-ingress and sap-egress qos policies. It is supported for both queues and policers. The user has the option of specifying **percent-rate** for **pir** and **cir** parameters. For **pir**, the range is 0.01 to 100.00, and for **cir**, the range is 0.00 to 100.00.

The rate can be also configured using the existing keyword **rate** in Kbps.

For queues, when the queue rate is in percent-rate either a local-limit or a port-limit can be applied.

When the local-limit is used the percent-rate is relative to the queue's parent scheduler rate or the agg-rate rate at egress, when the port-limit is used the percent-rate is relative to the rate of the port (including the ingress-rate/egress-rate setting) to which the queue is attached. port-limit is the default.

For policers, the percent-rate rate is always relative to the immediate parent root policer/arbiter rate or the FP capacity.

SAP Ingress QoS Policy:

```
*B:Dut-A>config>qos>sap-ingress# queue 1 percent-rate
- no percent-rate
- percent-rate <pir-percent> [cir <cir-percent>] [port-limit|local-limit]
- percent-rate <pir-percent> police [port-limit|local-limit]

<pir-percent> : [0.01..100.00]
<cir-percent> : [0.00..100.00]
<police> : keyword
<port-limit|local-*> : keyword

*B:Dut-A>config>qos>sap-ingress# policer 1 percent-rate
- no percent-rate
- percent-rate <pir-percent> [cir <cir-percent>]

<pir-percent> : [0.01..100.00]
<cir-percent> : [0.00..100.00]
```

SAP-Egress QoS Policy:

```
*B:Dut-A>config>qos>sap-egress# queue 1 percent-rate
- no percent-rate
- percent-rate <pir-percent> [cir <cir-percent>] [port-limit|local-limit]

<pir-percent> : [0.01..100.00]
<cir-percent> : [0.00..100.00]
<port-limit|local-*> : keyword

*B:Dut-A>config>qos>sap-egress# policer 1 percent-rate
- no percent-rate
- percent-rate <pir-percent> [cir <cir-percent>]

<pir-percent> : [0.01..100.00]
<cir-percent> : [0.00..100.00]
```

Service Egress QoS Policy

To create a service egress policy, you must define the following:

- A new policy ID value. The system will not dynamically assign a value.
- Specify the scope. A QoS policy must be defined as having either an *exclusive* scope for one-time use, or a *template* scope which enables its use with multiple SAPs.
- Include a description. The description provides a brief overview of policy features.

After the policy is created, the policy's behavior can be defined:

- Specify the forwarding class. The forwarding class name or names associated with the egress queue. The egress queue for the service traffic is selected based on the forwarding classes that are associated with the queue.
- A new queue ID value. The system will not dynamically assign a value.
- Define queue parameters. Queues support explicit and auto-expedite hardware queue scheduling, and parent virtual scheduler definition.

The following displays an egress QoS policy configuration:

```
A:ALA-7>config>qos# info
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
...
     sap-egress 105 create
            description "SAP egress policy"
            queue 1 create
            exit
            queue 2 create
            exit
            queue 3 expedite create
                parent test1
            exit
            fc af create
                queue 1
            exit
            fc ef create
                queue 2
            exit
        exit
...
#------------------------------------------
A:ALA-7>config>qos#
```

Service Egress QoS Queue

To create a service egress queue parameters, define the following:

- A new queue ID value. The system will not dynamically assign a value.
- Define queue parameters. Egress queues support explicit and auto-expedite hardware queue scheduling, and parent virtual scheduler definition.

The following displays an egress QoS policy configuration:

```
A:ALA-7>config>qos# info
----------------------------------------------
...
        sap-egress 105 create
            description "SAP egress policy"
            queue 1 create
                parent "scheduler-tier1"
            exit
            queue 2 create
            exit
            queue 3 expedite create
                parent "test1"
            exit
            fc af create
                queue 1
            exit
            fc ef create
            exit
        exit
...
----------------------------------------------
A:ALA-7>config>qos#
```

Egress Criteria Classification Directly to Policer

It is possible to classify traffic directly to a policer, independent of the policer/queue assigned to the traffic's forwarding class. This is supported at SAP egress by configuring a policer in the **action** statement within an **ip-criteria** or **ipv6-criteria** statement.

The policed traffic by default exits through one of the following methods:

- A queue in the **policer-output-queues queue** group that is automatically created on an access or hybrid port with the queue used that was chosen by the forwarding class definition in that queue group. If the forwarding class is modified in the **action** statement then the new forwarding class selects the queue to be used.
- A specific queue in a user configured queue group. For SAP egress, this requires the use of the **port-redirect-queue-group queue** parameter in the criteria **action** statement with the queue group name being specified when the egress QoS policy is applied to the SAP. For subscribers, the queue group to be used is selected using the inter-dest-id associated with the subscriber and configured as the **host-match dest** under the port access queue group configuration.
- A SAP queue configured within the SAP egress QoS policy.
- The queue to which the forwarding class for the traffic is mapped. This could be a queue group, SAP, or subscriber queue. This requires the use of the **use-fc-mapped-queue** parameter in the criteria **action** statement. If the forwarding class is modified in the **action** statement then new forwarding class selects the queue to be used.

The number of configuration combinations of a policer and one of the above methods is capped at 63 within a given SAP egress QoS policy. For two or more definitions to be counted as a single combination, their action statement must have the same policer ID, the same queue ID (if specified in either statement), the same **port-redirect-queue-group** (if specified in either statement) and the parameter **use-fc-mapped-queues** (if specified in either statement). The forwarding class and profile used are irrelevant when considering the number of combinations. For example, it is possible to configure 32 policers with traffic exiting queue 1 but then, only 31 of the same policers are exiting queue 2; this would use all 63 combinations. A resource is also allocated per FP where each combination configured corresponds to an egress bypass entry used in the FP per sap-instance or per subscriber-sap-sla instance which use the egress qos policy. The number of egress bypass entries available on an FP, together with the number allocated and the number free, can be seen using the following tools command.

```
A:PE# tools dump system-resources 1
Resource Manager info at 002 d 05/27/15 13:18:44.784:

Hardware Resource Usage for Slot #1, CardType iom3-xp, Cmplx #0:
                                |   Total   | Allocated |    Free
  ------------------------------|-----------|-----------|------------
...                             |           |           |
               Egress QoS Bypass |     262143|          1|     262142
```

This is supported on all FP2- and higher-based hardware, excluding when a HS-MDA is used. QPPB processing takes precedence over this feature.

This could be used, for example, when it is required that egress traffic with a DSCP value EF is to be policed instead of shaped in a queue on a given SAP. The traffic could be classified based on its DSCP value and directed to **policer 1** while the remainder of the customer's traffic is processed using egress **queue 1**. This is shown in Figure 14.

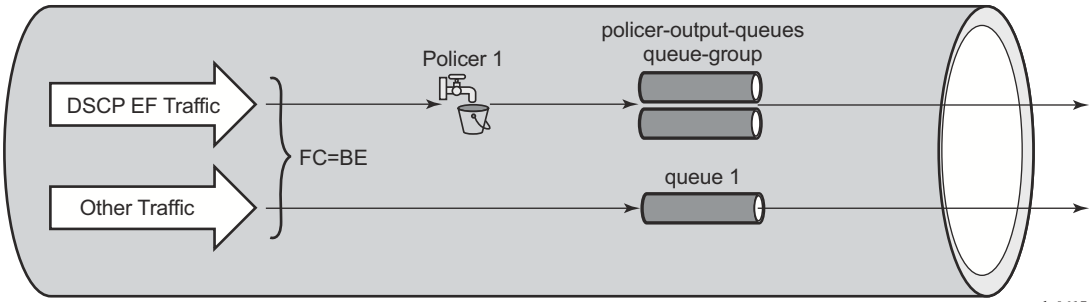

Figure 14: Egress SAP

The configuration would be as follows:

```
sap-egress 10 create
    queue 1 create
    exit
    policer 1 create
    exit
    ip-criteria
        entry 10 create
            match
                dscp ef
            exit
            action policer 1
        exit
    exit
exit
```

Service Ingress QoS Policy

To create an service ingress policy, define the following:

- A policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- Specify a default forwarding class for the policy. All packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class.
- Specify a default priority for all packets received on an ingress SAP using this policy.
- Define mappings from incoming packet contents to a forwarding class, and then, separately, from the forwarding class to queue.
- Define forwarding class parameters.
  - → Modify the **multicast-queue** default value to override the default multicast forwarding type queues mapping for **fc** *fc-name*.
  - → Modify the **unknown-queue** default value to override the default unknown unicast forwarding type queues mapping for **fc** *fc-name*.
  - → Modify the **broadcast-queue** default value to override the default broadcast forwarding type queues mapping for **fc** *fc-name*.
- Configure precedence value for the forwarding class or enqueuing priority when a packet is marked with an IP precedence value.
- Specify IP, IPv6 or MAC criteria. You can define IP, IPv6 and MAC-based SAP ingress policies to select the appropriate ingress queue and corresponding forwarding class for matched traffic.
- A SAP ingress policy is created with a template scope. The scope can be modified to exclusive for a special one-time use policy. Otherwise, the **template** scope enables the policy to be applied to multiple SAPs.

The following displays an service ingress policy configuration:

```
A:ALA-7>config>qos>sap-ingress# info
----------------------------------------------
...
        sap-ingress 100 create
            description "Used on VPN sap"
...
----------------------------------------------
A:ALA-7>config>qos>sap-ingress#
```

Service Ingress QoS Queue

To create service ingress queues parameters, define the following:

- A new queue ID value — The system will not dynamically assign a value.
- Queue parameters — Ingress queues support multipoint queues, explicit and auto-expedite hardware queue scheduling, and parent virtual scheduler definition.

The following displays an ingress queue configuration:

```
A:ALA-7>config>qos# info
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
...
        sap-ingress 100 create
            description "Used on VPN sap"
            queue 1 create
            exit
            queue 2 multipoint create
            exit
            queue 10 create
                parent VPN_be
                rate 11000
            exit
            queue 12 create
                parent VPN_priority
                rate 11000
            exit
            queue 13 create
                parent VPN_reserved
                rate 1
            exit
            queue 15 create
                parent VPN_video
                rate 1500 cir 1500
            exit
            queue 16 create
                parent VPN_voice
                rate 2500 cir 2500
            exit
            queue 17 create
                parent VPN_nc
                rate 100 cir 36
            exit
            queue 20 multipoint create
                parent VPN_be
                rate 11000
            exit
            queue 22 multipoint create
                parent VPN_priority
                rate 11000
            exit
            queue 23 multipoint create
                parent VPN_reserved
```

```
                rate 1
            exit
            queue 25 multipoint create
                parent VPN_video
                rate 1500 cir 1500
            exit
            queue 26 multipoint create
                parent VPN_voice
                rate 2500 cir 2500
            exit
            queue 27 multipoint create
                parent VPN_nc
                rate 100 cir 36
            exit
...
#--------------------------------------------
A:ALA-7>config>qos#
```

SAP Ingress Forwarding Class (FC)

The following displays a forwarding class and precedence configurations:

```
A:ALA-7>config>qos# info
#------------------------------------------
...
            fc af create
                queue 12
                broadcast-queue 22
                multicast-queue 22
                unknown-queue 22
            exit
            fc be create
                queue 10
                broadcast-queue 20
                multicast-queue 20
                unknown-queue 20
            exit
            fc ef create
                queue 13
                broadcast-queue 23
                multicast-queue 23
                unknown-queue 23
            exit
            fc h1 create
                queue 15
                broadcast-queue 25
                multicast-queue 25
                unknown-queue 25
            exit
            fc h2 create
                queue 16
                broadcast-queue 26
                multicast-queue 26
                unknown-queue 26
            exit
            fc nc create
                queue 17
                broadcast-queue 27
                multicast-queue 27
                unknown-queue 27
            exit
            prec 0 fc be
            prec 2 fc af
            prec 3 fc ef
            prec 5 fc h1
            prec 6 fc h2
            prec 7 fc nc
...
#------------------------------------------
A:ALA-7>config>qos#
```

Service Ingress IP Match Criteria

When specifying SAP ingress match criteria, only one match criteria type (IP/IPv6 or MAC) can be configured in the SAP ingress QoS policy.

The following displays an ingress IP criteria configuration:

```
A:ALA-7>config>qos# info
...
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
...
        sap-ingress 100 create
...
            ip-criteria
                entry 10 create
                    description "Entry 10-FC-AF"
                    match protocol 6
                        src-ip 10.10.10.103/24
                    exit
                    action fc af priority high
                exit
                entry 20 create
                    description "Entry 20-FC-BE"
                    match protocol 17
                        dst-port eq 255
                    exit
                    no action
                exit
            exit
        exit
..
#------------------------------------------
A:ALA-7>config>qos#
```

Service Ingress IPv6 Match Criteria

When specifying SAP ingress match criteria, only one match criteria type (IP/IPv6 or MAC) can be configured in the SAP ingress QoS policy.

The following displays an ingress IPv6 criteria configuration:

```
A:ALA-48>config>qos>sap-ingress# info
----------------------------------------------
            queue 1 create
            exit
            queue 11 multipoint create
            exit
            ip-criteria
            exit
            ipv6-criteria
                entry 10 create
                    description "IPv6 SAP-ingress policy"
                    match
                        src-ip ::/96
                        dst-ip 200::/7
                    exit
                    action fc be priority low
                exit
                entry 20 create
                    description "Entry 20-FC-AF"
                    match next-header tcp
                        src-port eq 500
                    exit
                    action fc af priority high
                exit
            exit
----------------------------------------------
A:ALA-48>config>qos>sap-ingress#
```

Service Ingress MAC Match Criteria

Both IP/IPv6 criteria and MAC criteria cannot be configured in the same SAP ingress QoS policy.

To configure service ingress policy MAC criteria, define the following:

- A new entry ID value. Entries must be explicitly created. The system will not dynamically assign entries or a value.
- The action to associate the forwarding class or enqueuing priority with a specific MAC criteria entry ID.
- A description. The description provides a brief overview of policy features.
- Match criteria for ingress SAP QoS policy. Optionally, specify an IP protocol to be used as an ingress SAP QoS policy match criterion.

The following displays an ingress MAC criteria configuration:

```
A:ALA-7>config>qos# info
...
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
...
        sap-ingress 101 create
...
            mac-criteria
                entry 10 create
                    description "Entry10-low prio"
                    match
                        dst-mac 04-67-ff-00-00-01 ff-ff-ff-ff-ff-ff
                        dot1p 7 7
                    exit
                    action fc be priority low
                exit
            exit
        exit
#------------------------------------------
A:ALA-7>config>qos#
```

Ingress Criteria Classification Directly to Policer

It is possible to classify traffic directly to a policer, independent of the policer/queue assigned to the traffic's forwarding class. This is supported at SAP ingress when using one of the following statements: ip-criteria, ipv6-critera or mac-criteria.

The standard mechanisms are still used to assign a forwarding class to the related traffic, and this forwarding class continues to be used for QOS processing at egress.

This is supported on all FP2 and higher based line cards. The use of explicitly configured broadcast, unknown, or multicast policers is not supported. QPPB processing takes precedence over this feature.

This could be used, for example, when it is required that ingress OAM traffic is not subject to the same QOS control as other customer traffic on a given SAP. The OAM traffic could be classified based on its source MAC address (for example, with an OUI of 00-xx-yy as configured below) and directed to policer 1 while the remainder of the customer's traffic is processed using ingress queue 1. This is shown in Figure 15.

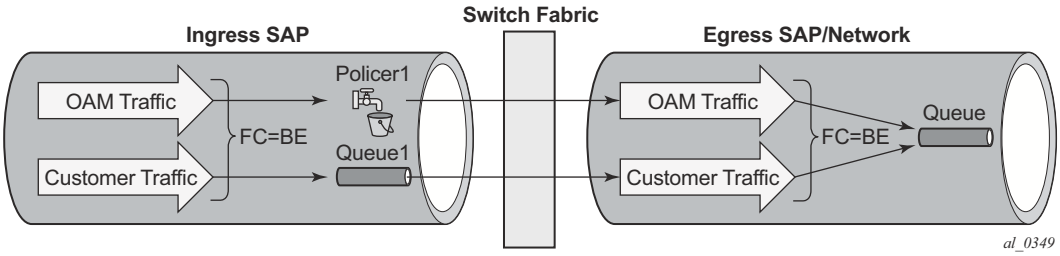

Figure 15: Ingress Criteria Classification Directly to Policer

The configuration would be as follows:

```
sap-ingress 10 create
    queue 1 create
    exit
    queue 11 multipoint create
    exit
    policer 1 create
    exit
    mac-criteria
        entry 10 create
            match
                src-mac 00-xx-yy-00-00-00 ff-ff-ff-00-00-00
            exit
            action policer 1
        exit
```

```
        exit
    exit
```

FC Mapping Based on EXP Bits

You can use the **lsp-exp** command to set your sap-ingress qos policy on Ethernet L2 SAPs to perform FC mapping based on EXP bits.

The **lsp-exp** option causes the forwarding class and drop priority of incoming traffic to be determined by the mapping result of the EXP bits in the top label.

The following example displays FC mapping based on EXP bits:

```
*A:Dut-T>config>qos>sap-ingress# info
----------------------------------------------
            queue 1 create
            exit
            queue 2 create
            exit
            queue 3 create
            exit
            queue 11 multipoint create
            exit
            fc "af" create
                queue 2
            exit
            fc "be" create
                queue 1
            exit
            fc "ef" create
                queue 3
            exit
            lsp-exp 0 fc "be" priority low
            lsp-exp 1 fc "af" priority high
            lsp-exp 2 fc "ef" priority low hsmda-counter-override 1
            lsp-exp 3 fc "ef" priority high hsmda-counter-override 2
```

VID Filters

VID filters extend the capability of current Ethernet ports with null or default SAP tag configuration to match and take action on VID tags. Service delimiting tags (for example qinq 1/1/1:10.20 or dot1q 1/1/1:10, where outer tag 10 and inner tags 20 are service delimiting) allow fine grain control of frame operations based on the VID tag. Service delimiting tags are exact match and are stripped from the frame as illustrated in Figure 16. Exact match or service delimiting tags do not require VID filters. VID filters can only be used to match on frame tags that are after the service delimiting tags.

With VID Filters operators can choose to match VID tags for up to two tags on ingress or egress or both.

- The outer-tag is the first tag in the packet that is carried transparently through the service.
- The inner-tag is the second tag in the packet that is carried transparently through the service.

VID filters add the capability to perform VID value filter policies on default tags (1/1/1:* or 1/1/1:x.*, or 1/1/1/:*.0), or null tags (1/1/1, 1/1/1:0 or 1/1/1:x.0). The matching is based on the port configuration and the SAP configuration.

QinQ tags are often referred to as the C-VID (Customer VID) and S-VID (service VID). The terms outer tag and inner tag allow flexibility without having to refer to C-TAG and an S-TAG explicitly. The position of inner and outer tags is relative to the port configuration and SAP configuration. Matching of tags is allowed for up to the first two tags on a frame. Since service delimiting tags may be 0, 1 or 2 tags.

The meaning of inner and outer has been designed to be consistent for egress and ingress when the number of non service delimiting tags is consistent. Service 1 in Figure 16 shows a conversion from qinq to a single dot1q example where there is one non-service delimiting tag on ingress and egress. Service 2 shows a symmetric example with two non-service delimiting tags (plus an additional tag for illustration) to two non-service delimiting tags on egress. Service 3 illustrates single non-service delimiting tags on ingress and to two tags with one non-service delimiting tag on ingress and egress.

SAP-ingress QoS setting allows for MAC-criteria type VID which uses the VID filter matching capabilities (see QoS and VID Filters on page 234).

A VID filter entry can be used as a debug or lawful intercept mirror source entry.

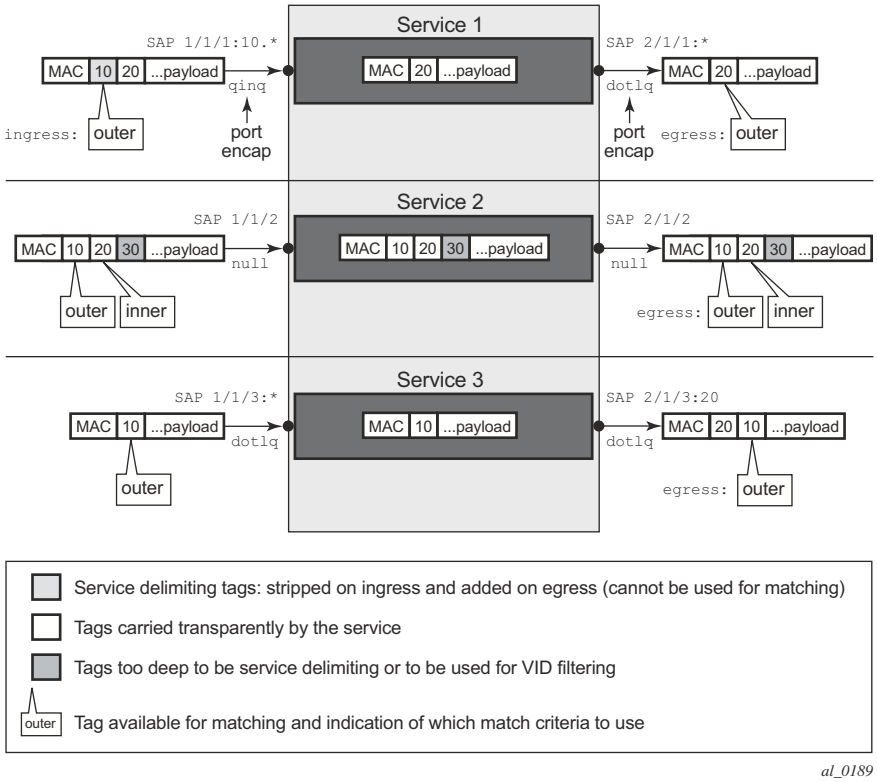

Figure 16: VID Filtering Examples

VID filters are available on Ethernet SAPs for Epipe, VPLS or I-VPLS including eth-tunnel and eth-ring services.

Arbitrary Bit Matching of VID Filters

In addition to matching an exact value, a VID filter mask allows masking any set of bits. The masking operation is ((value & vid-mask) = = (tag and vid-mask)). For example: A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6. VID filters allow explicit matching of VIDs and matching of any bit pattern within the VID tag.

When using VID filters on SAPs only VID filters are allowed on this SAP. Filters of type normal and ISID are not allowed.

An additional check for the “0” VID tag may be required when using certain wild card operations. For example frames with no tags on null encapsulated ports will match a value of 0 in outer tag and inner tag because there are no tags in the frame for matching. If a zero tag is possible but not desired it can be explicitly filtered using exact match on “0” prior to testing other bits for “0”.

Note that **configure>system>ethernet>new-qinq-untagged-sap** is a special QinQ function for single tagged QinQ frames with a null second tag. Using this in combination with VID filters is not recommended. Note that the outer-tag is the only tag available for filtering on egress for frames arriving from MPLS SDPs or from PBB services even though additional tags may be carried transparently.

QoS and VID Filters

On ingress VID filtering may also be used to set QoS on SAP ingress. The matching rules are the same as for VID filter but the action allows setting of the forwarding class.

For example, to set the forwarding class of all VIDs with 6 in the lower 3 bits of the VID a filter as illustrated below could be constructed and then ingress qos 5 could be applied to any SAP that requires the policy.

```
qos
        sap-ingress 5 create
            queue 1 create
            exit
            queue 11 multipoint create
            exit
            mac-criteria
                type vid
                entry 1 create
                    match frame-type ethernet-II
                        outer-tag 6 7
                    exit
                    action fc "af"
                exit
            exit
        exit
    exit
```

Port Group Configuration Example

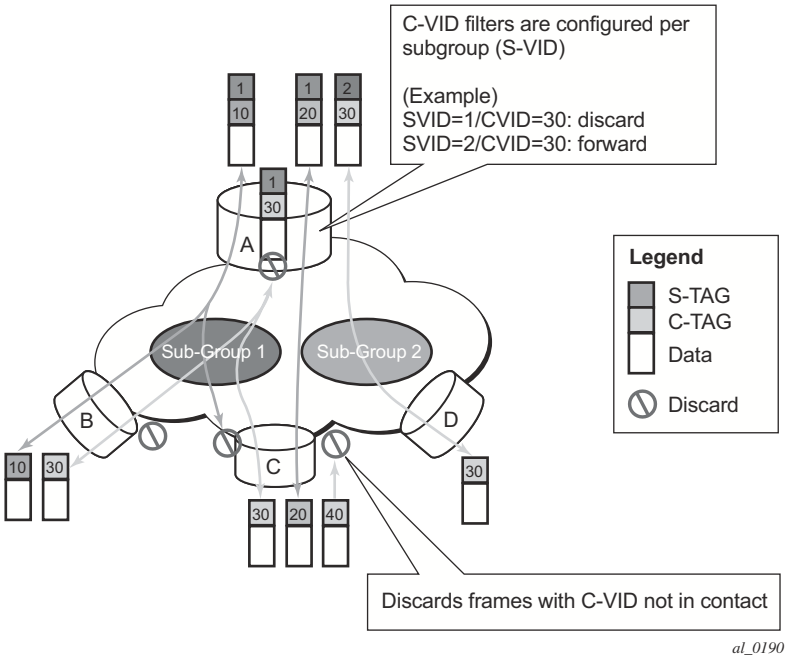

Figure 17: Port Groups

Figure 17 shows a customer use example where some VLANs are prevented from ingressing or egressing certain ports. In the example, port A sap 1/1/1:1.* would have a filter as shown below while port A sap 1/1/1:2.* would not.:

```
mac-filter 4 create
     default-action forward
            type vid
            entry 1 create
                match frame-type ethernet_II
                    outer-tag 30 4095
                exit
                action drop
            exit
        exit
```

Applying Service Ingress and Egress Policies

Apply SAP ingress and egress policies to the following service SAPs:

- Epipe
- IES
- VPLS
- VPRN

Refer to the Subscriber Services Overview section of the Services Guide for information about configuring service parameters.

Epipe

The following output displays an Epipe service configuration with SAP ingress policy 100 and SAP egress 105 applied to the SAP.

```
A:ALA-7>config>service# info
----------------------------------------------
        epipe 6 customer 6 vpn 6 create
            description "Distributed Epipe service to west coast"
            sap 1/1/10:010 create
                ingress
                    qos 100
                exit
                egress
                    qos 105
                exit
            exit
            spoke-sdp 2:6 create
                ingress
                    vc-label 6298
                exit
                egress
                    vc-label 6300
                exit
            exit
            no shutdown
        exit
----------------------------------------------
A:ALA-7>config>service#
```

IES

The following output displays an IES service configuration with SAP ingress policy 100 and SAP egress 105 applied to the SAP.

```
A:ALA-7>config>service# info
----------------------------------------------
        ies 88 customer 8 vpn 88 create
            interface "Sector A" create
                sap 1/1/1.2.2 create
                    ingress
                        qos 100
                    exit
                    egress
                        qos 105
                    exit
                exit
            exit
            no shutdown
        exit
----------------------------------------------
```

VPLS

The following output displays a VPLS service configuration with SAP ingress policy 100. The SAP egress policy 1 is applied to the SAP by default.

```
A:ALA-7>config>service# info
----------------------------------------------
        vpls 700 customer 7 vpn 700 create
            description "test"
            stp
                shutdown
            exit
            sap 1/1/9:010 create
                ingress
                    qos 100
                exit
            exit
            spoke-sdp 2:222 create
            exit
            mesh-sdp 2:700 create
            exit
            no shutdown
        exit
----------------------------------------------
A:ALA-7>config>service#
```

VPRN

The following output displays a VPRN service configuration.

```
A:ALA-7>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            ecmp 8
            autonomous-system 10000
            route-distinguisher 10001:1
            auto-bind-tunnel
                resolution-filter
                resolution-filter ldp
            vrf-target target:10001:1
            interface "to-ce1" create
                address 11.1.0.1/24
                sap 1/1/10:1 create
                    ingress
                        qos 100
                    exit
                    egress
                        qos 105
                    exit
                exit
            exit
            no shutdown
        exit
...
----------------------------------------------
A:ALA-7>config>service#
```

Service Management Tasks

This section discusses the following service management tasks:

Deleting QoS Policies

Every service SAP is associated, by default, with the appropriate egress or ingress policy (policy-id **1**). You can replace the default policy with a customer-configured policy, but you cannot entirely remove the policy from the SAP configuration. When you remove a non-default service egress or ingress policy, the association reverts to the default policy-id **1**.

A QoS policy cannot be deleted until it is removed from all SAPs where they are applied.

```
A:ALA-7>config>qos# no sap-ingress 100
MINOR: CLI SAP ingress policy "100" cannot be removed because it is in use.
A:ALA-7>config>qos#
```

Remove a QoS Policy from Service SAP(s)

The following Epipe and VPRN service output examples show that the SAP service egress and ingress reverted to policy-id "**1**" when the non-default policies were removed from the configuration.

```
A:ALA-104>config>service>epipe# info detail
----------------------------------------------
            description "Distributed Epipe service to west coast"
            service-mtu 1514
            sap 1/1/10:0 create
                no description
                no multi-service-site
                ingress
                    no scheduler-policy
                    qos 1
                exit
                egress
                    no scheduler-policy
                    qos 1
                exit
                no collect-stats
                no accounting-policy
```

```
                no shutdown
            exit
            spoke-sdp 2:6 vc-type ether create
                no shutdown
            exit
            no shutdown
-------------------------------------------
A:ALA-7>config>service>epipe#

A:ALA-7>config>service>vprn#
-------------------------------------------
...
        vprn 1 customer 1 create
            ecmp 8
            autonomous-system 10000
            route-distinguisher 10001:1
            auto-bind-tunnel
                resolution-filter
                resolution-filter ldp
            vrf-target target:10001:1
            interface "to-ce1" create
                address 11.1.0.1/24
                sap 1/1/10:1 create
                exit
            exit
            no shutdown
        exit
-------------------------------------------
A:ALA-7>config>service>vprn#
```

Copying and Overwriting QoS Policies

You can copy an existing service egress or ingress policy, rename it with a new policy ID value, or overwrite an existing policy ID. The `overwrite` option must be specified or an error occurs if the destination policy ID exists.

CLI Syntax: `config>qos# copy {sap-ingress | sap-egress}` *`source-policy-id dest-policy-id`* `[overwrite]`

The following output displays the copied policies:

```
A:ALA-7>config>qos# info
----------------------------------------------
...
exit
        sap-ingress 100 create
            description "Used on VPN sap"
            queue 1 create
            exit
            queue 2 multipoint create
            exit
            queue 10 create
                parent "VPN_be"
                rate 11000
            exit
...
        sap-ingress 101 create
            description "Used on VPN sap"
            queue 1 create
            exit
            queue 2 multipoint create
            exit
            queue 10 create
                parent "VPN_be"
                rate 11000
            exit
        sap-ingress 200 create
            description "Used on VPN sap"
            queue 1 create
            exit
            queue 2 multipoint create
            exit
            queue 10 create
                parent "VPN_be"
                rate 11000
            exit
...
----------------------------------------------
A:ALA-7>config>qos#
```

Remove a Policy from the QoS Configuration

CLI Syntax: `config>qos# no sap-ingress` *`policy-id`*

Example:
```
config>qos# no sap-ingress 100
config>qos# no sap-egress 1010
```

Editing QoS Policies

You can change QoS existing policies and entries. The changes are applied immediately to all services where this policy is applied. To prevent configuration errors copy the policy to a work area, make the edits, and then write over the original policy.

Queue Depth Monitoring

Queue depth monitoring gives more visibility to the operator of the queue depths being experienced on a set of queues when the traffic is bursty. The instantaneous depth of a queue can be seen using the **show pools** command, whereas queue depth monitoring shows the variation in queue depth over a period of time. It is applicable to SAP ingress unicast and multipoint queues and SAP egress queues, and for ingress and egress access and network queue group queues used by any service or network interfaces. The monitoring uses a polling mechanism by the line card CPU. Consequently, the results provided are statistical in nature. This is supported on FP2- and higher-based line cards.

An override (**monitor-depth**) is used to enable queue depth monitoring, which is configured under the SAP or queue group queue-overrides. There are show and clear commands, using the **queue-depth** parameter, for both service SAPs and port queue groups with associated MIB variables.

The configuration below gives an example of enabling the monitoring of the depth of queue 1 on an Epipe SAP.

```
epipe 1 customer 1 create
            sap 1/2/1 create
                egress
                    qos 10
                    queue-override
                        queue 1 create
                            monitor-depth
                        exit
                    exit
                exit
            exit
```

The queue depth can then be shown as follows:

```
*A:PE-1# show service id 1 sap 1/2/1 queue-depth

===============================================================================
Queue Depth Information (Ingress SAP)
===============================================================================
No Matching Entries
===============================================================================

===============================================================================
Queue Depth Information (Egress SAP)
===============================================================================
-------------------------------------------------------------------------------
Name                       : 1->1/2/1->1
MBS                        : Def

-------------------------------------------------------------------------------
Queue Depths (percentage)
-------------------------------------------------------------------------------
0%-10% 11%-20% 21%-30% 31%-40% 41%-50% 51%-60% 61%-70% 71%-80% 81%-90% 91%-100%
-------------------------------------------------------------------------------
68.21  3.64    3.43    3.47    3.86    3.22    3.86    2.87    3.78    3.66
-------------------------------------------------------------------------------
Average Elapsed Time      : 0d 00:11:48
Wghtd Avg Polling Interval: 99 ms
-------------------------------------------------------------------------------
===============================================================================
*A:PE-1#
```

The output shows the percentage of polls for each 10% range of queue depth. The output includes the name of the queue, its MBS configuration, the average elapsed time over which the depth was monitored (this is the elapsed time since the start of monitoring or the last clear), and the weighted average polling interval.

For example, in the above output, the queue depth was in the range of 51% to 60% for 3.22 percent of the polls, the polling was performed over an elapsed time of 11 minutes and 48 seconds, and the average polling interval was 99ms.

The monitoring is performed on the hardware queues corresponding to the configured queue. It is possible that the set of related hardware queues for a given configured queue changes over time. For example, when LAG ports are added or removed resulting in monitored hardware queues being added or removed. If the set of hardware queues for the configured queue changes, the system will only report occupancy information of all currently instantiated hardware queues, specifically, no attempt is made to keep historical occupancy information. The average polling interval is weighted based on the elapsed monitoring time of the individual hardware queues corresponding to the configured queue, and the elapsed monitoring time is averaged over the same set of hardware queues.

There is no specific limit on the number of queues that can be monitored but the amount of each line card CPU's resources allocated to the monitoring is bounded, consequently average polling interval will increase as more queues are monitored on the line card.

If the MBS of a queue is modified, the occupancy information is cleared and the elapsed timers reset to zero. Issuing a clear card will also clear this information. Note that packet drops caused at the pool level, rather than at the queue level, would result in lower queue depths being reported.