

Class Fair Hierarchical Policing for SAPs

In This Chapter

This section provides information to configure Class Fair Hierarchical Policing for SAPs.

Topics in this section include:

- [Applicability on page 1610](#)
- [Summary on page 1611](#)
- [Overview on page 1612](#)
- [Configuration on page 1623](#)
- [Conclusion on page 1652](#)

Applicability

The information in this note is applicable to all of the Alcatel-Lucent 7x50 platforms and is focused on the FP2 chipset, which is used in the IOM3-XP/IMMs and in the 7750 SRc-12/4. The configuration was tested on release 9.0R1. There are no specific pre-requisites for this configuration.

Summary

The Quality of Service (QoS) features of the 7x50 platforms provide traffic control with both shaping and policing.

Shaping is achieved using a queue; packets are placed on the queue and a scheduler removes packets from the queue at a given rate. This provides an upper bound to the traffic rate sent, thereby protecting downstream devices from bursts. However, shaping can introduce latency and jitter as packets are delayed in the queue. Packets can be dropped when the queue is full or statistically when weighted random early discard is applied. Configuration of shaping on the 7x50 is described in [QoS Architecture and Basic Operation on page 1679](#).

Policing is another mechanism for controlling traffic rates but it does not introduce latency/jitter. This is achieved using a token bucket mechanism which drops certain packets from the traffic. A common disadvantage of policing implementations is that they are usually applicable to a single level of traffic priority and have no way to fairly share capacity between multiple streams at the same priority level. Alcatel-Lucent's Class Fair Hierarchical Policing (CFHP) addresses these problems by implementing a four level prioritized policing hierarchy which also provides weighted fairness for traffic at a given priority.

Regardless of whether shaping or policing is being used, the preceding QoS classification and subsequent packet marking functionality is similar for both and is covered in more detail in [QoS Architecture and Basic Operation on page 1679](#).

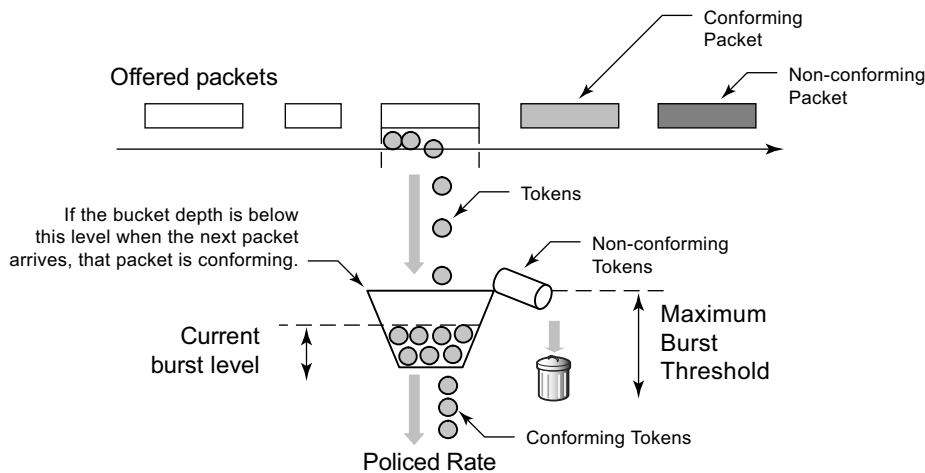
This note describes the configuration and operation of CFHP when applied to Service Access Points (SAPs). It is also possible to use CFHP for subscribers in a Triple Play Service Delivery Architecture (TPSDA) environment but it is beyond the scope of this note.

Overview

Policers

CFHP can be used both for ingress and egress QoS. The basic element is a policer which can apply both a committed information rate (CIR) and peak information rate (PIR) to a traffic flow (determined by the ingress classification). Traffic is directed to a policer by assigning a forwarding class (FC) to the policer.

To describe the operation of a policer we will use a token bucket model, this is shown in [Figure 244](#).



OSSG513

Figure 244: Policer Token Bucket Model

The policer is modeled by a bucket being filling with tokens which represent the bytes in the packets passing through the policer. The bucket drains at a given rate (the policed rate) and if the token (byte) arrival rate exceeds the drain rate then the bucket will fill. The bucket has a maximum depth, defined by a maximum burst threshold. If tokens for a packet arrive in the bucket when the current burst level of tokens is below the maximum burst threshold then the packet is considered to be conforming and all of its tokens are accepted into the bucket. If a packet's tokens arrive when the current burst level has exceeded the maximum burst threshold then none its tokens are accepted into the bucket and the packet is considered to be non-conforming (in the representation, these tokens over-flow into a waste bin).

Table 11 shows an example of the two possibilities.

Table 11: Burst Levels

Maximum burst threshold = 2000 tokens (bytes) Policed rate 2 Mbps = 250000 bytes/sec (250 tokens/ms)				
Arrival Time	Packet Size	Current Burst Level	Conforming Packet	New Burst Level
T_0	1024	1500	Yes	$1500 + 1024 = 2524$
$T_0 + 1ms$	128	$2524 - 250 = 2274$	No	2274

When the first packet arrives the current burst level is below the maximum burst threshold so it is conforming, however, when the second packet arrives the current burst level is above the maximum burst threshold so it is non-conforming.

An important aspect of the implementation of hierarchical policing is the ability of a policer bucket to have multiple burst thresholds. The tokens for each arriving packet are only compared against a single threshold relating to the characteristics of packet. These burst thresholds allow specific granular QoS control.

Policer Buckets

A policer uses up to 3 buckets depending on its configuration. A PIR bucket to control the traffic rate which is always used though its rate could be max, there can be an optional CIR bucket if a CIR rate is defined for dynamically profiling (in-profile/out-of-profile) packets, finally there may be a fair information rate (FIR) bucket used to maintain traffic fairness in a hierarchical policing scenario when multiple child policers are configured at the same parent priority level.

The PIR bucket is drained at the PIR rate and has two burst thresholds, one for high burst priority traffic (defined by the maximum burst size (MBS)) and a second for low burst priority traffic (defined by the MBS minus high-prio-only), see Figure 245. The traffic burst priority is determined at ingress by the configured priority of either high or low, and at the egress by the profile state of the packets (in-profile=high, out-of-profile=low). Note that by default all FCs are low burst priority. If a packet conforms at the PIR bucket (its tokens enter the bucket) then the packet is forwarded, otherwise the packet is discarded. Discarding logically results in the packet's tokens not being placed into the CIR, FIR or parent policer buckets.

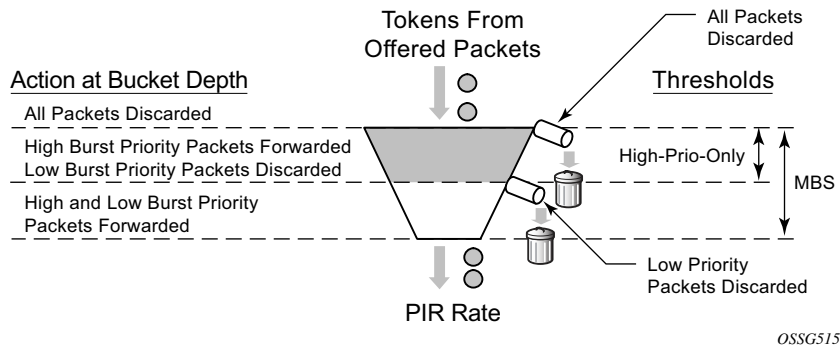


Figure 245: Peak Information Rate (PIR) Bucket

The CIR bucket is drained at the CIR rate and has one configurable burst threshold (defined by the committed burst size (CBS)). At the ingress, if the bucket level is below this threshold traffic is determined to be in-profile so the only action of the CIR bucket is to set the state of dynamically profiled packets to be either in-profile or out-of-profile. At the egress, re-profiling only affects Dot1P and DEI (Layer 2) egress marking (if the frame is double tagged, only the outer VLAN tag is remarked).

The CBS threshold is used when operating in color-blind mode, the profile of incoming packets is undefined and dynamically set based on the current burst level in the CIR bucket compared to the CBS threshold. It is also possible to operate (simultaneously) in color-aware mode, where the classification of incoming packets is used to explicitly determine whether a packet is in-profile or out-of-profile. For color-aware mode, the CIR bucket does not change the packet profile state.

In order to ensure that the overall amount of in-profile traffic takes into account both the explicit and dynamic in-profile packets, tokens from the explicit in-profile packets are allowed to fill the bucket above the CBS threshold. By doing this, dynamically profiled packets are only marked as in-profile after the token level representing dynamically in-profile and explicit in-profile packets have fallen below the CBS threshold (as the bucket drains). Note that explicitly marked out-of-profile packets remain out-of-profile, so the bottom of the bucket can be considered to be an implicit burst threshold for these packets. This is shown in [Figure 246](#).

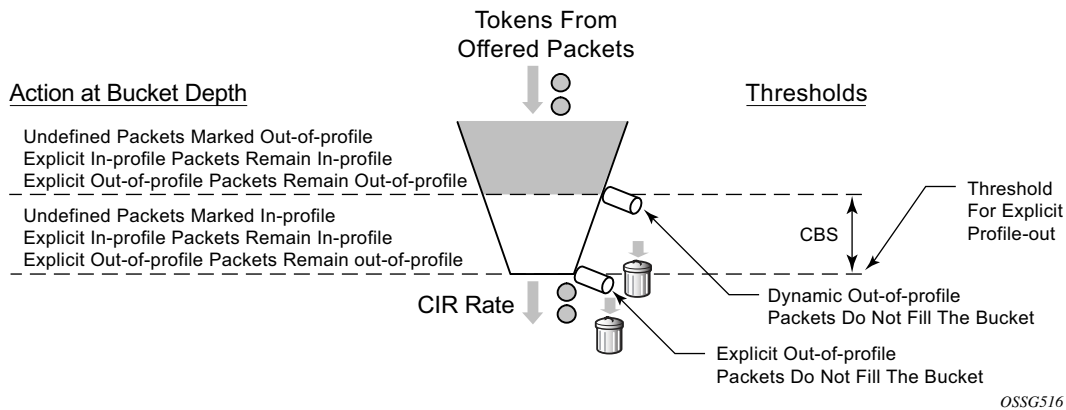
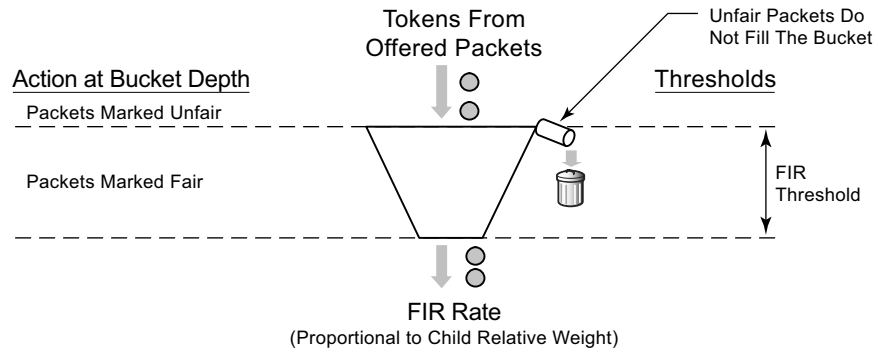


Figure 246: Committed Information Rate (CIR) Bucket

As the depths of the PIR and CIR buckets (MBS and CBS, respectively) are configured independently it is possible to have, for example, the CBS to be larger than the MBS (which is not possible for a queue). This could result in traffic being discarded because it is non-conforming at the PIR bucket but would have been conforming at the CIR bucket. Conversely, if the CBS is smaller than the MBS and the PIR=CIR traffic can be forwarded as out-of-profile, which would not be the case with a queue.

The FIR bucket is controlled by the system and is only used in hierarchical policing scenarios to determine a child's fair access to the available capacity at a parent priority level relative to other children at the same level. This bucket is only used when there is more than one child policer assigned to a given parent policer priority level. The drain rate of the FIR bucket is dynamically set proportionally to the weight configured for the child. This is shown in [Figure 247](#).



OSSG517

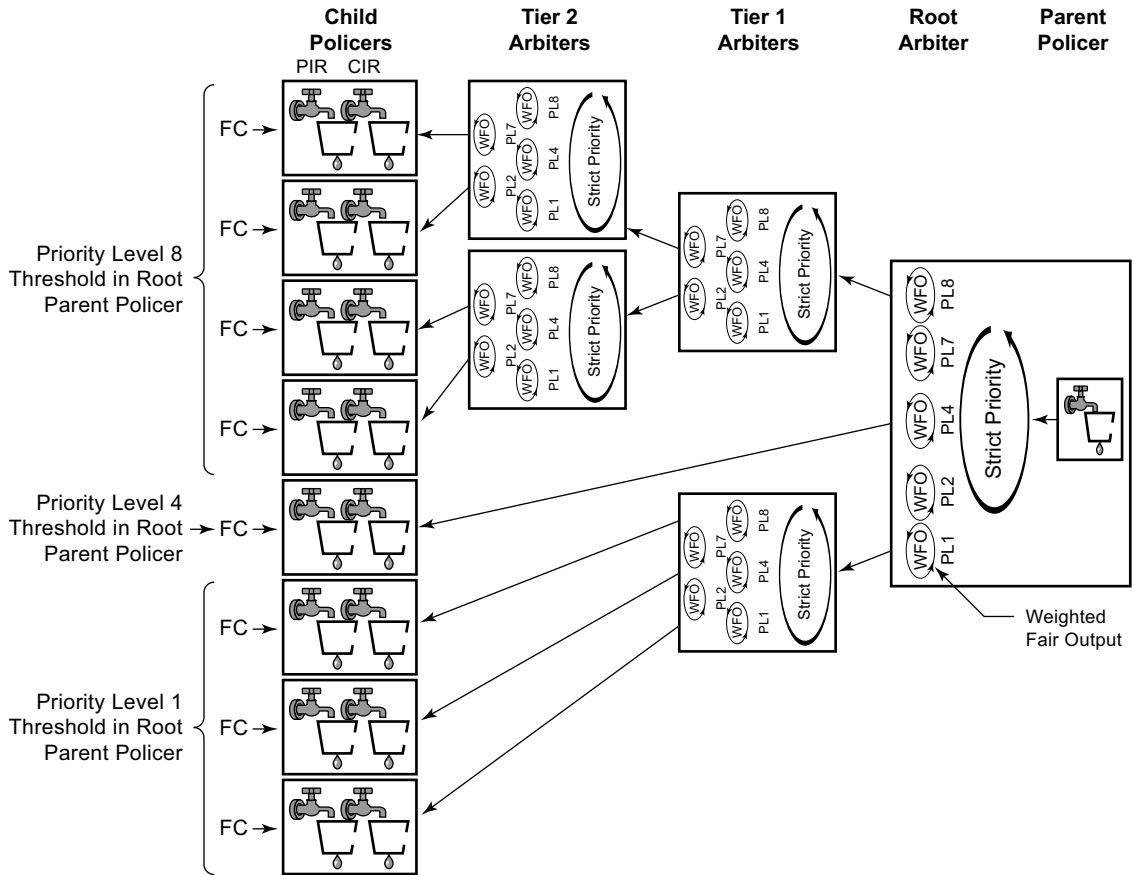
Figure 247: Fair Information Rate (FIR) Bucket

Hierarchical Policing

Policers can be used standalone or with a parent policer to provide hierarchical policing. Up to four stages can be configured in the hierarchy: the child policer, tier 1 and 2 intermediate arbiters, and a root arbiter (which is associated with the parent policer). The arbiters are logical entities that distribute bandwidth at a particular tier to their children in a priority level order, see [Figure 248](#).

This may result in the drain rates for the child policer buckets being modified, so each child policer PIR and CIR bucket has an administrative rate value (what it is configured to) and an operational rate value (the current operating rate) based on the bandwidth distribution by the parent arbiters.

Each stage in the hierarchy connects to its parent at a priority level and a weight. There are eight available priorities which are serviced in a strict order (8 to 1, highest to lowest, respectively). The weight is used to define relative fairness when multiple children are configured in the same priority level. Note that the child access to parent policer burst capacity is governed by the level at which the child ultimately connects into the root arbiter, not by its connection level at any intermediate arbiters.



OSSG518

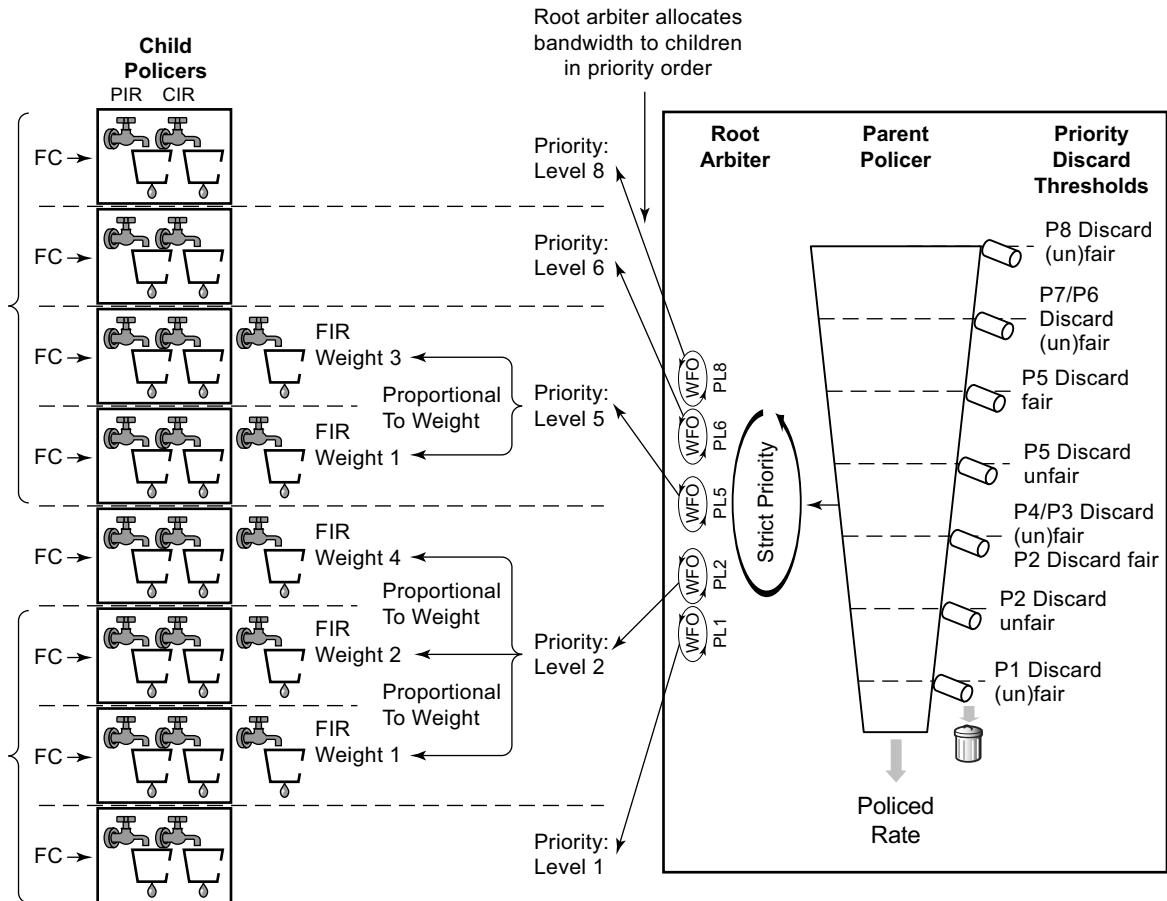
Figure 248: Policer and Arbiter Hierarchy

The final configuration aspect to consider is the parent policer, specifically its multiple thresholds and how they relate to the child policers. See [Figure 249](#).

There are 8 priority levels at the parent policer, each having an associated discard-fair and discard-unfair threshold.

The discard-fair threshold is the upper burst limit for all tokens (consequently, all packets) at the given priority, all traffic at a given priority level is discarded when its tokens arrive with this threshold being exceeded. The discard-fair thresholds enable prioritization at the parent policer by having the burst capacity for each priority threshold be larger (or equal) to those of lower priorities. For example, referring to [Figure 249](#), the priority 6 (P6) discard-fair threshold is larger than the priority 5 (P5) discard-fair threshold with the result that even if the priority 5 and below traffic is overloading the parent policer, the priority 6 traffic has burst capacity available in order to allow some of its packets to conform and get forwarded through the parent policer.

Note that if a packet is discarded at the parent policer, the discard needs to be reflected in the associated child policer, this is achieved by also logically removing the related tokens from the child policer buckets.



OSSG519

Figure 249: Parent Policer and Root Arbiter

Each priority also has a discard-unfair threshold which discards only unfair traffic of that priority, remembering that fair and unfair are determined by the FIR bucket based on the relative weights of the children.

By default, if there are no children configured at a given priority level then both its discard-fair and discard-unfair thresholds are set to zero bytes above the previous priority's discard-fair threshold.

Hierarchical Policing

If there is only a single child at a priority level, the discard-fair will be greater than the previous priority's discard-fair value (by an amount equal to the maximum of the min-thresh-separation and the mbs-contribution, see below) but the discard-unfair will be the same as the previous priority's discard-fair threshold (there is no need for a fairness function when there is only a single child at that priority).

If there is more than one child at a priority level, the discard-unfair threshold will be greater than the previous priority's discard-fair threshold by min-thresh-separation (see below) and the discard-fair threshold will be adjusted upwards by an amount equal to mbs-contribution minus min-thresh-separation.

The result can be summarized as follows:


- With no children at a priority level, the discard fair and unfair thresholds match the values of the previous priority.
- If there are at least two children at a priority level, the discard-unfair burst capacity equals min-thresh-separation.
- The burst capacity for a given priority level with at least one child equals the mbs-contribution, unless this is less than min-thresh-separation in which case the min-thresh-separation is used.

The burst tolerance for each threshold is its own burst capacity plus the sum of the burst capacities of all lower thresholds. Referring to [Figure 249](#), the total burst capacity for priority 6 is the sum of the burst capacities for priorities 1 to 6. Note that the burst for a given FC is normally controlled by the burst allowed at the child PIR threshold, not by the parent policer.

As the burst capacity at the parent policer for a given priority level can change when adding or removing children at lower priority levels, a parameter (fixed) is available per priority threshold which causes the discard-fair and discard-unfair thresholds to be non-zero and so greater than the previous priority's thresholds, calculated as above, even when there are no children at that priority level. An exception to this is when the mbs-contribution is set to zero with the fixed parameter configured, in which case both the discard unfair and fair for that priority level are set to zero bytes above the previous level's thresholds (which results in the corresponding traffic being dropped).

A specific configuration and associated show output is included below to highlight the different threshold options described above.

The QoS example shown in [Figure 250](#) is used to describe the configuration of CFHP.



OSSG520

Figure 250: Configuration Example

Five classes of services are accepted, each with a specific CIR and PIR. The data classes, bronze, silver and gold (L2/AF/L1), have a relative weighting of 50/25/25 at priority Level 2 of an intermediate arbiter which is constrained to 60Mbps. At the parent policer, the real time traffic (EF) is defined at level 5, with the data classes at Level 3 and a best effort class (BE) at Level 1. The overall traffic is constrained to 100Mbps at the parent policer. Only unicast traffic is policed in this example.

This example focuses on ingress policing, however, the configuration of policers, arbiters and the parent policer at the egress is almost identical to that at the ingress, the only difference being the particular statistics that can be collected.

There is a difference between ingress and egress policing in terms of how the ingress traffic accesses the switch fabric and the egress traffic access the port after it has been policed. In both cases, unicast access is enabled through a set of policer-output-queues, which are shared-queues at the ingress and queue-groups at the egress (at the egress, user defined queue-groups can be used). It is also possible to use a single service queue to access the egress port. Ingress multipoint traffic accesses the switch fabric using the Ingress Multicast Path Management (IMPM) queues.

Hierarchical Policing

This is shown in [Figure 251](#) on an IOM3-XP (other line cards have the same logic).




Figure 251: Post Policing Queues

The differences between the ingress and egress policing configuration will be high-lighted in the associated sections.

Configuration

To achieve the QoS shown in [Figure 250](#), configure a SAP-ingress QoS policy to define the child policers and a policer-control-policy to define the intermediate arbiter and the root arbiter/parent policer. As this example is for ingress, the unicast traffic will pass through a set of shared queues called policer-output-queues, which could be modified if required.

Policers

Policers control the CIR and PIR rates for each of the traffic classes and are defined in a SAP-ingress QoS policy. The focus here are parameters related to policing.

The configuration of a child (or standalone) policer is similar to that of a queue.

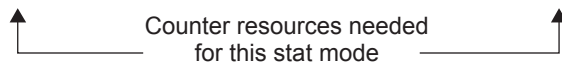
```
config>qos>sap-ingress# policer policer-id [create]
  description "description-string"
  adaptation-rule [pir {max | min | closest}] [cir {max | min | closest}]
  stat-mode {no-stats|minimal|offered-profile-no-cir|
             offered-priority-no-cir|offered-profile-cir|offered-priority-cir|
             offered-total-cir|offered-limited-profile-cir}
  rate {max | kilobits-per-second} [cir {max | kilobits-per-second}]
  percent-rate pir-percent [cir cir-percent]
  mbs size [bytes | kilobytes]
  cbs size [bytes | kilobytes]
  high-prio-only [default | percent-of-mbs]
  parent {root | arbiter-name} [level level] [weight weight-within-level]
  packet-byte-offset {add bytes | subtract bytes}
```

Parameters:

- **description** — This configures a text string, up to 80 characters, which can be used to describe the use of the policy.
- **adaptation-rule** — The hardware supports distinct values for the rates. This parameter tells the system how the rate configured should be mapped onto the possible hardware values. min results in the next higher hardware value being used, max results in the next lower hardware value being used and closest results in the closest available hardware value being used. As can be seen, it is possible to set the adaptation-rule independently for the CIR and PIR.
Default: closest
- **stat-mode** — This defines the traffic statistics collected by the policer, summarized in [Table 12](#).

Table 12: Policer stat-mode

stat-mode	Ingress	Egress
no-stats	0 Neither policer nor parent arbiter account are required.	0 Neither policer nor parent arbiter accounting are required.
minimal (default)	1 Basic policer accounting (default).	1 Basic policer accounting (default).
offered-profile-no-cir	2 All ingress packets are either in-profile or out-of-profile.	2 Accounting for egress offered profile is required. No visibility for CIR profile state output.
offered-priority-no-cir	2 Ingress packet burst priority accounting is the primary requirement.	N/A
offered-limited-profile-cir	3 Ingress color-aware profiling is in use but packets are not being classified as in-profile.	N/A
offered-profile-cir	4 Ingress color-aware profiling is in use and packets are undefined or classified as out-of-profile or in-profile.	4 Egress profile reclassification is performed.
offered-priority-cir	4 Ingress policer is used in color-blind mode and ingress packet priority and CIR state output accounting is needed.	N/A
offered-total-cir	2 Ingress priority and ingress profile accounting is not needed. CIR profiling is in use.	2 Offered profile visibility is not required (such as, all offered packets have the same profile) and CIR profiling is in use.



- rate and cir — The rate defines the PIR and the cir defines the CIR, both are in Kbps. The parameters rate and percent-rate are mutually exclusive and will overwrite each other when configured in the same policy.
Range: PIR=1 to 20,000,000 Kbps or max ; CIR=0 to 20,000,000 Kbps or max
Default: rate(PIR)=max ; cir=0
- percent-rate and cir — The percent-rate defines the PIR and the cir defines the CIR with their values being a percentage of the maximum policer rate of 20Gbps. The parameters rate and percent-rate are mutually exclusive and will overwrite each other when

configured in the same policy.

Range: pir-percent = [0.01..100.00]; cir-percent = [0.00..100.00]

Default: pir-percent = 100; cir-percent = 0.00

- mbs and cbs — The mbs defines the MBS for the PIR bucket and the cbs defines the CBS for the CIR bucket, both can be configured in bytes or kilobytes.
Note that the PIR MBS applies to high burst priority packets (these are packets whose classification match criteria is configured with priority high at the ingress and are in-profile packets at the egress).
Range: mbs=0 to 4194304 bytes; cbs=0 to 4194304 bytes
Note: mbs=0 prevents any traffic from being forwarded.
Default: mbs=10ms of traffic or 64KB if PIR=max; cbs=10ms of traffic or 64KB if CIR=max
- high-prio-only — This defines a second burst threshold within the PIR bucket to give a maximum burst size for low burst priority packets (these are packets whose classification match criteria is configured with priority low at the ingress and are out-of-profile packets at the egress). It is configured as a percentage of the MBS.
Default: 10%
- parent — This parameter is used when hierarchical policing is being performed and points to the parent arbiter (which could be the root arbiter or an intermediate arbiter), giving the level to which this policer connects to its parent arbiter and its relative weight compared to other children at the same level. Note that for a child policer to be associated with a parent, its stat-mode cannot be no-stats.
Range: level=1 to 8; weight=1 to 100
Default: level=1; weight=1
- packet-byte-offset — This changes the packet size used for accounting purposes, both in terms of the CIR and PIR rates and what is reported in the statistics. The change can either add or subtract a number of bytes. For example:
 - To have the policer work on Layer 2 frame size including inter-frame gap and preamble, add 20 bytes.
 - To have the policer work on IP packet size instead of the default layer 2 frame size, subtract the encapsulation overhead:
14 bytes L2 + 4bytes VLAN ID + 4 bytes FCS = 22 bytes
Range: add-bytes=0 to 31; sub-bytes=1 to 32
Default: add-bytes=0; sub-bytes=0

A FC must be assigned to the policer in order for the policer to be instantiated (allocating a hardware policer).

By default, any unicast traffic assigned to the FC at the ingress will be processed by the policer, non-unicast traffic would continue to use the multipoint queue. At the egress all traffic assigned to the FC is processed by the policer (as there is no distinction between unicast and non-unicast traffic at the egress).

Policers

If required, non-unicast traffic can be policed in IES/VPRN and VPLS services at the ingress (note: all Epipe traffic is treated as unicast). Within an IES/VPRN service, multicast traffic can be assigned to a specific ingress policer on a PIM enabled IP interface. When the service is VPLS, broadcast, unknown unicast and multicast traffic can be individually assigned to ingress policers. In each of these cases, the policers used could be separate from the unicast policer, resulting in the instantiation of additional hardware policers, or a single policer could be used for multiple traffic types (this differs from the queuing implementation where separate queue types are used for unicast and non-unicast traffic).

```
config>qos>sap-ingress>fc#  
  broadcast-policer <policer-id>  
  unknown-policer <policer-id>  
  multicast-policer <policer-id>
```

As mentioned above, the ingress policed unicast traffic passes through a set of shared-queues (policer-output-queues) to access the switch fabric with the multipoint traffic using the IMPM queues.

When policers are required at the egress, a SAP-egress policy is used. The configuration of the policers is almost identical to that used in the SAP-ingress policy, the only difference being the available stat-modes (as shown above).

At the egress, the policed traffic can also be directed to a specific queue-group (instead of the default policer-output-queues) and to a specific queue within that queue-group, as follows:

```
config>qos>sap-egress>fc# policer <policer-id> [group <queue-group-name> [queue  
<queueid>]]
```

It is also possible to direct the egress policed traffic to a single service queue if specific egress queuing is required, as follows:

```
config>qos>sap-egress>fc# policer <policer-id> queue <queue-id>
```

Multiple egress policers in a SAP-egress policy can use the same local queue and other forwarding classes can directly use the same local queue that is being used by policers.

Parent Policer and Arbiters

The parent policer and its associated root arbiter, together with the tier 1 and 2 arbiters, are configured within a policer-control-policy.

```
config>qos# policer-control-policy policy-name [create]
  description description-string
  root
    max-rate {kilobits-per-second | max}
    priority-mbs-thresholds
      min-thresh-separation size [bytes|kilobytes]
      priority level
      mbs-contribution size [bytes|kilobytes] [fixed]
  tier 1
    arbiter arbiter-name [create]
      description description-string
      rate {kilobits-per-second|max}
      parent root [level priority-level] [weight weight-within-level]
  tier 2
    arbiter arbiter-name [create]
      description description-string
      rate {kilobits-per-second | max}
      parent {root|arbiter-name} [level priority-level] [weight weight-within-level]
```

Parameters:

- **description** — This configures a text string, up to 80 characters, which can be used to describe the use of the policy.
- **root** — This section defines the configuration of the parent policer and the root arbiter.
 - **max-rate** — This defines the policed rate of the parent policer, the rate at which the bucket is drained. It is defined in Kbps with an option to use max, in which case the maximum possible rate is used.
Range: 1 to 20,000,000Kbps or max
Default: max
 - **priority-mbs-thresholds**
This section defines the thresholds used for the 8 priorities available in the parent policer.
 - **min-thresh-separation** — This defines the minimum separation between any two active thresholds in the parent policer in units of bytes or kilobytes.
It should be set to a value greater than the maximum packet size used for traffic passing through the policer. This ensures that a single packet arriving in the parent policer will not cause the depth of tokens to cross two burst thresholds, if this did happen it would result in the prioritization failing as a given priority level could be starved of burst capacity by a lower priority traffic.

Parent Policer and Arbiters

This parameter is also used as the burst capacity for each priority level's unfair packets.

Range: 0 to 4194304 bytes

Default: 1536 bytes

- mbs-contribution — This is normally used to define the amount of packet burst capacity required at the parent policer for a particular priority level with at least one child, keeping in mind that the total capacity is the sum of this plus that of all lower thresholds. The actual burst capacity used depends also on the setting of min-thresh-separation, as described earlier. This permits the tuning of the burst capacity at the parent for any children at a given priority level. A conservative setting would ensure that the burst at the parent policer for a given priority is the sum of the bursts of all children at that priority. Less conservative settings could use a lower value and assume some level of oversubscription.

The use of the fixed parameter causes both the fair and unfair discard thresholds to be non-zero even when there are no children assigned to this priority level (unless the mbs-contribution is set to zero).

Range: 0 to 4194304

Default: 8192 bytes

The relationship between these two parameters is shown in [Figure 252](#).

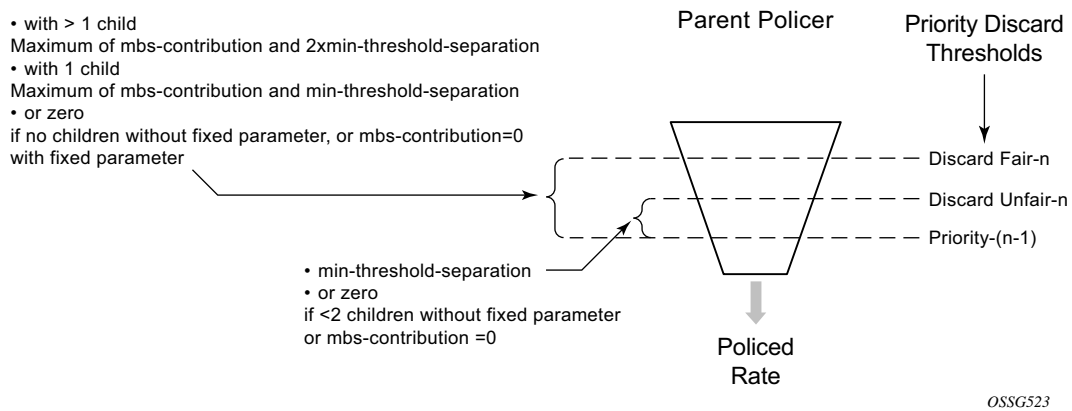


Figure 252: Parent Policer Thresholds

- tier

This section defines the configuration of any intermediate tier 1 or 2 arbiters.

→ arbiter

This specifies the name of the arbiter.

- description — This configures a text string, up to 80 characters, which can be used to describe the use of the policy.
- rate — This defines the rate of the arbiter, it is the maximum rate at which the arbiter will distribute burst capacity to its children. It is defined in Kbps with an option to use max, in which case the maximum possible rate is used.
Range: 1 to 20,000,000Kbps or max
Default: max
- parent — This parameter is used when hierarchical policing is being performed and points to the parent arbiter (which could be the root arbiter or a tier 1 arbiter), giving the level to which this arbiter connects to its parent arbiter and its relative weight compared to other children at the same level.
Range: level=1 to 8; weight=1 to 100
Default: level=1; weight=1

Access to Switch Fabric and Egress Port

After the traffic has been processed by the policers it must pass through a set of queues in order to access the switch fabric at the ingress or the port at the egress.

For the ingress unicast traffic, there is a set of shared-queues (one queue per FC for each possible switch fabric destination) called policer output queues. Note that only their queue characteristics can be configured, the FC to queue mapping is fixed. Also, the PIR/CIR rates only affect the packet scheduling, they do not alter the packet profile state. The details of shared-queues are beyond the scope of this note.

```
config>qos# shared-queue "policer-output-queues"
  description description-string
  fc <fc-name> [create]
    broadcast-queue <queue-id>
    multicast-queue <queue-id>
    queue <queue-id>
    unknown-queue <queue-id>
  queue queue-id [queue-type] [multipoint] [create]
    cbs percent
    mbs percent
    high-prio-only percent
    pool pool-name
    rate percent [cir percent]
```

Multipoint traffic uses the IMPM queues to access the switch fabric. For the egress to the port, either a queue-group or a single service queue is used. There is a default queue-group called policer-output-queues or a user configured queue-group can also be used.

As mentioned above, when a policer is assigned to a specific queue-group (default or user defined) it is optionally possible to configure explicitly the queue to be used. Within the queue-group it is also possible to redirect a FC for policed traffic to a specific queue, using the FC parameter. The preference of the FC to queue mapping is (in order, highest to lowest):

1. Explicitly configured in SAP-egress FC definition
2. Mapped using FC parameter within queue-group definition
3. Default is to use queue 1

```
config>qos>qgrps>egr# queue-group queue-group-name [create]
  description description-string
  queue queue-id [queue-type] [create]
    adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
    burst-limit size [bytes|kilobytes]
    cbs size-in-kbytes
    high-prio-only percent
    mbs size [bytes|kilobytes]
    parent scheduler-name [weight weight] [level level] [cir-weight cir-weight]
      [cir-level cir-level]
    percent-rate pir-percent [cir cir-percent]
```

```

pool pool-name
port-parent [weight weight] [level level] [cir-weight cir-weight]
             [cir-level cir-level]
rate pir-rate [cir cir-rate]
xp-specific
  wred-queue [policy slope-policy-name]
fc fc-name [create]
  queue queue-id

```

The default policer-output-queues queue-group consists of two queues; queue 1 being best-effort and queue 2 being expedite. The lowest four FCs (BE, L2, AF, L1) are assigned to queue 1 and the highest four queues (H2, EF, H1, NC) are assigned to queue 2. It may be important to change the queue 2 definition in the queue-group to have CIR=PIR when there are other best-effort queues using a non-zero CIR on the same egress port. This ensures that the policed traffic using queue 2 will be scheduled before any other best-effort within CIR traffic. It also results in the queue CBS being non-zero, allowing the queue 2 traffic access to reserved buffer space.

```

A:PE-1>config>qos# queue-group-templates egress queue-group "policer-output-queues"
A:PE-1>cfg>qos>qgrps>egr>qgrp# info detail

```

```

-----
description "Default egress policer output queues."
queue 1 best-effort create
  no parent
  no port-parent
  adaptation-rule pir closest cir closest
  rate max cir 0
  cbs default
  mbs default
  high-prio-only default
  no pool
  xp-specific
    no wred-queue
  exit
  no burst-limit
exit
queue 2 expedite create
  no parent
  no port-parent
  adaptation-rule pir closest cir closest
  rate max cir 0
  cbs default
  mbs default
  high-prio-only default
  no pool
  xp-specific
    no wred-queue
  exit
  no burst-limit
exit
fc af create
  queue 1
exit
fc be create
  queue 1
exit

```

Access to Switch Fabric and Egress Port

```
fc ef create
    queue 2
exit
fc h1 create
    queue 2
exit
fc h2 create
    queue 2
exit
fc l1 create
    queue 1
exit
fc l2 create
    queue 1
exit
fc nc create
    queue 2
exit
```

The remaining details of queue-groups are beyond the scope of this section.

Applying the SAP Ingress and Policer Control Policy

The SAP ingress policy and policer control policy are both applied under the associated SAP. After applying these, it is possible to override the configuration of specific policers and/or the policer control policy. This is shown below. The parameter values are the same as detailed for the policies, as above.

```
config>service><service>#
  sap sap-id [create]
    [ingress|egress]
      qos policy-id
        policer-control-policy policy-name
        policer-override
          policer policer-id [create]
            cbs size [bytes|kilobytes]
            mbs size [bytes|kilobytes]
            packet-byte-offset {add add-bytes | subtract sub-bytes}
            rate {rate | max} [cir {max | rate}
            percent-rate <pir-percent> [cir <cir-percent>]
            stat-mode stat-mode
          policer-control-override [create]
            max-rate {rate | max}
            priority-mbs-thresholds
              min-thresh-separation size [bytes | kilobytes]
              priority level
              mbs-contribution size [bytes | kilobytes]
```

The SAP ingress policy and policer control policy required for the configuration example in [Figure 250](#) is shown below.

```
#-----
echo "QoS Policy Configuration"
#-----
  qos
    policer-control-policy "cfhp-1" create
      root
        max-rate 100000
      exit
      tier 1
        arbiter "a3" create
          parent "root" level 3
          rate 60000
        exit
      exit
    exit
  sap-ingress 10 create
    queue 1 create
    exit
    queue 11 multipoint create
    exit
  policer 1 create
    stat-mode offered-total-cir
    parent "root"
    rate 100000
```

Applying the SAP Ingress and Policer Control Policy

```
        high-prio-only 0
    exit
    policer 2 create
        stat-mode offered-total-cir
        parent "a3" level 2 weight 50
        rate 60000 cir 20000
        high-prio-only 0
    exit
    policer 3 create
        stat-mode offered-total-cir
        parent "a3" level 2 weight 25
        rate 60000 cir 20000
        high-prio-only 0
    exit
    policer 4 create
        stat-mode offered-total-cir
        parent "a3" level 2 weight 25
        rate 60000 cir 20000
        high-prio-only 0
    exit
    policer 5 create
        stat-mode offered-total-cir
        parent "root" level 5
        rate 10000 cir 10000
        high-prio-only 0
    exit
    fc "af" create
        policer 3
    exit
    fc "be" create
        policer 1
    exit
    fc "ef" create
        policer 5
    exit
    fc "l1" create
        policer 4
    exit
    fc "l2" create
        policer 2
    exit
    dot1p 1 fc "be"
    dot1p 2 fc "l2"
    dot1p 3 fc "af"
    dot1p 4 fc "l1"
    dot1p 5 fc "ef"
exit
```

Traffic is classified based on dot1p values, each of which is assigned to an individual FC which in turn is assigned to a policer. The policer rates are configured as required for the example with an appropriate stat-mode. Default values are used for the policer burst thresholds. As all FCs are low burst priority by default, the high-prio-only has been set to zero in order to allow the traffic to use all of the MBS available at the PIR bucket.

Policers 2, 3 and 4 are parented to the arbiter “a3” with the required weights and at a single level (Level 2). In this example it does not matter which level of “a3” is used to parent these policers,

the important aspect is the level at which “a3” is parented to the root. Consequently, these policers use the Level 3 parent policer thresholds (not the level they are parented on a “a3” not Level 2). Arbiter “a3” has a rate of 60Mbps so that its children cannot exceed this rate (except up to the burst tolerances).

Policers 1 and 5 are directly parented to the root arbiter, together with tier 1 arbiter “a3”.

The total capacity for the 5 traffic streams is constrained to 100Mbps by the parent policer, again with the default burst tolerances at the root arbiter.

The SAP-ingress and policer-control-policies are applied to a SAP within an Epipe.

```
#-----  
echo "Service Configuration"  
#-----  
service  
  epipe 1 customer 1 create  
    sap 1/1/3:1 create  
      ingress  
        policer-control-policy "cfhp-1"  
        qos 10  
      exit  
    exit  
  sap 1/1/4:1 create  
  exit  
no shutdown  
exit  
exit
```

Applying the SAP Ingress and Policer Control Policy

The following configuration is used to highlight the relative thresholds in the parent policer when a priority level has 0, 1 or 2 associated children, both with and without using the fixed parameter.

```
-----
echo "QoS Policy Configuration"
#-----
  qos
    policer-control-policy "cfhp-2" create
      root
        max-rate 100000
        priority-mbs-thresholds
          min-thresh-separation 256 bytes
          priority 1
            mbs-contribution 1 kilobytes
          exit
          priority 2
            mbs-contribution 1 kilobytes
          exit
          priority 3
            mbs-contribution 1 kilobytes
          exit
          priority 4
            mbs-contribution 1 kilobytes fixed
          exit
          priority 5
            mbs-contribution 1 kilobytes fixed
          exit
          priority 6
            mbs-contribution 1 kilobytes fixed
          exit
        exit
      exit
    exit
  sap-ingress 20 create
    queue 1 create
    exit
    queue 11 multipoint create
    exit
    policer 1 create
      parent "root" level 2
    exit
    policer 2 create
      parent "root" level 3
    exit
    policer 3 create
      parent "root" level 3
    exit
    policer 4 create
      parent "root" level 5
    exit
    policer 5 create
      parent "root" level 6
    exit
    policer 6 create
      parent "root" level 6
    exit
  fc "af" create
```

```

        policer 3
    exit
    fc "be" create
        policer 1
    exit
    fc "ef" create
        policer 6
    exit
    fc "h2" create
        policer 5
    exit
    fc "l1" create
        policer 4
    exit
    fc "l2" create
        policer 2
    exit
    exit
#-----
echo "Service Configuration"
#-----
    service
        epipe 2 customer 1 create
            sap 1/1/3:2 create
                ingress
                    policer-control-policy "cfhp-2"
                    qos 20
                exit
            exit
        sap 1/1/4:2 create
        exit
        no shutdown
    exit

```

A policer-control-policy can also be applied under a multi-service site (MSS) so that the hierarchical policing applies to traffic on multiple SAPs, potentially from different services. The MSS can only be assigned to a port, which could be a LAG, but it is not possible to assign an MSS to a card. When MSS are used, policer overrides are not supported.

```

config>service><service>#
    service
        customer customer-id [create]
            multi-service-site customer-site-name [create]
                assignment port port-id
                egress
                    policer-control-policy name
                ingress
                    policer-control-policy name
        service-type
            sap sap-id
            multi-service-site customer-site-name
            ingress
                qos policy-id
            egress
                qos policy-id

```

Show Output

After configuring the example as described in the previous section, steady state traffic was sent through the Epipe to overload each of the policers and the show output below was collected. This output focuses on the policer and arbiter details.

The following shows the policers on the SAP and their current state.

```
A:PE-1# show qos policer sap 1/1/3:1
=====
Policer Information (Summary), Slot 1
=====
```

Name Direction	FC-Maps	MBS CBS	HP-Only Depth	A.PIR O.PIR	A.CIR O.CIR	O.FIR
1->1/1/3:1->1 Ingress	be	124 KB 0 KB	0 KB 82	100000 30000	0 0	30000
1->1/1/3:1->2 Ingress	12	76 KB 25 KB	0 KB 77846	60000 30000	20000 20000	30000
1->1/1/3:1->3 Ingress	af	76 KB 25 KB	0 KB 77824	60000 15000	20000 15000	15000
1->1/1/3:1->4 Ingress	11	76 KB 25 KB	0 KB 77868	60000 15000	20000 15000	15000
1->1/1/3:1->5 Ingress	ef	12800 B 12800 B	0 KB 12834	10000 10000	10000 10000	10000

```
=====
A:PE-1#
```

The output above shows the configured values for the policers, e.g. PIR and CIR, together with their operational (current) state, such as PIR, CIR and FIR. The depth of each of the PIR buckets is also shown.

The detailed state of each policer can be seen by adding the parameter detail. The following is the output for policer 3.

```
A:PE-1# show qos policer sap 1/1/3:1 ingress detail
...
=====
Policer Info (1->1/1/3:1->3), Slot 1
=====
```

Policer Name	: 1->1/1/3:1->3	Fwding Plane	: 1
Direction	: Ingress		
FC-Map	: af		
Depth PIR	: 77842 Bytes	Depth CIR	: 25618 Bytes
Depth FIR	: 77842 Bytes		
MBS	: 76 KB	CBS	: 25 KB
Hi Prio Only	: 0 KB	Pkt Byte Offset	: 0

Class Fair Hierarchical Policing for SAPs

```
Admin PIR      : 60000 Kbps      Admin CIR      : 20000 Kbps
Oper PIR      : 15000 Kbps      Oper CIR      : 15000 Kbps
Oper FIR      : 15000 Kbps
Stat Mode     : offered-total-cir
PIR Adaption  : closest        CIR Adaption   : closest
Parent Arbiter Name: a3
```

Arbiter Member Information

```
Offered Rate   : 45800 Kbps
Level          : 2              Weight         : 25
Parent PIR     : 15000 Kbps    Parent FIR     : 15000 Kbps
Consumed      : 15000 Kbps
```

```
A:PE-1#
```

Notice that the above output shows the depth of the PIR, CIR and FIR buckets together with their operational rates. This can be used to explain the operation of the policers in this example and is discussed later in this section.

The stat-mode of offered-total-cir configured on policer 3 results in these statistics being collected.

```
A:PE-1# show service id 1 sap 1/1/3:1 stats
```

```
=====
...
-----
Sap per Policer stats
-----
```

	Packets	Octets
Ingress Policer 1 (Stats mode: offered-total-cir)		
Off. All	: 2690893	172217152
Dro. InProf	: 0	0
Dro. OutProf	: 967465	61917760
For. InProf	: 0	0
For. OutProf	: 1723428	110299392
Ingress Policer 2 (Stats mode: offered-total-cir)		
Off. All	: 2690988	172223232
Dro. InProf	: 0	0
Dro. OutProf	: 909492	58207488
For. InProf	: 1178507	75424448
For. OutProf	: 602989	38591296

```
...
```

Show Output

The following output is included for reference and shows the statistics which are collected for each of the ingress and egress stat-modes.

```
PE-1# show service id 2 sap 1/1/1:2 stats
...
-----
Sap per Policer stats
-----
                Packets                Octets

Ingress Policer 1 (Stats mode: no-stats)

Ingress Policer 2 (Stats mode: minimal)
Off. All          : 0                    0
For. All          : 0                    0
Dro. All          : 0                    0

Ingress Policer 3 (Stats mode: offered-profile-no-cir)
Off. InProf       : 0                    0
Off. OutProf      : 0                    0
For. InProf       : 0                    0
For. OutProf      : 0                    0
Dro. InProf       : 0                    0
Dro. OutProf      : 0                    0

Ingress Policer 4 (Stats mode: offered-priority-no-cir)
Off. HiPrio       : 0                    0
Off. LowPrio      : 0                    0
For. HiPrio       : 0                    0
For. LoPrio       : 0                    0
Dro. HiPrio       : 0                    0
Dro. LowPrio      : 0                    0

Ingress Policer 5 (Stats mode: offered-profile-cir)
Off. InProf       : 0                    0
Off. OutProf      : 0                    0
Off. Uncolor      : 0                    0
For. InProf       : 0                    0
For. OutProf      : 0                    0
Dro. InProf       : 0                    0
Dro. OutProf      : 0                    0

Ingress Policer 6 (Stats mode: offered-priority-cir)
Off. HiPrio       : 0                    0
Off. LowPrio      : 0                    0
For. InProf       : 0                    0
For. OutProf      : 0                    0
Dro. InProf       : 0                    0
Dro. OutProf      : 0                    0

Ingress Policer 7 (Stats mode: offered-total-cir)
Off. All          : 0                    0
For. InProf       : 0                    0
For. OutProf      : 0                    0
Dro. InProf       : 0                    0
Dro. OutProf      : 0                    0

Ingress Policer 8 (Stats mode: offered-limited-profile-cir)
```


Class Fair Hierarchical Policing for SAPs

Off. OutProf	: 0	0
Off. Uncolor	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0
Dro. InProf	: 0	0
Dro. OutProf	: 0	0

Egress Policer 1 (Stats mode: no-stats)

Egress Policer 2 (Stats mode: minimal)		
Off. All	: 0	0
For. All	: 0	0
Dro. All	: 0	0

Egress Policer 3 (Stats mode: offered-profile-no-cir)

Off. InProf	: 0	0
Off. OutProf	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0
Dro. InProf	: 0	0
Dro. OutProf	: 0	0

Egress Policer 4 (Stats mode: offered-profile-cir)

Off. InProf	: 0	0
Off. OutProf	: 0	0
Off. Uncolor	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0
Dro. InProf	: 0	0
Dro. OutProf	: 0	0

Egress Policer 5 (Stats mode: offered-total-cir)

Off. All	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0
Dro. InProf	: 0	0
Dro. OutProf	: 0	0

=====

Show Output

It is possible to show the policer-control-policy details and the SAPs with which it is associated, as shown here.

```
A:PE-1# show qos policer-control-policy cfhp-1
=====
QoS Policer Control Policy
=====
Policy-Name       : cfhp-1
Description       : (Not Specified)
Min Threshold Sep : Def

-----
Priority MBS Thresholds
-----
Priority          MBS Contribution
-----
1                 none
2                 none
3                 none
4                 none
5                 none
6                 none
7                 none
8                 none

-----
Tier/Arbiter          Lvl/Wt      Rate      Parent
-----
   root                N/A        100000    None
1 a3                   3/1        60000     root

=====
A:PE-1# show qos policer-control-policy "cfhp-1" association
=====
QoS Policer Control Policy
=====
Policy-Name       : cfhp-1
Description       : (Not Specified)

-----
Associations
-----
Service-Id        : 1 (Epipe)          Customer-Id      : 1
- SAP : 1/1/3:1 (Ing)

=====
A:PE-1
```

The following command shows the policer hierarchy, including the child policers and their relationship to the intermediate arbiter (a3) and the root arbiter. It can be used to monitor the status of the child policers in the hierarchy. The output shows the assigned, offered and consumed capacity for each policer.

```
A:PE-1# show qos policer-hierarchy sap 1/1/3:1
=====
Policer Hierarchy - Sap 1/1/3:1
=====
Ingress Policer Control Policy : cfhp-1
Egress Policer Control Policy :
-----
root (Ing)
|
| slot(1)
|
|--(A) : a3 (Sap 1/1/3:1)
|
| |--(P) : Policer 1->1/1/3:1->4
| |
| | | [Level 2 Weight 25]
| | | Assigned PIR:15000      Offered:45800
| | | Consumed:15000
| |
| | | Assigned FIR:15000
| |
| |--(P) : Policer 1->1/1/3:1->3
| |
| | | [Level 2 Weight 25]
| | | Assigned PIR:15000      Offered:45800
| | | Consumed:15000
| |
| | | Assigned FIR:15000
| |
| |--(P) : Policer 1->1/1/3:1->2
| |
| | | [Level 2 Weight 50]
| | | Assigned PIR:30000      Offered:45800
| | | Consumed:30000
| |
| | | Assigned FIR:30000
| |
|--(P) : Policer 1->1/1/3:1->5
|
| | [Level 5 Weight 1]
| | Assigned PIR:10000      Offered:10000
| | Consumed:10000
|
| | Assigned FIR:10000
|
|--(P) : Policer 1->1/1/3:1->1
|
| | [Level 1 Weight 1]
| | Assigned PIR:30000      Offered:45800
| | Consumed:30000
|
| | Assigned FIR:30000
```

Show Output

```
root (Egr)
|
No Active Members Found on slot 1
=====
A:PE-1#
```

The complete information about the policer hierarchy can be seen by adding the detail parameter, as shown below, with alternative parameters to select more specific information.

- root-detail — Rates, depth and thresholds for the root arbiter.
- thresholds — CBS, MBS and high-prio-only thresholds with associated rates of child policers.
- priority-info — Discard-fair and discard-unfair thresholds, with number of associated children, for each of the root priority levels.
- depth — Parent policer and child PIR buckets depth, with PIR and FIR rate information.
- arbiter — Specific information of a given arbiter.
- port — For use with LAGs in different line cards or using adapt-qos link.

The output adds a good representation of the root arbiter thresholds, indicating the priority levels, discard-unfair and discard-fair thresholds, and how many child policers are associated with each level. It also includes the current depth of the child policer PIR buckets and the parent policer bucket.

```
A:PE-1# show qos policer-hierarchy sap 1/1/3:1 detail
=====
Policer Hierarchy - Sap 1/1/3:1
=====
Ingress Policer Control Policy : cfhp-1
Egress Policer Control Policy :
-----
Legend :
(*) real-time dynamic value
(w) Wire rates
-----
root (Ing)
|
| slot(1)
|   MaxPIR:100000
|   ConsumedByChildren:100000
|   OperPIR:100000      OperFIR:100000
|
|   DepthPIR:8111 bytes
| Priority 8
|   Oper Thresh Unfair:17408      Oper Thresh Fair:25600
|   Association count:0
| Priority 7
|   Oper Thresh Unfair:17408      Oper Thresh Fair:25600
|   Association count:0
| Priority 6
|   Oper Thresh Unfair:17408      Oper Thresh Fair:25600
```

Class Fair Hierarchical Policing for SAPs

```

| Association count:0
| Priority 5
| Oper Thresh Unfair:17408      Oper Thresh Fair:25600
| Association count:1
| Priority 4
| Oper Thresh Unfair:9728      Oper Thresh Fair:17408
| Association count:0
| Priority 3
| Oper Thresh Unfair:9728      Oper Thresh Fair:17408
| Association count:3
| Priority 2
| Oper Thresh Unfair:0         Oper Thresh Fair:8192
| Association count:0
| Priority 1
| Oper Thresh Unfair:0         Oper Thresh Fair:8192
| Association count:1
|
|--(A) : a3 (Sap 1/1/3:1)
| | MaxPIR:60000
| | ConsumedByChildren:60000
| | OperPIR:60000      OperFIR:60000
| |
| | [Level 3 Weight 1]
| | Assigned PIR:60000      Offered:60000
| | Consumed:60000
| |
| | Assigned FIR:60000
| |
|--(P) : Policer 1->1/1/3:1->4
| | | MaxPIR:60000      MaxCIR:20000
| | | CBS:25600        MBS:77824
| | | HiPrio:0
| | | Depth:77876
| | |
| | | OperPIR:15000      OperCIR:15000
| | | OperFIR:15000
| | | PacketByteOffset:0
| | | StatMode: offered-total-cir
| | |
| | | [Level 2 Weight 25]
| | | Assigned PIR:15000      Offered:45800
| | | Consumed:15000
| | |
| | | Assigned FIR:15000
| | |
|--(P) : Policer 1->1/1/3:1->3
| | | MaxPIR:60000      MaxCIR:20000
| | | CBS:25600        MBS:77824
| | | HiPrio:0
| | | Depth:77834
| | |
| | | OperPIR:15000      OperCIR:15000
| | | OperFIR:15000
| | | PacketByteOffset:0
| | | StatMode: offered-total-cir
| | |
| | | [Level 2 Weight 25]
| | | Assigned PIR:15000      Offered:45800
| | | Consumed:15000

```

Show Output

```
| | |
| | | Assigned FIR:15000
| | |
| | | --(P) : Policer 1->1/1/3:1->2
| | | MaxPIR:60000 MaxCIR:20000
| | | CBS:25600 MBS:77824
| | | HiPrio:0
| | | Depth:77848
| | |
| | | OperFIR:30000 OperCIR:20000
| | | OperFIR:30000
| | | PacketByteOffset:0
| | | StatMode: offered-total-cir
| | |
| | | [Level 2 Weight 50]
| | | Assigned PIR:30000 Offered:45800
| | | Consumed:30000
| | |
| | | Assigned FIR:30000
| | |
| | | --(P) : Policer 1->1/1/3:1->5
| | | MaxPIR:10000 MaxCIR:10000
| | | CBS:12800 MBS:12800
| | | HiPrio:0
| | | Depth:12854
| | |
| | | OperFIR:10000 OperCIR:10000
| | | OperFIR:10000
| | | PacketByteOffset:0
| | | StatMode: offered-total-cir
| | |
| | | [Level 5 Weight 1]
| | | Assigned PIR:10000 Offered:10000
| | | Consumed:10000
| | |
| | | Assigned FIR:10000
| | |
| | | --(P) : Policer 1->1/1/3:1->1
| | | MaxPIR:100000 MaxCIR:0
| | | CBS:0 MBS:126976
| | | HiPrio:0
| | | Depth:135
| | |
| | | OperFIR:30000 OperCIR:0
| | | OperFIR:30000
| | | PacketByteOffset:0
| | | StatMode: offered-total-cir
| | |
| | | [Level 1 Weight 1]
| | | Assigned PIR:30000 Offered:45800
| | | Consumed:30000
| | |
| | | Assigned FIR:30000
```

```
root (Egr)
|
No Active Members Found on slot 1
```

```
=====
A:PE-1#
```

The output above gives the depth of the parent policer, which can be used with the output below to explain the operation of the policing in this example.

```
A:PE-1# show qos policer sap 1/1/3:1 detail | match expression "Slot | Bytes | Kbps"
Policer Info (1->1/1/3:1->1), Slot 1
Depth PIR      : 153 Bytes          Depth CIR      : 0 Bytes
Depth FIR      : 153 Bytes
Admin PIR      : 100000 Kbps        Admin CIR      : 0 Kbps
Oper PIR       : 30000 Kbps        Oper CIR       : 0 Kbps
Oper FIR       : 30000 Kbps
Offered Rate   : 45800 Kbps
Parent PIR     : 30000 Kbps        Parent FIR     : 30000 Kbps
Consumed       : 30000 Kbps
Policer Info (1->1/1/3:1->2), Slot 1
Depth PIR      : 77828 Bytes        Depth CIR      : 25624 Bytes
Depth FIR      : 77828 Bytes
Admin PIR      : 60000 Kbps        Admin CIR      : 20000 Kbps
Oper PIR       : 30000 Kbps        Oper CIR       : 20000 Kbps
Oper FIR       : 30000 Kbps
Offered Rate   : 45800 Kbps
Parent PIR     : 30000 Kbps        Parent FIR     : 30000 Kbps
Consumed       : 30000 Kbps
Policer Info (1->1/1/3:1->3), Slot 1
Depth PIR      : 77858 Bytes        Depth CIR      : 25634 Bytes
Depth FIR      : 77858 Bytes
Admin PIR      : 60000 Kbps        Admin CIR      : 20000 Kbps
Oper PIR       : 15000 Kbps        Oper CIR       : 15000 Kbps
Oper FIR       : 15000 Kbps
Offered Rate   : 45800 Kbps
Parent PIR     : 15000 Kbps        Parent FIR     : 15000 Kbps
Consumed       : 15000 Kbps
Policer Info (1->1/1/3:1->4), Slot 1
Depth PIR      : 77838 Bytes        Depth CIR      : 25614 Bytes
Depth FIR      : 77838 Bytes
Admin PIR      : 60000 Kbps        Admin CIR      : 20000 Kbps
Oper PIR       : 15000 Kbps        Oper CIR       : 15000 Kbps
Oper FIR       : 15000 Kbps
Offered Rate   : 45800 Kbps
Parent PIR     : 15000 Kbps        Parent FIR     : 15000 Kbps
Consumed       : 15000 Kbps
Policer Info (1->1/1/3:1->5), Slot 1
Depth PIR      : 12814 Bytes        Depth CIR      : 12814 Bytes
Depth FIR      : 12814 Bytes
Admin PIR      : 10000 Kbps        Admin CIR      : 10000 Kbps
Oper PIR       : 10000 Kbps        Oper CIR       : 10000 Kbps
Oper FIR       : 10000 Kbps
Offered Rate   : 10000 Kbps
Parent PIR     : 10000 Kbps        Parent FIR     : 10000 Kbps
Consumed       : 10000 Kbps
A:PE-1#
```

From the output above, it can be seen that the offered rate for policers 1-4 is 45800Kbps, in fact it is the same for policer 5 but this is capped at the admin PIR rate, 10000Kbps.

Show Output

The depth of the parent policer is only 8111 bytes, so this is not causing any discarding of priority 2-5 traffic at the parent policer as their discard thresholds are all above this value. Therefore the drops in policers 2-5 are all occurring in the child policers.

Policer 5 is consuming all of its operational capacity (PIR, CIR and FIR), and it can be seen that the level of the PIR bucket is 12814 bytes, which is slightly above its MBS of 12800 bytes. The level of the PIR bucket will oscillate around the MBS value as tokens are added to exceed the threshold (causing discards) then the draining reduces the level to just below the threshold (allowing forwarding).

Policers 2-4 are functioning in the same way as policer 5, as can be seen from their PIR bucket levels (levels are 77828 bytes with MBS of 77824), resulting in the PIR buckets constraining the rates of the traffic through these policers. This is happening because the arbiter "a3" is distributing its 60000Kbps in the configured ratio to these policers, which changes the operational PIR to 30000Kbps for policer 2 and 15000Kbps for policers 3 and 4, all being below the offered traffic rate. A similar effect can be seen with the CIR rates and bucket depths, as the operational CIR rate of policer 2 has reached its administrative value with those of policer 3 and 4 being constrained by the operational PIR. The CIR bucket depths are just above the CBS, again this will oscillate causing traffic to both in-profile and out-of-profile. As this is steady state traffic, the operational FIR rates for these policers have settled to match their operational PIR rates.

Policer 1 is also discarding traffic at the PIR bucket but it is also discarding traffic at the parent policer. This can be seen by the fact that policer 1 PIR depth is nowhere near its MBS whereas the parent policer level is just below the priority 1 discard-fair threshold. The level of the parent policer bucket will oscillate around this threshold causing policer 1 traffic to be discarded, which in turn is reflected back into the level of tokens in the policer 1 PIR bucket.

As this example is based on ingress unicast policing, the traffic exits the policers and then accesses the switch fabric using a set of shared-queue (policer-output-queues). The parameters for these queues can be seen using the following **show** command.

```
A:PE-1# show qos shared-queue "policer-output-queues" detail
=====
QoS Shared Queue Policy
=====
-----
Shared Queue Policy (policer-output-queues)
-----
Policy          : policer-output-queues
Description     : Default Policer Output Shared Queue Policy
-----
Queue CIR      PIR      CBS      MBS      HiPrio  Multipoint Pool-Name
-----
1      0       100      1       50      10     FALSE
2      25      100      3       50      10     FALSE
3      25      100      10      50      10     FALSE
4      25      100      3       25      10     FALSE
5      100     100      10      50      10     FALSE
6      100     100      10      50      10     FALSE
7      10      100      3       25      10     FALSE
```


Class Fair Hierarchical Policing for SAPs

```

8      10      100      3      25      10      FALSE
9      0       100      1      50      10      TRUE
10     25      100      3      50      10      TRUE
11     25      100      10     50      10      TRUE
12     25      100      3      25      10      TRUE
13     100     100      10     50      10      TRUE
14     100     100      10     50      10      TRUE
15     10      100      3      25      10      TRUE
16     10      100      3      25      10      TRUE

```

```

-----
FC      UCastQ      MCastQ      BCastQ      UnknownQ
-----
be      1           9           9           9
12      2           10          10          10
af      3           11          11          11
l1      4           12          12          12
h2      5           13          13          13
ef      6           14          14          14
h1      7           15          15          15
nc      8           16          16          16

```

Associations

```

-----
Service : 1          SAP : 1/1/3:1
=====

```

```
A:PE-1#
```

For egress policing, policed traffic can access the exit port by a queue-group, the default being called policer-output-queues. The following shows the parameters for these queues.

```
A:PE-1# show qos queue-group "policer-output-queues" detail
```

```
=====
QoS Queue-Group Ingress
=====
```

```
QoS Queue-Group Egress
=====
```

```
-----
QoS Queue Group
-----
```

```
Group-Name      : policer-output-queues
Description     : Default egress policer output queues.

```

```
-----
Q  CIR Admin PIR Admin CBS      HiPrio PIR Lvl/Wt      Parent      BurstLimit(B)
   CIR Rule  PIR Rule  MBS      CIR Lvl/Wt      Wred-Queue      Slope
   Named-Buffer Pool
-----
1  0          max      def      def      1/1          None          default
   closest   closest  def      0/1          disabled      default
   (not-assigned)
2  0          max      def      def      1/1          None          default
   closest   closest  def      0/1          disabled      default

```

Show Output

(not-assigned)

```
=====
Queue Group Ports (access)
=====
```

Port	Sched Pol	Acctg Pol	Stats	Description
1/1/3		0	No	
1/1/4		0	No	

```
=====
Queue Group Ports (network)
=====
```

Port	Sched Pol	Acctg Pol	Stats	Description
------	-----------	-----------	-------	-------------

No Matching Entries

```
=====
Queue Group Sap FC Maps
=====
```

Sap Policy	FC Name	Queue Id
------------	---------	----------

No Matching Entries

```
=====
A:PE-1#
```

The following output shows the relative thresholds in the parent policer when a priority level has 0, 1 or 2 associated children, both with and without using the fixed parameter.

```
A:PE-1# show qos policer-hierarchy sap 1/1/3:2 ingress priority-info
```

```
=====
Policer Hierarchy - Sap 1/1/3:2
=====
```

```
Ingress Policer Control Policy : cfhp-2
=====
```

```
root (Ing)
|
| slot(1)
|  Priority 8
|    Oper Thresh Unfair:4352      Oper Thresh Fair:5120
|    Association count:0
|  Priority 7
|    Oper Thresh Unfair:4352      Oper Thresh Fair:5120
|    Association count:0
|  Priority 6
|    Oper Thresh Unfair:4352      Oper Thresh Fair:5120
|    Association count:2 fixed
|  Priority 5
|    Oper Thresh Unfair:3328      Oper Thresh Fair:4096
|    Association count:1 fixed
|  Priority 4
|    Oper Thresh Unfair:2304      Oper Thresh Fair:3072
|    Association count:0 fixed
|  Priority 3
```

Class Fair Hierarchical Policing for SAPs

```
| Oper Thresh Unfair:1280      Oper Thresh Fair:2048
| Association count:2
| Priority 2
| Oper Thresh Unfair:0        Oper Thresh Fair:1024
| Association count:1
| Priority 1
| Oper Thresh Unfair:0        Oper Thresh Fair:0
| Association count:0
```

=====
A:PE-1#

Where

- Priority Level 1 has no children so both its fair and unfair thresholds are 0.
- Priority Level 2 has one child so its unfair threshold is 0 and its fair threshold is at the configured mbs-contribution [1024 bytes] (given that this is larger than the min-thresh-separation).
- Priority Level 3 has two children so its unfair threshold is equal to the min-thresh-separation plus the fair threshold of priority 2 [256+1024=1280 bytes]. Its fair threshold is effectively the mbs-contribution plus the fair threshold of priority 2 [1024+1024=2048 bytes] (given that the mbs-contribution is larger than 2x min-thresh-separation).
- Priorities 4, 5 and 6 have the fixed parameter configured. Even though priority 4 has no children, priority 5 has only one child and priority 6 has two children, all three priorities have the same incremental values for their unfair and fair discard threshold. This result in
 - Priority 4's unfair threshold being equal to the min-thresh-separation plus the fair threshold of priority 3 [256+2048=2304 bytes]. Its fair threshold is effectively the mbs-contribution plus the fair threshold of priority 3 [1024+2048=3072 bytes] (given that the mbs-contribution is larger than 2x min-thresh-separation).
 - Priority 5's unfair threshold being equal to the min-thresh-separation plus the fair threshold of priority 4 [256+3072=3328 bytes]. Its fair threshold is effectively the mbs-contribution plus the fair threshold of priority 4 [1024+3072=4096 bytes] (given that the mbs-contribution is larger than 2x min-thresh-separation).
 - Priority 6's unfair threshold being equal to the min-thresh-separation plus the fair threshold of priority 5 [256+4096=4352 bytes]. Its fair threshold is effectively the mbs-contribution plus the fair threshold of priority 5 [1024+4096=5120 bytes] (given that the mbs-contribution is larger than 2x min-thresh-separation).

Note that the above parameter values were chosen to exactly match available hardware values to simplify the output.

Conclusion

This note has described the configuration of Class Fair Hierarchical Policing for SAPs. This hardware policing provides low latency ingress and egress prioritized traffic control with the ability to provide fairness between child policers at the same parent policer priority level.