



Configuring Hierarchical Modular QoS

Hierarchical QoS (H-QoS) is a QoS model that enables you to specify QoS behavior at multiple levels of hierarchy. This chapter provides information about this feature and the different steps involved in configuring it.

This chapter covers the following topics:

- [Overview of Hierarchical Modular QoS, on page 1](#)
- [Restrictions for Configuring H-QoS, on page 2](#)
- [Configuring Hierarchical Queuing, on page 3](#)

Overview of Hierarchical Modular QoS

Hierarchical QoS (H-QoS) allows you to specify QoS behavior at multiple policy levels, which provides a high degree of granularity in traffic management.

H-QoS is applied on the router interface using nested traffic policies. The first level of traffic policy, the parent traffic policy, is used for controlling the traffic at the main interface or sub-interface level. The second level of traffic policy, the child traffic policy, is used for more control over a specific traffic stream or class. The child traffic policy, is a previously defined traffic policy, that is referenced within the parent traffic policy using the **service-policy** command.

Two-level H-QoS is supported on both ingress and egress directions on all line cards and on physical or bundle main interfaces and sub-interfaces.

Three-level Hierarchical QoS (H-QoS) enables enforcement of class/service, group/ Ethernet Flow Point (EFP), and port level SLAs. You can apply regular two-level egress H-QoS policies on the sub-interfaces to achieve class and EFP SLAs at child and parent levels. In addition, you can apply a port shaper policy on the main interface to achieve an aggregated port level SLA in a 1+2 H-QoS or three-level H-QoS model.

An important point to note is that before Release 6.6.25 (where the three-level H-QoS capability was introduced), when you applied class-default shaper on a main interface, it was enforced *only* on the traffic going through the main interface. With three-level H-QoS, a class default shaper that is applied on the main interface is considered as a port shaper and enforced on *all* traffic going out of that physical port. The advantage of three-level H-QoS is that the parent shaper on the sub-interfaces is allowed to oversubscribe, thus enabling best effort sharing of the aggregate port shaper at the third level.

Restrictions for Configuring H-QoS

The following restrictions are applicable while configuring H-QoS:

1. The parent traffic policy only supports the traffic class of type class-default.
2. The parent traffic policy only supports the class-action **shape** and no other queuing action can be configured in it.
3. While configuring on the router, it is mandatory that the priority class must have traffic shaper in the child traffic policy.
4. The sum of the bandwidth of the child policies must be less than the parent policy's traffic shaper.
5. For congestion avoidance and management, the traffic shaper in the parent traffic policy calculates the queue limit and drop priority.
6. PBTS feature does not work when the H-QoS profile is enabled. This is due to TCAM limitations.
7. A maximum of 896 bundle sub-interfaces are only supported in the system, even if there are no QoS policies applied. This is due to an internal LAG_ID resource consumption in HQoS profile mode for bundle sub-interfaces with or without QoS policies being applied.
8. A maximum of 4 priority levels are only supported in HQoS profile mode unlike the default mode where 7-priority levels are supported. The restriction also applies to physical and bundle main interface policies where 7-level priorities were previously used in non-H-QoS profile mode.
9. Bandwidth and Bandwidth remaining configurations are not supported simultaneously within the same policy-map. If a class has bandwidth (CIR), other classes must also have only bandwidth configuration. If a class-map has bandwidth remaining percent/ratio (EIR), other classes should also have only the bandwidth remaining configuration. Shaping is applied on any class.
10. Priority classes must have rate limit configuration by using a Shaping configuration. The effective shaper value is taken as priority bandwidth reservation. Sum of priority bandwidth reservations across all sub-interfaces and main interfaces must not exceed the network interface (NIF) port speed. This is to avoid over-subscription of priority traffic across the network interface port.
Rates of non-priority classes and parent shaping can be over-subscribed.
11. The granularity of bandwidth or bandwidth remaining ration (BRR) is 1:64 as compared to 1:4096 in non-hqos mode. So, there could be accuracy differences in bandwidth performance based on the values used.
12. Filtering for egress IPv4 and IPv6 multicast traffic is not supported if H-QoS is configured on the router.

The following restrictions are applicable while configuring three-level H-QoS:

- There is no support for bandwidth action at the EFP parent level. All EFP/sub-interface policies get a fair share of the port shaper.
- Three-level H-QoS does not apply to ingress policies or to egress marking policies.
- Executing **clear qos counters** on the main interface clears only the main interface policy statistics. Use the "all" option to clear all sub-interface statistics or alternately, clear the sub-interface policy statistics individually.

- Main interface policy statistics do not reflect the sub-interface packet / byte counters, although the port shaper is enforced on all logical ports for a given physical interface. The sub-interface policy-map statistics reflect the transmitted and dropped packet/byte count post-port shaper enforcement.

Configuring Hierarchical Queuing

Before you configure H-QoS, you must enable the H-QoS profile on the router. After enabling H-QoS profile, reload the router, as shown in the following configuration.

```
admin
hw-module location all reload
Router# configure
Router(config)# hw-module profile qos hqos-enable
Router(config)# commit
Router# admin
sysadmin-vm:0_RP0# hw-module location all reload
```

The steps that are involved in configuring hierarchical queuing are as follows:

1. Configure a class-map.
2. Configure a child traffic policy using the class-map that was configured in the previous step.
3. Configure a parent traffic policy and add the child traffic policy in it.

The parent traffic policy is the H-QoS traffic policy and it can be applied on physical or bundle main interfaces and sub-interfaces.

Configuration Example

Configuration of a class-map is as follows:

```
Router# configure
Router(config)# class-map match-any tc2
Router(config-cmap)# match traffic-class 1
Router(config-cmap)# end-class-map
Router(config)# commit
```

Configuration of a child traffic policy is as follows:

```
Router# configure
Router(config)# policy-map child
Router(config-pmap)# class tc2
Router(config-pmap-c)# shape average percent 20
Router(config-pmap-c)# exit
Router(config-pmap)# class class-default
Router(config-pmap-c)# shape average percent 1
Router(config-pmap)# end-policy-map
Router(config)# commit
```

Configuration of a parent traffic policy is as follows:

```
Router# configure
Router(config)# policy-map parent
Router(config-pmap)# class class-default
Router(config-pmap-c)# service-policy child
```

```
Router(config-pmap-c)# shape average percent 50
Router(config-pmap)# end-policy-map
Router(config)# commit
```

Running Configuration

```
/* Configuration of a Class-map */
class-map match-any tc2
  match traffic-class 1
end-class-map
!
/* Configuration of a Child Traffic Policy */
policy-map child
  class tc2
    shape average percent 20
  !
  class class-default
    shape average percent 1
  !
end-policy-map
!
/* Configuration of a Parent Traffic Policy */
policy-map parent
  class class-default
    service-policy child
    shape average percent 50
  !
end-policy-map
!
```

Applying the Parent Traffic Policy on a Main Interface

```
Router# configure
Router(config)# interface TenGigE 0/0/0/10
Router(config-int)# service-policy output parent
Router(config-int)# commit
```

Applying the Parent Traffic Policy on a Sub-interface

```
Router# configure
Router(config)# interface TenGigE 0/0/0/10.1
Router(config-int)# service-policy output parent
Router(config-int)# commit
```

Verification

Verify if the H-QoS traffic policy is applied correctly on the interface using the commands **show qos interface interface-name output**. In the following example, the **Level1 Class** gives information about the class-map that is associated with the parent traffic policy and the **Level2 Class** gives information about the class-maps that are associated with the child traffic policy.

```
RP/0/RP0/CPU0:ios#show qos interface ten0/0/0/10 output
```

```
NOTE:- Configured values are displayed within parentheses
Interface TenGigE0/0/0/10 ifh 0x1e0 -- output policy
NPU Id:                                0
Total number of classes:                 3
Interface Bandwidth:                     10000000 kbps
```

```

VOQ Base:                               1136
Accounting Type:                         Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class                             = class-default
Queue Max. BW.                           = no max (50 %)
Queue Min. BW.                           = 0 kbps (default)
Inverse Weight / Weight                   = 0 / (BWR not configured)
  Level2 Class                             = tc2
  Egressq Queue ID                        = 1138 (LP queue)
  Queue Max. BW.                          = 1020015 kbps (20 %)
  Queue Min. BW.                          = 0 kbps (default)
  Inverse Weight / Weight                  = 1 / (BWR not configured)
  Guaranteed service rate                  = 1000000 kbps
  TailDrop Threshold                       = 1253376 bytes / 10 ms (default)
  WRED not configured for this class
  Level2 Class                             = class-default
  Egressq Queue ID                        = 1136 (Default LP queue)
  Queue Max. BW.                          = 50625 kbps (1 %)
  Queue Min. BW.                          = 0 kbps (default)
  Inverse Weight / Weight                  = 1 / (BWR not configured)
  Guaranteed service rate                  = 50000 kbps
  TailDrop Threshold                       = 62720 bytes / 10 ms (default)
  WRED not configured for this class

```

The statistics for the packets that have matched the different traffic classes of the parent and child traffic policies can be viewed using the command **show policy-map interface *interface-name* output**. Also, this command also shows the number of packets that are transmitted or dropped when the specified action is applied on the packets that have matched the respective traffic class.

```

Router# show policy-map interface ten0/0/0/10 output

TenGigE0/0/0/10 output: parent
Class class-default
  Classification statistics                (packets/bytes)    (rate - kbps)
  Matched                                : 2313578823/296138089344    8494665
  Transmitted                             : 232805738/29799134464     854465
  Total Dropped                           : 2080773085/266338954880   7640200
Policy child Class tc2
  Classification statistics                (packets/bytes)    (rate - kbps)
  Matched                                : 2313578823/296138089344    8494665
  Transmitted                             : 232805738/29799134464     854465
  Total Dropped                           : 2080773085/266338954880   7640200
  Queueing statistics
  Queue ID                                : 1138
  Taildropped(packets/bytes)              : 2080773085/266338954880
Policy child Class class-default
  Classification statistics                (packets/bytes)    (rate - kbps)
  Matched                                : 0/0                       0
  Transmitted                             : 0/0                       0
  Total Dropped                           : 0/0                       0
  Queueing statistics
  Queue ID                                : 1136
  Taildropped(packets/bytes)              : 0/0

```

When using hierarchical policers, there is no independent set of hardware counters to store the parent policer statistics. Instead, parent policer statistics are manipulated in the software to be the sum of all child policers under the same policy-map.

This is shown in the following example where two streams of traffic, with CoS value of 1 and 2 are sent at a speed of 3.5 Gbps each.

```

/*Hierarchical Policy Map Configuration*/
=====
Router# show running-config policy-map Hingress
policy-map Hingress
  class class-default
    service-policy ingress
    police rate 5 gbps peak-rate 9 gbps
  !
!
end-policy-map
!
/*Ingress Policy Map Configuration*/
=====
Router#show running-config policy-map ingress
policy-map ingress
  class cos1
    set traffic-class 1
    police rate 5 gbps
  !
!
  class cos2
    set traffic-class 2
    police rate 5 gbps
  !
!
  class class-default
  !
end-policy-map
!
/*Policy Map applied at TenGigE0/0/0/6.100 Interface*/
=====
Router#show policy-map interface tenGigE 0/0/0/6.100 input

TenGigE0/0/0/6.100 input: Hingress

Class class-default
  Classification statistics          (packets/bytes)    (rate - kbps)
  Matched                          :          856717937/109659895936      6683676
  Transmitted                       :          856717937/109659895936      6683676
  Total Dropped                     :                   0/0              0
  Policing statistics              (packets/bytes)    (rate - kbps)
  Policed(conform)                 :          856717937/109659895936      6683674
  Policed(exceed)                   :                   0/0              0
  Policed(violate)                  :                   0/0              0
  Policed and dropped               :                   0/0

Policy ingress Class cos1
  Classification statistics          (packets/bytes)    (rate - kbps)
  Matched                          :          437826303/56041766784      3341838
  Transmitted                       :          437826303/56041766784      3341838
  Total Dropped                     :                   0/0              0
  Policing statistics              (packets/bytes)    (rate - kbps)
  Policed(conform)                 :          437826303/56041766784      3341838
  Policed(exceed)                   :                   0/0              0
  Policed(violate)                  :                   0/0              0
  Policed and dropped               :                   0/0
  Policed and dropped(parent policer) : 0/0

Policy ingress Class cos2
  Classification statistics          (packets/bytes)    (rate - kbps)
  Matched                          :          418891634/53618129152      3341838
  Transmitted                       :          418891634/53618129152      3341838
  Total Dropped                     :                   0/0              0
  Policing statistics              (packets/bytes)    (rate - kbps)

```

```

    Policed(conform)      :          418891634/53618129152      3341838
    Policed(exceed)      :              0/0                  0
    Policed(violate)     :              0/0                  0
    Policed and dropped  :              0/0
    Policed and dropped(parent policer) : 0/0

```

```

Policy ingress Class class-default
Classification statistics      (packets/bytes)      (rate - kbps)
    Matched                    :              0/0              0
    Transmitted                 :              0/0              0
    Total Dropped               :              0/0              0
Policy Bag Stats time: 0
Policy Bag Stats time: 0

```

Configuration Example for Three-Level H-QoS

To configure three-level H-QoS:

1. Configure the port shaper or EFP group shaper.
2. Configure the two-level H-QoS policy with an EFP parent shaper and class or service level actions
3. Enable the port or EFP group shaper on the main interface to become the root policy.
4. Enable the two-level H-QoS policy on each EFP instance, thus providing a three-level hierarchy of service, EFP and EFP group or port SLAs.

A sample configuration for three-level H-QoS is as follows.

```

policy-map port_shaper
  class class-default
    shape average 6 gbps
  !
end-policy-map
!

policy-map efp_policy
  class class-default
    service-policy efp_policy_child
    shape average 4 gbps
  !
end-policy-map

!

policy-map efp_policy_child
  class tc1
    shape average 50 mbps
    priority level 1
  !
  class tc2
    bandwidth percent 50
  !
  class tc3
    bandwidth percent 30
  !
  class class-default
  !
end-policy-map
!

interface TenGigE0/5/0/4
  service-policy output port_shaper

```

```
!  
  
interface TenGigE0/5/0/4.1  
  service-policy output efp_policy  
  encapsulation dot1q 11  
!  
  
interface TenGigE0/5/0/4.2  
  service-policy output efp_policy  
  encapsulation dot1q 12  
!
```

Verification

Run the **show policy-map interface** command in XR EXEC mode to display the packet/byte count and rate post port shaper enforcement on each of the sub-interface/EFPP policies.