

Configuring Quality of Service

2

This chapter provides information for configuring Quality of Service (QoS) on your ERX edge router. QoS provides the service that allows an ERX system to distinguish traffic with strict timing requirements from traffic that can tolerate delay, jitter, and loss.

Topic	Page
Overview	2-1
References	2-4
Configuration Tasks	2-5
Traffic Classes	2-6
Traffic Class Groups	2-7
Queue Profiles	2-8
Scheduler Profiles	2-10
QoS Profiles	2-13
QoS Profile Attachments	2-17
QoS Port-Type Profile Attachments	2-18
Strict Priority Scheduling	2-18
Rate Shaping	2-19
Monitoring QoS	2-19

Overview

QoS is a suite of features that configure queuing and scheduling on the forwarding path of your ERX system. QoS provides a level of predictability and control beyond the current *best-effort* service. Your ERX system provides best-effort data delivery by default. Packets not

assigned to a specific traffic class are carried in the best-effort traffic class. Best-effort service provides packet transmission with no guarantee of results. For information on traffic classes, see *Traffic Classes* later in this chapter.

QoS, as developed by Juniper Networks for your ERX system, conforms to the IETF Differentiated Services (DiffServ) model (RFCs 2597 and 2598). DiffServ networks classify packets into one of a small number of aggregated flows or traffic classes. The Juniper Networks QoS architecture extends DiffServ to support edge features such as high-density queuing.

The ERX system supports:

- The IETF architecture for differentiated services
- The assured forwarding (AF) per-hop-behavior (PHB) groups
- The expedited forwarding (EF) per-hop-behavior (PHB) groups

See *References* for the list of related RFCs.

The ERX architecture extends the IETF model with edge features, such as dense queuing.

The ERX system supports configurable queuing and configurable scheduling. It has an ASIC scheduler that supports thousands of queues in a hierarchical round robin (HRR) scheduler. This allows your system to allocate separate queues for each forwarding interface. Separate queues enable fair access to buffers and bandwidth for each subscriber connected to the system.

Allocating queues per interface allows an Internet service provider (ISP) to shape an individual subscriber's traffic flows to specified rates independent of the underlying layer 2 network type.

Your ERX system supports QoS on the 5 Gbps, 10 Gbps and the 40 Gbps fabric boards in the system. It supports egress line module functions only on ASIC-based line modules.

For an illustration of traffic flow through your ERX system, see Figure 2-1.

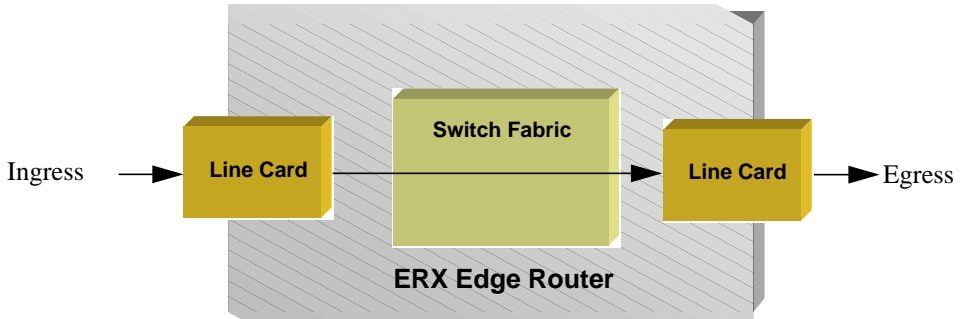


Figure 2-1 Traffic flow through an ERX system

Features

The major QoS features Juniper Networks provides are described in Table 2-1.

Table 2-1 QoS features

Feature	Description
Best effort	Default traffic class for packets being forwarded across the device. Packets that are not assigned to a specific traffic class in a policy are assigned to the best-effort traffic class.
Differentiated services	<ul style="list-style-type: none"> Assured forwarding (AF) – See RFC 2597. Expedited forwarding (EF) – See RFC 2598.
QoS port-type profile	QoS profile that is automatically attached to ports of the corresponding type if you do not explicitly attach a QoS profile.
QoS profile	A collection of QoS commands that specify queue profiles and scheduler profiles in combination with interface types
Queue profile	A template that specifies the buffering and dropping behavior of an egress queue
Rate shaping	Mechanism that throttles the rate at which an interface can transmit packets <i>Note: Rate shaping as presented in policy management in previous releases is deprecated and converted to QoS profiles and scheduler profiles.</i>
Scheduler profile	Configures the bandwidth at which queues drain as a function of relative weight and shaping rate
Traffic class	A set of resources on the data path of a router that is reserved for use by packets assigned to the traffic class <i>Note: The Juniper Networks QoS architecture supports up to eight traffic classes.</i>

Table 2-1 QoS features (continued)

Feature	Description
Traffic class group	A separate hierarchy of scheduler nodes and queues over a port, used for strict-priority queues. Each traffic class belongs to only one group (or to no group). There can only be one group.

Terms

Table 2-2 provides information on the major terms used in discussing QoS.

Table 2-2 QoS terminology used in this chapter

Term	Meaning
Best effort	Network forwards as many packets as possible in as reasonable a time as possible. This is the default per-hop behavior (PHB) for packet transmission.
QoS port-type profile	Supplies the QoS information for forwarding interfaces stacked above ports of the associated interface type
Latency	The delay in the transmission of a packet through a network from beginning to end
Proprietary QoS MIB	Supported on your ERX system
QoS profile attachment	Applies the rules in the QoS profile to a specific interface and all interfaces
Queue	First-in-first-out (FIFO) set of buffers that control packets on the data path
Rate shaping	Allows you to throttle a queue to a specified rate
Strict priority	Designates the traffic class (queue) that receives top priority for transmission of its packets through a port
Weight	Specifies the relative weight for queues in the traffic class
HRR	Hierarchical round robin. Allocates bandwidth to queues in proportion to their weights.

References

For more information about QoS, see the following resources:

- RFC 3260 – New Terminology and Clarifications for Diffserv (April 2002)
- RFC 3246 – An Expedited Forwarding PHB (Per-Hop Behavior) (March 2002)

- RFC 2998 – A Framework for Integrated Services Operation over Diffserv Networks (November 2000)
- RFC 2990 – Next Steps for the IP QoS Architecture (November 2000)
- RFC 2698 – A Two Rate Three Color Marker (September 1999)
- RFC 2597 – Assured Forwarding PHB Group (June 1999)
- RFC 2475 – An Architecture for Differentiated Services (December 1998)
- RFC 2474 – Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers (December 1998)

Configuration Tasks

To configure QoS, perform the following tasks:

- 1 Create a traffic class.

You can create up to eight traffic classes.

- 2 (Optional) Create a traffic class group.

You can add multiple traffic classes to the group. Groups are required only for strict priority queues.



Note: Only one group is supported. All strict-priority traffic classes should belong to the group.

- 3 Create a queue profile, if you want to configure non-default buffer management.

The ERX system supports up to 16 queue profiles.

- 4 Create a scheduler profile.

The ERX system supports up to 1,000 scheduler profiles.

- 5 Create a QoS profile.

A QoS profile references queue profiles and scheduler profiles.

- 6 Attach the QoS profile to one or more interface(s), or specify the profile as a QoS port-type profile for a given interface type.

Traffic Classes

A traffic class is a system-wide collection of resources configured to provide a defined level of service to packets assigned to the traffic class. The resources consist of buffers, queues, and bandwidth.

A traffic class corresponds to what the IETF DiffServ working group calls a traffic class in RFC 2597 – *Assured Forwarding PHB Group* (June 1999).

Traffic classes are global to the system. Packets are:

- Classified into a traffic class on ingress or egress
- Queued on fabric queues that are specific to the traffic class
- Queued on the egress line module on queues that are specific to the traffic class

Input policies classify packets into the traffic class; the fabric carries the packets to an egress line module in a fabric queue that is specific to the traffic class; and the packets are placed into traffic class-specific queues on the egress line module.

Best-Effort Forwarding

Your ERX system automatically defines a default traffic class called *best-effort*. When no other configuration is specified, it is assumed that packets belong to this aggregate. You cannot delete this traffic class. Packets that are not classified into a traffic class are assigned to the best-effort traffic class. Packets that arrive at an egress line module that has no queues allocated for their traffic class are also assigned to the best-effort traffic class.

Configuring a Traffic Class

To configure a traffic class, you need to name the class. You also have the option of setting its fabric weight and fabric strict priority.

fabric-strict-priority

- Use to specify strict priority scheduling for queues in the traffic class.
- If multiple traffic classes are strict priority, the weight determines which class gets more bandwidth.
- Example

```
host1(config-traffic-class)#fabric-strict-priority
```
- Use the **no** version to delete the strict priority setting.

fabric-weight

- Use to specify the relative weight for queues in the traffic class *in the fabric*.
- The weight value is in the range 1–63.
- Zero is not a legal weight.
- The default weight is 8.
- Example

```
host1(config-traffic-class)#fabric-weight 8
```
- Use the **no** version to set the fabric to the default weight value.



Note: *Weight controls the bandwidth management of fabric queues associated with the traffic class. It does not control the weight of egress queues associated with the traffic class.*



Note: *The **fabric-weight** command currently works only with ERX-1440 edge routers.*

traffic-class

- Use to configure a traffic class in the ERX system.
- The traffic class name can be up to 32 characters with no spaces.
- Your system currently supports up to eight global traffic classes.
- Each traffic class can appear in only one traffic class group.
- This command accesses Traffic Class Configuration mode.
- Example

```
host1(config)#traffic class low-loss1
host1(config-traffic-class)#
```
- Use the **no** version to delete a specified traffic class.

Traffic Class Groups

Each traffic class can belong either to only one traffic class group or to no group. A traffic class group can contain multiple traffic classes.



Note: *In this release, you can configure only one group.*

Configuring Traffic Class Groups

To configure a traffic class group, use the following commands to name the group and specify the traffic classes that belong to the group. A strict-priority traffic class should belong to the group.

traffic-class

- Use to specify a traffic class that belongs to the traffic class group.
- The traffic class group is used in strict priority applications.

- Use in Traffic Class Group Configuration mode.
- Example

```
host1(config-traffic-class-group)#traffic-class
low-latency-group
```
- Use the **no** version to delete a traffic class from the selected traffic-class group.

traffic-class-group

- Use to configure a traffic class group.
- Use to access Traffic Class Group Configuration mode from which you can add or delete the traffic classes that are members of the group.
- Example

```
host1(config)#traffic-class-group strict-priority
low-latency-group
host1(config-traffic-class-group)#
```
- Use the **no** version to remove the selected traffic-class group.

Queue Profiles

A queue is a first-in-first-out (FIFO) set of buffers that buffer packets on the data path. The Juniper Networks QoS architecture associates queues with a traffic class-interface pair. For example, if you create 4,000 IP interfaces and configure each interface with four traffic classes, then 16,000 queues are created.

Your ERX system dynamically manages the shared memory on the egress line card to provide a good balance between sharing the memory between queues and protecting an individual queue's claim on its fair share of the egress memory.

When the egress packet memory is in high demand and aggregate utilization of the 32 MB is high, queue lengths are set to lengths that strictly partition egress memory into per-queue memory sections. This conservative buffer management strategy reserves a fair share of buffers for each queue, so that aggressive bandwidth consumers cannot starve out moderate traffic consumers by allocating all the shared memory resource.

When the egress packet memory is in low demand, a more liberal buffer management strategy is used to allow the queues that are active more access to the shared memory resource.

The system dynamically varies queue lengths for all queues as the real time demand on the egress packet memory changes.

You can configure limits to prevent the system from setting queue lengths too low, or too high.

Configuring Queue Profiles

You can configure only egress queues. A queue profile specifies the buffering and dropping behavior of a set of egress queues.

To configure a queue profile, you need to name the profile and set the queue's length and buffer weight. You can set the queue lengths to a minimum and maximum length. Setting a minimum length allows you to oversubscribe buffer memory. Setting a maximum length allows you to limit the buffering in queues. If you do not set the lengths, the ERX system varies the queue length dynamically between a minimum of 1 KB and a maximum of 7 MB.



Note: You can configure only 16 queue profiles on a system.

buffer-weight

- Use to set the buffer weight of the queue.
- The buffer weight range is 1–63; the default is 8.
- Queues with a buffer weight of 16 are twice as long as queues with a buffer weight of 8.
- Example

```
host1(config-queue)#buffer-weight 8
```

- Use the **no** version to remove the buffer weight for a queue.

committed-length **conformed-length** **exceeded-length**

- Use to set minimum or maximum constraints on queue lengths.
- You can set minimum and maximum constraints.
 - › minimum length: range 0–16MB; default is 0 (no minimum)
 - › maximum length: range 0–1GB; default is 0 (no maximum)
- The **committed-length** command sets a minimum or maximum queue length for committed packets. The color for committed packets is green. See *Color-Based Thresholding below*.
- The **conformed-length** command sets a minimum or maximum queue length for conformed packets. The color for conformed packets is yellow. See *Color-Based Thresholding*.
- The **exceeded-length** command sets a queue minimum or maximum length for exceeded packets. The color for exceeded packets is red. See *Color-Based Thresholding*.

- Example

```
host1(config-queue)#committed length 8 10
```

- Use the **no** version to remove constraints on the queue length.

conformed-action
exceeded-fraction

- Use to set the conformed and exceeded drop thresholds as a percentage of the committed threshold.
 - › exceeded fraction: range is 0–100; default is 25
 - › conformed fraction: range is 0–100; default is 50
- Example

```
host1(config-queue)#exceeded-fraction 25
```
- Use the **no** version to return the fraction to its default setting.

queue-profile

- Use to configure a queue profile.
- Use to enter Queue Configuration mode.
- Example

```
host1(config)#queue-profile qp-low-latency
host1(config-queue)#
```
- Use the **no** version to remove the queue profile.

Color-Based Thresholding

When the queue fills above the exceeded threshold, red packets are dropped, but yellow and green are still queued. When the queue fills above the conformed drop threshold, only green packets are queued.

Packets within the system are tagged with a drop precedence:

- Committed – green
- Conformed – yellow
- Exceeded – red



Note: *The exceeded and conformed thresholds vary in proportion to the dynamic queue length.*

Scheduler Profiles

The egress line module scheduler is a hierarchical round-robin (HRR) scheduler. Figure 2-2 is an example of a QoS scheduler’s hierarchy.




Figure 2-2 QoS scheduler hierarchy

The following list describes Figure 2-2:

- The queues feeding a physical port are organized in a hierarchy.
- The scheduler selects a first-level node based on relative weights and rate shaping rates. See *Rate Shaping* later in this chapter.
- The scheduler then selects a second-level node from the group of nodes that are stacked above the first-level node that was selected. This selection is also based on relative weights and rate shaping rates.
- Finally, a queue is selected from the group of queues stacked above the second-level node.

The scheduler supports relative weighting and rate shaping on all three levels of the hierarchy: first-level node, second-level node, and queue. The bandwidth delivered from a given node or queue is a function of both the relative weight and the shaping rate:

- The weight provides minimum bandwidth to the node or queue.
- The shaping rate provides maximum bandwidth to the node or queue.

When the scheduler is congested, the relative weights determine which queue can claim bandwidth. When the scheduler is not congested, the shaping rates determine the claim to bandwidth.



Note: For ERX ASIC cards, strict priority is supported only for a single first-level scheduler node.

Configuring Scheduler Profiles

To configure a scheduler profile, you need to name the profile and specify the shaping rate and weight. You can also set the scheduler node or queue to strict priority.



Note: If you configured traffic shaping through traffic shape profiles on your ERX system in earlier ERX software releases, with this release “traffic shaping” is replaced with the “rate shaping” feature which is configured when you configure a scheduler profile.

scheduler-profile

- Use to configure a scheduler profile.
- Your ERX system supports up to 1,000 scheduler profiles.
- Use to enter Scheduler Profile Configuration mode.
- Example


```
host1(config)#scheduler-profile sp-1mbs
host1(config-scheduler-profile)#
```
- Use the **no** version to remove the scheduler profile.

shaping-rate

- Use to set the shaping rate of the scheduler node or queue in bits per second.
- The rate is in bits per second. Any positive number is acceptable, but your ERX system rounds to the next higher 8 Kbps.
- Shaping rate range is 64,000–1,000,000,000 bps; default is no shaping rate.
- Burst is the catch up number associated with the shaper.
- Burst rate range is 8192–32,767 bps. The default is 32 Kbps.
- Example


```
host1(config-scheduler-profile)#shaping-rate 32
```
- Use the **no** version to delete the shaping rate.

strict-priority

- Use to set strict-priority scheduling for the scheduler node.
- Example


```
host1(config-scheduler-profile)#strict-priority
```
- Use the **no** version to delete the strict priority setting.

weight

- Use to set the hierarchical round robin (HRR) weight of the scheduler node or queue.
- The weight value is in the range 1–63 inclusive. The default weight is 8.
- Example


```
host1(config-scheduler-profile)#weight 8
```
- Use the **no** version to set the weight setting to the default weight.

QoS Profiles

A QoS profile is a collection of QoS commands that specify queue profiles and scheduler profiles in combination with interface types. A QoS profile specifies the queue and scheduler configuration for a given subtree of the interface hierarchy. The QoS profile controls the way scheduler nodes and queues are bound to the interfaces above its attachment point in the interface hierarchy.

A QoS profile is attached to the interface at the base of the subtree hierarchy. For example, a QoS profile attached to an ATM port specifies queuing attributes for interfaces of all types that are stacked over the port.



Note: QoS commands affect only ASIC cards.

Configuring QoS Profiles

To configure a QoS profile, you need to name the profile and also name the traffic class and/or the queue profile, scheduler profile, or traffic class group that belong to the QoS profile. Each command begins with a keyword that designates an interface type. Table 2-3 lists the interface type and the commands they can be used with.

Table 2-3 Interface types and commands

Interface type	queue	node	group
atm	x	x	x
atm-vc	x	x	
cbf	x	x	
ethernet	x	x	x
fr-vc	x	x	
ip	x	x	
ip-tunnel	x	x	
l2tp-tunnel	x	x	
serial	x	x	x

Table 2-3 Interface types and commands (continued)

Interface type	queue	node	group
server-port	x	x	x
vlan	x	x	

group

- Use to configure a group node for each interface of the specified type.
- The group defaults to *no group*. In this release, your ERX system supports only one named traffic-class group above a given port.
- Each traffic class can belong to only one traffic-class group.
- With this command, you can also associate a scheduler profile with the QoS profile.
- Example

```
host1(config-qos-profile)#atm group groupA scheduler-profile scheduler1
```

- Use the **no** version to remove this rule from the QoS profile.

node

- Use to configure a scheduler node for each interface of the specified type.
- The optional scheduler profile supplies a relative weight and potentially a shaping rate to be applied at the scheduler node.



Note: For ERX ASIC cards, a scheduler profile cannot be associated with a port-type interface unless the strict-priority group is also specified.

- Example

```
host1(config-qos-profile)#ip node scheduler-profile scheduler1 group strict-priority
```

- Use the **no** version to remove this rule from the QoS profile.

queue

- Use to configure a queue for each interface in the specified traffic class.
- The optional scheduler profile supplies a relative weight and potentially a shaping rate to be applied at the queue.

The optional queue profile supplies threshold information for the queue, if the system defaults are not appropriate.

- Each queue traffic class can appear in only one traffic class group.
- Example

```
host1(config-qos-profile)#atm queue traffic-class strict-priority
```

- Use the **no** version to remove this rule from the QoS profile.

qos-profile

- Use to create a QoS profile on your system.
- Use to enter QoS Profile Configuration mode.
- Example


```
host1(config)#qos-profile qosp-vc-queuing
host1(config-qos-profile)#
```
- Use the **no** version to remove the QoS profile.

Example

The following **show** command displays information on a QoS profile. The example illustrates the components of a QoS profile.

```
host1#show qos-profile
```

```
atm-default:
interface  rule      traffic      scheduler   queue      drop      t-class
  type     type      class      profile     profile    profile   group
-----
ip         node      default
atm-vc     node      default
cbf        node      default
ip         queue    best-effort default     default    default
atm        queue    best-effort default     default    default
atm-vc     queue    best-effort default     default    default
cbf        queue    best-effort default     default    default
```

Configuring QoS for ATM Interfaces

ATM modules present a higher degree of complexity than other interfaces, due to the presence of a separate buffer pool and ATM cell scheduler that reschedules cells onto the ATM network based on ATM QoS configuration.

Table 2-4 defines the acronyms used in this section.

Table 2-4 Acronyms

Acronym	Term
CBR	constant bit rate
PCR	peak cell rate
SAR	segmentation and reassembly
UBR	unspecified bit rate
VBR	variable bit rate
VC	virtual circuit
VP	virtual path

By default, the ATM SAR uses a per-VC scheduler. The SAR requires per-VC queues to implement any non-UBR ATM QoS. Non-UBR rates are *specified* bit rates, such as CBR, ntVBR, UBR with PCR, and VP shaping. Frame scheduler configurations that provide a single per-VC scheduler node at level 1 of the scheduler hierarchy are compatible with the default SAR scheduler. This means that the QoS profile must include an **atm-vc node** subcommand, and must not include **atm queue**, or **atm group** commands.

```
host1(config)#qos-profile athens
host1(config-qos-profile)#atm-vc node scheduler-profile
scheduler2 group strict-priority
```

The ERX system supports the **qos-mode-port** command in ATM major Interface Configuration mode.

```
host1(config)#interface atm 0/1
host1(config-if)#qos-mode-port
```

This command configures the SAR for per-port queuing and is mutually exclusive with any non-UBR ATM QoS service. In port-mode queuing, the SAR manages its buffer pool to contain a very small number of frames at any moment in time, just the minimum required to maintain line rate forwarding onto the ATM port. Instead of buffering up to 32 MB of frames and implementing per-VC round robin among thousands of VCs, the SAR buffers only 4-12 KB of frames, and the total VC count in the SAR at any moment in time is small enough so that the effect of the cell scheduler is minimized. In port-mode queuing the SAR buffer pool and scheduler are effectively eliminated from the traffic model, so that the packet scheduler dominates the sequencing of frames into the SAR, and the SAR introduces only a minor second-order effect of interleaving cells from a small number of VCs. Any frame scheduler configuration can be supported while the SAR is configured in port-mode.

For ERX ATM modules, configuring any port of an ATM card in **qos-mode-port** puts all ports for that module into port-mode queuing. Port mode queuing is mutually exclusive with configuration of any non-UBR ATM service on any VC on that module.

qos-mode-port

- Use to configure an ATM port for per-port queuing.
- This command runs in Interface Configuration mode.
- This command excludes non-UBR ATM QoS service on any VC on the ATM module.
- Example

```
host1(config)#interface atm 0/1
host1(config-if)#qos-mode-port
```
- Use the **no** version to remove per-port queuing on the ATM port.

QoS Profile Attachments

You can attach a QoS profile to an interface at the base of an interface hierarchy. The sequence of commands for doing this is:

```
host1(config)#interface interfaceID
host1(config-if)#qos-profile qosProfileName
```

interface

- Use to create an interface of a particular interface type. See Table 2-3.
- Use to enter Interface Configuration mode.
- Example

```
host1(config)#interface atm 1.0/1
```
- Use the **no** version to remove the interface.

qos-profile

- Use to attach a QoS profile to a given interface.
- Interface types below the attachment point cannot be referenced in the QoS profile.
- Example

```
host1(config-if)#qos-profile qosp-vc-queuing
```
- Use the **no** version to remove the QoS profile from a given interface.

QoS Port-Type Profile Attachments

A QoS port-type profile supplies the QoS information for all forwarding interfaces stacked above all ports of the associated interface type. If you explicitly attach a QoS profile to a port, however, that profile overrides all the commands in the QoS port-type profile. The QoS profile associates queue profiles and scheduler profiles with interface types, and it applies to all interfaces stacked above ports of the associated type.

qos-port-type-profile

- Use to associate a QoS profile with all the ports of a given interface type.
- The interface type can be any one of the following: atm, serial, ethernet, or server.
- A profile attached to a port must specify a queue for each forwarding interface type in the best-effort traffic class.
- Example

```
host1(config)#qos-port-type-profile atm qos-profile  
strict-priority
```

- There is no **no** version for this command.

Strict Priority Scheduling

Your ERX system can configure one or more strict priority queues per interface. Strict priority scheduling is implemented with a special strict priority scheduler node that is stacked directly above the port. Queues stacked on top of the strict priority scheduler node always get the bandwidth before any other queues.

Only one node at the first scheduler level can be configured as strict priority. If any node or queue above the strict priority node has packets, it is scheduled next. If multiple queues above the strict priority node have packets, the hierarchical round-robin (HRR) algorithm selects which strict priority queue is scheduled next, based on the HRR algorithm.

Example

```
host1(config-qos-profile)#atm scheduler-profile  
strictSchedulerProfile group strict priority
```

Rate Shaping

Rate shaping is a service that throttles the rate at which queues transmit packets. Rate shaping is TCP friendly, that is, it buffers packets that are above the rate, rather than dropping them. Your ERX system supports 64,000 rate shapers per line module. Shaping rates are multiples of 8 Kbps.



Note: With this release of ERX software, rate shaping is configured via the scheduler profile. See *Configuring Scheduler Profiles* earlier in this chapter.

Monitoring QoS

You can monitor the elements and profiles supported by QoS. To do so, run the following commands.

show fabric-queue

- Use to display forwarded and dropped statistics for the fabric.
- If you do not specify one of the keywords (**traffic-class**, **egress-slot**, or **detail**), this command displays general data about the fabric queue.
- Example

```
host1#show fabric-queue
```

traffic class	egress slot	type	forwarded packets	forwarded bytes	dropped packets	dropped bytes
best-effort	all	committed	0	0	0	0
best-effort	all	conformed	0	0	0	0
best-effort	all	exceeded	0	0	0	0

show ip interface

- Use to display QoS parameters on a particular interface.
- Example

```
host1#show ip interface atm 2/0.1
ATM2/0.1 line protocol Atm1483 is up, ip is up
Network Protocols: IP
Internet address is 90.120.1.1/255.255.0.0
Broadcast address is 255.255.255.255
Operational MTU = 9180 Administrative MTU = 0
Operational speed = 155520000 Administrative speed = 0
Discontinuity Time = 722186
Router advertisement = disabled
Proxy Arp = disabled
Administrative debounce-time = disabled
Operational debounce-time = disabled
```

```
Access routing = disabled
Multipath mode = hashed
```

```
In Received Packets 2, Bytes 256
  Unicast Packets 2, Bytes 256
  Multicast Packets 0, Bytes 0
In Policed Packets 0, Bytes 0
In Error Packets 0
In Invalid Source Address Packets 0
In Discarded Packets 0
Out Forwarded Packets 2, Bytes 256
  Unicast Packets 2, Bytes 256
  Multicast Routed Packets 0, Bytes 0
Out Scheduler Dropped Packets 0, Bytes 0
Out Policed Packets 0, Bytes 0
Out Discarded Packets 0
```

```
queue 0: traffic class best-effort, bound to ip ATM2/0.1
```

```
  Queue length 0 Bytes
  Forwarded packets 0, Bytes 0
  Dropped committed packets 0, Bytes 0
  Dropped conformed packets 0, Bytes 0
  Dropped exceeded packets 0, Bytes 0
```

```
queue 1: traffic class tcl, bound to ip ATM2/0.1
```

```
  Queue length 0 Bytes
  Forwarded packets 0, Bytes 0
  Dropped committed packets 0, Bytes 0
  Dropped conformed packets 0, Bytes 0
  Dropped exceeded packets 0, Bytes 0
```

```
queue 2: traffic class tcs4, bound to ip ATM2/0.1
```

```
  Queue length 0 Bytes
  Forwarded packets 0, Bytes 0
  Dropped committed packets 0, Bytes 0
  Dropped conformed packets 0, Bytes 0
  Dropped exceeded packets 0, Bytes 0
```

```
queue 3: traffic class tcs5, bound to ip ATM2/0.1
```

```
  Queue length 0 Bytes
  Forwarded packets 0, Bytes 0
  Dropped committed packets 0, Bytes 0
  Dropped conformed packets 0, Bytes 0
  Dropped exceeded packets 0, Bytes 0
```

show qos interface-hierarchy

- Use to display brief or default (normal) information about a queue profile.
- If you do not specify the queue profile name, data for all queue profiles is displayed.
- The brief format contains the name of the queue profile and the reference count. The reference count is the number of times the queue profile is referenced by a QoS profile.
- The normal format (no brief option) displays one line per traffic class. This is the default format. The format contains the name of the traffic class, the reference count, the weight, and information on whether or not strict priority is enabled or disabled.
- Field descriptions
 - › committed length – greater queue length than the length of the conformed or exceeded length
 - › conformed length – a queue length that is less than the committed length but greater than the exceeded length
 - › exceeded length – a queue length less than the conformed length which is less than the committed length
 - › conformed fraction – percentage of the total queue that can be occupied before dropping conformed packets
 - › exceeded fraction – percentage of the total queue that can be occupied before dropping exceeded packets
 - › buffer weight – the weight of the queue
- Example 1

```
host1#show queue-profile brief
queue-profile default referenced
```

- Example 2

This is the default (normal) format.

```
host1#show queue-profile
committed   conformed   exceeded   fraction:
queue      length:    length:    length:    conformed,  buffer
profile   min, max  min, max  min, max  exceeded    weight
-----   -
default  0, <none> 0, <none> 0, <none> 50, 25      8
```

show qos-port-type-profile

- Use to display the default (normal) information about a QoS port-type profile.
- If you do not specify the profile name, data for all interface types is displayed.
- The normal format contains a list of all the qos-port-type-profile commands as they have been entered. This is the default format.
- Example

```
host1#show qos-port-type-profile
default-port-profile Ethernet qos-profile ethernet-default
default-port-profile Atm qos-profile atm-default
```

```
default-port-profile HDLC qos-profile serial-default
default-port-profile ServerPort qos-profile server-default
```

show qos-profile

- Use to display brief or default (normal) information about a QoS profile.
- If you do not specify the QoS profile name, data for all QoS profiles is displayed.
- The brief format contains the name of the QoS profile and the reference count. The reference count is the number of times the QoS profile is referenced by an interface or protocol profile.
- The normal format (no brief or detail option) contains the name of the QoS profile, and the configuration parameters. This is the default format.
- The detailed format contains the name of the QoS profile and the reference count.
- Example 1

```
host1#show qos-profile brief
qos-profile atm-default referenced
qos-profile serial-default referenced
qos-profile ethernet-default referenced
qos-profile server-default referenced
```

- Example 2
This is the default (normal) format.

```
host1#show qos-profile
atm-default:
interface rule traffic scheduler queue drop t-class
  type type class profile profile profile group
-----
ip node default
atm-vc node default
cbf node default
ip queue best-effort default default default
atm queue best-effort default default default
atm-vc queue best-effort default default default
cbf queue best-effort default default default

serial-default:
interface rule traffic scheduler queue drop t-class
  type type class profile profile profile group
-----
ip node default
fr-vc node default
cbf node default
ip queue best-effort default default default
serial queue best-effort default default default
cbf queue best-effort default default default
```

```

ethernet-default:
interface  rule      traffic  scheduler  queue  drop  t-class
  type    type      class    profile    profile profile profile
-----  -
ip        node          default
vlan     node          default
ip       queue    best-effort  default    default  default
ethernet queue    best-effort  default    default  default

server-default:
interface  rule      traffic  scheduler  queue  drop  t-class
  type    type      class    profile    profile profile profile
-----  -
ip        node          default
l2tp-tunnel node          default
ip-tunnel node          default
ip       queue    best-effort  default    default  default
svrPort  queue    best-effort  default    default  default
l2tp-tunnel queue    best-effort  default    default  default
ip-tunnel queue    best-effort  default    default  default

```

show scheduler-profile

- Use to display brief or default (normal) information about a scheduler profile.
- If you do not specify the scheduler profile name, data for all scheduler profiles is displayed.
- The brief format contains the name of the scheduler profile and the reference count. The reference count is the number of times the traffic class is referenced by a QoS profile.
- The normal format (no brief or detail option) contains the name of the scheduler profile, the shaping rate, the weight, and information on whether or not strict priority is enabled or disabled. This is the default format.
- The detailed format contains the name of the scheduler profile, the shaping rate, the weight, information on whether or not strict priority is enabled or disabled, and the name of each QoS profile that references the scheduler profile.
- Example 1

```

host1#show scheduler-profile brief
scheduler-profile default referenced

```

- Example 2

This is the default (normal) format.

```

host1#show scheduler-profile
  shaping    strict
scheduler   rate      burst   weight  priority
-----
default     <none>  32767   8       no
  
```

show traffic-class

- Use to display brief or default (normal) information about a traffic class.
- If you do not specify the traffic class name, data for all traffic classes is displayed.
- The brief format contains the name of the traffic class and the reference count. The reference count is the number of times the traffic class is referenced by a QoS profile.
- The normal format (no brief or detail option) contains the name of the traffic class, the weight, and information on whether or not strict priority is enabled or disabled. This is the default format.
- The detailed format contains the name of the traffic class, the weight, information on whether or not strict priority is enabled or disabled, the name of each QoS profile that references the traffic class, and the name of each policy list that references the traffic class.
- Field descriptions
 - › referenced – QoS profiles that reference this traffic class
 - › fabric weight – weight of the queue in the fabric
 - › fabric strict priority – setting strict priority queues in the fabric

- Example 1

```

host1#show traffic class brief
traffic-class best-effort referenced
traffic-class tc1 referenced
traffic-class tc2
traffic-class tc3
traffic-class tcs4 referenced
traffic-class tcs5 referenced
  
```

- Example 2

This is the default (normal) format.

```

host1>show traffic-class
  fabric
traffic      fabric  strict
class  weight  priority
-----
best-effort    8      no
best-effort    8      no
tc1            8      no
tc2            8      no
  
```

tc3	8	no
tcs4	8	yes
tcs5	8	yes

show traffic-class-group

- Use to display the name of a traffic class group and the classes in the group.
- Example

```
host1#show traffic-class-group  
traffic-class-group sg  
  traffic-class tcs4  
  traffic-class tcs5
```

