Chapter 1

Quality of Service Overview

The quality of service (QoS) feature enables your E-series router to distinguish traffic with strict timing requirements from traffic that can tolerate delay, jitter, and loss.

QoS topics are discussed in the following sections:

QoS on the E-series Router Overview

QoS is a suite of features that configure queuing and scheduling on the forwarding path of the E-series router. QoS provides a level of predictability and control beyond the best-effort delivery that the router provides by default. Best-effort service provides packet transmission with no assurance of reliability, delay, jitter, or throughput.

QoS as developed for E-series routers 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 for which you can configure different QoS characteristics. The Juniper Networks QoS architecture extends DiffServ to support edge features such as high-density queuing.

The E-series router supports:

- IETF architecture for differentiated services
- Assured forwarding per-hop-behavior (PHB) groups
- Expedited forwarding PHB groups

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

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.

Related Topics

- For a list of related RFCs, see *Configuring QoS on the E-series Router* on page 8

QoS Audience

This topic collection contains configuration information for two types of QoS users: QoS administrators and QoS clients.

QoS administrators are responsible for implementing a QoS queuing architecture by defining drop profiles, queue profiles, scheduler profiles, QoS profiles, and QoS parameter definitions.

QoS clients are responsible for configuring services for individual subscribers by creating parameter instances. The parameter instances that QoS clients can create depend on the settings defined in parameter definitions by the QoS administrator.

Related Topics

- For information about QoS users and QoS parameters, see *QoS Parameter Audience* on page 222

QoS Platform Considerations

QoS is supported on all E-series line modules except for the ES2 10G Uplink LM.

Figure 1 shows the traffic flow through the router.

Figure 1: Traffic Flow Through an E-series Router

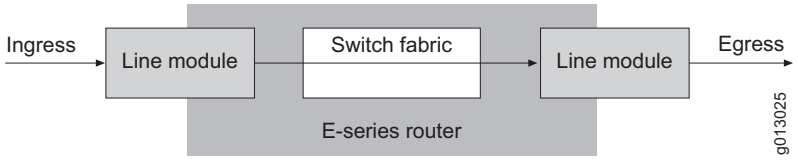

For information about the modules supported on E-series routers:

- See the *ERX Module Guide* for modules supported on ERX-7xx models, ERX-14xx models, and the ERX-310 router.
- See the *E120 and E320 Module Guide* for modules supported on the E120 router and the E320 router.

Interface Specifiers

The majority of the configuration task examples in this topic collection use the *slot*/*port* format to specify an interface. However, the interface specifier format that you use depends on the router that you are using.

For ERX-7xx models, ERX-14xx models, and ERX-310 routers, use the *slot*/*port* format. For example, the following command specifies an ATM interface on slot 0, port 1 of an ERX-7xx model, ERX-14xx model, or ERX-310 router.

```
host1(config)#interface gigabitEthernet 0/1
```

For E120 and E320 routers, use the *slot*/*adapter*/*port* format, which includes an identifier for the bay in which the I/O adapter (IOA) resides. In the software, adapter 0 identifies the right IOA bay (E120 router) and the upper IOA bay (E320 router); adapter 1 identifies the left IOA bay (E120 router) and the lower IOA bay (E320 router). For example, the following command specifies a10-Gigabit Ethernet interface on slot 5, adapter 0, port 0 of an E320 router.

```
host1(config)#interface tenGigabitEthernet 5/0/0
```

Related Topics

- For more information about supported interface types and specifiers on E-series routers, see *Interface Types and Specifiers* in *JUNOSe Command Reference Guide, About This Guide*

QoS Terms

Table 4 defines terms used in this discussion of QoS.

Table 4: QoS Terminology

Term	Description
Assured rate	Bandwidth guaranteed until the node below in the scheduler hierarchy is oversubscribed.
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.
Best-effort queue	For a logical interface, the queue associated with the best-effort traffic class for that logical interface,
Best-effort scheduler node	The scheduler node associated with a logical interface and traffic class group pair, and where the traffic class group contains the best-effort traffic class. Also known as best-effort node.

Table 4: QoS Terminology (continued)

Term	Description
CDV	Cell delay variation. Measures the difference between a cell's expected and actual transfer delay. Determines the amount of jitter.
CDVT	Cell delay variation tolerance. Specifies the acceptable tolerance of CDV (jitter).
Effective weight	The result of a weight or an assured rate. Users configure the scheduler node by specifying either an assured rate or a weight within a scheduler profile. An assured rate, in bits per second, is translated into a weight. The resultant weight is referred to as an effective weight.
Group node	A scheduler node associated with a {port interface, traffic-class group} pair. Because the logical interface is the port, only one such scheduler node can exist for each traffic-class group above the port. This node aggregates all traffic for traffic classes in the group.
HAR	Hierarchical assured rate. Dynamically adjusts bandwidth for scheduler nodes.
HRR	Hierarchical round-robin. Allocates bandwidth to queues in proportion to their weights.
Latency	Delay in the transmission of a packet through a network from beginning to end.
Proprietary QoS Management Information Base (MIB)	Supported on the E-series router.
Queue	First-in-first-out (FIFO) set of buffers that control packets on the data path.
QoS port-type profile	Supplies the QoS information for forwarding interfaces stacked above ports of the associated interface type.
QoS profile attachment	Applies the rules in the QoS profile to a specific interface.
Rate shaping	Allows you to throttle a queue to a specified rate.
RED	Random early detection congestion avoidance technique.
Scheduler hierarchy	A hierarchical, tree-like arrangement of scheduler nodes and queues. The router supports up to three levels of scheduler nodes stacked above a port. The port scheduler is at level 0, with two levels of scheduler nodes at levels 1 and 2. A final level of queues is stacked above the nodes.
Scheduler node	An element within the hierarchical scheduler that implements bandwidth controls for a group of queues. Queues are stacked above scheduler nodes in a hierarchy. The root node is associated with a channel or physical port.
Shaping rate	Bandwidth in a queue or node can be throttled to a specified rate.
Shared shaper constituent	All nodes and queues that are associated with a logical interface that is being shared shaped are considered potential constituents of the shared shaper.
Weight	Specifies the relative weight for queues in the traffic class.
WRED	Weighted random early detection congestion avoidance technique.

QoS Features

Table 5 describes the major QoS features supported on the E-series router.

Table 5: 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 are assigned to the best-effort traffic class.
Differentiated services	■ Assured forwarding—See RFC 2597. ■ Expedited forwarding—See RFC 2598.
Drop profile	Template that specifies active queue management in the form of WRED behavior of an egress queue.
Port shaping	Shapes the aggregate traffic through a port or channel to a rate that is less than the line or port rate.
QoS parameters	Creates a queuing architecture without the numeric subscriber rates and weights in scheduler profiles. You then use the same QoS and scheduler profiles across all subscribers who use the same services but at different bandwidths, reducing the total number of QoS profiles and scheduler profiles required.
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	Collection of QoS commands that specify queue profiles, drop profiles, scheduler profiles, and statistics profiles in combination with interface types.
Queue profile	Template that specifies the buffering and tail-dropping behavior of an egress queue.
Rate shaping	Mechanism that throttles the rate at which an interface can transmit packets. **Note:** Rate shaping as presented in policy management in releases before JUNOSe Release 4.0 is deprecated and converted to QoS profiles and scheduler profiles.
Relative strict-priority scheduling	Provides strict-priority scheduling within a shaped aggregate rate. For example, it lets you provide 1 Mbps of aggregate bandwidth to a subscriber, with up to 500 Kbps of the bandwidth for low-latency traffic. If there is no strict-priority traffic, the low-latency traffic can use up to the full aggregate rate of 1 Mbps.
Scheduler profile	Configures the bandwidth at which queues drain as a function of relative weight, assured rate, and shaping rate.
Shared rate shaping	Mechanism for shaping a logical interface's aggregate traffic to a rate when the traffic for that logical interface is queued through more than one scheduler hierarchy.
Statistics profile	Template that specifies rate statistics and event-gathering characteristics.
Strict-priority scheduling	Designates the traffic class (queue) that receives top priority for transmission of its packets through a port. It is implemented with a special strict-priority scheduler node that is stacked directly above the port.

Table 5: QoS Features (continued)

Feature	Description
Traffic class	A chassis-wide grouping of queues and buffers that support transmission of a designated set of traffic across the chassis, from ingress line module, through the switch fabric, and onto the egress line module. The router supports up to eight traffic classes, and therefore up to eight queues per logical interface.
Traffic-class group	Separate hierarchy of scheduler nodes and queues over a port. A traffic-class group uses one level of the scheduler hierarchy, level 1. Traffic classes belong to the default group unless they are specifically assigned to a named group. All queues are stacked in a single scheduler hierarchy above the physical port. When you configure a traffic class inside a group, its queues are stacked separately. The most common reason for creating separate scheduler hierarchies is to implement strict priority scheduling for all queues in the group. The router supports up to four traffic-class groups. A traffic class cannot belong to more than one group.
WRED	Signals end-to-end protocols such as TCP that the router is becoming congested along a particular egress path. The intent is to trigger TCP congestion avoidance in a random set of TCP flows before congestion becomes severe and causes tail dropping on a large number of flows.

Configuring QoS on the E-series Router

Several of the tasks for configuring QoS on your E-series router are optional.

To configure QoS on your E-series router:

1. Create and configure a traffic class.

   See *Chapter 2, Defining Service Levels with Traffic Classes and Traffic-Class Groups*.

2. (Optional) Create one or more traffic-class groups.

   See *Chapter 2, Defining Service Levels with Traffic Classes and Traffic-Class Groups*.

3. (Optional) To configure nondefault buffer management, create a queue profile.

   See *Chapter 3, Configuring Queue Profiles for Buffer Management*.

4. (Optional) To configure RED or WRED, create a drop profile.

   See *Chapter 4, Configuring Dropping Behavior with RED and WRED*.

5. (Optional) To gather rate statistics, create a statistics profile.

   See *Chapter 5, Gathering Statistics for Rates and Events in the Queue*.

6. Configure a scheduler hierarchy with a scheduler profile.

   See *Chapter 6, QoS Scheduler Hierarchy Overview*.

7. (Optional) Configure shaping:

   - Configure shaping and shared shaping using the scheduler profile.

     See *Chapter 7, Configuring Rates and Weights in the Scheduler Hierarchy*, *Chapter 10, Configuring Simple Shared Shaping of Traffic*, and *Chapter 12, Configuring Compound Shared Shaping of Traffic*.

   - Configure shaping rates independent of the QoS profile and scheduler profile using QoS parameters.

     See *Chapter 24, Configuring a QoS Parameter*.

8. Create a QoS profile. QoS profiles reference queue, drop, statistics, and scheduler profiles.

   See *Chapter 16, Configuring and Attaching QoS Profiles to an Interface*.

9. Attach the QoS profile to one or more interfaces, or specify the profile as a QoS port-type profile for a given interface type.

   See *Chapter 16, Configuring and Attaching QoS Profiles to an Interface*.

QoS References

For more information about QoS, see the following resources:

- RFC 2474—Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers (December 1998)
- RFC 2475—An Architecture for Differentiated Services (December 1998)
- RFC 2597—Assured Forwarding PHB Group (June 1999)
- RFC 2598—An Expedited Forwarding PHB (June 1999)
- RFC 2698—A Two Rate Three Color Marker (September 1999)
- RFC 2990—Next Steps for the IP QoS Architecture (November 2000)
- RFC 2998—A Framework for Integrated Services Operation over Diffserv Networks (November 2000)
- RFC 3246—An Expedited Forwarding PHB (Per-Hop Behavior) (March 2002)

- RFC 3260—New Terminology and Clarifications for Diffserv (April 2002)
- DSL Forum Technical Report (TR)-059—DSL Evolution - Architecture Requirements for the Support of QoS-Enabled IP Services
- Floyd, S., and Jacobson, V. Random Early Detection for Congestion Avoidance. *IEEE/ACM Transactions on Networking* 1(4), August 1993