

Queuing and Buffer Management Overview

A queue is a set of first-in, first-out (FIFO) buffers that buffer packets on the data path. QoS associates queues with a traffic class/interface pair. For example, if you create 4000 IP interfaces and configure each interface with four traffic classes, then 16,000 queues are created. For specific information about the maximum number of QoS queues supported, see *JUNOS Release Notes, Appendix A, System Maximums*.

The E Series router dynamically manages the shared memory on egress line modules to provide a good balance between sharing the memory among queues and protecting an individual queue's claim on its fair share of the egress memory.

When egress packet memory is in high demand and aggregate utilization of the packet memory 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 high bandwidth consumers cannot starve out moderate traffic consumers by allocating all the shared memory resource for themselves.

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

The router 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 router from setting queue lengths too low or too high.

Static Oversubscription

The router uses static oversubscription to vary queue thresholds based on the number of queues currently configured, which is relatively static. Static oversubscription is based on the assumption that, when a few queues are configured, many of the queues are likely to be active at the same time. When a large number of queues are configured, fewer queues are likely to be active at the same time.

When few queues are configured, buffer memory is strictly partitioned between queues to ensure that buffers are available for all queues. As the number of configured queues increases, buffer memory is increasingly oversubscribed to allow more buffer sharing. Reserving buffer space for all queues when many are expected to be idle is unnecessary and wasteful.

Dynamic Oversubscription

The router uses dynamic oversubscription to vary queue thresholds based on the amount of egress buffer memory in use. The router divides egress buffer memory into eight regions.

The size of the region depends on the ASIC type. For more information, see *Memory Requirements for Queue and Buffers*.

When buffer memory is in low demand, queues are given large amounts of buffer memory. As the demand for buffer memory increases, queues are given progressively smaller amounts of buffer memory.

Color-Based Thresholding

Packets within the router are tagged with a drop precedence:

- Committed—Green
- Conformed—Yellow
- Exceeded—Red

When the queue fills above the exceeded threshold, the router drops red packets, but still queues yellow and green packets. When the queue fills above the conformed drop threshold, the router queues only green packets.



NOTE: All color-based thresholds vary in proportion to the dynamic queue length.

- Related Topics**
- Configuring Queue Profiles to Manage Buffers and Thresholds
 - Guidelines for Managing Queue Thresholds
 - Guidelines for Managing Buffers
 - RED and WRED Overview

Published: 2010-01-12