

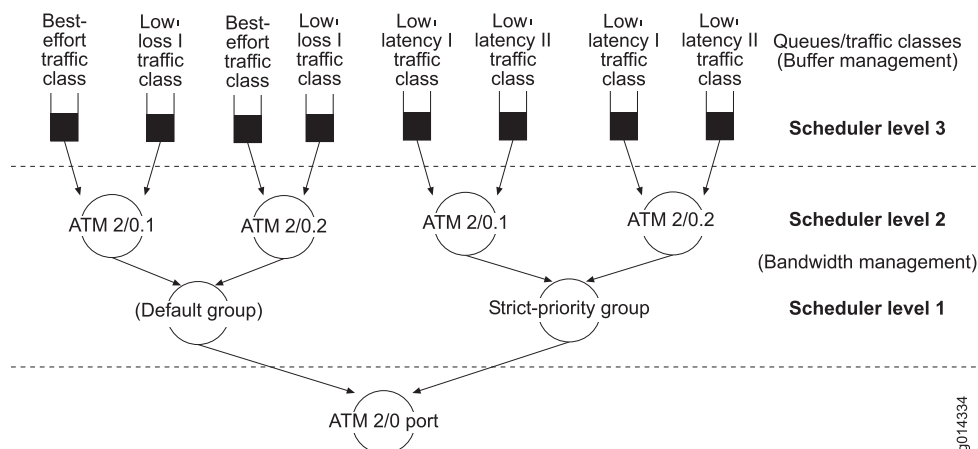
Scheduler Hierarchy Overview

The egress line module scheduler is an HRR scheduler. Figure 1 on page 1 is an example of a QoS scheduler's hierarchy.

As shown in Figure 1 on page 1, the queues feeding a physical port are organized in a hierarchy. At each level in the hierarchy, the scheduler uses shaping rates, hierarchical or assured rates, and relative weights to determine the allocated bandwidth:

- The scheduler selects a first-level node based on the allocated bandwidth.
- The scheduler then selects a second-level node from the group of nodes that are stacked above the selected first-level node. This selection is also based on the allocated bandwidth.
- Finally, the scheduler selects a queue from the group of queues stacked above the second-level node.

Figure 1: QoS Scheduler Hierarchy



Shaping Rates, Assured Rates, and Relative Weights in a Scheduler Hierarchy

The scheduler supports hierarchical and static assured rates, relative weights, and shaping rates 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 the shaping rate and either the assured rate or relative weight:

- When the scheduler is not congested, the shaping rates determine which node or queue can claim the bandwidth. The shaping rate specifies the maximum bandwidth to the node or queue.
- When the scheduler is congested, either the hierarchical or static assured rate or the weight specifies the minimum bandwidth.
 - If the scheduler is configured to use a static assured rate and the assured rate is other than none (the default), it is used to determine the allocated bandwidth, and the weight setting is ignored. If the assured rate is zero, the weight setting is used to determine the bandwidth.

The static assured rate specifies the desired bandwidth. This rate is guaranteed until the bandwidth becomes oversubscribed.

- If the scheduler is configured to use hierarchical assured rate, the scheduler dynamically adjusts the amount of allocated bandwidth for service delivery based on the sum of the assured rates of all child nodes and queues.
- The assured rate also specifies that if bandwidth is over- or undersubscribed, all adjustments are made in proportion to the original assured-rate specification.

For example, if Node A is configured to receive 40 Mbps and Node B receives 20 Mbps, any available bandwidth above the subscribed total of 60 Mbps would be allocated to the two nodes at the same 2-to-1 ratio. Similarly, if the bandwidth were oversubscribed and only 30 Mbps were available, this amount would also be allocated to the two nodes at the 2-to-1 ratio, with Node A getting 20 Mbps and Node B getting 10 Mbps.



NOTE: For E-series ASIC modules, strict priority is supported only for a single first-level scheduler node.

When determining the shaping rate, the system includes all bytes in Layer 2 encapsulations. The packets that are included in the rate depend on the Layer 2 node that is specified in the QoS profile. For example, the shaping rate for an Ethernet node includes bytes from the Ethernet and VLAN encapsulations.

Related Topics

- Static and Hierarchical Assured Rate Overview
- Rate Shaping and Port Shaping Overview
- Shared Shaping Overview
- Configuring a Scheduler Hierarchy