

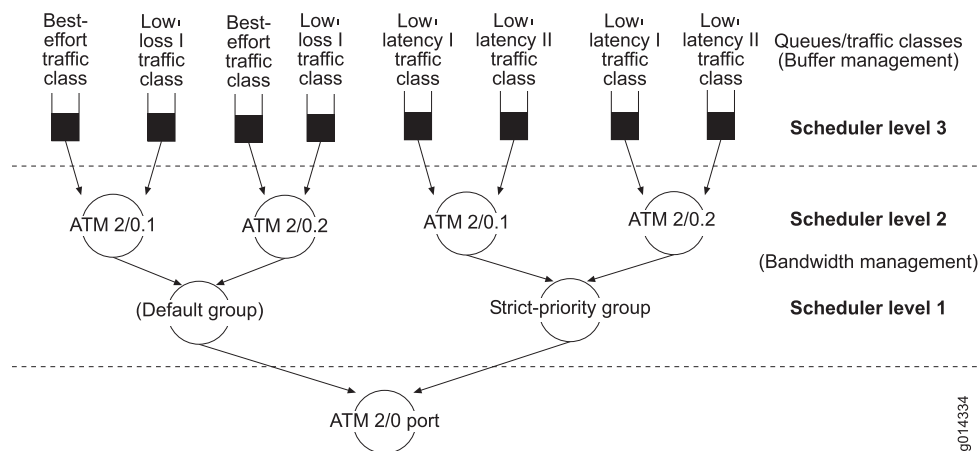
Strict-Priority and Relative Strict-Priority Scheduling Overview

You 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 bandwidth before other queues.

You can configure only one node at the first scheduler level 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 HRR algorithm selects which strict-priority queue is scheduled next.

Figure 1 illustrates an example of a QoS scheduler's hierarchy.

Figure 1: Sample Strict-Priority Scheduling Hierarchy



One strict priority traffic-class group is called the auto-strict-priority group. The scheduler nodes and queues in the auto-strict-priority group receive strict-priority scheduling. If multiple queues above the strict-priority node have packets, the HRR algorithm selects which strict-priority queue is scheduled next.



NOTE: If you configured traffic shaping through traffic shape profiles in JUNOS releases before Release 4.0, traffic shaping is replaced with the rate-shaping feature, which is configured when you configure a scheduler profile.

Relative Strict-Priority Scheduling Overview

Relative strict-priority scheduling provides strict-priority scheduling within a shaped aggregate rate. For example, it allows you to 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.

Relative strict priority differs from true strict priority in that it can implement the aggregate shaping rate for both strict and nonstrict traffic. With true strict priority, you can shape the nonstrict or the strict traffic separately, but you cannot shape the aggregate to a single rate.

The best application of relative strict priority is on Ethernet, where you can shape the aggregate for each VLAN to a specified rate, and provision a strict and nonstrict queue for each VLAN above the shaped VLAN node.

To use relative strict priority, you configure strict-priority queues above the VC or VLAN scheduler node, thereby providing for strict-priority scheduling of the queues within the VC or VLAN. You configure relative strict priority without using QoS traffic-class groups, which causes strict-priority queues to appear in the same scheduler hierarchy as the nonstrict queues.

Relative strict priority provides low latency only if you undersubscribe the port by shaping all VCs on the port so that the sum of the shaping rates is less than the port rate. The port will not become congested, and the latency caused by the round-robin behavior of both the HRR and cell schedulers is nominal. In these undersubscribed conditions, the latency of a strict-priority queue within each VC is calculated as if the VC were draining onto a wire with bandwidth equal to the shaped rate.

Relative strict priority is carried out in the HRR scheduler on E Series ASIC line modules.

- Related Topics**
- Comparison of True Strict Priority with Relative Strict Priority Scheduling
 - Configuring Strict-Priority Scheduling
 - Configuring Relative Strict-Priority Scheduling for Aggregate Shaping Rates

Published: 2010-03-21