

Understanding CoS Schedulers

You use schedulers to define the properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the drop profiles associated with the queue.

You associate the schedulers with forwarding classes by means of scheduler maps. You can then associate each scheduler map with an interface, thereby configuring the queues, packet schedulers, and tail drop processes that operate according to this mapping.

- Default Schedulers on page 1
- Transmission Rate on page 2
- Scheduler Buffer Size on page 2
- Priority Scheduling on page 3
- Scheduler Drop-Profile Maps on page 3
- Scheduler Maps on page 3

Default Schedulers

Each forwarding class has an associated scheduler priority. Only two forwarding classes, best-effort (queue0) and network-control (queue7) are used in the default configuration.



NOTE: On EX 8200 series switches three forwarding classes—best-effort (queue0), multicast best-effort (queue2), and network-control (queue7)—are used in the default configuration.

By default, the best-effort forwarding class (queue 0) receives 95 percent of the bandwidth and buffer space for the output link, and the network-control forwarding class (queue 7) receives 5 percent. The default drop profile causes the buffer to fill completely and then to discard all incoming packets until it has free space.



NOTE: On EX 8200 series switches, by default, the best-effort forwarding class (queue 0) receives 75 percent of the bandwidth, the multicast best-effort forwarding class (queue 2) receives 20 percent of the bandwidth and buffer space for the output link, and the network-control forwarding class (queue 7) receives 5 percent.

The expedited-forwarding and assured-forwarding classes have no scheduler because no resources are assigned to queue 5 and queue 1, by default. However, you can manually configure resources for the expedited-forwarding and assured-forwarding classes.

Also by default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues. When a forwarding class does not fully use the allocated transmission bandwidth, the remaining bandwidth can be used by other forwarding classes if they receive a larger amount of offered load than their allocated bandwidth allows.

Transmission Rate

The transmission-rate control determines the actual traffic bandwidth from each forwarding class you configure. The rate is specified in bits per second. Each queue is allocated some portion of the bandwidth of the outgoing interface.

This bandwidth amount can be a fixed value, such as 1 megabit per second (Mbps), a percentage of the total available bandwidth, or the rest of the available bandwidth. You can allow transmission bandwidth to exceed the configured rate if additional bandwidth is available from other queues. In case of congestion, configured amount of transmission rate is guaranteed for the queue. This property allows you to ensure that each queue receives the amount of bandwidth appropriate to its level of service.

Scheduler Buffer Size

To control congestion at the output stage, you can configure the delay-buffer bandwidth. The delay-buffer bandwidth provides packet buffer space to absorb burst traffic up to the specified duration of delay. Once the specified delay buffer becomes full, packets with 100 percent drop probability are dropped from the tail of the buffer.

The default scheduler transmission rate for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent of the total available bandwidth. The default buffer-size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent of the total available buffer.



NOTE: On EX 8200 series switches, the default scheduler transmission rates for queues 0 through 7 are 75, 0, 20, 0, 0, 0, 0, and 5 percent of the total available bandwidth. The default buffer-size percentages for queues 0 through 7 are 75, 0, 20, 0, 0, 0, 0, and 5 percent of the total available buffer.

For each scheduler, you can configure the buffer size as one of the following:

- A percentage of the total buffer.
- The remaining buffer available. The remainder is the buffer percentage that is not assigned to other queues. For example, if you assign 40 percent of the delay buffer to queue 0, allow queue 2 to keep the default allotment of 20 percent, allow queue 7 to keep the default allotment of 5 percent, and assign the remainder to queue 3, then queue 3 uses approximately 35 percent of the delay buffer.

Priority Scheduling

Priority scheduling determines the order in which an output interface transmits traffic from the queues, thus ensuring that queues containing important traffic are provided better access to the outgoing interface.

Priority scheduling is accomplished through a procedure in which the scheduler examines the priority of the queue. JUNOS software supports two levels of transmission priority:

- **Low**—The scheduler determines if the individual queue is within its defined bandwidth profile. This binary decision, which is reevaluated on a regular time cycle, compares the amount of data transmitted by the queue against the amount of bandwidth allocated to it by the scheduler. When the transmitted amount is less than the allocated amount, the queue is considered to be in profile. A queue is out of profile when its transmitted amount is larger than its allocated amount. Out of profile queue will be transmitted only if bandwidth is available. Otherwise, it will be buffered.

A queue from the set is selected based on the shaped deficit weighted round robin (SDWRR) algorithm, which operates within the set.

- **Strict-high**—Strict-high priority queue receives preferential treatment over low priority queue. Unlimited bandwidth is assigned to strict-high priority queue. Queues are scheduled according to the queue number, starting with the highest queue 7, with decreasing priority down through queue 0. Traffic in higher queue numbers is always scheduled prior to traffic in lower queue numbers. In other words, in case of two high priority queues, the queue with higher queue number is processed first.

Packets in low priority queues are transmitted only when strict-high priority queues are empty.

Scheduler Drop-Profile Maps

Drop-profile maps associate drop profiles with a scheduler. Drop-profile map sets the drop profile for a specific packet loss priority (PLP) and protocol type. The inputs for the drop-profile map are the PLP and the protocol type. The output is the drop profile.

Scheduler Maps

A scheduler map associates a specified forwarding class with a scheduler configuration. After configuring a scheduler, you must include it in a scheduler map and then associate the scheduler map with an output interface.

EX-series switches allow you to associate up to four user-defined scheduler maps with interfaces.

- Related Topics**
- Understanding JUNOS CoS Components for EX-series Switches
 - Example: Configuring CoS on EX-series Switches

- Defining CoS Schedulers (CLI Procedure)
- Defining CoS Schedulers (J-Web Procedure)