

3) Policing and Remarking Best Practices

There is little reason to forward unwanted traffic only to police and drop it at a downstream node. Therefore, it is recommended to police traffic flows as close to their sources as possible.

Whenever supported, markdown should be done according to standards-based rules, such as RFC 2597, The Assured Forwarding PHB. For example, excess traffic marked to AFx1 should be marked down to AFx2 (or AFx3 whenever dual-rate policing—such as defined in RFC 2698—is supported). Following such markdowns, congestion management policies, such as DSCP-based Weighted Random Early Detection (WRED), should be configured to drop AFx3 more aggressively than AFx2, which in turn should be dropped more aggressively than AFx1.

4) Queuing and Dropping Best Practices

Business-critical applications require service guarantees regardless of network conditions. The only way to provide service guarantees is to enable queuing at any and every node that has the potential for congestion.

In addition, because each application class has unique service level requirements, each should optimally be assigned a dedicated queue. In such a manner, specific bandwidth allocations and dropping policies can be assigned to each discrete application class to meet its distinctive QoS requirements. Otherwise, if multiple application classes are assigned into a common queuing bucket, the administrator no longer can control if bandwidth resources are being shared among these application classes according to their individual requirements.

At a minimum, however, the following standards-based queuing behaviors should be supported:

- Real-time queue(s)-to support an RFC 3246 Expedite Forwarding service
- Guaranteed-bandwidth queue(s)-to support RFC 2597 Assured Forwarding services
- Default queue-to support an RFC 2474 Default Forwarding service
- Bandwidth-constrained queue-to support an RFC 3662 “Scavenger” service

Cisco offers design recommendations for each of these types of queues. These queuing best practices are illustrated in Figure 1.

The **Real-Time Queue** corresponds to the RFC 3246 EF PHB. The amount of bandwidth assigned to the real-time queue is usually variable. However, if the majority of bandwidth is provisioned with strict-priority queuing (which is effectively a first-in, first-out [FIFO] queue), the overall effect is a dampening of QoS functionality. Remember the goal of convergence is to enable voice, video, and data applications to transparently coexist on a single network. When real-time applications dominate a link, non-real-time applications fluctuate significantly in their response times, destroying the transparency of the converged network.

Cisco has done extensive testing and has found that a significant decrease in non-real-time application response times occurs when real-time traffic exceeds one-third of link bandwidth capacity. In fact, both testing and customer deployments have shown that a general best queuing practice is to **limit the amount of strict-priority queuing to 33% of link bandwidth capacity**. This strict priority queuing recommendation is a conservative and safe design ratio for merging real-time applications with data applications.

Finally, WRED—or any similar congestion avoidance mechanism—should never be enabled on the strict-priority queue. Traffic assigned to this queue is often highly drop sensitive; therefore, early dropping should never be induced on these flows.

At least one queue should be provisioned as an **Assured Forwarding Queue**. Per RFC 2597, up to four queues can be provisioned with this service:

- AF Class 1-AF11, AF12, AF13
- AF Class 2-AF21, AF22, AF23
- AF Class 3-AF31, AF32, AF33
- AF Class 4-AF41, AF42, AF43

These queues should have bandwidth guarantees that correspond with the application class requirements of the traffic assigned to it.

In addition, DSCP-based WRED should be enabled on these queues, such that traffic marked AFx3 is (statistically) dropped sooner and more often than AFx2, which in turn is (statistically) dropped more aggressively than AFx1.

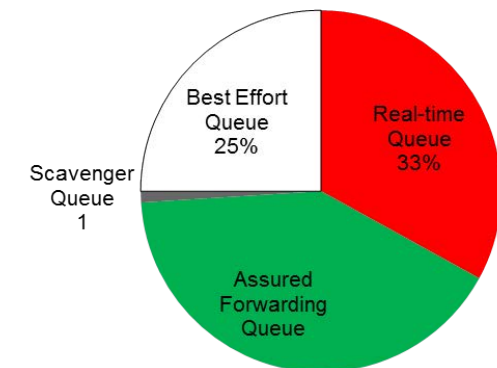
The **Best Effort Queue** is the default treatment for all traffic that has not been explicitly assigned to another queue. Only if an application has been selected for preferential/deferential treatment is it removed from the default class. Because most enterprises have several thousand applications running over their networks, adequate bandwidth must be provisioned for this class as a whole to handle the sheer number and volume of applications that default to it. Therefore, Cisco recommends **provisioning at least 25% of link bandwidth for the default Best Effort class**.

In addition, WRED is recommended to be enabled on the default class to improve throughput and reduce TCP synchronization. Because all traffic destined to this class is to be marked to the same DSCP value (of 0), there is no “weight” component to the WRED dropping decision, and therefore the congestion algorithm is effectively random early detect (RED).

Whenever the **Scavenger Queue** is enabled, it should be assigned a **minimal amount of bandwidth, such as 1%** (or whatever the minimal bandwidth allocation that the platform supports).

WRED is not required on the Scavenger class queue because traffic assigned to this queue has no implied “good-faith” service guarantee or expectation. Therefore, there is little to gain by adding this feature and it may even be wasteful of router CPU resources.

Figure 1 Queuing Best Practices



For more details, see:

http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSIntro_40.html

And the Cisco Press Book: **End-to-End QoS Network Design** (Second Edition)-Chapter 11