

Shared Shaping Overview

In the JUNOS software QoS implementation, you configure a traffic-class group to create a separate scheduler hierarchy. Traffic classes in a traffic-class group are queued through a scheduler hierarchy dedicated to that group. QoS supports up to five user-configurable, named traffic-class groups. Traffic classes that do not belong to any named group belong to the default traffic-class group. With the factory default configuration, the best-effort traffic class is in the default traffic-class group.

Shared shaping is a mechanism for shaping a logical interface's aggregate traffic to a rate when the traffic for that logical interface is queued through more than one scheduler hierarchy. For example, a service provider can configure QoS for voice, video, and data traffic on a single ATM VC. The video traffic and the voice traffic are placed in separate scheduler hierarchies from the data traffic to provision the low latency that is required for voice traffic and the higher bandwidth that is required for video traffic.

In this scenario, the data traffic needs to be dynamically shaped so that its rate matches the bandwidth available after the voice and video bandwidth requirements are met. When less voice and video traffic is being forwarded, then the data traffic can expand to fill the line rate.

When determining a shared shaping rate, the system includes all bytes in Layer 2 encapsulations. The packets that are included in the rate depend on the node specified. For example, rates for an Ethernet node include the Ethernet and VLAN encapsulations.

Shared shaping is typically enabled on the access-facing line module, but you can enable the feature for any interface type recognized by QoS, on any line module and any E Series Broadband Services Routers.

- Related Topics**
- Simple Shared Shaping Overview
 - Compound Shared Shaping Overview

Published: 2010-03-24