

Simple Shared Shaping Overview

Simple shared shaping shapes the best-effort node or queue associated with a logical interface to a shared rate.

Bandwidth Allocation for Simple Shared Shaping

Once per second, the simple shared shaper calculates the combined rate of the voice and video queues for the logical interface, and shapes the best-effort queue for the data traffic to the shared rate minus the video and voice queue rates. The bandwidth for the voice and video queues is determined by the configuration of the hierarchical scheduler. The shared shaper does not actively manage the video and voice queues.

Simple Shared Shaping on the Best-Effort Scheduler Node

If you have a second traffic class for data in addition to the best-effort data traffic class, configure shared shaping on the best-effort scheduler node. In this scenario, two weighted queues are stacked above the best-effort scheduler node, one for the best-effort traffic class and the other for the second data traffic class. If you configure the shared-shaping rate on the best-effort queue, then the shared shaper can have a tendency to starve the best-effort queue in favor of the second data queue. If you instead configure the shared-shaping rate on the best-effort node, the hierarchical scheduler allocates bandwidth between multiple data queues based on their relative weight and assured rate.

If you are configuring VP shared shaping, configure shared shaping on the best-effort scheduler node for the VP. Shaping the best-effort scheduler node for the VP has the effect of shaping all the VC best-effort queues for that VP. This enables you to retain the advantages of per-VC queuing in the hierarchical scheduler.

If you are configuring VC shared shaping and the SAR is operating in low-CDV mode, we recommend you configure the shared-shaping rate on the best-effort scheduler node for the VP or VC. The router sets the SAR shaper for the VC or VP to match the shared-shaping rate on VC and VP nodes in the hierarchical scheduler; this is usually the desired behavior. A shared shaper configured on the best-effort queue does not trigger the matching shaper in the SAR.

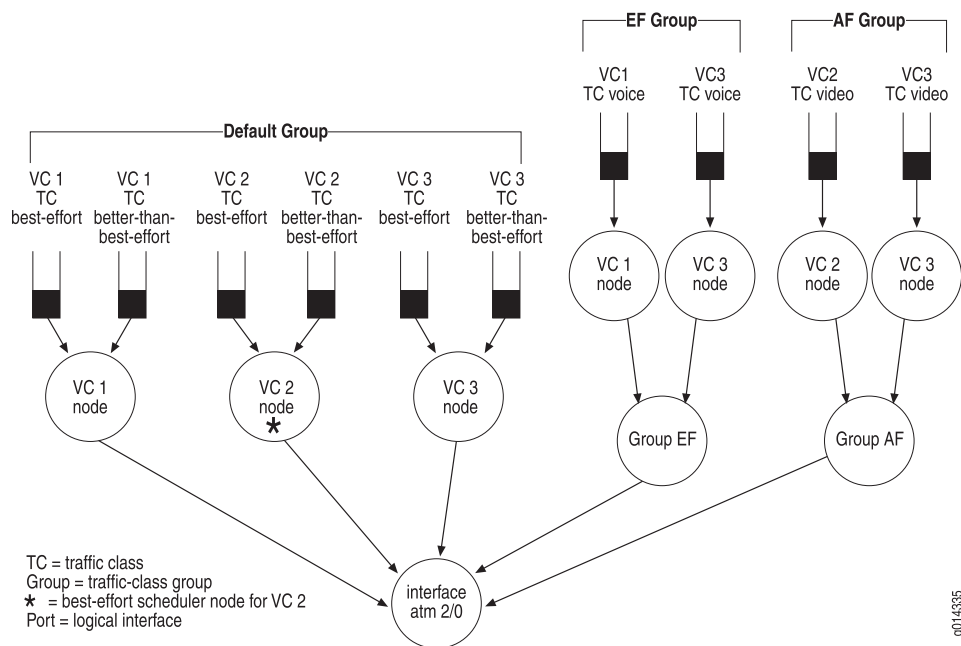
Simple Shared Shaping for Triple-Play Networks

Simple shared shaping enables you to shape the logical interface to a single rate for triple-play networks.

In Figure 1 on page 2, the AF traffic-class group contains the video traffic class. The EF traffic-class group contains the voice traffic class. The best-effort and better-than-best-effort traffic classes remain outside any traffic-class group. Because the voice, video, and data queues are stacked in separate scheduler hierarchies, you must use the shared shaper to shape the logical interface aggregate to a single rate.

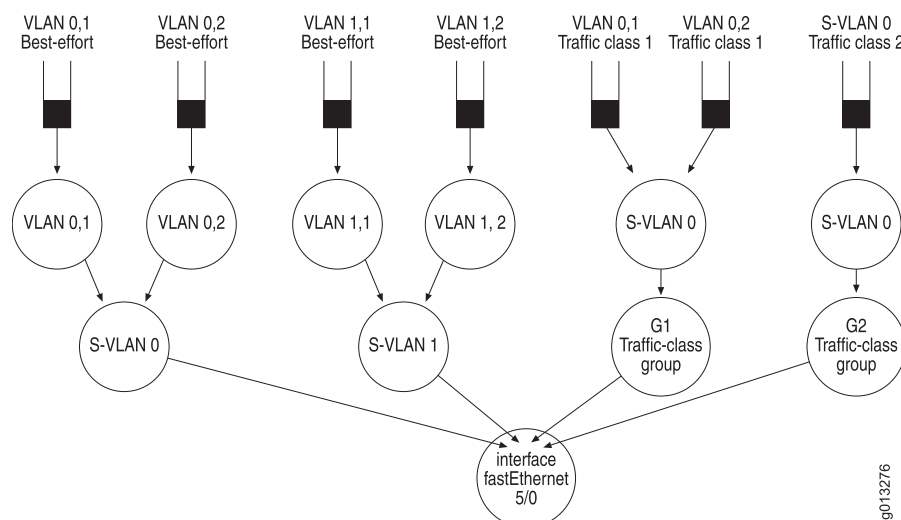
In this example, VC 1 is configured for voice and data. VC 2 is configured for data and video. VC 3 is configured for data, voice, and video. The shared shaper is configured on the best-effort node or queue for VC 1; the corresponding voice queue for VC 1 shares the configured rate.

Figure 1: Simple Shared Shaping over ATM



In a typical triple-play network configuration over Ethernet, individual subscribers are represented on the B-RAS by VLANs and DSLAMs by SVLANs. Figure 2 on page 2 illustrates how to shape the subscriber aggregate of voice, video, and data to a single rate in Ethernet.

Figure 2: Simple Shared Shaping over Ethernet



- Related Topics**
- For a list of shared shaper terms, see Shared Shaping Overview
 - Configuring Simple Shared Shaping
 - Constituent Selection for Shared Shaping Overview