

Shadow Nodes and Scheduler Behavior

You can configure *shadow nodes* when you want to explicitly set the queues at the required scheduler level for any line module with the EFA, EFA2, FFA, or TFA hardware. Shadow nodes enable you to specify the weight and the shaping rate of the added node. Shadow nodes can also conserve scheduler node resources.

You define the shadow node by referencing the shadow node in the QoS profile. Like phantom nodes, the router creates shadow nodes only when the additional node is required to meet the proper queue level.

The router creates shadow nodes after all the nodes and group nodes are created, and only when a node of the same interface type has existed in the same group of the scheduler hierarchy.

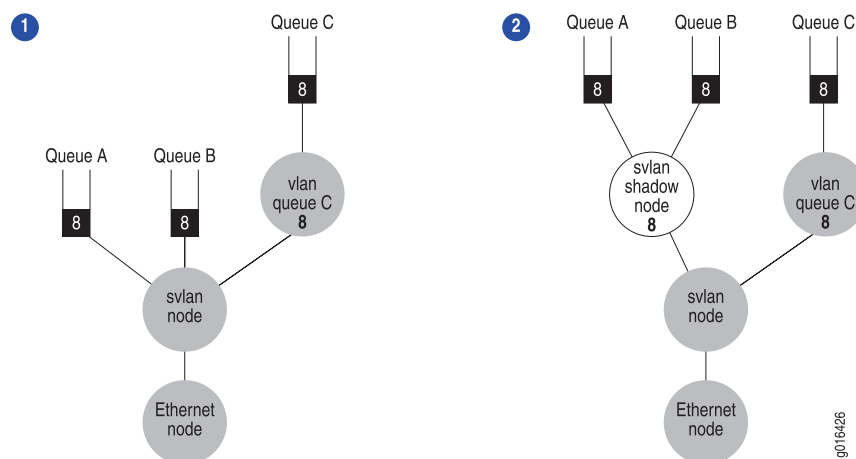
Shadow nodes can be configured for all interface types available for nodes.



NOTE: Shadow nodes ignore any shared-shaping rates in a scheduler profile.

Figure 1 on page 1 compares a scheduler hierarchy with and without shadow nodes.

Figure 1: Shadow Nodes



Unlike phantom nodes, shadow nodes can alter the behavior of the scheduler.

The first scheduler hierarchy in Figure 1 on page 1 shows VLAN interfaces A, B and C stacked above the same S-VLAN interface. Interfaces A and B have the same scheduler hierarchy (referencing qos-profile AB) and have a VLAN queue stacked directly above the S-VLAN node. In this case, VLAN interfaces A, B and C share the same 33 percent bandwidth available to the S-VLAN node.

Interface C has a VLAN queue stacked above a VLAN node and the S-VLAN node (referencing qos-profile C).

Specifying a shadow node forces the VLAN queue to the proper scheduler level. The second scheduler hierarchy in Figure 1 on page 1 shows the shadow node that is applied after QoS profile AB-shadow is assigned to interfaces A and B. As a result, interfaces A and B have 25 percent of the S-VLAN bandwidth and interface C has 50 percent of the S-VLAN bandwidth.

The S-VLAN shadow node uses the same scheduler profile as the queue.

To provide interfaces A and B with the proper weight, configure the weight of the shadow node to the sum of its queue weight. You can use hierarchical parameter instances and weight expressions to configure an appropriate weight.

- Related Topics**
- For a list of interface types supported for shadow nodes, see [Supported Interface Types for QoS Profiles](#)
 - For more information about hierarchical parameters, see [Hierarchical QoS Parameters Overview](#)