Applying QoS in the Campus

Need for QoS in the Campus

This topic describes the need for QoS measures in a campus network.

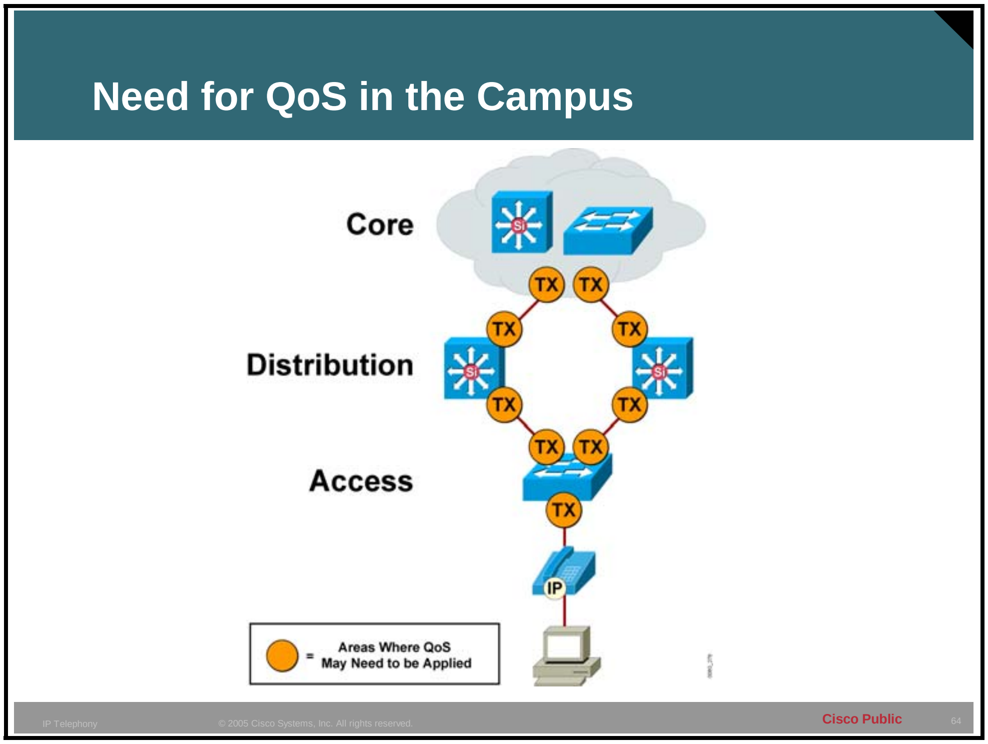

The bursty nature of data networks, combined with the high volume of smaller TCP packets, causes transmit buffers to fill to capacity in high-speed campus networks. If an output buffer fills, ingress interfaces are unable to place new flow traffic into the output buffer. When the ingress buffer is filled, which can happen quickly, packet drops occur. Typically, these drops are more than a single packet in any given flow. TCP packet loss causes retransmissions to occur, further aggravating the problem. Voice packet loss results in voice clipping and skips.

The figure illustrates the areas of a campus network where QoS is a concern. These areas are as follows:

- **Access ports:** Access-level ports that connect to IP Phones or gateways and the distribution switches
- **Distribution ports:** Uplink ports that connect to access switches and core switches
- **Core ports:** Uplink ports that connect the distribution switches to the core campus switches

VoIP traffic is sensitive to both delayed packets and dropped packets. Even though a campus may be using gigabit Ethernet trunks, which have extremely fast serialization times, delay and dropped packets may still be an issue because of the type of packets, type of applications, and number of packets. However, drops *always* adversely affect voice quality in the campus. Using multiple queues on transmit interfaces is the only way to eliminate the potential for dropped traffic caused by buffers operating at 100 percent capacity. By separating voice (which is sensitive to delays and drops) into its own queue, you can prevent flows from being dropped at the ingress interface even if data flows fill up the data transmit buffer. If video is present, it should also have its own queue because it has similar delay characteristics.

Caution It is critical that you verify that flow control is disabled when enabling QoS (multiple queues) on Catalyst switches. Flow control interferes with the configured queuing behavior by acting on the ports before activation of queuing. Flow control is disabled by default.

Marking Control and Management Traffic

Marking control and management traffic is critical to VoIP network performance. Control and management traffic should be marked for preferential dispatching throughout the campus. This topic describes marking control and management traffic as an approach to the management of QoS in a campus network.

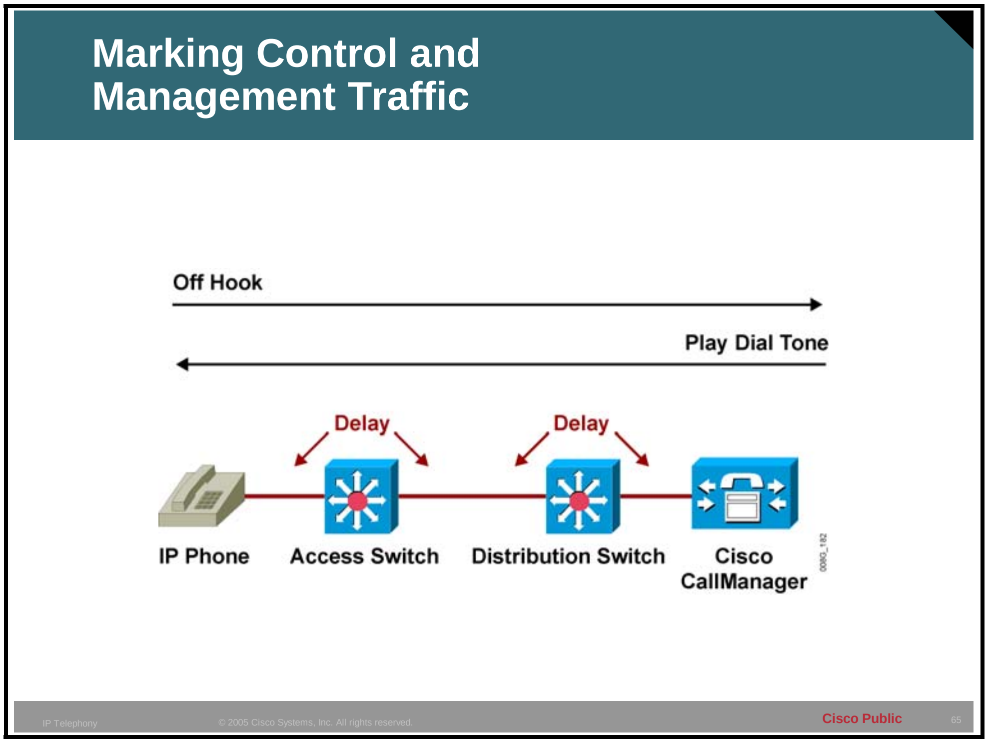

In networks with high traffic loads, managing the delivery of control traffic is critical to ensuring a positive user experience with VoIP. Delay to Dial-Tone (DTT) time is an example of this type of delivery management. When a Cisco IP Phone goes off hook, it “asks” Cisco CallManager for instructions. Cisco CallManager then instructs the Cisco IP Phone to play a dial tone. If this management and control traffic is dropped or delayed within the network, the user experience is adversely affected. The same logic applies to all signaling traffic for gateways and telephones.

.

Configuring a Voice VLAN

```
Router# configure terminal
Router(config)# interface fastethernet 5/1
Router(config-if)# switchport voice vlan 101
Router(config-if)# exit
```

  Cisco Public

The figure shows how to configure Fast Ethernet port 5/1 to send Cisco Discovery Protocol (CDP) packets that tell the Cisco IP Phone to use VLAN 101 as the voice VLAN.

Verifying the Configuration

```
Router# show interfaces fastethernet 5/1 switchport
Name: Fa5/1
Switchport: Enabled
Administrative Mode: access
Operational Mode: access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: off
Access Mode VLAN: 100
Voice VLAN: 101
Trunking Native Mode VLAN: 1 (default)
Administrative private-vlan host-association: none
Administrative private-vlan mapping: 900 ((Inactive)) 901
((Inactive))
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL
```

  Cisco Public

This example shows how to verify the configuration of Fast Ethernet port 5/1.

Switchwide Queuing

Configuration for switchwide queuing depends upon the type of interface used. The two types of interfaces used are:

- **Receive interface:** The recommended configuration for the receive interface is one standard FIFO queue.
- **Transmit interface:** The recommended configuration for the transmit interface is two standard queues with a single threshold (2Q1T). This method is preferred over other methods because it places traffic with specific CoS in one queue (as configured) and all other traffic in the second queue, thus simplifying configuration. The threshold indicates that traffic is dropped when the buffer capacity reaches 100 percent. Scheduling is done on a WRR basis. Admission to the queues is based on the 802.1p CoS value and is user-configurable in pairs. If you enable QoS, but do not modify the CoS-to-transmit queue mappings, switch performance could be affected because all traffic goes to queue 1.

Note 802.1p is a Layer 2 mechanism for classifying traffic using 3 bits. 802.1p adds 16 bits to the Layer 2 header.

IP Phone Port Queuing

In Cisco Catalyst software Release 5.5.1, the Catalyst 4000 line does not offer any advanced IP Phone-queuing features. Because of this, the Catalyst 4000 depends on the default CoS marking and enforcement of the Cisco IP Phone. Other access switches, such as the Catalyst 3500, offer more advanced IP port phone-queuing options.

Uplink Interface to the Distribution Switch

You do not need to configure special queuing or scheduling commands on the Catalyst 4000 side of the link (from the access layer Catalyst 4000 to the distribution layer Catalyst 6000). Queuing is automatically enabled after QoS has been enabled and the classification and queue admission have been configured.