Campus QoS Design

The Case for QoS in Campus Networks

The primary role of QoS in campus networks is not to control latency or jitter (as it is in the WAN/VPN), but to manage packet loss. In GE/10GE campus networks, it takes only a few milliseconds of congestion to cause instantaneous buffer overruns resulting in packet drops. Rich media applications—particularly HD video applications—are extremely sensitive to packet drops, to the point where even 1 packet dropped in 10,000 is discernible by the end-user.

Classification, marking, policing, queuing, and congestion avoidance are therefore critical QoS functions that are optimally performed within the campus network.

Four QoS design principles that apply to campus QoS deployments include:

- Always perform QoS in hardware rather than software when a choice exists.
- Classify and mark applications as close to their sources as technically and administratively feasible.
- Police unwanted traffic flows as close to their sources as possible.
- Enable queuing policies at every node where the potential for congestion exists.

Campus QoS Design Considerations

There are several considerations that impact QoS designs within the campus:

- Global Default QoS Setting
- Trust States and Conditional Trust
- Per-Port QoS., Per-VLAN QoS, Per-Port/Per-VLAN QoS
- Ingress QoS Models
- Egress QoS Models
- EtherChannel QoS
- QoS Roles in a campus
- AutoQoS

Global Default QoS Setting

On some platforms QoS is globally disabled by default (such as the Cisco Catalyst 2960/3650/3750). A fundamental first step is to globally enable QoS on these platforms.

Trust States

A switch port that is set to trust will accept and preserve either Layer 2 or Layer 3 packet markings. There are four static trust states with which a switch port may be configured:

- Untrusted—The default state with QoS enabled
- Trust CoS—Accepts Layer 2 802.1P CoS markings
- Trust IP Precedence—Accepts Layer 3 IP Precedence markings; largely deprecated
- Trust DSCP—Accepts Layer 3 DSCP markings; this is the most granular and flexible static state and thus the most utilized static trust state in campus networks

Conditional Trust

Trust may also be extended dynamically, provided a successful condition has been met. In Cisco campus networks this condition is a successful Cisco Discovery Protocol (CDP) negotiation between the access switch and the endpoints. Endpoints that can be extended conditional trust by Cisco Catalyst switches include Cisco IP phones, Cisco TelePresence Systems, Cisco IP Surveillance Cameras, and Cisco Digital Media Players. Conditional trust operation is shown in Figure 1.

Figure 1 Conditional Trust Operation

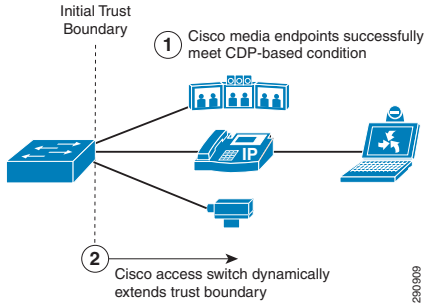

Per-Port QoS

When a QoS policy is applied on a per-port basis, it is attached to a specific physical switch port and is active on all traffic received on that specific port (only). QoS policies are applied on a per-port basis by default. Figure 2 illustrates port-based QoS.

Figure 2 Port-Based QoS

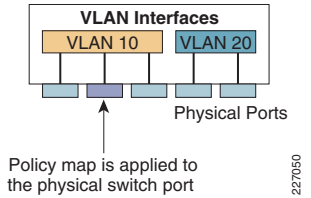

Per-VLAN QoS

When a QoS policy is applied on a per-VLAN basis, it is attached to a logical VLAN interface and is active on all traffic received on all ports that are currently assigned to the VLAN. Figure 3 illustrates VLAN-based QoS.

Figure 3 VLAN-Based QoS

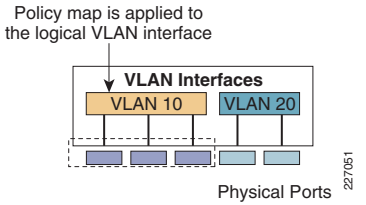

Per-Port/Per-VLAN QoS

When a QoS policy is applied on a Per-Port/Per-VLAN basis, it is attached to specific VLAN on a trunked port and is active on all traffic received from that specific VLAN from that specific trunked port (only). Figure 4 illustrates Per-Port/Per-VLAN-based QoS.

Figure 4 Per-Port-Per-VLAN-Based QoS

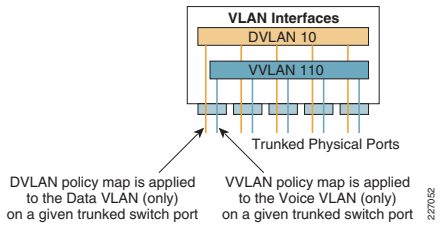

Ingress QoS Models

There are many options for an administrator to choose from for ingress QoS models, as shown in Figure 5.

Figure 5 Ingress QoS Models

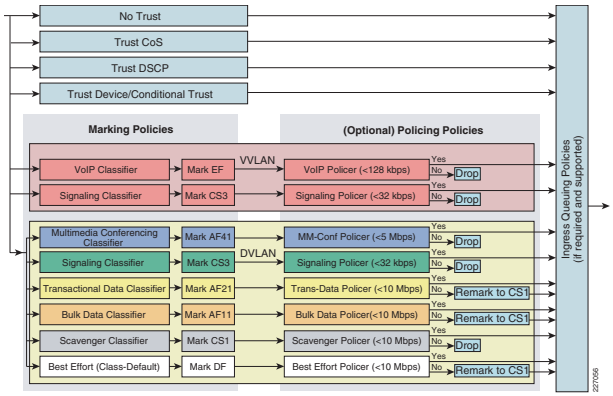

The three most utilized ingress QoS models for campus networks are:

- Trust DSCP Model
- Conditional Trust Model
- Service Policy Models

Combinations of these may be used at the same time

Egress QoS Models

Cisco Catalyst switches perform queuing in hardware and as such are limited to a fixed number of queues. The nomenclature used to describe these queuing structures is 1PxQyT, where:

- 1**P** represents a strict priority queue
- x**Q** represents x-number of non-priority queues
- y**T** represents y-number of drop-thresholds per non-priority queue

No fewer than four hardware queues would be required to support QoS policies in the campus; the following queues would be considered a minimum:

- Realtime queue (RFC 3246 EF PHB)
- Guaranteed bandwidth queue (RFC 2597 AF PHB)
- Default queue (RFC 2474 DF PHB)
- Bandwidth constrained queue (RFC 3662 PDB or "scavenger"service)

Additionally, the following bandwidth allocations are recommended for these queues:

- Realtime queue should not exceed 33% BW
- Default queue should be at least 25% BW
- Bulk/scavenger queue should not exceed 5% BW

Given these minimum queuing requirements and bandwidth recommendations, the following application classes can be mapped to the respective queues:

- Voice, broadcast video, and realtime interactive may be mapped to the realtime queue (per RFC 4594)
- Network/internetwork control, signaling, network management, multimedia conferencing, multimedia streaming, and transactional data can be mapped to the guaranteed bandwidth queue. Congestion avoidance mechanisms, such as WRED, can be enabled on this class; furthermore, if configurable drop thresholds are supported on the platform, these may be enabled to provide intra-queue QoS to these application classes, in the respective order they are listed (such that control plane protocols receive the highest level of QoS within a given queue)
- Bulk data and scavenger traffic can be mapped to the bandwidth-constrained queue and congestion avoidance mechanisms can be enabled on this class. If configurable drop thresholds are supported on the platform, enabling them provides intra-queue QoS to drop scavenger traffic ahead of bulk data
- Best effort traffic can be mapped to the default queue; congestion avoidance mechanisms can be enabled on this class

An egress queuing example based on these design considerations is shown in Figure 6.

Figure 6 An Egress Queuing Example Model

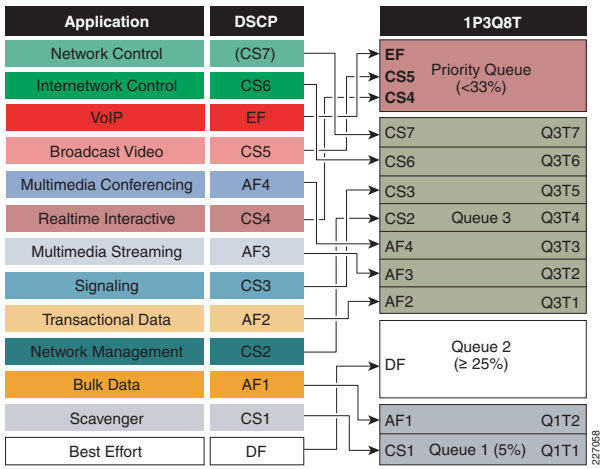

EtherChannel QoS

On some platforms ingress QoS policies (such as DSCP trust) are applied on the logical Port-Channel interface; however, on all platforms egress QoS policies (such as queuing policies) are always applied to the physical port-member interfaces.

QoS Roles in a Campus

Access edge switch ports have the most variation in QoS policy roles and these will vary depending on the type of endpoint to which these are connecting.

For all switch-to-switch links the only QoS policies that are required are DSCP-trust (on ingress) and queuing (on egress). QoS roles in a campus network are shown in Figure 7.

Figure 7 Campus Port QoS Roles

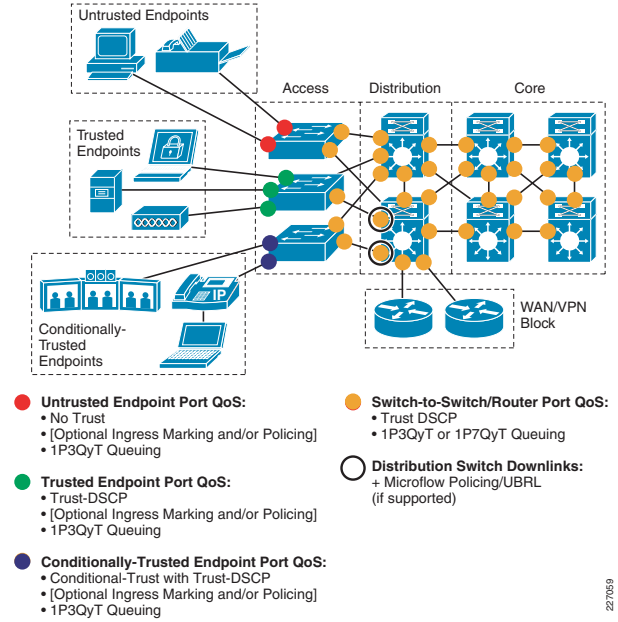

AutoQoS

On some Catalyst switching platforms Cisco has already updated and expanded the functionality of its AutoQoS feature to automatically provision QoS best practice designs for voice, IP-based video applications (such as IP Video Surveillance, Cisco TelePresence, conferencing applications, and streaming video applications), as well as for multiple types of data applications.

On these switch platforms, an administrator can automatically provision these best practice designs via a single interface-level command that corresponds to the endpoint to which the switch port is connecting.

For more details, see Campus QoS Design 4.0:
http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSCampus_40.html
And the Cisco Press book: **End-to-End QoS Network Design** (Second Edition)-Chapter 13