



CHAPTER 2

Medianet Campus QoS Design 4.0

Overview

The case for Quality of Service (QoS) in WANs/VPNs is largely self-evident because of the relatively low-speed bandwidth links at these Places-in-the-Network (PINs), as compared to Gigabit/Ten Gigabit campus networks, where the need for QoS is sometimes overlooked or even challenged. This is sometimes due to network administrators equating QoS with queuing policies only; whereas, the QoS toolset extends considerably beyond just queuing tools. Classification, marking, and policing are all important QoS functions that are optimally performed within the campus network, particularly at the access layer ingress edge (access edge).

Five strategic QoS design principles discussed in [Chapter 1, “Enterprise Medianet Quality of Service Design 4.0—Overview”](#) are relevant when deploying QoS in the campus:

- **Always perform QoS in hardware rather than software when a choice exists.** Cisco IOS routers perform QoS in software. This places additional demands on the CPU, depending on the complexity and functionality of the policy. Cisco Catalyst switches, on the other hand, perform QoS in dedicated hardware Application-Specific Integrated Circuits (ASICs) and as such do not tax their main CPUs to administer QoS policies. You can therefore apply complex QoS policies at Gigabit/Ten Gigabit line rates in these switches.
- **Classify and mark applications as close to their sources as technically and administratively feasible.** This principle promotes end-to-end Differentiated Services/Per-Hop Behaviors. Sometimes endpoints can be trusted to set Class of Service (CoS) or Differentiated Services Code Point (DSCP) markings correctly, but this is not always recommended as users could easily abuse provisioned QoS policies if permitted to mark their own traffic. For example, if DSCP Expedited Forwarding (EF) received priority services throughout the enterprise, a user could easily configure the NIC on a PC to mark all traffic to DSCP EF, thus hijacking network priority queues to service their non-real time traffic. Such abuse could easily ruin the service quality of real time applications (like VoIP) throughout the enterprise.
- **Police unwanted traffic flows as close to their sources as possible.** There is little sense in forwarding unwanted traffic only to police and drop it at a subsequent node. This is especially the case when the unwanted traffic is the result of Denial of Service (DoS) or worm attacks. Such attacks can cause network outages by overwhelming network device processors with traffic.
- **Enable queuing policies at every node where the potential for congestion exists,** regardless of how rarely this in fact may occur. This principle applies to campus edge and interswitch links, where oversubscription ratios create the potential for congestion. There is simply no other way to guarantee service levels than by enabling queuing wherever a potential speed mismatch exists.

- **Protect the control plane and data plane** by enabling control plane policing (on platforms supporting this feature) as well as data plane policing (scavenger class QoS) on campus network switches to mitigate and constrain network attacks.

However, before these strategic QoS design principles can be translated into platform-specific configuration recommendations, a few additional campus-specific considerations need to be taken into account and are discussed below.

Medianet Campus QoS Design Considerations

There are several considerations unique to the campus that factor into QoS designs, including:

- [Internal DSCP](#)
- [Trust States and Operation](#)
- [Trust Boundaries](#)
- [Port-Based, VLAN-Based, and Per-Port/Per-VLAN-Based QoS](#)
- [EtherChannel QoS](#)
- [Campus QoS Models](#)
- [Medianet Campus Port QoS Roles](#)
- [AutoQoS](#)
- [Smartport Macros](#)
- [Control Plane Policing](#)

These are discussed in the following sections.

Internal DSCP

For the most part, Cisco Catalyst switches perform QoS operations by assigning each packet (where “packet” is being used loosely in this chapter to describe Layer 2 frames as well as Layer 3 packets) an internal DSCP value (which is sometimes referred to as a “QoS label”, but is not to be confused with an MPLS label). This internal DSCP value is used to determine if a packet is to be remarked or policed or to which queue it is to be assigned or whether it should be dropped. The internal DSCP value may—or may not be—the same as the actual DSCP value of an IP (IPv4 or IPv6) packet; furthermore, an internal DSCP value is generated even for non-IP protocols (such as Layer 2 protocols like Spanning Tree as well as non-IP Layer 3 protocols like IPX).

The manner in which the internal DSCP value is generated for a packet depends on the trust state of the port on which the packet was received, which is described next.

Trust States and Operation

There are four (static) trust states with which a switch port can be configured:

- **Untrusted**—A port in this trust state disregards any and all Layer 2 or Layer 3 markings that a packet may have and generates an internal DSCP value of 0 (by default, unless explicitly overridden by the `[mls] qos cos` interface configuration command) for all received packets. This port trust state can be enabled with the interface configuration command `no [mls] qos trust`.

**Note**

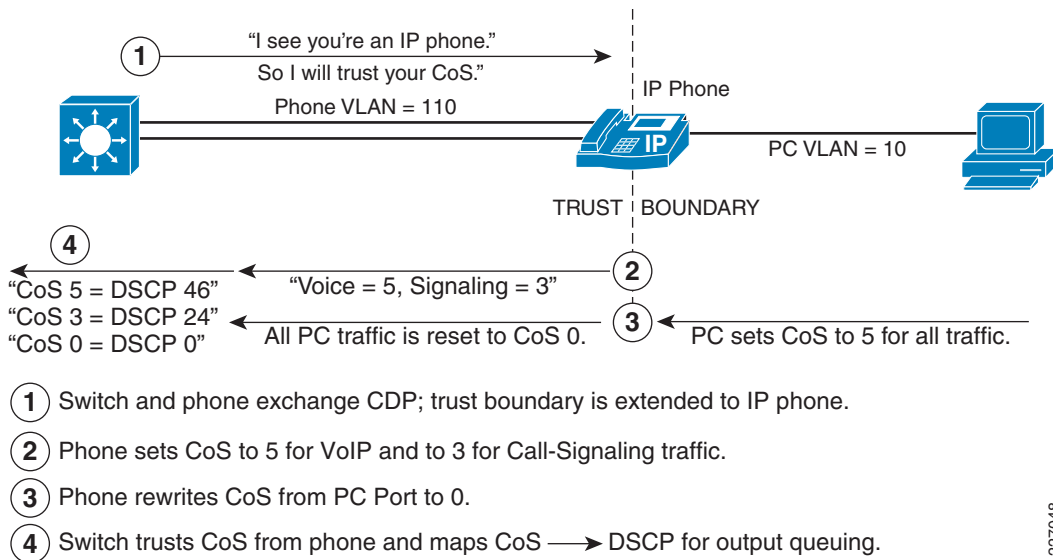
Cisco switches—with the exception of the 4500/4900 family—use the **mls** prefix for these QoS commands, whereas the 4500/4900 family omits this prefix. Otherwise, these commands are compatible across Cisco Catalyst 2960, 2975, 3560, 3750, 4500, 4900, and 6500 series platforms.

- **Trust CoS**—A port in this trust state accepts the 802.1p CoS marking of a 802.1Q tagged packet and use this value—in conjunction with the CoS-to-DSCP mapping table—to calculate an internal DSCP value for the packet. The default CoS-to-DSCP mapping table multiplies each CoS value by a factor of 8 to calculate the default internal DSCP (for example, CoS 1 maps to DSCP 8, CoS 2 maps to DSCP 16, and so on); however, the default CoS-to-DSCP mapping table can be modified with the **[mls] qos map cos-dscp** global configuration command (for example to map CoS 5 to the non-default DSCP value of EF [46]). In the case of an untagged packet (such as a packet received from the native VLAN) the default Internal DSCP value of 0 is applied. This port trust state can be enabled with the interface configuration command **[mls] qos trust cos**.
- **Trust IP Precedence**—A port in this trust state accepts the IP Precedence (IPP) marking of a packet (that is, the first three bits of the IPv4 or IPv6 Type of Service byte) and uses this value—in conjunction with the IP Precedence-to-DSCP mapping table—to calculate an internal DSCP value for the packet. The default IPP-to-DSCP mapping table multiplies each IPP value by a factor of 8 to calculate the default internal DSCP (for example, IPP 1 maps to DSCP 8, IPP 2 maps to DSCP 16, and so on); however, the default IPP-to-DSCP mapping table can be modified with the **[mls] qos map ip-prec-dscp** global configuration command (for example to map IPP 5 to the non-default DSCP value of EF [46]). In the case of a non-IP packet (such as an IPX packet) the default Internal DSCP value of 0 is applied. This port trust state can be enabled with the interface configuration command **[mls] qos trust ip-precedence**; however, it should be noted that this trust state is a legacy state, having been relegated by the trust DSCP state.
- **Trust DSCP**—A port in this trust state accepts the DSCP marking of a packet and sets the internal DSCP value to match. In the case of a non-IP packet (such as an IPX packet), the default internal DSCP value of 0 is applied. This port trust state can be enabled with the interface configuration command **[mls] qos trust dscp**.

**Note**

While the preceding serves to summarize these port trust states and operations, more complex options and scenarios also exist, as illustrated in [Figure 2-15](#).

In addition to the four static trust states described above, Cisco Catalyst switches also support a dynamic trust state, where the applied trust state for a port can dynamically toggle, depending on a successful endpoint identification exchange and the configured endpoint trust policy. This feature is referred to as conditional trust and automates user mobility for Cisco IP telephony deployments. Conditional trust operation is illustrated in [Figure 2-1](#).

Figure 2-1 Conditional Trust Operation

- 1 Switch and phone exchange CDP; trust boundary is extended to IP phone.
- 2 Phone sets CoS to 5 for VoIP and to 3 for Call-Signaling traffic.
- 3 Phone rewrites CoS from PC Port to 0.
- 4 Switch trusts CoS from phone and maps CoS → DSCP for output queuing.

227048

The sequence shown in [Figure 2-1](#) is:

1. The Cisco Catalyst access switch and Cisco Unified IP phone exchange Cisco Discovery Protocol (CDP) information; after a successful exchange, the switch recognizes that the endpoint is an IP phone and—in accordance with the switch port's configured policy—can extend trust to it.
2. The Cisco IP phone sets CoS to 5 for VoIP and to 3 for call signaling traffic.
3. The Cisco IP phone rewrites CoS from PC to 0.
4. The switch trusts CoS from phone and maps CoS-to-DSCP to generate internal DSCP values for all incoming packets.

**Note**

CDP is a lightweight, proprietary protocol engineered to perform neighbor discovery and as such was never engineered to be used as a security authentication protocol. Therefore, CDP should not be viewed or relied on as secure, as it can easily be spoofed.

The dynamic conditional trust state for Cisco Unified IP phones can be enabled with the interface configuration command **[mls] qos trust device cisco-phone**. Additionally, newer medianet devices, such as Cisco TelePresence Systems and IP Video Surveillance cameras, can also support conditional trust (on certain platforms with the latest versions of software); these devices use the **cts** and **cisco-camera** keywords, respectively

Regardless of how the Internal DSCP is generated—either by one of the four static port trust states or by the dynamic conditional trust state—it is important to note that as the packet exits the switch (unless explicitly overridden, as discussed in the following paragraph) the Catalyst switch sets the exiting IP packet's DSCP value to the final computed internal DSCP value. If trunking is enabled on the exiting switch port, the exiting packet's CoS value is also similarly set, but only to the first three bits of the final computed internal DSCP value.

If an administrator does not want the internal DSCP to overwrite the packet's ingress DSCP value, they can utilize the DSCP transparency feature, which is enabled by the **no mls qos rewrite ip dscp** global configuration command. When the DSCP transparency feature is enabled, the packet always has the same DSCP value on egress as it had on ingress, regardless of any internal QoS operations performed on the packet.

**Note**

The DSCP transparency is supported on all switching platforms discussed in this chapter, with the exception of the Catalyst 4500/4900 family.

Trust Boundaries

Having reviewed the internal DSCP concept and trust state operations, the administrator needs to consider where to enforce the trust boundary, i.e., the network edge at which packets are trusted (or not).

In line with the strategic QoS classification principle mentioned at the outset of this chapter, **the trust boundary should be set as close to the endpoints as technically and administratively feasible.**

The reason for the “administratively feasible” caveat within this design recommendation is that, while many endpoints (including user PCs) technically support the ability to mark traffic on their NICs, allowing a blanket trust of such markings could easily facilitate network abuse, as users could simply mark all their traffic with Expedited Forwarding, which would allow them to hijack network priority services for their non-realtime traffic and thus ruin the service quality of real time applications throughout the enterprise.

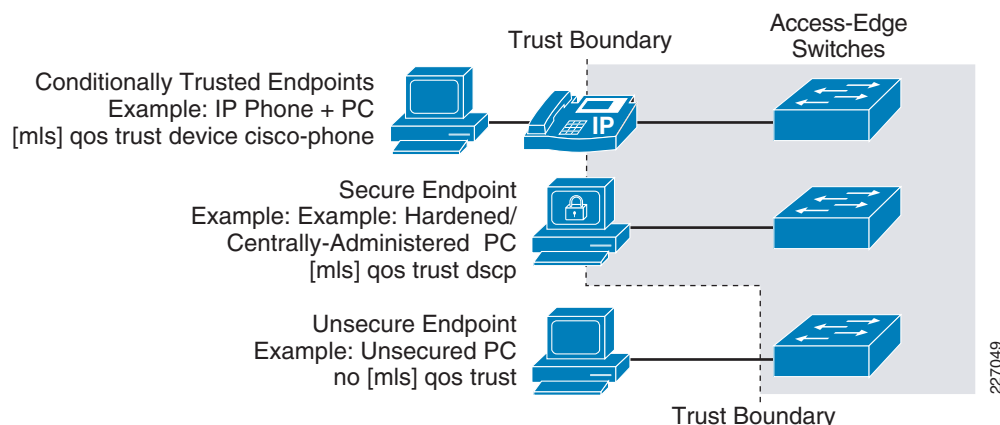
Thus, for many years it was advocated to not trust traffic from user PCs. However, more recently, various secure endpoint software has been released, such as Cisco Security Agent, that allows for PC markings to be centrally administered and enforced. Such centrally-administered software—along with quarantine VLANs for PCs that do not have such software installed—presents the option for network administrators to trust such secure endpoint PCs.

Therefore, from a trust perspective, there are three main categories of endpoints:

- Conditionally trusted endpoints—These include Cisco Unified IP phones as well as Cisco TelePresence systems.
- Trusted endpoints—These can include secure endpoint PCs and servers, IP video surveillance (IPVS) units, IP conferencing stations, wireless access points, analog and videoconferencing gateways, and other similar devices.
- Untrusted endpoints—Unsecure PCs and devices

The optimal trust boundaries and configuration commands for each of these categories of endpoints are illustrated in [Figure 2-2](#).

Figure 2-2 Optimal Trust Boundaries




Port-Based, VLAN-Based, and Per-Port/Per-VLAN-Based QoS

QoS classification (including trust), marking, and policing policies on Cisco Catalyst switches can be applied in one of three ways:


- Port-based 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 2-3](#) illustrates port-based QoS.

Figure 2-3 Port-Based QoS




- VLAN-based 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. Applying QoS policies on a per-VLAN basis requires the `[mls] qos vlan-based` interface command. [Figure 2-4](#) illustrates VLAN-based QoS.

Figure 2-4 VLAN-Based QoS



- Per-port/per-VLAN-based 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). Per-port/per-VLAN QoS is not supported on all platforms and the configuration commands are platform-specific, and as such is discussed on a per-platform basis later in this chapter. [Figure 2-5](#) illustrates per-port/per-VLAN-based QoS.

Figure 2-5 Per-Port/Per-VLAN-Based QoS

These application options allow for efficiency and granularity. For example, marking policies may be more efficiently scaled when applied on a per-VLAN basis. On the other hand, policies requiring policing granularity are best performed on a per-port/per-VLAN basis. Specifically, if an administrator wanted to police VoIP traffic from IP phones to a maximum of 128 kbps from each IP phone, this would could best be achieved by deploying a per-port/per-VLAN policing policy applied to the VVLAN on a given port. A per-port policer would not be sufficient, unless additional classification criteria was provided to specifically identify traffic from the IP phone only; neither could a per-VLAN policy be used, as this would police the aggregate traffic from all ports belonging to the VVLAN to 128 kbps.

EtherChannel QoS

Another case where logical versus physical interfaces has a bearing on QoS design is when provisioning QoS over EtherChannel interfaces. Multiple Gigabit Ethernet or 10-Gigabit Ethernet ports can be logically bundled into a single EtherChannel interface (which is also known as a PortChannel interface, as this is how it appears in the configuration syntax). From a Layer 2/Layer 3 standpoint, these bundled interfaces are represented-and function-as a single interface.

Two important considerations that should be kept in mind when deploying QoS policies on EtherChannel interfaces:

- The first EtherChannel QoS design consideration relates to load-balancing. Depending upon the platform, load balancing on the port-channel group can be done in various ways—by source IP address, by destination IP address, by source MAC address, by destination MAC address, by source and destination IP address, or by source and destination MAC address. It should be noted that EtherChannel technology does not take into account the bandwidth of each flow. Instead, it relies on the statistical probability that, given a large number of flows of relatively equal bandwidths, the load is equally distributed across the links of the port-channel group. However, this may not always be true. In general, it is recommended to load-balance based on the source-and-destination IP address, as this allows for statistically-superior load-distribution. And when loads are balanced in this manner, packets belonging to a single flow will retain their packet order.
- The second EtherChannel QoS design consideration is that EtherChannel technology does not take into account any QoS configuration on the individual Gigabit Ethernet links. Again, it relies on the statistical probability that, given a large number of flows with different QoS markings, the load of those individual flows is equally distributed across the links of the port-channel group. Given a failover situation in which one of the links of an EtherChannel group fails, the sessions crossing that link would be re-allocated across the remaining links. Since EtherChannel technology has no awareness of QoS markings, it could easily re-allocate more real-time flows across any one of the

links than the link is configured to accommodate. This could result in degraded real-time services. Incidentally, this scenario could also occur in a non-failover situation. Therefore, caution should be used when deciding to utilize EtherChannel technology versus a single higher-speed uplink port.

When configuring QoS policies over EtherChannel interfaces, the policies must often (but not always) be split two ways:

- Ingress policies, such as trust or marking and/or policing policies, are attached to the (logical) PortChannel interface. For example **[mls] qos trust dscp** or **service-policy input** commands would be applied to the PortChannel interface; this is the case for the Catalyst 4500/4500-E and 6500-6500-E series switches. An exception to this is the Catalyst 2960-G/S, 2975-GS, 3560-G/E/X, and 3750-G/E/X family of switches, which require ingress trust/classification/marketing/policing policies to be identically-configured on each and every Etherchannel physical port-member interface.
- Egress queuing policies are applied directly on the (physical) interfaces that compose the EtherChannel bundle. As queuing policies and commands vary by platform and/or linecard, these must be configured according to the platform-specific queuing sections outlined later in this design chapter. This requirement applies to all Catalyst switches discussed in this design chapter.

Therefore, as there is some slight per-platform variation in EtherChannel QoS configuration, a design example is included within each platform-family.

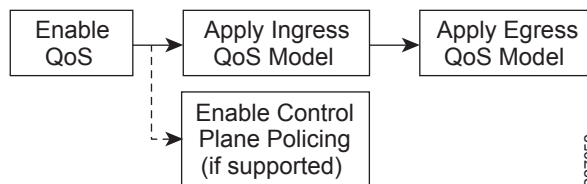
Campus QoS Models

Generally speaking, there are four main steps to deploying QoS models in the campus:

1. Enable QoS.
2. Apply an ingress QoS model, to assign trust or to explicitly classify and mark flows, to (optionally) police flows and to enable ingress queuing (if required).
3. Apply an egress QoS model, to assign flows to transmit queues, to enable dropping policies and egress policing (if supported and required).
4. Enable control plane policing (on platforms that support this feature).

These campus QoS deployment steps are illustrated in [Figure 2-6](#) and are discussed in additional detail in the following sections.

Figure 2-6 *Campus QoS Deployment Steps*



Ingress QoS Models

The ingress QoS model applies either a port trust state or an explicit classification and marking policy to the switch ports (or VLANs, in the case of VLAN-based QoS), as well as optional ingress policers and ingress queuing (as required and supported).

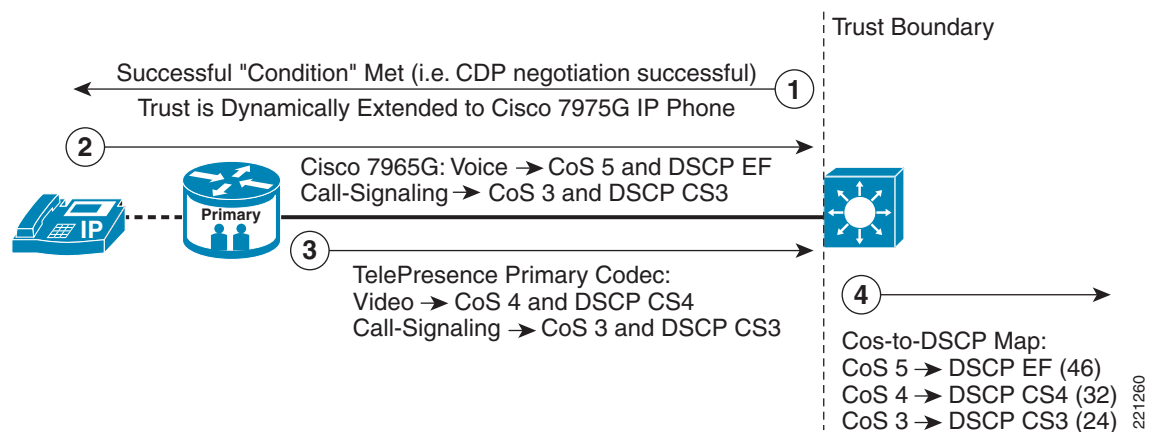
To begin with, the administrator needs to consider what application classes are present at the campus access edge (in the ingress direction) and whether these application classes are sourced from trusted or untrusted endpoints. As previously discussed, if PC endpoint markings are secured and centrally administered, then endpoint PCs can also be considered trusted endpoints; however, in most deployment scenarios this is not the case, and as such PCs are considered as untrusted endpoints for the remainder of this chapter.

Not every application class in the Cisco-modified RFC 4594-based model, shown in [Figure 1-9](#) in [Chapter 1, “Enterprise Medianet Quality of Service Design 4.0—Overview”](#), is present in the ingress direction at the campus access edge and as such, do not need to be provisioned for at this node. Specifically, network control traffic should never be received from endpoints, and as such, this class is not needed at the campus access edge. A similar case can be made for OAM traffic, as this traffic is primarily generated by network devices and is collected by management stations, which are typically in a data center or a network control center (and not the campus in general). Also, broadcast video and multimedia streaming traffic would originate from data center servers and would be unidirectional to campus endpoints (and should not be sourced from campus endpoints); therefore, these classes also would not need to be provisioned at the campus access edge.

That being said, of the remaining classes, consideration has to be given to which are sourced from trusted versus untrusted endpoints. Voice traffic is primarily sourced from Cisco IP telephony devices residing in the voice VLAN (VVLAN), and as such can be trusted (optimally, by conditional trust policies to facilitate user mobility, as illustrated in [Figure 2-1](#)). On the other hand, voice traffic may also be sourced from PC soft-phone applications, like Cisco Unified Personal Communicator (CUPC). However, because such applications share the same UDP port range as multimedia conferencing traffic (specifically, UDP/RTP ports 16384-32767), from a campus access edge classification standpoint, this renders soft-phone VoIP streams virtually indistinguishable from multimedia conferencing streams (unless NBAR technologies are used at the campus access edge). Unless soft-phone VoIP can be definitively distinguished from multimedia conferencing flows, it is simpler and safer to classify and mark the UDP port range of 16384-32767 as multimedia conferencing flows (AF4), as the alternative could allow multimedia conferencing flows to be admitted into strict priority queues intended (and capacity planned) for VoIP-only.

Realtime interactive flows may be sourced from Cisco TelePresence systems, which—like other Cisco IP telephony products—reside in the VVLAN and can be trusted to mark their own traffic, as shown in [Figure 2-7](#). Cisco TelePresence systems can be configured with either static or conditional trust policies.

Figure 2-7 Cisco TelePresence Conditional Trust Operation



221260

At the campus edge, signaling traffic may be sourced from both trusted endpoints (such as Cisco IP phones or Cisco TelePresence systems) or from untrusted endpoints (in the case of soft-phone applications running on PC endpoints, like CUPC). Therefore, both cases need to be accounted for with access edge policies.

Data applications, whether transactional, bulk, or best effort, are typically sourced from untrusted PC endpoints, as are scavenger applications.

These campus access edge endpoint marking and trust categories are summarized in [Figure 2-8](#).

Figure 2-8 *Campus Ingress Edge Marking and Trust Categories*

Application	PHB	Application Examples	Present at Campus Access-Edge (Ingress)?	Trusted Endpoint?	Untrusted Endpoint?
Network Control	CS6	EIGRP, OSPF, HSRP, IKE			
VoIP	EF	Cisco IP Phones	Yes	Trusted	
Broadcast Video		Cisco IPVS, Enterprise TV	Yes	Trusted	
Realtime Interactive	CS4	Cisco TelePresence	Yes	Trusted	
Multimedia Conferencing	AF4	Cisco CUPC, WebEx	Yes		Untrusted
Multimedia Streaming	AF3	Cisco DMS, IP/TV			
Signaling	CS3	SCCP, SIP, H.323	Yes	Trusted	Untrusted
Transactional Data	AF2	ERP Apps, CRM Apps	Yes		Untrusted
OAM	CS2	SNMP, SSH, Syslog			
Bulk Data	AF1	Email, FTP, Backups	Yes		Untrusted
Best Effort	DF	Default Class	Yes		Untrusted
Scavenger	CS1	YouTube, Gaming, P2P	Yes		Untrusted

227054

Traffic sourced from untrusted endpoints requires explicit classification and marking policies. While the number of applications assigned to the five (non-default) untrusted campus access edge application classes shown in [Figure 2-9](#) is virtually limitless—and is a function of the business objectives of the enterprise, as well as the technical proficiency of the network administrators—only a relatively few applications are used in this design chapter to illustrate these design concepts. Specifically, multimedia conferencing applications are sourced from the DVLAN to/from UDP ports 16384-32767. Signaling applications are limited to Skinny Call Control Protocol (SCCP) on TCP ports 2000-2002 and Session Initiation Protocol (SIP) on TCP/UDP ports 5060 and 5061. HTTPs are classified as a transactional data application (as the use of a secure transport implies a transaction). Additionally, an sample Enterprise Resource Planning (ERP) application, namely Oracle, is likewise classified as a transactional data application. FTP and email applications are classified as bulk data, as are PC-backup applications, such as Connected Backup for PC. Various peer-to-peer media sharing applications, such as iTunes, BitTorrent, and Kazaa, are classified as scavenger, as are gaming applications like Microsoft and Yahoo online gaming services. These applications classes, along with their classification criteria, are summarized in [Figure 2-9](#).

Figure 2-9 Untrusted Application Classification Examples

Application-Class	Application/Protocol	TCP	UDP	Port/Port-Range
Multimedia Conferencing	CUPC	TCP		16384-32767
Signaling	SCCP	TCP		2000-2002
Signaling	SIP	TCP	UDP	5060-5061
Transactional Data	HTTPS	TCP		443
Transactional Data	Oracle-SQL *NET	TCP	UDP	1521
Transactional Data	Oracle	TCP	UDP	1526
Transactional Data	Oracle	TCP	UDP	1575
Transactional Data	Oracle	TCP	UDP	1630
Bulk Data	FTP	TCP		20-21
Bulk Data	SSH/SFTP	TCP		22
Bulk Data	SMTP	TCP		25
Bulk Data	Secure SMTP	TCP		465
Bulk Data	IMAP	TCP		143
Bulk Data	Secure IMAP	TCP		993
Bulk Data	POP3	TCP		110
Bulk Data	Secure POP3	TCP		995
Bulk Data	Connected PC Backup	TCP		1914
Scavenger	Kazaa	TCP	UDP	1214
Scavenger	Microsoft DirectX Gaming	TCP	UDP	2300-2400
Scavenger	Apple iTunes Music Sharing	TCP	UDP	3689
Scavenger	BitTorrent	TCP		6881-6999
Scavenger	Yahoo Games	TCP		11999
Scavenger	MSN Gaming Zone	TCP	UDP	28800-6999

227055

**Note**

It is important to note that the list of TCP/UDP ports for applications shown in [Figure 2-9](#) is merely an example list and is not to be taken as an application port list reference. Some application ports are not included in the list above (to simplify the examples that follow); additionally, many applications add or change ports with incremental software revisions (and this list will not be maintained or updated to reflect such revisions).

In addition to explicit marking policies, optional policing policies may also be implemented on the campus access ingress edges to meter and manage flows. For example, voice flows could be policed to 128 kbps, while remaining traffic from the VVLAN (which would for the most part be signaling traffic, with a negligible amount of management traffic) could be policed to 32 kbps. Both VVLAN policers could be configured to drop violating flows, as VoIP and signaling traffic is well-defined and behaved, and traffic bursts in excess of these rates would indicate a network violation or abuse.

**Note**

Policing VoIP to 128 kbps is adequate to support G.711, G.722 and G.729 VoIP codecs. However, other VoIP codecs may require additional bandwidth, such as the Cisco Wideband (L16) Codec, which requires 256 kbps + network overhead (for a 320 kbps total). In such cases, the VoIP policers need to be provisioned accordingly.

In the DVLAN, multimedia conferencing flows come in various resolutions and quality. For example, 384 kbps or 768 kbps H.323 video conferencing streams can be policed at 500 kbps and 1 Mbps, respectively. Higher quality streams, such as 720p or 1080p H.264 streams, can be policed at (approximately) 2 Mbps and 5 Mbps, respectively (depending on motion-handling algorithms and other factors).

Data plane policing policies (discussed in [QoS for Security Best Practices](#) in Chapter 1, “Enterprise Medianet Quality of Service Design 4.0—Overview”) can be applied to monitor transactional data, bulk data, and best effort flows, such that these flows are metered, with violations being remarked to either an increased Drop Precedence within a given AF class (such as AF12, AF22, AF32, or AF42 or even to AF13, AF23, AF33, or AF43 in the case of dual-rate policers) or to CS1. What is important is that these packets are not dropped on ingress. For example, each of these classes can be policed to remark at 10 Mbps.

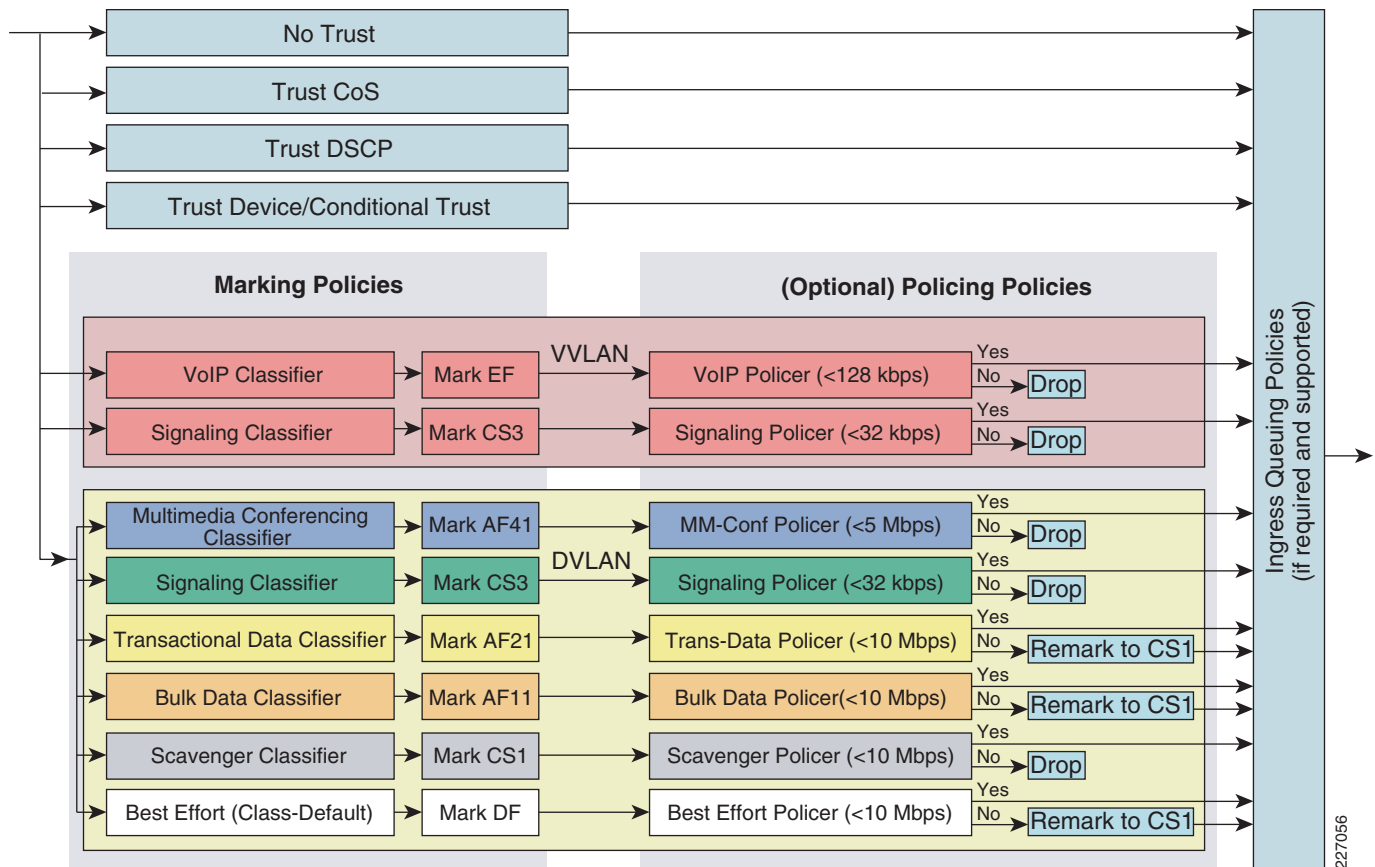
**Note**

This data plane policing rate (of 10 Mbps) is an **example** value. Such values could vary from enterprise to enterprise, even from department to department within an enterprise. The key is to set data plane policing rates such that approximately 95% of traffic flows for a given application class fall below the metered rate. For the sake of simplicity, a data plane policing rate of 10 Mbps is used for these application classes within this chapter.

Finally, a scavenger class can also be implemented to meter “less than best effort” flows—such as peer-to-peer media sharing applications or gaming applications. Such flows could also be policed to 10 Mbps (which is still only 1% of a GE link’s capacity), but with a more severe penalty for violations, namely dropping rather than remarking.

Once all ports have been set to trust or classify and mark (and optionally police) traffic, then ingress queuing policies may be applied (on platforms that require and support this feature). Ingress queuing details are discussed in the relevant platform-specific sections of this chapter.

[Figure 2-10](#) summarizes these campus ingress QoS model examples.

Figure 2-10 Campus Ingress QoS Example Models

It bears repeating that not every application class described here needs to be provisioned for at the access edge. For example, if multimedia conferencing applications are not widely deployed or utilized, then this class (along with the DVLAN signaling class) need not be provisioned at the access edge. Similarly, administrators may choose to simplify their data plane provisioning models, such that rather than explicitly provisioning transactional data, bulk data, and best effort classes, these could be provisioned as an aggregate best effort class (and marked as DF and optionally policed at an aggregate policing level). Likewise, explicitly provisioning a scavenger class is completely optional. Nonetheless, full examples, as described, are shown in this design chapter to provide template configurations which may be simplified as needed (or alternatively, expanded on).

Once ingress traffic has been trusted, classified, and (optionally) policed at the campus access edge, then the ingress QoS model for all campus inter-switch links can be set to trust the DSCP markings of all incoming packets.

Egress QoS Models

Egress QoS models primarily deal with queuing and dropping policies (although additional egress QoS features—such as egress policing—are supported on some platforms). As discussed in the previous chapter, critical media applications require service guarantees regardless of network conditions. The only way to provide service guarantees is to enable queuing at any node that has the potential for congestion, regardless of how rarely this may actually occur. This principle applies not only to campus-to-WAN/VPN edges, where speed mismatches are most pronounced, but also to campus

inter-switch links, where oversubscription ratios create the potential for instantaneous congestion. There is simply no other way to guarantee service levels other than by enabling queuing wherever a speed mismatch exists.

Additionally, because each medianet application class has unique service level requirements, each should optimally be assigned a dedicated queue. However, on platforms bounded by a limited number of hardware queues, no fewer than four queues would be required to support medianet QoS policies in the campus; specifically the following queues would be considered a minimum:


- Realtime queue (to support a RFC 3246 EF PHB service)
- Guaranteed bandwidth queue (to support RFC 2597 AF PHB services)
- Default queue (to support a RFC 2474 DF service)
- Bandwidth constrained queue (to support a RFC 3662 scavenger service)

Additionally, given the queuing best practice guidelines outlined in the previous chapter, the following bandwidth allocations are recommended for these queues:

- Realtime queue should not exceed 33% of the link's bandwidth.
- Default queue should be at least 25% of the link's bandwidth.
- Bulk/scavenger queue should not exceed 5% of the link's bandwidth.

These campus queuing bandwidth allocation recommendations are illustrated in [Figure 2-11](#).

Figure 2-11 *Campus Queuing Bandwidth Recommendations*



On some platforms, not only bandwidth allocations may be tuned, but also buffer allocations. Per-queue buffer allocations can be *directly* proportional to per-queue bandwidth allocations (for example, the buffer allocations for the best effort queue may be set to 25% to match the bandwidth allocation for this queue) or these can be *indirectly* proportional (for example, a strict priority queue which is being serviced in real-time would likely not need a corresponding 33% buffer allocation; whereas a bandwidth-constrained queue would benefit from deeper buffers to offset its minimal bandwidth allocation). Tuning buffer allocations is less impactful than tuning bandwidth allocations alone, but serves to complement the scheduling policies. Thus, in this design chapter—wherever possible—the strict-priority and less-than-best-effort queues are tuned to be indirectly proportional to their bandwidth allocations, while all other non-priority preferential queues are tuned to be directly proportional to their bandwidth allocations.

Given these minimum queuing requirements and bandwidth and buffer allocation 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 (i.e., selective dropping tools), 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, these may be enabled to provide inter-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.

Obviously, if more queues are supported these should be leveraged to give more granular bandwidth guarantees to these respective application classes. Nonetheless, the general application class hierarchy is to provision realtime applications (such as voice, broadcast video and realtime interactive) in a strict priority queue, followed by control plane protocols (including network/internetwork control, signaling [which is control plane traffic for the voice/video infrastructure] and network management), followed by guaranteed bandwidth, non-realtime applications (including multimedia conferencing, multimedia streaming, and transactional data), followed by the default best effort class, followed by bulk data and scavenger applications. Maintaining such an application class hierarchy serves to ensure consistent per-hop behaviors (PHBs).

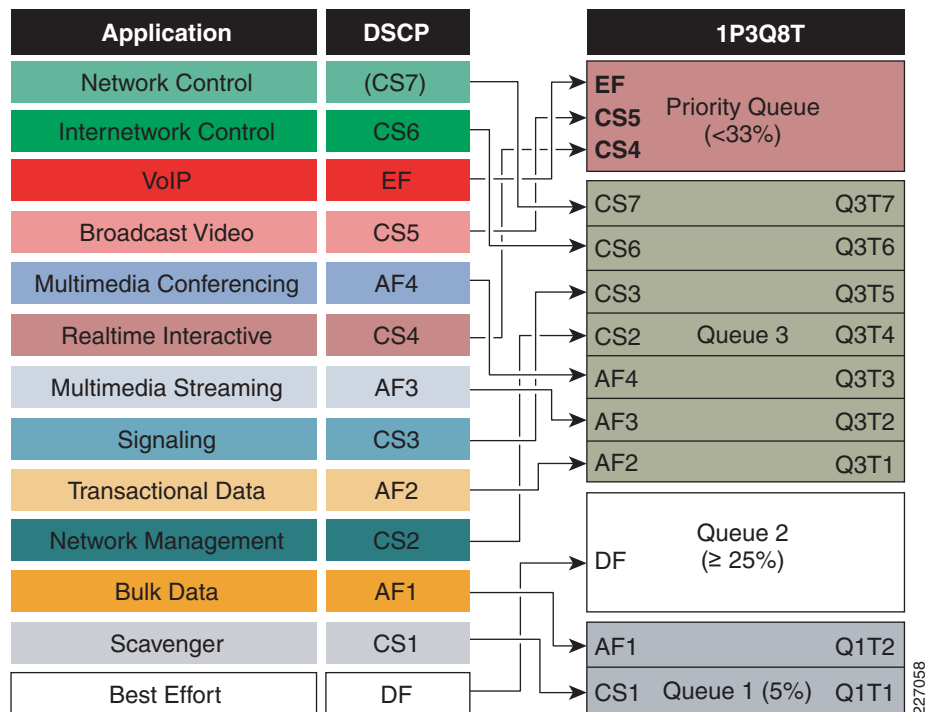
Some platforms provide DSCP-to-queue mapping functionality, whereas others (such as some Catalyst 6500 linecards) are limited to CoS-to-queue mapping functionality only. In both cases, it is the value of the internal DSCP that decides the transmit queue to which the packet is assigned; but in the case of CoS-to-queue mapping, internal DSCP values are assigned to queues in blocks of eight. For example, if CoS 1 is mapped to queue 1 (Q1), this means that internal DSCP values 8 through 15 are assigned to Q1; if CoS 2 is assigned to queue 2, this means that internal DSCP values 16-23 are assigned to Q2; if CoS 3 is mapped to queue 3, this means that internal DSCP values 24-31 are assigned to Q3, and so on. Essentially, CoS-to-queue mapping assigns the internal DSCP value that corresponds to the (CoS value * 8), along with the following seven internal DSCP values, to a given queue.

In some CoS-to-queue mapping scenarios, certain application classes may not be distinguishable from one another (due to the limited marking granularity of the 3-bit 802.1Q/p CoS model) and as such need to be assigned to the same queues. For example, since realtime interactive traffic (CS4/32) and multimedia conferencing traffic (AF41/34) share the same CoS value (of 4), these could not be mapped to different queues within a CoS-to-queue mapping model. Such considerations are discussed in more detail in the platform-specific sections of this chapter.

In contrast, with DSCP-to-queue mapping, discrete DSCP values can be mapped to specific queues, allowing for better queuing-policy granularity.

A campus egress QoS model example for a platform that supports DSCP-to-queue mapping with a 1P3Q8T queuing structure is shown in [Figure 2-12](#).

Figure 2-12 Campus Egress QoS Model Example



Medianet Campus Port QoS Roles

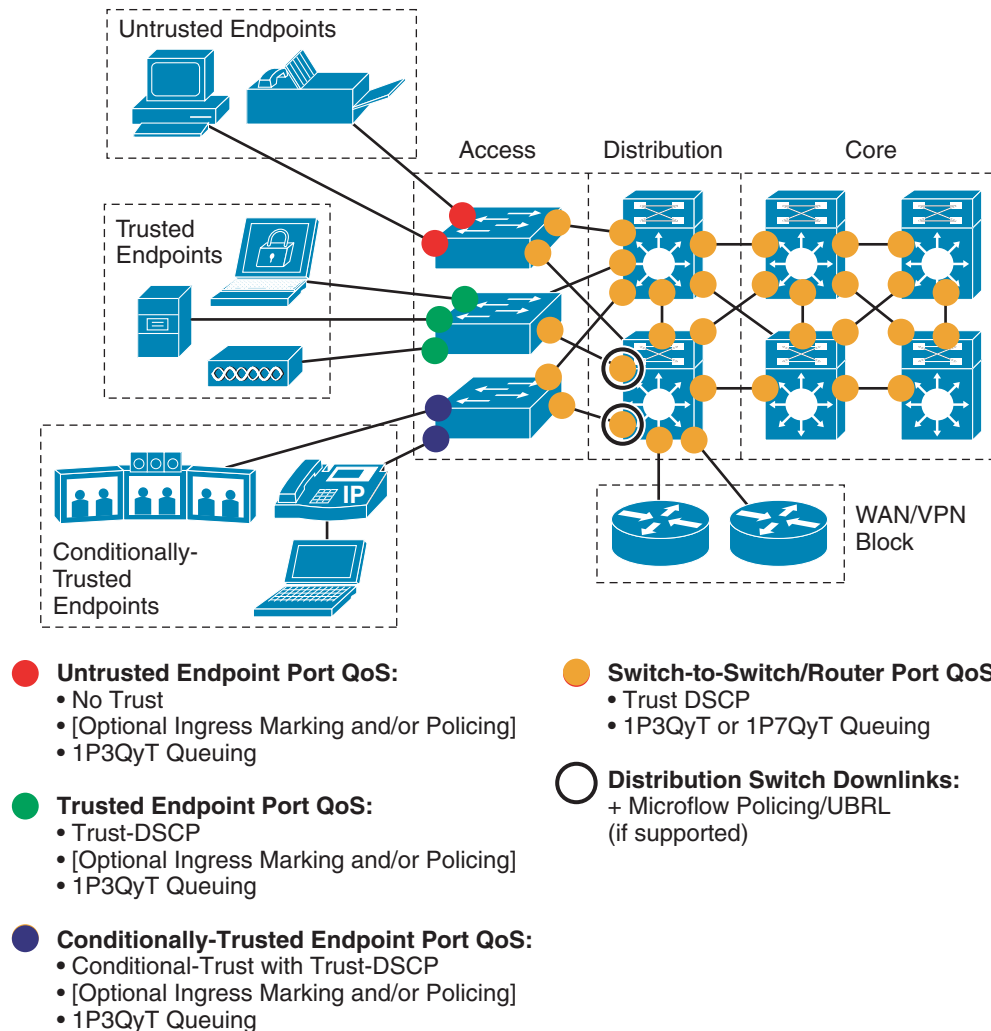
The policy elements discussed thus far can be grouped into roles that various switch ports serve within the medianet campus architecture, such as:

- Switch ports connecting to untrusted endpoints:
 - Endpoint examples include (unsecured/unmanaged) PCs, PDAs, printers, or other devices.
 - Trust **should** be disabled on these ports.
 - Optional ingress marking or policing policies (such as data plane policing policies) **may** be configured on these ports.
 - Ingress queuing policies (if supported and if required due to oversubscription scenarios, such as switch stacks) **may** be configured on these ports.
 - Egress queuing policies that support (at a minimum) 1P3QyT queuing should be configured on these ports, preferably with DSCP-to-queue mapping.
- Switch ports connecting to trusted endpoints:
 - Endpoint examples include secure/centrally-managed PCs and servers, IP video surveillance (IPVS) units, IP conferencing stations, wireless access points, analog and videoconferencing gateways, and similar other devices.
 - Static trust policies **should** be configured on these ports, preferably DSCP-trust for maximum classification and marking granularity.
 - Optional ingress marking or policing policies (such as data plane policing policies) **may** be configured on these ports.

- Ingress queuing policies (if supported and if required due to oversubscription scenarios, such as switch stacks) **may** be configured on these ports.
 - Egress queuing policies that support (at a minimum) 1P3QyT queuing **should** be configured on these ports, preferably with DSCP-to-queue mapping.
- Switch ports connecting to conditionally-trusted endpoints:
 - Endpoint examples include Cisco IP phones and Cisco TelePresence systems.
 - Conditional trust policies **should** be configured on these ports, preferably in conjunction with DSCP-trust extension, for maximum classification and marking granularity.
 - Optional ingress marking or policing policies (such as data plane policing policies) **may** be configured on these ports.
 - Ingress queuing policies (if supported and if required due to oversubscription scenarios, such as switch stacks) **may** be configured on these ports.
 - Egress queuing policies that support (at a minimum) 1P3QyT queuing **should** be configured on these ports, preferably with DSCP-to-queue mapping.
- Switch ports connecting to switch (or router) ports:
 - Access/distribution uplinks/downlinks; distribution/core uplinks/downlinks; core links; and campus-to-WAN/VPN-edge links
 - Static trust policies **should** be configured on these ports, preferably DSCP-trust for maximum classification and marking granularity.
 - Optional ingress marking or policing policies (such as data plane policing policies) **may** be configured on these ports.
 - Egress queuing policies that support (at a minimum) 1P3QyT queuing **should** be configured on these ports, preferably with DSCP-to-queue mapping. However, switch platforms/linecards that support 1P7QyT queuing are preferred at the distribution and core layers for increased queuing granularity at these aggregation layers.
 - Distribution downlinks (to the access layer) may be configured with microflow policing or User-Based Rate Limiting (UBRL) to provide a potential second line of policing defense for the medianet campus network.

Figure 2-13 shows these switch port QoS roles within a medianet campus architecture.

Figure 2-13 Medianet Campus Port QoS Roles



227059

AutoQoS

Due to the complexity of some QoS policies, coupled with the large number of ports on typical Catalyst switches, QoS deployment can often become unwieldy. One option is to make liberal use of the **interface range** configuration command to deploy policies to multiple interfaces at once. Another option is to use Automatic QoS (AutoQoS), if applicable. Yet another option is to use Smartport macros (which is discussed in the following section).

To address customer demands for simplification of QoS deployment, Cisco developed the AutoQoS feature. AutoQoS is an intelligent macro that allows an administrator to enter one or two simple AutoQoS commands to enable all the recommended QoS settings for one (or more) applications. In its first release, AutoQoS-VoIP, AutoQoS provisioned all the recommended QoS settings for IP telephony deployments for a specific switch port interface.

AutoQoS-VoIP for Catalyst switches supports three modes of operation, all of which are preceded by the **auto qos voip** interface configuration command:

- **cisco-phone**—This mode is intended for switch ports that may be connected to PCs or Cisco IP phones and sets the port to a conditional trust state, as well as configures mapping and queuing policies for QoS for VoIP.
- **cisco-softphone**—This mode is intended for switch ports that may be connected to PCs running Cisco IP Communicator or similar soft-phone software, and polices VoIP and signaling traffic, as well as configures mapping and queuing policies for QoS for VoIP (note that this feature is not supported on the Catalyst 4500 series of switches).
- **trust**—This mode is intended for switch ports that are within the trusted-boundary (such as inter-switch links, including uplinks and downlinks) or switch ports that are connecting to trusted endpoints, and sets the port to a static trust-dscp state, as well as configures mapping and queuing policies for QoS for VoIP.

**Note**

AutoQoS-VoIP is not supported on the Catalyst 4500-E/4900M series switches.

For additional details on AutoQoS-VoIP and the platform-specific commands and settings that it generates, refer to the respective platform's AutoQoS-VoIP documentation:

- Catalyst 2960 AutoQoS-VoIP Documentation:
http://www.cisco.com/en/US/docs/switches/lan/catalyst2960/software/release/12.2_50_se/configuration/guide/swqos.html#wp1231112
- Catalyst 3560/3750 AutoQoS-VoIP Documentation:
http://www.cisco.com/en/US/docs/switches/lan/catalyst3560/software/release/12.2_50_se/configuration/guide/swqos.html#wp1231112
- Catalyst 3560-E/3750-E AutoQoS-VoIP Documentation:
http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_50_se/configuration/guide/swqos.html#wp1231112
- Catalyst 4500 “Classic Supervisor” AutoQoS-VoIP Documentation:
<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/50sg/configuration/guide/qos.html#wp1583167>
- Catalyst 6500 AutoQoS-VoIP Documentation:
http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/auto_qos.html

Some may naturally ask, why should I read this lengthy and complex QoS design document when I have AutoQoS-VoIP? AutoQoS-VoIP is an excellent tool for customers who want to enable QoS for VoIP (only), that have basic QoS needs, or do not have the time or desire to do more with QoS.

However, it is important to remember how AutoQoS developed. AutoQoS features are the result of Cisco QoS feature development coupled with Cisco QoS design guides based on large-scale lab testing. AutoQoS-VoIP is the product of the first QoS Solution Reference Network Design (SRND) guide (published in 1999) and the AutoQoS-VoIP feature has not been significantly updated since. Therefore, if the business requirement for QoS is for IP Telephony **only**, then AutoQoS would be an excellent tool to expedite the QoS deployment.

However, as of August 2010, an updated version of AutoQoS was released for the Catalyst 2960-G/S, 2975-GS, 3560-G/E/X, and 3750-G/E/X family of switches (with IOS release 12.2(55)SE). This release was directly based on the recommendations put forward in this design chapter to support medianet applications; in fact, the new keyword and name for this version of AutoQoS is AutoQoS-SRND4 (taken from Solution Reference Network Design guide version 4, which is the Cisco name for this design chapter). AutoQoS-SRND4 is the fastest and most accurate method to deploy the recommended QoS

designs to support rich media applications across this family of switches. Details on this feature, as well as the complete configurations produced, are presented in [Cisco Catalyst 2960-G/S, 2975-GS, 3560-G/E/X, and 3750-G/E/X QoS Design](#).

**Note**

It should be mentioned that—at the time of writing—there are initiatives to update AutoQoS to also support medianet applications on other switching platforms. As these become available, details will be added to this design chapter.

Smartport Macros

Smartports macros provide static (and on some platforms, dynamic) configurations to port or VLAN interfaces. With Smartport macros, longer configuration snippets can be deployed with a single command, with some configuration parameters modified dynamically (such as VLAN IDs and IP addresses).

Certain Smartport macros are pre-defined, or built in, within Catalyst IOS switch software, such as macros that configure ports to connect to Cisco IP phones (which includes the configuration and execution of AutoQoS-VoIP on the switchport), Cisco Catalyst switches, Cisco routers, and Cisco wireless access points (among other devices).

Additionally, Smartport macros can be deployed on event triggers. The most common event triggers are based on CDP messages received from connected devices. The detection of a device invokes a CDP event trigger, such as a Cisco IP phone, Cisco switch, Cisco router, or Cisco wireless access point.

Finally, Smartport macros can be custom defined, such that an administrator can assign a Smartport macro name to a custom configuration snippet and apply the macro statically or have it triggered dynamically by an event.

For additional information on Smartport macros refer to the respective platform Smartport Macro documentation.

- Catalyst 2960 Smartport Macros Documentation:
http://www.cisco.com/en/US/docs/switches/lan/catalyst2960/software/release/12.2_50_se/configuration/guide/swmacro.html
- Catalyst 2975 Smartport Macros Documentation:
http://www.cisco.com/en/US/docs/switches/lan/catalyst2975/software/release/12.2_46_ex/configuration/guide/swmacro.html
- Catalyst 3560G/3750G Smartport Macros Documentation:
http://www.cisco.com/en/US/docs/switches/lan/catalyst3560/software/release/12.2_50_se/configuration/guide/swmacro.html
- Catalyst 3560-E/3750-E Smartport Macros Documentation:
http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_50_se/configuration/guide/swmacro.html
- Catalyst 4500 Smartport Macros Documentation:
<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/50sg/configuration/guide/macro.html>
- Catalyst 6500 Smartport Macros Documentation:
<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/smrtport.html>

Control Plane Policing

Control plane policing (CoPP) is a security infrastructure feature available on Catalyst 4500 and 6500 Series switches running Cisco IOS that allows the configuration of QoS policies that rate limit the traffic handled by the main CPU of the switch. This protects the control plane of the switch from direct DoS attacks and reconnaissance activity.

CoPP protects Catalyst 4500 and 6500 switches by allowing the definition and enforcement of QoS policies that regulate the traffic processed by the main switch CPU (route or switch processor). With CoPP, these QoS policies are configured to permit, block, or rate limit the packets handled by the main CPU.

Packets handled by the main CPU, referred to as control plane traffic, typically include:

- Routing protocols
- Packets destined to the local IP address of the router
- Packets from network management protocols, such as SNMP
- Interactive access protocols, such as SSH, and telnet
- Other protocols, such as ICMP, or IP options, might also require handling by the switch CPU
- Layer 2 packets such as BPDU, CDP, DOT1X, etc.

CoPP leverages the modular QoS command line interface (MQC) for its QoS policy configuration. MQC allows the classification of traffic into classes and lets you define and apply distinct QoS policies to separately rate limit the traffic in each class. MQC lets you divide the traffic destined to the CPU into multiple classes based on different criteria. For example, four traffic classes could be defined based on relative importance: critical, normal, undesirable, and default. After the traffic classes are defined, a QoS policy can be defined and enforced for each class according to importance. The QoS policies in each class can be configured to permit all packets, drop all packets, or drop only those packets exceeding a specific rate limit.

**Note**

The number of control plane classes is not limited to four, but should be chosen based on local network requirements, security policies, and a thorough analysis of the baseline traffic.

CoPP comes into play right after the switching or the routing decision and before traffic is forwarded to the control plane. When CoPP is enabled the sequence of events (at a high level) is:

1. A packet enters the switch configured with CoPP on the ingress port.
2. The port performs any applicable input port and QoS services.
3. The packet gets forwarded to the switch CPU.
4. The switch CPU makes a routing or a switching decision, determining whether or not the packet is destined for the control plane.
5. Packets destined for the control plane are processed by CoPP and are dropped or delivered to the control plane according to each traffic class policy. Packets that have other destinations are forwarded normally.

The Catalyst 4500 and Catalyst 6500 Series switches implement CoPP similarly; however, CoPP has been enhanced on both platforms to leverage the benefits of their hardware architectures, and as a result each platform provides unique features. Therefore, the CoPP implementations on Catalyst 4500 and Catalyst 6500 Series switches are discussed in platform-specific detail in their respective sections within this chapter. Nonetheless, some general guidelines to deploying CoPP are common to both platforms.

Defining CoPP Traffic Classes

Developing a CoPP policy starts with the classification of the control plane traffic. To that end, the control plane traffic needs to be first identified and separated into different class maps.

The Catalyst 4500 Series switches provides a macro which automatically generates a collection of class maps for common Layer 3 and Layer 2 control plane traffic. While very useful, these predefined class maps might not include all the necessary traffic classes reaching the control plane and as a result they might need to be complemented with other user-defined class maps. The Catalyst 6500 Series switches do not provide a configuration macro. Therefore, all class maps need to be defined by the user.

This section presents a classification template that can be used as a model when implementing CoPP on Catalyst 4500 and Catalyst 6500 Series switches. This template presents a realistic classification, where traffic is grouped based on its relative importance and protocol type. The template uses nine different classes, which provide great granularity and make it suitable for real-world environments. It is important to note that, even though you can use this template as a reference, the actual number and type of classes needed for a given network can differ and should be selected based on local requirements, security policies, and a thorough analysis of baseline traffic.

This CoPP template defines these nine traffic classes:

- **Border Gateway Protocol (BGP)**—This class defines traffic that is crucial to maintaining neighbor relationships for BGP routing protocol, such as BGP keepalives and routing updates. Maintaining BGP routing protocol is crucial to maintaining connectivity within a network or to an ISP. Sites that are not running BGP would not use this class.
- **Interior Gateway Protocol (IGP)**—This class defines traffic that is crucial to maintaining IGP routing protocols, such as Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), and Routing Information Protocol (RIP). Maintaining IGP routing protocols is crucial to maintaining connectivity within a network.
- **Interactive Management**—This class defines interactive traffic that is required for day-to-day network operations. This class would include light volume traffic used for remote network access and management. For example, telnet, Secure Shell (SSH), Network Time Protocol (NTP), Simple Network Management Protocol (SNMP), and Terminal Access Controller Access Control System (TACACS).
- **File Management**—This class defines high volume traffic used for software image and configuration maintenance. This class would include traffic generated for remote file transfer, for example Trivial File Transfer Protocol (TFTP) and File Transfer Protocol (FTP).
- **Reporting**—This class defines traffic used for generating network performance statistics for reporting. This class would include traffic required for using Cisco IOS IP Service Level Agreements (SLAs) to generate ICMP with different DSCP settings in order to report on response times within different QoS data classes.
- **Monitoring**—This class defines traffic used for monitoring a router. This kind of traffic should be permitted but should never be allowed to pose a risk to the router. With CoPP, this traffic can be permitted but limited to a low rate. Examples would include packets generated by ICMP echo requests (ping and trace route).
- **Critical Applications**—This class defines application traffic that is crucial to a specific network. The protocols that might be included in this class include generic routing encapsulation (GRE), Hot Standby Router Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), Dynamic Host Configuration Protocol (DHCP), IPsec, and multicast traffic.
- **Undesirable**—This explicitly identifies unwanted or malicious traffic that should be dropped and denied access to the RP. For example, this class could contain packets from a well-known worm. This class is particularly useful when specific traffic destined to the router should always be denied

rather than be placed into a default category. Explicitly denying traffic allows you to collect rough statistics on this traffic using **show** commands and thereby offers some insight into the rate of denied traffic.

- **Default**—This class defines all remaining traffic destined to the route processor (RP) that does not match any other class. MQC provides the default class so you can specify how to treat traffic that is not explicitly associated with any other user-defined classes. It is desirable to give such traffic access to the RP, but at a highly reduced rate. With a default classification in place, statistics can be monitored to determine the rate of otherwise unidentified traffic destined to the control plane. After this traffic is identified, further analysis can be performed to classify it. If needed, the other CoPP policy entries can be updated to account for this traffic.

**Note**

On Catalyst 6500 Supervisors 32 and 720 the default class (class-default) is the only traffic class that matches both IP and non-IP packets.

Deploying CoPP Policies

Because CoPP filters traffic, it is critical to gain an adequate level of understanding about the legitimate traffic destined to the RP prior to deployment. CoPP policies built without proper understanding of the protocols, devices, or required traffic rates involved can block critical traffic, which has the potential of creating a DoS condition. Determining the exact traffic profile needed to build the CoPP policies might be difficult in some networks.

The following steps employ a conservative methodology that facilitates the process of designing and deploying CoPP. This methodology uses iterative ACL configurations to help identify and to incrementally filter traffic.

To deploy CoPP, it is recommended that you perform these steps:

Step 1 Determine the classification scheme for your network.

Identify the known protocols that access the RP and divide them into categories using the most useful criteria for your specific network. In the case of the Catalyst 4500 Series switch, you can take advantage of the system predefined classes and chose to combine them with your own classes. In the case of Catalyst 6500 there are no predefined classes, so you need to define all the classes. As an example of classification, the nine categories template presented earlier in this section (BGP, IGP, interactive management, file management, reporting, critical applications, undesirable, and default) use a combination of relative importance and traffic type. Select a scheme suited to your specific network, which might require a larger or smaller number of classes.

Step 2 Define classification access lists.

Configure each ACL to permit all known protocols in its class that require access to the RP. At this point, each ACL entry should have both source and destination addresses set to any. In addition, the ACL for the default class should be configured with a single entry, permit ip any any. This matches traffic not explicitly permitted by entries in the other ACLs. After the ACLs have been configured, create a class map for each class defined in Step 1, including one for the default class. Then assign each ACL to its corresponding class map.

**Note**

In this step you should create a separate class map for the default class, rather than using the class default available in some platforms. Creating a separate class map and assigning a **permit ip any any** ACL allows you to identify traffic not yet classified as part of another class.

Each class map should then be associated with a policy map that permits all traffic, regardless of classification. The policy for each class should be set as conform-action transmit exceed-action transmit.

Step 3 Review the identified traffic and adjust the classification.

Ideally, the classification performed in Step 1 identified all required traffic destined to the router. However, realistically, not all required traffic is identified prior to deployment and the **permit ip any any** entry in the default class ACL logs a number of packet matches. Some form of analysis is required to determine the exact nature of the unclassified packets. For example, you can use the **show access-lists** command to see the entries in the ACLs that are in use and to identify any additional traffic sent to the RP. However, to analyze the unclassified traffic you can use one of these techniques:

- General ACL classification as described in Characterizing and Tracing Packet Floods Using Cisco Routers, which is available at http://www.cisco.com/en/US/tech/tk59/technologies_tech_note09186a0080149ad6.shtml
- Packet analyzers

When traffic has been properly identified, adjust the class configuration accordingly. Remove the ACL entries for those protocols that are not used. Add a **permit any any** entry for each protocol just identified.

Step 4 Restrict a macro range of source addresses.

Refine the classification ACLs by only allowing the full range of the allocated CIDR block to be permitted as the source address. For example, if the network has been allocated 172.68.0.0/16, then permit source addresses from 172.68.0.0/16 where applicable.

This step provides data points for devices or users from outside the CIDR block that might be accessing the equipment. An external BGP (eBGP) peer requires an exception because the permitted source addresses for the session lies outside the CIDR block. This phase might be left on for a few days to collect data for the next phase of narrowing the ACL entries.

Step 5 Narrow the ACL permit statements to authorized source addresses.

Increasingly limit the source address in the classification ACLs to only permit sources that communicate with the RP. For example, only known network management stations should be permitted to access the SNMP ports on a router.

Step 6 Refine CoPP policies by implementing rate limiting.

Use the **show policy-map control-plane** command to collect data about the actual policies in place. Analyze the packet count and rate information and develop a rate limiting policy accordingly. At this point, you might decide to remove the class map and ACL used for the classification of default traffic. If so, you should also replace the previously defined policy for the default class by the class default policy.

A tested and validated set of CoPP rates are presented in [Table 2-1](#). It is important to note that the values presented here are solely for illustration purposes, as every environment has different baselines.

Table 2-1 Example Control Plane Policing Rate Limits and Actions

Traffic Class	Rate (bps)	Conform Action	Exceed Action
Border Gateway Protocol	4,000,000	Transmit	Drop
Interior Gateway Protocol	300,000	Transmit	Drop
Interactive Management	500,000	Transmit	Drop
File management	6,000,000	Transmit	Drop

Table 2-1 Example Control Plane Policing Rate Limits and Actions

Traffic Class	Rate (bps)	Conform Action	Exceed Action
Monitoring	900,000	Transmit	Drop
Critical applications	900,000	Transmit	Drop
Undesirable	32,000	Drop	Drop
Default	500,000	Transmit	Drop

This CoPP classification template, deployment model, and rate limits are used in the Catalyst 4500 and 6500 CoPP configuration examples later in this chapter.

Cisco Catalyst 2960-G/S, 2975-GS, 3560-G/E/X, and 3750-G/E/X QoS Design

The Cisco Catalyst 2960-G/S, 2975-GS, 3560-G/E/X, and 3750-G/E/X family of switches all support the (previously discussed) minimum requirements for medianet switches, including Gigabit Ethernet support, as well as supporting a strict priority hardware queue with at least three additional hardware queues.

The specific switch hardware configurations that meet these requirements are shown below, by switch family.

- Catalyst 2960-G series switches include:
 - Cisco Catalyst 2960G-8TC-L—8 Ethernet 10/100/1000 ports
 - Cisco Catalyst 2960G-24TC-L—24 Ethernet 10/100/1000 ports
 - Cisco Catalyst 2960G-48-TC-L—48 Ethernet 10/100/1000 ports
- Catalyst 2960-S series switches include:
 - Cisco Catalyst 2960S-48FPD-L—48 Ethernet 10/100/1000 PoE+ ports and 2 Ten Gigabit Ethernet SFP+ or 2 One Gigabit Ethernet SFP ports
 - Cisco Catalyst 2960S-48LPD-L—48 Ethernet 10/100/1000 PoE+ ports and 2 Ten Gigabit Ethernet SFP+ or 2 One Gigabit Ethernet SFP ports
 - Cisco Catalyst 2960S-24PD-L—24 Ethernet 10/100/1000 PoE+ ports and 2 Ten Gigabit Ethernet SFP+ or 2 One Gigabit Ethernet SFP ports
 - Cisco Catalyst 2960S-48TD-L—48 Ethernet 10/100/1000 ports and 2 Ten Gigabit Ethernet SFP+ or 2 One Gigabit Ethernet SFP ports
 - Cisco Catalyst 2960S-24TD-L—24 Ethernet 10/100/1000 ports and 2 Ten Gigabit Ethernet SFP+ or 2 One Gigabit Ethernet SFP ports
 - Cisco Catalyst 2960S-48FPS-L—48 Ethernet 10/100/1000 PoE+ ports and 4 One Gigabit Ethernet SFP ports
 - Cisco Catalyst 2960S-48LPS-L—48 Ethernet 10/100/1000 PoE+ ports and 4 One Gigabit Ethernet SFP ports
 - Cisco Catalyst 2960S-24PS-L—24 Ethernet 10/100/1000 PoE+ ports and 4 One Gigabit Ethernet SFP ports

- Cisco Catalyst 2960S-48TS-L—48 Ethernet 10/100/1000 ports and 4 One Gigabit Ethernet SFP ports
 - Cisco Catalyst 2960S-24TS-L—24 Ethernet 10/100/1000 ports and 4 One Gigabit Ethernet SFP ports
- Catalyst 2975-GS series switches include:
 - Cisco Catalyst 2975GS-48PS-L—48 Ethernet 10/100/1000 PoE ports and 4 SFP ports
- Catalyst 3560-G series switches include:
 - Cisco Catalyst 3560G-24TS—24 Ethernet 10/100/1000 ports and 4 SFP-based Gigabit Ethernet ports
 - Cisco Catalyst 3560G-48TS—48 Ethernet 10/100/1000 ports and 4 SFP-based Gigabit Ethernet ports
 - Cisco Catalyst 3560G-24PS—24 Ethernet 10/100/1000 ports with PoE and 4 SFP-based Gigabit Ethernet ports
 - Cisco Catalyst 3560G-48PS—48 Ethernet 10/100/1000 ports with PoE and 4 SFP-based Gigabit Ethernet ports
- Catalyst 3650-E series switches include:
 - Cisco Catalyst 3560E-24TD—24 Ethernet 10/100/1000 ports and 2 X2 10 Gigabit Ethernet uplinks
 - Cisco Catalyst 3560E-24PD—24 Ethernet 10/100/1000 ports with PoE and 2 X2 10 Gigabit Ethernet uplinks
 - Cisco Catalyst 3560E-48TD—48 Ethernet 10/100/1000 ports and 2 X2 10 Gigabit Ethernet uplinks
 - Cisco Catalyst 3560E-48PD—48 Ethernet 10/100/1000 ports with PoE and 2 X2 10 Gigabit Ethernet uplinks
 - Cisco Catalyst 3560E-48PD-F—48 Ethernet 10/100/1000 ports with 15.4W PoE on all 48 ports and 2 X2 10 Gigabit Ethernet uplinks
 - Cisco Catalyst 3560E-12D—12 X2 10 Gigabit Ethernet ports
 - Cisco Catalyst 3560E-12SD—12 SFP Gigabit Ethernet ports and 2 X2 10 Gigabit Ethernet ports
- Catalyst 3560-X series switches include:
 - WS-C3560X-24T-L—24 Ethernet 10/100/1000 ports and modular 4x Gigabit Ethernet or 2x Ten-Gigabit Ethernet SFP ports
 - WS-C3560X-24P—24 Ethernet 10/100/1000 ports with PoE+ and modular 4x Gigabit Ethernet or 2x Ten-Gigabit Ethernet SFP ports
 - WS-C3560X-48T—48 Ethernet 10/100/1000 ports and modular 4x Gigabit Ethernet or 2x Ten-Gigabit Ethernet SFP ports
 - WS-C3560X-48P—48 Ethernet 10/100/1000 ports with PoE+ and modular 4x Gigabit Ethernet or 2x Ten-Gigabit Ethernet SFP ports
 - WS-C3560X-48PF—48 Ethernet 10/100/1000 ports with PoE+ and modular 4x Gigabit Ethernet or 2x Ten-Gigabit Ethernet SFP ports and 1100W power supply
- Catalyst 3750-G series switches include:
 - Cisco Catalyst 3750G-24TS-1U—24 Ethernet 10/100/1000 ports and four SFP uplinks
 - Cisco Catalyst 3750G-24PS—24 Ethernet 10/100/1000 ports with IEEE 802.3af and Cisco pre-standard PoE and four SFP uplinks

- Cisco Catalyst 3750G-48TS—48 Ethernet 10/100/1000 ports and four SFP uplinks
- Cisco Catalyst 3750G-48PS—48 Ethernet 10/100/1000 ports with IEEE 802.3af and Cisco pre-standard PoE and four SFP uplinks
- Cisco Catalyst 3750G-24WS—24 Ethernet 10/100/1000 ports with IEEE 802.3af, Cisco prestandard PoE and two SFP uplinks and an integrated wireless LAN controller
- Catalyst 3750-E series switches include:
 - Cisco Catalyst 3750E-24TD—24 Ethernet 10/100/1000 ports and 2 X2 10 Gigabit Ethernet uplinks
 - Cisco Catalyst 3750E-24PD—24 Ethernet 10/100/1000 ports with PoE and 2 X2 10 Gigabit Ethernet uplinks
 - Cisco Catalyst 3750E-48TD—48 Ethernet 10/100/1000 ports and 2 X2 10 Gigabit Ethernet uplinks
 - Cisco Catalyst 3750E-48PD—48 Ethernet 10/100/1000 ports with PoE and 2 X2 10 Gigabit Ethernet uplinks
 - Cisco Catalyst 3750E-48PD-F—48 Ethernet 10/100/1000 ports with > 15.4 watts PoE on all 48 ports and 2 X2 10 Gigabit Ethernet uplinks
- Cisco Catalyst 3750-X series switches include:
 - WS-C3750X-24T-L—24 Ethernet 10/100/1000 ports and modular 4x Gigabit Ethernet or 2x Ten-Gigabit Ethernet SFP ports
 - WS-C3750X-24P—24 Ethernet 10/100/1000 ports with PoE+ and modular 4x Gigabit Ethernet or 2x Ten-Gigabit Ethernet SFP ports
 - WS-C3750X-48T—48 Ethernet 10/100/1000 ports and modular 4x Gigabit Ethernet or 2x Ten-Gigabit Ethernet SFP ports
 - WS-C3750X-48P—48 Ethernet 10/100/1000 ports with PoE+ and modular 4x Gigabit Ethernet or 2x Ten-Gigabit Ethernet SFP ports
 - WS-C3750X-48PF—48 Ethernet 10/100/1000 ports with PoE+ and modular 4x Gigabit Ethernet or 2x Ten-Gigabit Ethernet SFP ports and 1100W power supply

**Note**

These are the current shipping hardware configurations for these switching families at the time of writing. Additional configuration options may be added over time. As long as future hardware configuration options include the minimum requirements for medianet campus switches (namely, the support of Gigabit interfaces, along with a strict priority hardware queue and at least three additional non-priority hardware queues), these can also be deployed across medianet campus network infrastructures according to the guidelines presented in this chapter.

At a high-level, the major differences between these switch product families are as follows: the Catalyst 2960-G, 2960-S, and 2975-GS are Layer 2-only switches, while the 3560-G, 3560-E, 3560-X, and 3750-G, 3750-E, and 3750-X support Layer 2/Layer 3 multilayer switch feature sets. Additionally, the Catalyst 2960-G, 3560-G, 3560-E, and 3560-X are standalone switches, while (some models of) the Catalyst 2960-S, the Catalyst 2975-GS, 3750-G, 3750-E, and 3750-X are stackable switches. The Catalyst 2975-GS and 3750-G support stacking with Cisco StackWise technology, while the 3750-E and 3750-X use StackWise Plus technology; however, the models of the 2960-S family that support stacking do so using FlexStack technology. All of these Catalyst switches support a dual-counter-rotating ring, which effectively serves as the switching backplane; these rings are internal for non-stackable switches, but external (via special cables) for stackable switches. These rings operate at 16 Gbps each (for a total

switching capacity of 32 Gbps) for the Catalyst 2960-G and 2975-GS series, 20 Gbps each (for a total switching capacity of 40 Gbps) for the 2960-S series and at 32 Gbps each (for a total switching capacity of 64 Gbps) for the 3750-G, 3750-E, and 3750-X series switches.

**Note**

For additional product-specific details, refer to the product data sheets for each switch product family.

These major feature and functionality differences between these switch product families are summarized in [Table 2-2](#).

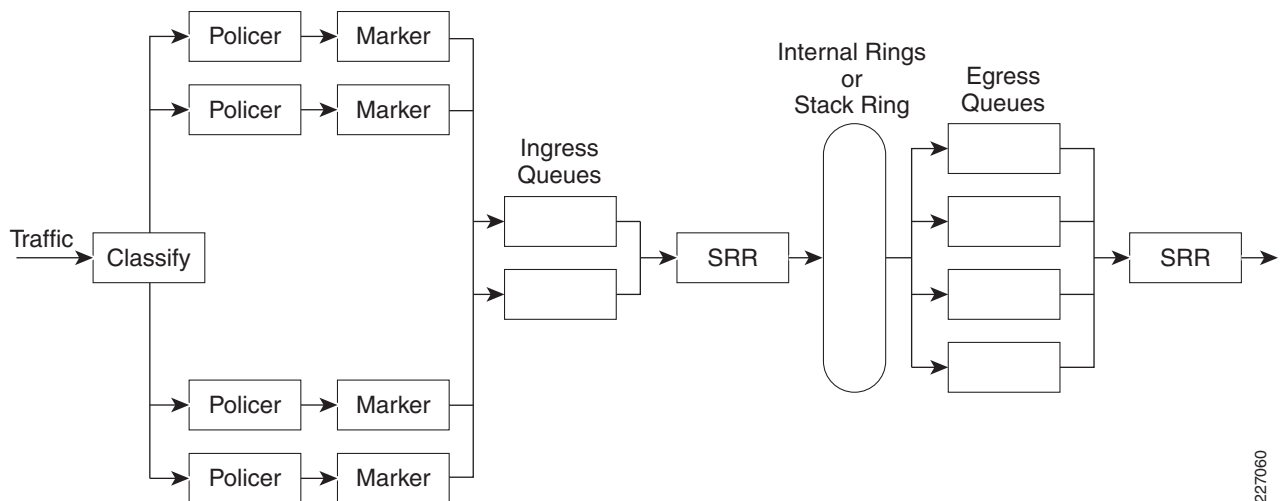
Table 2-2 Cisco Catalyst 2960-G/S, 2975-GS, 3560-G/E/X, and 3750-G/E/X—Major Feature and Functionality Matrix

Switch	Layer 2-Only Switch	Layer 2/Layer 3 Multilayer Switch	Stackable?	Stacking Technology	Total Switching Capacity
Catalyst 2960-G	Yes				32 Gbps
Catalyst 2960-S	Yes		Some Models	FlexStack	40 Gbps
Catalyst 2975-GS	Yes		Yes	StackWise	32 Gbps
Catalyst 3560-G		Yes			32 Gbps
Catalyst 3560-E		Yes			64 Gbps
Catalyst 3560-X		Yes			64 Gbps
Catalyst 3750-G		Yes	Yes	StackWise	32 Gbps
Catalyst 3750-E		Yes	Yes	StackWise Plus	64 Gbps
Catalyst 3750-X		Yes	Yes	StackWise Plus	64 Gbps

While these switches have some major feature and functionality differences, their QoS feature set and command syntax are virtually identical, with a few minor differences, as is discussed in the following section.

Platform-Specific QoS Considerations

Cisco Catalyst 2960-G/S, 2975-GS, 3560-G/E/X, and 3750-G/E/X have virtually identical QoS feature sets, and as such are discussed collectively; additionally, for brevity, these switches are collectively referred to as the Catalyst 3750-E, except when discussing switch-specific differences. The complete QoS model for the Catalyst 3750-E is shown in [Figure 2-14](#).

Figure 2-14 Catalyst 3750-E QoS Model

Traffic is classified on ingress, based on trust-states, access-lists, or class-maps. Marking or policing policies can be applied to physical switch ports or—on multilayer switch platforms—to Switch Virtual Interfaces (SVIs), which allows for per-VLAN or per-port/per-VLAN policies.

Because the total inbound bandwidth of all ports can exceed the bandwidth of the stack or internal ring, ingress queues are located after the packet is classified, policed, and marked and before packets are forwarded into the switch fabric (i.e., the internal or stack rings). Because multiple ingress ports can simultaneously send packets to an egress port (such as an uplink port) and cause congestion, outbound queues are located after the stack or internal rings. The queuing scheduler is Shared Round Robin (SRR), and the dropping algorithm is Weighted Tail Drop (WTD), both of which are discussed in more detail in [Queuing Models](#).

Relating to QoS, these key switch-specific differences exist:

- The Catalyst 2960-G/S and 2975-GS do not support multilayer switching and as such do not correspondingly support per-VLAN or per-port/per-VLAN policies.
- The Catalyst 2960-G and 2975-GS can only police to a minimum rate of 1 Mbps; all other platforms within this switch product family can police to a minimum rate of 8 kbps (with the exception of the 2960-S, which-although it can be configured to police at 8 kbps-can only police at a minimum rate of 16 kbps).
- The Catalyst 2960-S does not support ingress queuing.
- The Catalyst 2960-S does not support a “class-default” class-map.
- Only the Catalyst 3650-E/X, 3750-E/X support IPv6 QoS.
- Only the Catalyst 3650-E/X and 3750-E/X support policing on 10 Gigabit Ethernet interfaces.
- Only the Catalyst 3650-E/X and 3750-E/X support SRR shaping weights on 10 Gigabit Ethernet interfaces (SRR shaping weights are discussed in more detail in [Queuing Models](#)).

Other than these key exceptions, the following commands and configurations work across these switch platforms (unless explicitly noted otherwise).

Enabling QoS

On all the switching platforms discussed in this chapter (with the exception of the Catalyst 4500-E/4900M) QoS needs to be explicitly enabled, as it is disabled by default. This is a critical first step to deploying QoS on these platforms. If this small—but important—step is overlooked, this can lead to frustration in troubleshooting QoS problems; this is because the switch software accepts QoS commands and even displays these within the switch configuration, but none of the QoS commands are active until the **mls qos** global command is enabled, as shown in [Example 2-1](#).

Example 2-1 Enabling QoS on a Catalyst 3750-E

```
C3750-E(config)#mls qos
```

This configuration can be verified with the command:

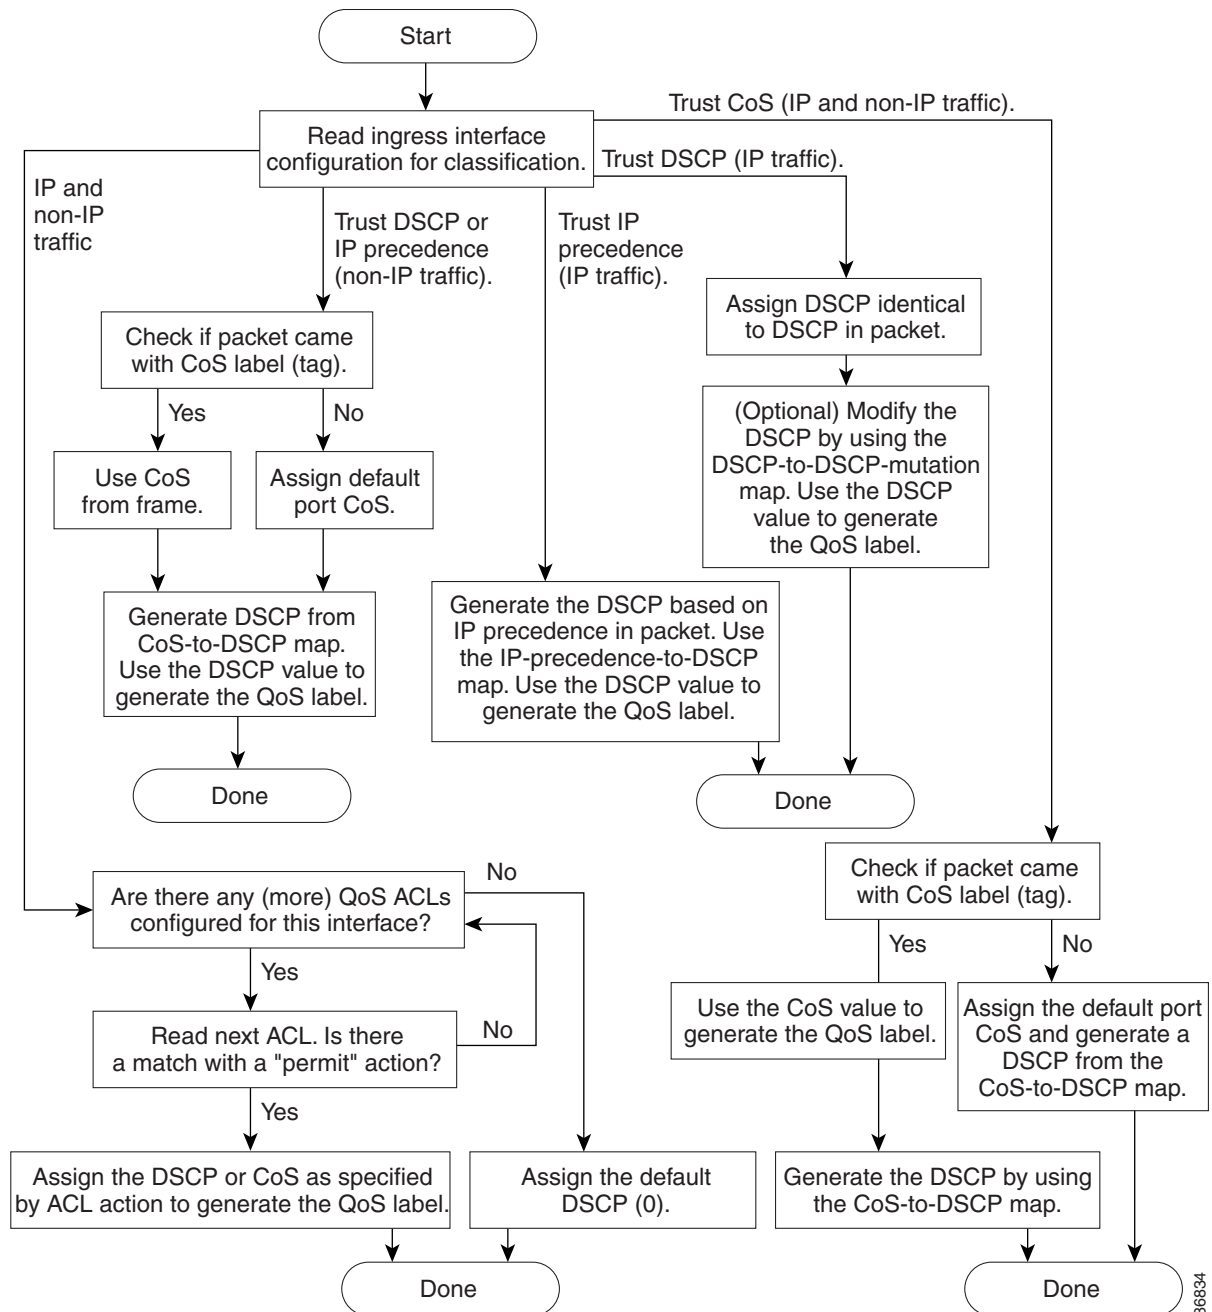
- **show mls qos** (as shown in [Example 2-2](#))

Example 2-2 Verifying Global QoS on a Catalyst 3750-E—show mls qos

```
C3750-E#show mls qos  
QoS is enabled  
QoS ip packet dscp rewrite is enabled  
  
C3750-E#
```

Trust Models

The Catalyst 3750-E switch ports can be configured to statically trust CoS, DSCP, and IP Precedence (although this is considered to be relegated by DSCP-trust) or to dynamically and conditionally trust Cisco IP phones. By default, with QoS enabled, all ports are set to an untrusted state. The complete port trust classification flowchart for the Catalyst 3750-E switch product family is shown in [Figure 2-15](#).

Figure 2-15 Catalyst 3750-E Port Trust Classification Flowchart

86834

Trust-CoS Model

A Catalyst 3750-E switch port can be configured to trust CoS by configuring the interface with the **mls qos trust cos** command. However, if an interface is set to trust CoS, then it by default calculates a packet's internal DSCP to be the incoming packet's (CoS value * 8). While this may be suitable for most markings, this default mapping may not be suitable for VoIP, as VoIP is usually marked CoS 5, which

would map by default to DSCP 40 (and not 46, which is the EF PHB as defined by RFC 3246). Therefore, if an interface is set to trust CoS, then the default CoS-to-DSCP mapping table should be modified such that CoS 5 maps to DSCP 46, as shown in [Example 2-3](#).

Example 2-3 Configuring Trust CoS and CoS-to-DSCP Mapping Modification on a Catalyst 3750-E

```
C3750-E(config)#mls qos map cos-dscp 0 8 16 24 32 46 48 56
! CoS 5 (the sixth CoS value, starting from 0) is mapped to 46
C3750-E(config)#interface GigabitEthernet 1/0/1
C3750-E(config-if)#mls qos trust cos
! The interface is set to statically trust CoS
```

This configuration can be verified with the commands:

- **show mls qos map cos-dscp** (as shown in [Example 2-4](#))
- **show mls qos interface** (as shown in [Example 2-5](#))

Example 2-4 Verifying Global CoS-to-DSCP Mapping Modifications on a Catalyst 3750-E—show mls qos map cos-dscp

```
C3750-E#show mls qos map cos-dscp
Cos-dscp map:
    cos:   0   1   2   3   4   5   6   7
-----
    dscp:  0  8 16 24 32 46 48 56
```

```
C3750-E#
```

In [Example 2-4](#), the CoS-to-DSCP mapping value for CoS 5 has been modified from the default mapping of 40 (CoS 5 * 8) to 46 (to match the recommendation from RFC 3246 that realtime applications be marked DSCP 46/EF).

Example 2-5 Verifying Interface Trust Settings on a Catalyst 3750-E—show mls qos interface

```
C3750-E#show mls qos interface GigabitEthernet 1/0/1
GigabitEthernet1/0/1
trust state: trust cos
trust mode: trust cos
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
qos mode: port-based

C3750-E#
```

In [Example 2-5](#), the port trust mode is set to trust CoS and the current (static) state of the interface is likewise set to trust CoS.

Trust-DSCP Model

Because of the additional granularity of DSCP versus QoS markings, it is generally recommended to trust DSCP rather than CoS (everything else being held equal). A Catalyst 3750-E switch port can be configured to trust DSCP with the **mls qos trust dscp** interface command, as shown in [Example 2-6](#).

Example 2-6 Configuring Trust-DSCP on a Catalyst 3750-E

```
C3750-E(config)#interface GigabitEthernet 1/0/1
C3750-E(config-if)#mls qos trust dscp
! The interface is set to statically trust DSCP
```

This configuration can be verified with the command:

- **show mls qos interface** (as shown in [Example 2-7](#))

Example 2-7 Verifying Interface Trust Settings on a Catalyst 3750-E—show mls qos interface

```
C3750-E#show mls qos interface GigabitEthernet 1/0/1
GigabitEthernet1/0/1
trust state: trust dscp
trust mode: trust dscp
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
qos mode: port-based

C3750-E#
```

In [Example 2-7](#), the port trust mode is set to trust DSCP and the current (static) state of the interface is likewise set to trust DSCP.

Conditional-Trust Model

In addition to configuring switch ports to statically trust endpoints, the Catalyst 3750-E family supports dynamic, conditional trust with the **mls qos trust device** interface command, which can be configured with the **cisco-phone** keyword to extend trust to Cisco IP phones, after these have been verified via a CDP-negotiation. Additionally, the type of trust to be extended must be specified (either CoS or DSCP). When configuring conditional trust to Cisco IP Phones, it is recommended to dynamically extend CoS-Trust, as Cisco IP Phones can only remark PC QoS markings at Layer 2 (CoS) and not at Layer 3 (DSCP). For other endpoints that do not have this remarking limitation, it is recommended to dynamically extend DSCP-trust (over CoS-trust), not only because DSCP has greater marking granularity, but also because the type of trust configured on the **ingress** switch port on a Catalyst 3750-E family of switches ultimately determines the type of queuing policies that are applied on the **egress** switch port. Specifically, if an ingress switch port is configured to trust-CoS—whether this is configured statically or dynamically (in conjunction with the **mls qos trust device** interface command)—a CoS-to-queue mapping determines the (ingress and) egress queuing policy. Conversely, if an ingress switch port is configured to trust-DSCP—whether this is configured statically or dynamically—a DSCP-to-queue mapping determines the (ingress and) egress queuing policy. Since DSCP-to-queue mapping has more granular policy options, it is the preferred way to assign packets to queues and as such depends on the ingress switch port being set to trust DSCP.

An example of a dynamic, conditional trust policy that is set to extend CoS-trust to CDP-verified Cisco IP phones is shown in [Example 2-8](#).

Example 2-8 Configuring (CoS-mode) Conditional Trust on a Catalyst 3750-E

```
C3750-E(config)#mls qos map cos-dscp 0 8 16 24 32 46 48 56
! CoS 5 (the sixth CoS value, starting from 0) is mapped to 46
C3750-E(config)#interface GigabitEthernet 1/0/1
C3750-E(config-if)# switchport access vlan 10
C3750-E(config-if)# switchport voice vlan 110
```

```

C3750-E(config-if)# spanning-tree portfast
C3750-E(config-if)# mls qos trust device cisco-phone
! The interface is set to conditionally-trust Cisco IP Phones
C3750-E(config-if)# mls qos trust cos
! CoS-trust will be dynamically extended to Cisco IP Phones

```

This configuration can be verified with the command:

- **show mls qos interface** (as shown in [Example 2-9](#))

Example 2-9 Verifying Interface Trust Settings on a Catalyst 3750-E—show mls qos interface

```

C3750-E#show mls qos interface GigabitEthernet 1/0/1
GigabitEthernet1/0/1
trust state: trust cos
trust mode: trust cos
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: cisco-phone
qos mode: port-based

C3750-E#

```

In [Example 2-9](#), the trust device feature has been enabled, with the trusted device being specified as a **cisco-phone**. The port trust mode—that is, the mode of trust (CoS | DSCP | IP Precedence) that is extended dynamically to the IP phone—is set to trust CoS. Similarly, the current (dynamic) trust state of the interface is likewise set to trust CoS. This is because there is a Cisco IP phone currently connected to the switch port; if this IP phone is removed from the switch port, the trust state of the interface toggles to “not trusted”.

Marking Models

The Catalyst 3750-E family of switches supports two main marking models:

- Per-port marking model—This is the only option on Catalyst 2960 and 2975 series switches, as these do not support multilayer switching (and therefore do not support SVI interfaces and per-VLAN policies).
- Per-VLAN marking model—This model is supported on the Catalyst 3560G, 3750G, 3560-E, and 3750-E series switches.

Each model is detailed in the following sections.

Per-Port Marking Model

The per-port marking model (based on [Figure 2-10](#)) matches VoIP and signaling traffic from the VVLAN by matching on DSCP EF and CS3, respectively. Multimedia conferencing traffic from the DVLAN is matched by UDP/RTP ports 16384-32767. Signaling traffic is matched on SCCP ports (TCP 2000-2002), as well as on SIP ports (TCP/UDP 5060-5061). Other transactional data traffic, bulk data, and scavenger traffic are matched on various ports (outlined in [Figure 2-9](#)). The service policy is applied to an interface range, along with (DSCP-mode) conditional trust, as shown in [Example 2-10](#).

Example 2-10 Per-Port Marking Configuration Example on a Catalyst 3750-E

```
! This first section configures IP access-lists to match applications
```

```
C3750-E(config)#ip access-list extended MULTIMEDIA-CONFERENCING
C3750-E(config-ext-nacl)# remark RTP
C3750-E(config-ext-nacl)# permit udp any any range 16384 32767

C3750-E(config)#ip access-list extended SIGNALING
C3750-E(config-ext-nacl)# remark SCCP
C3750-E(config-ext-nacl)# permit tcp any any range 2000 2002
C3750-E(config-ext-nacl)# remark SIP
C3750-E(config-ext-nacl)# permit tcp any any range 5060 5061
C3750-E(config-ext-nacl)# permit udp any any range 5060 5061

C3750-E(config)#ip access-list extended TRANSACTIONAL-DATA
C3750-E(config-ext-nacl)# remark HTTPS
C3750-E(config-ext-nacl)# permit tcp any any eq 443
C3750-E(config-ext-nacl)# remark ORACLE-SQL*NET
C3750-E(config-ext-nacl)# permit tcp any any eq 1521
C3750-E(config-ext-nacl)# permit udp any any eq 1521
C3750-E(config-ext-nacl)# remark ORACLE
C3750-E(config-ext-nacl)# permit tcp any any eq 1526
C3750-E(config-ext-nacl)# permit udp any any eq 1526
C3750-E(config-ext-nacl)# permit tcp any any eq 1575
C3750-E(config-ext-nacl)# permit udp any any eq 1575
C3750-E(config-ext-nacl)# permit tcp any any eq 1630
C3750-E(config-ext-nacl)# permit udp any any eq 1526

C3750-E(config)#ip access-list extended BULK-DATA
C3750-E(config-ext-nacl)# remark FTP
C3750-E(config-ext-nacl)# permit tcp any any eq ftp
C3750-E(config-ext-nacl)# permit tcp any any eq ftp-data
C3750-E(config-ext-nacl)# remark SSH/SFTP
C3750-E(config-ext-nacl)# permit tcp any any eq 22
C3750-E(config-ext-nacl)# remark SMTP/SECURE SMTP
C3750-E(config-ext-nacl)# permit tcp any any eq smtp
C3750-E(config-ext-nacl)# permit tcp any any eq 465
C3750-E(config-ext-nacl)# remark IMAP/SECURE IMAP
C3750-E(config-ext-nacl)# permit tcp any any eq 143
C3750-E(config-ext-nacl)# permit tcp any any eq 993
C3750-E(config-ext-nacl)# remark POP3/SECURE POP3
C3750-E(config-ext-nacl)# permit tcp any any eq pop3
C3750-E(config-ext-nacl)# permit tcp any any eq 995
C3750-E(config-ext-nacl)# remark CONNECTED PC BACKUP
C3750-E(config-ext-nacl)# permit tcp any eq 1914 any

C3750-E(config)#ip access-list extended SCAVENGER
C3750-E(config-ext-nacl)# remark KAZAA
C3750-E(config-ext-nacl)# permit tcp any any eq 1214
C3750-E(config-ext-nacl)# permit udp any any eq 1214
C3750-E(config-ext-nacl)# remark MICROSOFT DIRECT X GAMING
C3750-E(config-ext-nacl)# permit tcp any any range 2300 2400
C3750-E(config-ext-nacl)# permit udp any any range 2300 2400
C3750-E(config-ext-nacl)# remark APPLE ITUNES MUSIC SHARING
C3750-E(config-ext-nacl)# permit tcp any any eq 3689
C3750-E(config-ext-nacl)# permit udp any any eq 3689
C3750-E(config-ext-nacl)# remark BITTORRENT
C3750-E(config-ext-nacl)# permit tcp any any range 6881 6999
C3750-E(config-ext-nacl)# remark YAHOO GAMES
C3750-E(config-ext-nacl)# permit tcp any any eq 11999
C3750-E(config-ext-nacl)# remark MSN GAMING ZONE
C3750-E(config-ext-nacl)# permit tcp any any range 28800 29100

C3750-E(config)#ip access-list extended DEFAULT
C3750-E(config-ext-nacl)# remark EXPLICIT CLASS-DEFAULT
C3750-E(config-ext-nacl)# permit ip any any
```

```

! This section configures the class-maps
C3750-E(config-cmap)#class-map match-all VVLAN-VOIP
C3750-E(config-cmap)# match ip dscp ef
! VoIP is trusted (from the VVLAN)

C3750-E(config-cmap)#class-map match-all VVLAN-SIGNALING
C3750-E(config-cmap)# match ip dscp cs3
! Signaling is trusted (from the VVLAN)

C3750-E(config-cmap)#class-map match-all MULTIMEDIA-CONFERENCING
C3750-E(config-cmap)# match access-group name MULTIMEDIA-CONFERENCING
! Associates MULTIMEDIA-CONFERENCING access-list with class-map

C3750-E(config-cmap)#class-map match-all SIGNALING
C3750-E(config-cmap)# match access-group name SIGNALING
! Associates SIGNALING access-list with class-map

C3750-E(config-cmap)#class-map match-all TRANSACTIONAL-DATA
C3750-E(config-cmap)# match access-group name TRANSACTIONAL-DATA
! Associates TRANSACTIONAL-DATA access-list with class-map

C3750-E(config-cmap)#class-map match-all BULK-DATA
C3750-E(config-cmap)# match access-group name BULK-DATA
! Associates BULK-DATA access-list with class-map

C3750-E(config-cmap)#class-map match-all SCAVENGER
C3750-E(config-cmap)# match access-group name SCAVENGER
! Associates SCAVENGER access-list with class-map

C3750-E(config-cmap)#class-map match-all DEFAULT
C3750-E(config-cmap)# match access-group name DEFAULT
! Associates DEFAULT access-list with class-map

! This section configures the Per-Port ingress marking policy-map
C3750-E(config-cmap)#policy-map PER-PORT-MARKING
C3750-E(config-pmap)# class VVLAN-VOIP
C3750-E(config-pmap-c)# set dscp ef
! VoIP is marked EF (see note below)
C3750-E(config-pmap-c)# class VVLAN-SIGNALING
C3750-E(config-pmap-c)# set dscp cs3
! Signaling (from the VVLAN) is marked CS3 (see note below)
C3750-E(config-pmap-c)# class MULTIMEDIA-CONFERENCING
C3750-E(config-pmap-c)# set dscp af41
! Multimedia-conferencing is marked AF41
C3750-E(config-pmap-c)# class SIGNALING
C3750-E(config-pmap-c)# set dscp cs3
! Signaling (from the DVLAN) is marked CS3
C3750-E(config-pmap-c)# class TRANSACTIONAL-DATA
C3750-E(config-pmap-c)# set dscp af21
! Transactional Data is marked AF21
C3750-E(config-pmap-c)# class BULK-DATA
C3750-E(config-pmap-c)# set dscp af11
! Bulk Data is marked AF11
C3750-E(config-pmap-c)# class SCAVENGER
C3750-E(config-pmap-c)# set dscp cs1
! Scavenger traffic is marked CS1
C3750-E(config-pmap-c)# class DEFAULT
C3750-E(config-pmap-c)# set dscp default
! An explicit class-default marks all other IP traffic to 0 (see note)

! This section attaches the service-policy to the interface(s)
C3750-E(config)#interface range GigabitEthernet 1/0/1-48

```



```

C3750-E(config-if-range)# switchport access vlan 10
C3750-E(config-if-range)# switchport voice vlan 110
C3750-E(config-if-range)# spanning-tree portfast
C3750-E(config-if-range)# mls qos trust device cisco-phone
! The interface is set to conditionally-trust Cisco IP Phones
C3750-E(config-if-range)# mls qos trust cos
! CoS-trust will be dynamically extended to Cisco IP Phones
C3750-E(config-if-range)# service-policy input PER-PORT-MARKING
! Attaches the Per-Port Marking policy to the interface(s)

```

**Note**

While the Catalyst 3750-E MQC syntax includes an implicit class-default, any policy actions assigned to this class are not enforced. Therefore, an explicit class DEFAULT is configured in [Example 2-10](#) to enforce a marking/remarking policy to DSCP 0 for all other IP traffic.

**Note**

An explicit marking command (**set dscp**) is used even for trusted application classes (like VVLAN-VOIP and VVLAN-SIGNALING) rather than a **trust** policy-map action. This is because a trust statement in a policy map requires multiple hardware entries and, as such, might be too large to fit into the available QoS hardware memory, triggering an error when the policy map is applied to a port. The use of an explicit (but seemingly redundant) explicit marking command actually improves the policy efficiency from a hardware perspective.

This configuration can be verified with the commands:

- **show mls qos interface** (as shown in [Example 2-11](#))
- **show class-map** (as shown in [Example 2-12](#))
- **show policy-map** (as shown in [Example 2-13](#))
- **show policy-map interface** (as shown in [Example 2-14](#))

Example 2-11 Verifying Interface Trust and Policy Settings on a Catalyst 3750-E—show mls qos interface

```

C3750-E#show mls qos interface GigabitEthernet 1/0/1
GigabitEthernet1/0/1
Attached policy-map for Ingress: PER-PORT-MARKING
trust state: trust cos
trust mode: trust cos
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: cisco-phone
qos mode: port-based

C3750-E#

```

In [Example 2-11](#), CoS-mode conditional trust has been applied to the interface (which allows the port to dynamically extend CoS-trust to the Cisco IP phone, such that VVLAN-VoIP and VVLAN-Signaling traffic can be matched on CoS 5 and 3, respectively). Additionally the PER-PORT-MARKING service policy has been attached to the interface to classify both VVLAN and DVLAN traffic.

Example 2-12 Verifying Class Maps on a Catalyst 3750-E—show class-map

```

C3750-E#show class-map

```

```

Class Map match-any class-default (id 0)
  Match any

Class Map match-all BULK-DATA (id 6)
  Match access-group name BULK-DATA

Class Map match-all VVLAN-SIGNALING (id 2)
  Match ip dscp cs3 (24)

Class Map match-all MULTIMEDIA-CONFERENCING (id 3)
  Match access-group name MULTIMEDIA-CONFERENCING

Class Map match-all DEFAULT (id 8)
  Match access-group name DEFAULT

Class Map match-all SCAVENGER (id 7)
  Match access-group name SCAVENGER

Class Map match-all SIGNALING (id 4)
  Match access-group name SIGNALING

Class Map match-all VVLAN-VOIP (id 1)
  Match ip dscp ef (46)

Class Map match-all TRANSACTIONAL-DATA (id 5)
  Match access-group name TRANSACTIONAL-DATA

C3750-E#

```

Example 2-13 Verifying Policy Maps on a Catalyst 3750-E—show policy-map

```

C3750-E#show policy-map
Policy Map PER-PORT-MARKING
  Class VVLAN-VOIP
    set dscp ef
  Class VVLAN-SIGNALING
    set dscp cs3
  Class MULTIMEDIA-CONFERENCING
    set dscp af41
  Class SIGNALING
    set dscp cs3
  Class TRANSACTIONAL-DATA
    set dscp af21
  Class BULK-DATA
    set dscp af11
  Class SCAVENGER
    set dscp cs1
  Class DEFAULT
    set dscp default

C3750-E#

```

Example 2-14 Verifying Service Policies on a Catalyst 3750-E—show policy-map interface

```

C3750-E#show policy-map interface GigabitEthernet 1/0/1
GigabitEthernet1/0/1

Service-policy input: PER-PORT-MARKING

Class-map: VVLAN-VOIP (match-all)

```

```

    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: ip dscp ef (46)

Class-map: VVLAN-SIGNALING (match-all)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: ip dscp cs3 (24)

Class-map: MULTIMEDIA-CONFERENCING (match-all)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: access-group name MULTIMEDIA-CONFERENCING

Class-map: SIGNALING (match-all)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: access-group name SIGNALING

Class-map: TRANSACTIONAL-DATA (match-all)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: access-group name TRANSACTIONAL-DATA

Class-map: BULK-DATA (match-all)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: access-group name BULK-DATA

Class-map: SCAVENGER (match-all)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: access-group name SCAVENGER

Class-map: DEFAULT (match-all)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: access-group name DEFAULT

Class-map: class-default (match-any)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: any
        0 packets, 0 bytes
        5 minute rate 0 bps
C3750-E#

```

As shown in [Example 2-14](#), unlike the **show policy-map interface** outputs on IOS routers, the corresponding command on the Catalyst 3750-E series of switches does not dynamically increment packet, byte, drop, and bps counters.

Per-VLAN Marking Model

An alternative approach for deploying marking policies on the Catalyst 3560/3750 platforms is to deploy these on a per-VLAN basis. In order to do so, the interfaces belonging to the VLANs need to be configured with the **mls qos vlan-based** interface command. Additionally, the policy-map can be simplified/broken-apart, as applicable to each VLAN. Adapting the previous example to a VLAN-based marking policing allows for the VVLAN-based policy map to be reduced to only three explicit classes:

VoIP, signaling, and the explicit default class. Similarly, the DVLAN-based policy map is reduced to six explicit classes: multimedia conferencing, signaling, transactional data, bulk data, scavenger, and the explicit default class. A per-VLAN marking model is shown in [Example 2-15](#).

**Note**

As the access lists and class maps are identical to [Example 2-14](#), these are omitted for brevity in this—and in following—examples for this switch platform family.

Example 2-15 Per-VLAN Marking Configuration Example on a Catalyst 3750-E

```

! This section configures the ingress marking policy-map for the VVLAN
C3750-E(config)#policy-map VVLAN-MARKING
C3750-E(config-pmap)# class VVLAN-VOIP
C3750-E(config-pmap-c)# set dscp ef
! VoIP is trusted (from the VVLAN)
C3750-E(config-pmap-c)# class VVLAN-SIGNALING
C3750-E(config-pmap-c)# set dscp cs3
! Signaling is trusted (from the VVLAN)
C3750-E(config-pmap-c)# class DEFAULT
C3750-E(config-pmap-c)# set dscp default
! An explicit DEFAULT class marks all other VVLAN IP traffic to DF

! This section configures the ingress marking policy-map for the DVLAN
C3750-E(config)#policy-map DVLAN-MARKING
C3750-E(config-pmap)# class MULTIMEDIA-CONFERENCING
C3750-E(config-pmap-c)# set dscp af41
! Multimedia-conferencing is marked AF41
C3750-E(config-pmap-c)# class SIGNALING
C3750-E(config-pmap-c)# set dscp cs3
! Signaling (from the DVLAN) is marked CS3
C3750-E(config-pmap-c)# class TRANSACTIONAL-DATA
C3750-E(config-pmap-c)# set dscp af21
! Transactional Data is marked AF21
C3750-E(config-pmap-c)# class BULK-DATA
C3750-E(config-pmap-c)# set dscp af11
! Bulk Data is marked AF11
C3750-E(config-pmap-c)# class SCAVENGER
C3750-E(config-pmap-c)# set dscp cs1
! Scavenger traffic is marked CS1
C3750-E(config-pmap-c)# class DEFAULT
C3750-E(config-pmap-c)# set dscp default
! An explicit DEFAULT class marks all other DVLAN IP traffic to DF

! This section configures the interface(s) for conditional trust
! and enables VLAN-based QoS
C3750-E(config)#interface range GigabitEthernet 1/0/1-48
C3750-E(config-if-range)# switchport access vlan 10
C3750-E(config-if-range)# switchport voice vlan 110
C3750-E(config-if-range)# spanning-tree portfast
C3750-E(config-if-range)# mls qos trust device cisco-phone
! The interface is set to conditionally-trust Cisco IP Phones
C3750-E(config-if-range)# mls qos vlan-based
! Enables VLAN-based QoS on the interface(s)

! This section attaches the DVLAN policy to the DVLAN interface
C3750-E(config)#interface Vlan 10
C3750-E(config-if)# description DVLAN
C3750-E(config-if)# service-policy input DVLAN-MARKING

```

```

! Attaches the DVLAN Per-VLAN Marking policy to the DVLAN interface

! This section attaches the VVLAN policy to the VVLAN interface
C3750-E(config)#interface Vlan 110
C3750-E(config-if)# description VVLAN
C3750-E(config-if)# service-policy input VVLAN-MARKING
! Attaches the VVLAN Per-VLAN Marking policy to the VVLAN interface

```

This configuration can be verified with the commands:

- **show mls qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Policing Models

The Catalyst 3750-E family of switches support 256 policers per hardware ASIC. These switches share 2, 4, 6, 8, or 24 ports per ASIC (this depends on the platform and hardware configuration). The number of ASICs for a specific switch can be verified by using the **show platform port-asic version** verification command. Additionally, the specific switch ports associated with each ASIC can further be identified by the **show platform pm platform-block** verification command (in the ASIC column).

As a reminder, these policing caveats apply to these switches:

- The Catalyst 2960 and 2975 can only police to a minimum rate of 1 Mbps; all other platforms within this switch-product family can police to a minimum rate of 8 kbps.
- Only the Catalyst 3650-E and 3750-E support policing on 10 Gigabit Ethernet interfaces.

The Catalyst 3750-E family of switches supports these ingress policing models:

- Per-port policing model—This model (which is the only option on Catalyst 2960 and 2975 series switches—as these do not support multilayer switching and therefore do not support SVI interfaces and per-VLAN policies) attaches policers to physical switch port interfaces.
- Per-VLAN policing model—This model (which is supported on the Catalyst 3560G, 3750G, 3560-E, and 3750-E series switches) attaches policers to logical VLAN interfaces. However, there is an inherent limitation with this policing model: it only supports a single-aggregate policer per VLAN and—since the number of ports associated with a VLAN is dynamic and variable—thus is quite restricted in overall policing effectiveness. Therefore, it is generally recommended to use the per-port/per-VLAN policing model instead, as it offers more discrete policing options.
- Per-port/per-VLAN policing model—This model (which is supported on the Catalyst 3560G, 3750G, 3560-E, and 3750-E series switches) attaches policers to discrete VLANs traversing a single switch trunk interface.

The per-port and per-port/per-VLAN policing models for the Catalyst 3750-E family of switches are detailed in the following sections.

Per-Port Policing Model

The per-port policing model is quite similar to the per-port marking model, except that the policy action includes a policing function—in some cases to drop, in others to remark. As shown in [Figure 2-10](#), the VoIP and signaling traffic from the VVLAN can be policed to drop at 128 kbps and 32 kbps, respectively (as any excessive traffic matching this criteria would be indicative of network abuse). Similarly, the

multimedia conferencing, signaling, and scavenger traffic from the DVLAN can be policed to drop. On the other hand, data plane policing policies can be applied to transactional, bulk, and best effort data traffic, such that these flows are subject to being remarked (but not dropped at the ingress edge) when severely out-of-profile. Remarketing is performed by configuring a policed-DSCP map with the global configuration command **mls qos map policed-dscp**, which specifies which DSCP values are subject to remarketing if out-of-profile and what value these should be remarked as (which in the case of data plane policing/scavenger class QoS policies, this value is CS1/DSCP 8). A per-port policing model for a Catalyst 3750-E is shown in [Example 2-16](#).

Example 2-16 Per-Port Policing Configuration Example on a Catalyst 3750-E

```
! This section configures the global policed-DSCP markdown map
C3750-E(config)#mls qos map policed-dscp 0 10 18 to 8
! DSCP 0 (DF), 10 (AF11) and 18 (AF21) are marked down to 8 (CS1)
! if found to be in excess of their (respective) policing rates

! This section configures the Per-Port policing policy-map
C3750-E(config)#policy-map PER-PORT-POLICING
C3750-E(config-pmap)# class VVLAN-VOIP
C3750-E(config-pmap-c)# set dscp ef
C3750-E(config-pmap-c)# police 128k 8000 exceed-action drop
! VoIP is marked EF and policed to drop at 128 kbps
C3750-E(config-pmap-c)# class VVLAN-SIGNALING
C3750-E(config-pmap-c)# set dscp cs3
C3750-E(config-pmap-c)# police 32k 8000 exceed-action drop
! (VVLAN) Signaling is marked CS3 and policed to drop at 32 kbps
C3750-E(config-pmap-c)# class MULTIMEDIA-CONFERENCING
C3750-E(config-pmap-c)# set dscp af41
C3750-E(config-pmap-c)# police 5m 8000 exceed-action drop
! Multimedia-conferencing is marked AF41 and policed to drop at 5 Mbps
C3750-E(config-pmap-c)# class SIGNALING
C3750-E(config-pmap-c)# set dscp cs3
C3750-E(config-pmap-c)# police 32k 8000 exceed-action drop
! (DVLAN) Signaling is marked CS3 and policed to drop at 32 kbps
C3750-E(config-pmap-c)# class TRANSACTIONAL-DATA
C3750-E(config-pmap-c)# set dscp af21
C3750-E(config-pmap-c)# police 10m 8000 exceed-action policed-dscp-transmit
! Trans-data is marked AF21 and policed to remark (to CS1) at 10 Mbps
C3750-E(config-pmap-c)# class BULK-DATA
C3750-E(config-pmap-c)# set dscp af11
C3750-E(config-pmap-c)# police 10m 8000 exceed-action policed-dscp-transmit
! Bulk-data is marked AF11 and policed to remark (to CS1) at 10 Mbps
C3750-E(config-pmap-c)# class SCAVENGER
C3750-E(config-pmap-c)# set dscp cs1
C3750-E(config-pmap-c)# police 10m 8000 exceed-action drop
! Scavenger traffic is marked CS1 and policed to drop at 10 Mbps
C3750-E(config-pmap-c)# class DEFAULT
C3750-E(config-pmap-c)# set dscp default
C3750-E(config-pmap-c)# police 10m 8000 exceed-action policed-dscp-transmit
! An explicit default class marks all other IP traffic to DF
! and polices all other IP traffic to remark (to CS1) at 10 Mbps

! This section attaches the service-policy to the interface(s)
C3750-E(config)#interface range GigabitEthernet 1/0/1-48
C3750-E(config-if-range)# switchport access vlan 10
C3750-E(config-if-range)# switchport voice vlan 110
C3750-E(config-if-range)# spanning-tree portfast
C3750-E(config-if-range)# mls qos trust device cisco-phone
! The interface is set to conditionally-trust Cisco IP Phones
```

```
C3750-E(config-if-range)# mls qos trust cos
! CoS-trust will be dynamically extended to Cisco IP Phones
C3750-E(config-if-range)# service-policy input PER-PORT-POLICING
! Attaches the Per-Port Policing policy to the interface(s)
```

**Note**

Catalyst 3750-G software allows for policing rates to be entered using the postfixes **k** (for kilobits), **m** (for megabits), and **g** (for gigabits), as shown in [Example 2-16](#). Additionally, decimal points are allowed in conjunction with these postfixes; for example, a rate of 10.5 Mbps could be entered with the policy-map command **police 10.5m**. While these policing rates are converted to their full bps values within the configuration, it makes the entering of these rate more user-friendly and less error prone (as could easily be the case when having to enter up to 10 zeros to define the policing rate).

This configuration can be verified with the commands:

- **show mls qos maps policed-dscp** (as shown in [Example 2-17](#))
- **show mls qos interface**
- **show mls qos interface interface x/y policers** (as shown in [Example 2-18](#))
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Example 2-17 Verifying Global Policing Markdown Mappings on a Catalyst 3750-E—show mls qos maps policed-dscp

```
C3750-E#show mls qos maps policed-dscp
Policed-dscp map:
d1 : d2 0 1 2 3 4 5 6 7 8 9
-----
0 : 08 01 02 03 04 05 06 07 08 09
1 : 08 11 12 13 14 15 16 17 08 19
2 : 20 21 22 23 24 25 26 27 28 29
3 : 30 31 32 33 34 35 36 37 38 39
4 : 40 41 42 43 44 45 46 47 48 49
5 : 50 51 52 53 54 55 56 57 58 59
6 : 60 61 62 63
```

C3750-E#

In [Example 2-17](#), the policing DSCP-markdown mapping is shown. The first digit of the DSCP value of a packet offered to a policer is shown along the Y-axis of the table; the second digit of the DSCP value of a packet offered to a policer is shown along the X-axis of the table. For example, the DSCP value for the transactional data application class (AF21/18) is found in the row d1=1 and column d2=8. And, as shown, packets with this offered DSCP value (along with DF/0 and AF11/10) are remarked to CS1 (08) if found to be in excess of the policing rate.

Example 2-18 Verifying Interface Policers on a Catalyst 3750-E—show mls qos interface interface x/y policers

```
C3750-E#show mls qos interface GigabitEthernet 1/0/1 policers
GigabitEthernet1/0/1
policy-map=PER-PORT-POLICING

type=Single, id=1 rate=128000, qlimit=8000, drop=1
```

```

type=Single, id=2 rate=32000, qlimit=8000, drop=1
type=Single, id=3 rate=5000000, qlimit=8000, drop=1
type=Single, id=4 rate=32000, qlimit=8000, drop=1
type=Single, id=5 rate=10000000, qlimit=8000, drop=0
type=Single, id=6 rate=10000000, qlimit=8000, drop=0
type=Single, id=7 rate=10000000, qlimit=8000, drop=1
type=Single, id=8 rate=10000000, qlimit=8000, drop=0
C3750-E#

```

In [Example 2-18](#), the interface policers for GigabitEthernet 1/0/1 are shown, including the policing rates, burst, and drop-function values (drop=1 means that exceeding-traffic is dropped, while drop=0 value means that exceeding-traffic is not dropped, but remarked).

Per-Port/Per-VLAN Policing Model

An alternative—and more discrete—approach for deploying policing policies on the Catalyst 3560/3750 platforms is to deploy these on a per-port/per-VLAN basis, which (on this family of switch platforms) requires the use of hierarchical QoS policies, also known as nested QoS policies.

The first step is to configure a class-map that defines the switch port(s) to which the policers are attached. Then one or more per-port policers need to be defined (according to the various levels of policing rates or exceeding-actions required); these policers reference the previously-defined class map that specifies the switch port(s) are policed. These per-port policers comprise the “child” policy maps in the hierarchy.

Following this, “parent” policy maps are configured that combine the various per-port policers for the various classes of traffic for a given VLAN. Each of these parent policy maps reference child policies that implement the per-port policing functions. Finally, these parent policy maps are applied to the VLAN SVI interfaces.

In [Example 2-18](#), a class map (VLAN-10/110-PORTS) defines the ports on which the policers are enforced, specifically the ports belonging to DVLAN 10 and VVLAN 110 (which in this example equates to Gigabit Ethernet ports 1/0/1 through 1/0/48). Then a series of per-port policers (child policy maps) are defined, one each for 128 kbps (with a dropping action), 32 kbps (with a dropping action), 5 Mbps (with a dropping action), 10 Mbps (with a dropping action), and 10 Mbps (with a remarking action). Following this, a parent policy map for the VVLAN references the child policy maps to police VoIP to 128 kbps, (VLAN) signaling to 32 kbps, and all other VVLAN IP traffic to 32 kbps.

Similarly, a parent policy map for the DVLAN references the child policy maps to police multimedia conferencing to 5 Mbps, (DVLAN) signaling to 32 kbps, and scavenger traffic to 10 Mbps. However, data plane policing (scavenger class QoS) policies are applied to the transactional data, bulk data, and the (explicitly defined) best effort class to police these (respectively) to 10 Mbps with a remarking action and not a dropping action.

As in the previous example, remarking is performed by configuring a policed DSCP map with the global configuration command **mls qos map policed-dscp**, which specifies which DSCP values are subject to remarking if out-of-profile and what value these should be remarked as (which in the case of data plane policing/scavenger class QoS policies, this value is CS1/DSCP 8). The switch ports have VLAN-based QoS enabled on them and the parent service policies are applied to the VLAN SVI interfaces for the DVLAN and the VVLAN.

Example 2-19 Per-Port/Per-VLAN Policing Configuration Example on a Catalyst 3750-E

```

! This section configures the global policed-DSCP markdown map
C3750-E(config)#mls qos map policed-dscp 0 10 18 to 8
! DSCP 0 (DF), 10 (AF11) and 18 (AF21) are marked down to 8 (CS1)
! if found to be in excess of their (respective) policing rates

! This section configures the class-map of switch ports
! that the Per-Port/Per-VLAN policing actions will be enforced on
C3750-E(config)#class-map match-all VLAN-10/110-PORTS
C3750-E(config-cmap)# match input-interface GigabitEthernet1/0/1 - GigabitEthernet1/0/48

! This section configures Per-Port policers
C3750-E(config-pmap)#policy-map PER-PORT-POLICER-128K-DROP
C3750-E(config-pmap)# class VLAN-10/110-PORTS
C3750-E(config-pmap-c)# police 128k 8000 exceed-action drop
! This policy-map configures a 128 kbps dropping policer for the ports

C3750-E(config-pmap)#policy-map PER-PORT-POLICER-32K-DROP-VVLAN
C3750-E(config-pmap)# class VLAN-10/110-PORTS
C3750-E(config-pmap-c)# police 32k 8000 exceed-action drop
! This policy-map configures a 32 kbps dropping policer for the ports
! but only for use on the VVLAN (see note below)

C3750-E(config-pmap)#policy-map PER-PORT-POLICER-32K-DROP-DVLAN
C3750-E(config-pmap)# class VLAN-10/110-PORTS
C3750-E(config-pmap-c)# police 32k 8000 exceed-action drop
! This policy-map configures a 32 kbps dropping policer for the ports
! but only for use on the DVLAN (see note below)

C3750-E(config-pmap)#policy-map PER-PORT-POLICER-5M-DROP
C3750-E(config-pmap)# class VLAN-10/110-PORTS
C3750-E(config-pmap-c)# police 5m 8000 exceed-action drop
! This policy-map configures a 5 Mbps dropping policer for the ports

C3750-E(config-pmap)#policy-map PER-PORT-POLICER-10M-DROP
C3750-E(config-pmap)# class VLAN-10/110-PORTS
C3750-E(config-pmap-c)# police 10m 8000 exceed-action drop
! This policy-map configures a 10 Mbps dropping policer for the ports

C3750-E(config-pmap)#policy-map PER-PORT-POLICER-10M-REMARK
C3750-E(config-pmap)# class VLAN-10/110-PORTS
C3750-E(config-pmap-c)# police 10m 8000 exceed-action policed-dscp-transmit
! This policy-map configures a 10 Mbps remarking policer for the ports

! This section combines the Per-Port policers for the VVLAN policy-map
C3750-E(config-pmap)#policy-map VVLAN-POLICERS
C3750-E(config-pmap)# class VVLAN-VOIP
C3750-E(config-pmap-c)# set dscp ef
C3750-E(config-pmap-c)# service-policy PER-PORT-POLICER-128K-DROP
! VoIP is marked to EF and
! (via a nested service-policy) is policed to drop at 128 kbps
C3750-E(config-pmap-c)# class VVLAN-SIGNALING
C3750-E(config-pmap-c)# set dscp cs3
C3750-E(config-pmap-c)# service-policy PER-PORT-POLICER-32K-DROP-VVLAN
! (VVLAN) Signaling is marked to CS3 and
! (via a nested service-policy) is policed to drop at 32 kbps
C3750-E(config-pmap-c)# class DEFAULT
C3750-E(config-pmap-c)# set dscp default

```

```

C3750-E(config-pmap-c)# service-policy PER-PORT-POLICER-32K-DROP-VVLAN
! An explicit default class marks all other VVLAN IP traffic to DF
! and (via a nested service-policy) polices to drop at 32 kbps
C3750-E(config-pmap-c)#exit
C3750-E(config-pmap)#
C3750-E(config-pmap)#

! This section combines the Per-Port policers for the DVLAN policy-map
C3750-E(config-pmap)#policy-map DVLAN-POLICERS
C3750-E(config-pmap)# class MULTIMEDIA-CONFERENCING
C3750-E(config-pmap-c)# set dscp af41
C3750-E(config-pmap-c)# service-policy PER-PORT-POLICER-5M-DROP
! Multimedia-conferencing is marked AF41 and
! (via a nested service-policy) is policed to drop at 5 Mbps
C3750-E(config-pmap-c)# class SIGNALING
C3750-E(config-pmap-c)# set dscp cs3
C3750-E(config-pmap-c)# service-policy PER-PORT-POLICER-32K-DROP-DVLAN
! (DVLAN) Signaling is marked CS3 and
! (via a nested service-policy) is policed to drop at 32 kbps
C3750-E(config-pmap-c)# class TRANSACTIONAL-DATA
C3750-E(config-pmap-c)# set dscp af21
C3750-E(config-pmap-c)# service-policy PER-PORT-POLICER-10M-REMARK
! Transactional-data is marked AF21 and (via a nested service-policy)
! is policed to remark (to CS1) at 10 Mbps
C3750-E(config-pmap-c)# class BULK-DATA
C3750-E(config-pmap-c)# set dscp af11
C3750-E(config-pmap-c)# service-policy PER-PORT-POLICER-10M-REMARK
! Bulk-data is marked AF11 and (via a nested service-policy)
! is policed to remark (to CS1) at 10 Mbps
C3750-E(config-pmap-c)# class SCAVENGER
C3750-E(config-pmap-c)# set dscp cs1
C3750-E(config-pmap-c)# service-policy PER-PORT-POLICER-10M-DROP
! Scavenger traffic is marked CS1 and (via a nested service-policy)
! is policed to drop at 10 Mbps
C3750-E(config-pmap-c)# class DEFAULT
C3750-E(config-pmap-c)# set dscp default
C3750-E(config-pmap-c)# service-policy PER-PORT-POLICER-10M-REMARK
! An explicit default class marks all other DVLAN IP traffic to DF and
! (via a nested service-policy) polices to remark (to CS1) at 10 Mbps

! This section configures the interfaces for conditional trust
! and enables VLAN-based QoS
C3750-E(config)#interface range GigabitEthernet 1/0/1-48
C3750-E(config-if-range)# switchport access vlan 10
C3750-E(config-if-range)# switchport voice vlan 110
C3750-E(config-if-range)# spanning-tree portfast
C3750-E(config-if-range)# mls qos trust device cisco-phone
! The interface is set to conditionally-trust Cisco IP Phones
C3750-E(config-if-range)# mls qos vlan-based
! Enables VLAN-based QoS on the interface(s)

! This section attaches the DVLAN policers to the DVLAN interface
C3750-E(config)#interface Vlan 10
C3750-E(config-if)# description DVLAN
C3750-E(config-if)# service-policy input DVLAN-POLICERS
! Attaches the DVLAN Per-VLAN Policing policy to the DVLAN interface

! This section attaches the VVLAN policers to the VVLAN interface

```

```
C3750-E(config)#interface Vlan 110
C3750-E(config-if)# description VVLAN
C3750-E(config-if)# service-policy input VVLAN-POLICERS
! Attaches the VVLAN Per-VLAN Policing policy to the VVLAN interface
```

**Note**

On Catalyst 3750-E switches, a policer cannot be attached to both a port and a SVI; separate policers must be configured for these different types of interfaces.

**Note**

On Catalyst 3750-E switches, a nested/child policy map can only be referenced by one parent service policy. Therefore, separate (child) policers are configured in [Example 2-19](#) for the signaling classes (one each for the DVLAN-POLICER parent policy map and another for the VVLAN-POLICER parent policy map).

**Note**

It is important to note that on Catalyst 3750G and 3750-E switches, when you enable VLAN-based QoS and configure a hierarchical policy map in a switch stack, these automatic actions occur when the stack configuration changes:

- When a new stack master is selected, the stack master re-enables and reconfigures these features on all applicable interfaces on the stack master.
- When a stack member is added, the stack master re-enables and reconfigures these features on all applicable ports on the stack member.
- When you merge switch stacks, the new stack master re-enables and reconfigures these features on the switches in the new stack.
- When the switch stack divides into two or more switch stacks, the stack master in each switch stack re-enables and reconfigures these features on all applicable interfaces on the stack members, including the stack master.

This configuration can be verified with the commands:

- **show mls qos maps policed-dscp**
- **show mls qos interface**
- **show mls qos interface *interface x/y* policers**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Queuing Models

As shown in [Figure 2-14](#), on the Catalyst 3750-E switch-family platforms, because the total inbound bandwidth of all ports can exceed the bandwidth of the stack or internal ring, ingress queues are located after the packet is classified, policed, and marked and before packets are forwarded into the switch fabric. Additionally, because multiple ingress ports can simultaneously send packets to an egress port and cause congestion, outbound queues are located after the stack or internal ring.

Both the ingress and egress queues are serviced by a Shaped Round Robin (SRR) scheduling algorithm. SRR can be configured in two modes, shaped or sharing.

In shaped mode, the egress queues are guaranteed a percentage of the bandwidth and they are rate limited to that amount. Shaped traffic does not use more than the allocated bandwidth even if the link is idle. Shaping provides a more even flow of traffic over time and reduces the peaks and valleys of bursty traffic. With shaping, the absolute value of each weight is used to compute the bandwidth available for the queues. SRR shaping is configured with the **srr-queue bandwidth shape** interface command.

In shared mode, the ingress or egress queues share the bandwidth among them according to the configured weights. The bandwidth is guaranteed at this level but not limited to it. For example, if a queue is empty and no longer requires a share of the link, the remaining queues can expand into the unused bandwidth and share it among them. With sharing, the ratio of the weights controls the frequency of dequeuing; the absolute values are meaningless. SRR sharing is configured with the **srr-queue bandwidth share** interface command.

Furthermore, both the ingress and egress queuing structures support the enabling of a single priority queue, or expedite queue, as it corresponds to the EF PHB. An ingress or egress queue operating as an expedite queue is fully serviced ahead of all other queues until empty. After the priority queue has been fully serviced, the scheduler services the non-priority queues, which are configured in either shaped or shared SRR modes. A strict priority queue is enabled with the **priority-queue** interface command.

With respect to scheduling hierarchy in the Catalyst 3750-E family of switches, shaped mode overrides shared mode and priority mode overrides both shaped and shared modes.

Additionally, the Catalyst 3750-E family of switches supports the weighted tail drop (WTD) congestion avoidance mechanism. WTD is implemented on queues to manage the queue lengths and to provide drop preferences for different traffic classifications. As a packet is enqueued to a particular ingress or egress queue, WTD uses the frame's assigned internal DSCP to subject it to different drop thresholds. If the threshold is exceeded for a given internal DSCP value (in other words, the space available in the destination queue is less than the size of the packet), the switch drops the packet. Each queue has three threshold values. The internal DSCP determines which of the three threshold values is subjected to the frame. Of the three thresholds, two are configurable (explicit) and one is not (implicit), as this last threshold corresponds to the tail of the queue (100% limit).

Packets are mapped to queues and thresholds on the Catalyst 3750-E by either CoS-to-queue/threshold or DSCP-to-queue/threshold mappings. The mapping used directly corresponds to whether the packet was configured to trust CoS on ingress or to trust DSCP on ingress (untrusted packets are simply assigned to the default queue).

Ingress Queuing 1P1Q3T Model

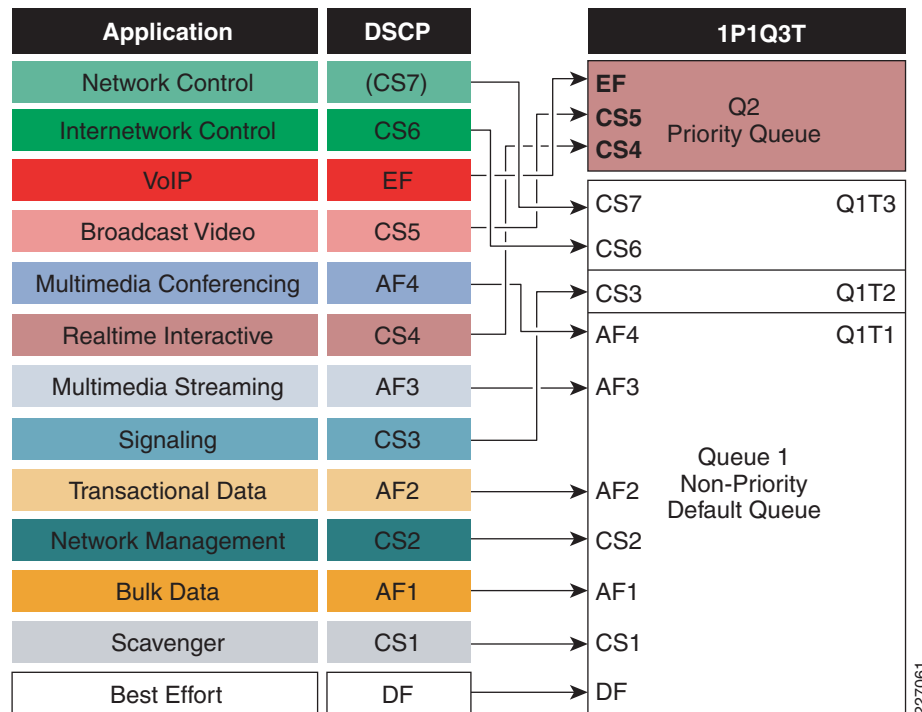
As the Catalyst 3750-E switch platforms have architectures based on oversubscription, they have been engineered to guarantee QoS by protecting critical traffic trying to access the backplane/stack-ring via ingress queuing. Ingress queuing on this platform can be configured as 2Q3T or 1P1Q3T, with the latter being the recommended configuration (as it supports the RFC 3246 EF PHB).

1P1Q3T ingress queuing is configured by explicitly enabling Q2 as a priority queue and assigning it a bandwidth allocation, such as 30%. Next, an SRR weight can be assigned to the non-priority queue, which in this case would be 70%. The buffer allocations can be tuned such that Q1 gets 90% of the buffers, while Q2 (the PQ) gets only 10%; since the PQ is serviced in realtime, it is generally more efficient to provision fewer buffers to it and more to the non-priority queues. After this, WTD thresholds can be defined on Q1 to provide inter-queue QoS; specifically, Q1T1 can be explicitly set at 80% queue depth and Q1T2 can be explicitly set at 90% queue depth (while Q3 remains implicitly set at 100% queue depth).

With the queues and thresholds set, then VoIP (EF), broadcast video (CS5), and realtime interactive (CS4) traffic can be mapped to the strict priority ingress queue. All other traffic classes can be mapped to the default (non-priority) ingress queue. However, drop preference can be given to control plane

traffic, such that network control (CS7) and internetwork control (CS6) traffic is mapped to the highest WTD threshold (Q1T3); additionally, signaling (CS3) traffic can be mapped to the middle WTD threshold (Q1T2). All other flows would be mapped to Q1T1. These 1P1Q3T ingress queuing mappings for the Catalyst 3750-E are shown in Figure 2-16.

Figure 2-16 Catalyst 3750-E 1P1Q3T Ingress Queuing Model



The corresponding configuration for 1P1Q3T ingress queuing on the Catalyst 3750-E is shown in Example 2-20.

Example 2-20 1P1Q3T Ingress Queuing Configuration Example on a Catalyst 3750-E

```

! This section configures the ingress queues
C3750-E(config)#mls qos srr-queue input priority-queue 2 bandwidth 30
! Q2 is enabled as a strict-priority ingress queue with 30% BW
C3750-E(config)#mls qos srr-queue input bandwidth 70 30
! Q1 is assigned 70% BW via SRR shared weights
! Q2 SRR shared weight is ignored (as it has been configured as a PQ)
C3750-E(config)# mls qos srr-queue input buffers 90 10
! Q1 is assigned 90% of queuing buffers and Q2 (PQ) is assigned 10%
C3750-E(config)#mls qos srr-queue input threshold 1 80 90
! Q1 thresholds are configured at 80% (Q1T1) and 90% (Q1T2)
! Q1T3 is implicitly set at 100% (the tail of the queue)
! Q2 thresholds are all set (by default) to 100% (the tail of Q2)

! This section configures ingress CoS-to-Queue mappings (if required)
C3750-E(config)#mls qos srr-queue input cos-map queue 1 threshold 1 0 1 2
! CoS values 0, 1 and 2 are mapped to Q1T1
C3750-E(config)#mls qos srr-queue input cos-map queue 1 threshold 2 3
! CoS value 3 is mapped to ingress Q1T2
C3750-E(config)#mls qos srr-queue input cos-map queue 1 threshold 3 6 7

```

```

! CoS values 6 and 7 are mapped to ingress Q1T3
C3750-E(config)#mls qos srr-queue input cos-map queue 2 threshold 1 4 5
! CoS values 4 and 5 are mapped to ingress Q2 (the PQ)

! This section configures ingress DSCP-to-Queue Mappings
C3750-E(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 0 8 10 12 14
! DSCP DF, CS1 and AF1 are mapped to ingress Q1T1
C3750-E(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 16 18 20 22
! DSCP CS2 and AF2 are mapped to ingress Q1T1
C3750-E(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 26 28 30 34 36 38
! DSCP AF3 and AF4 are mapped to ingress Q1T1
C3750-E(config)#mls qos srr-queue input dscp-map queue 1 threshold 2 24
! DSCP CS3 is mapped to ingress Q1T2
C3750-E(config)#mls qos srr-queue input dscp-map queue 1 threshold 3 48 56
! DSCP CS6 and CS7 are mapped to ingress Q1T3 (the tail of Q1)
C3750-E(config)#mls qos srr-queue input dscp-map queue 2 threshold 3 32 40 46
! DSCP CS4, CS5 and EF are mapped to ingress Q2T3 (the tail of the PQ)

```

**Note**

CoS-to-queue mappings are only required if some switch ports are configured to trust-CoS on ingress. In which case, the CoS-to-DSCP map should also be modified to map CoS 5 to DSCP EF (as shown in [Example 2-3](#)). Additionally, it should be noted that due to the limited granularity of CoS-to-queue mapping, it is not possible to assign multimedia conferencing (AF4) and realtime interactive (CS4) traffic into separate queues (as both share the same CoS value of 4); nor is it possible to assign signaling (CS3) and multimedia streaming (AF3) traffic into separate queue thresholds (as both share the same CoS value of 3).

**Note**

Non-standard DSCP-to-queue mappings are not shown in the configurations in this chapter for the sake of simplicity.

This configuration can be verified with the commands:

- **show mls qos input-queue** (shown in [Example 2-21](#))
- **show mls qos maps cos-input-q** (shown in [Example 2-22](#))
- **show mls qos maps dscp-input-q** (shown in [Example 2-23](#))

Example 2-21 Verifying Ingress Queuing on a Catalyst 3750-E—show mls qos input-queue

```

C3750-E#show mls qos input-queue
Queue      :      1      2
-----
buffers    :    90      10
bandwidth  :    70      30
priority   :      0      30
threshold1 :    80     100
threshold2 :    90     100
C3750-E#

```

[Example 2-21](#) shows that ingress queuing buffers and bandwidth have been allocated between Q1 and Q2 by a 70:30 split, respectively. Also, that Q2 has been enabled as a strict-priority queue with a 30% maximum bandwidth guarantee. Q1T1 and Q1T2 thresholds have been set to 80% and 90%, but all Q2 thresholds are at 100%.

Example 2-22 Verifying Ingress Queue Mapping on a Catalyst 3750-E—show mls qos maps cos-input-q

```
C3750-E#show mls qos maps cos-input-q
Cos-inputq-threshold map:
cos:  0    1    2    3    4    5    6    7
-----
queue-threshold: 1-1 1-1 1-1 1-2 2-3 2-3 1-3 1-3

C3750-E#
```

Example 2-22 shows the ingress CoS-to-queue mappings. Specifically, CoS values 1 and 2 have been mapped to Q1T1, CoS 3 has been mapped to Q1T2, CoS values 4 and 5 have been mapped to Q2T1 (the PQ), and CoS values 6 and 7 have been mapped to Q1T3.

Example 2-23 Verifying Ingress Queue Mapping on a Catalyst 3750-E—show mls qos maps dscp-input-q

```
C3750-E#show mls qos maps dscp-input-q
Dscp-inputq-threshold map:
d1 :d2   0    1    2    3    4    5    6    7    8    9
-----
0 : 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01
1 : 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01
2 : 01-01 01-01 01-01 01-01 01-02 01-01 01-01 01-01 01-01 01-01
3 : 01-01 01-01 02-03 01-01 01-01 01-01 01-01 01-01 01-01 01-01
4 : 02-03 02-01 02-01 02-01 02-01 02-01 02-03 02-01 01-03 01-01
5 : 01-01 01-01 01-01 01-01 01-01 01-01 01-03 01-01 01-01 01-01
6 : 01-01 01-01 01-01 01-01

C3750-E#
```

Example 2-23 shows the ingress DSCP-to-queue mappings. The first digit of the DSCP value of a packet is shown along the Y-axis of the table; the second digit of the DSCP value of a packet is shown along the X-axis of the table. The mapping table corresponds to Figure 2-16. It can be noted that CS4 (DSCP 32), CS5 (DSCP 40), and EF (DSCP 46) are all mapped to Q2 (the PQ). It should also be noted that internal DSCP values 40 through 47 are mapped to Q2 by default, which is why the table shows additional values being mapped to this queue.

Egress Queuing 1P3Q3T Model

Egress queuing on the Catalyst 3750-E family of switches can be configured as 4Q3T or 1P3Q3T, with the latter being the recommended configuration (as it supports the RFC 3246 EF PHB).

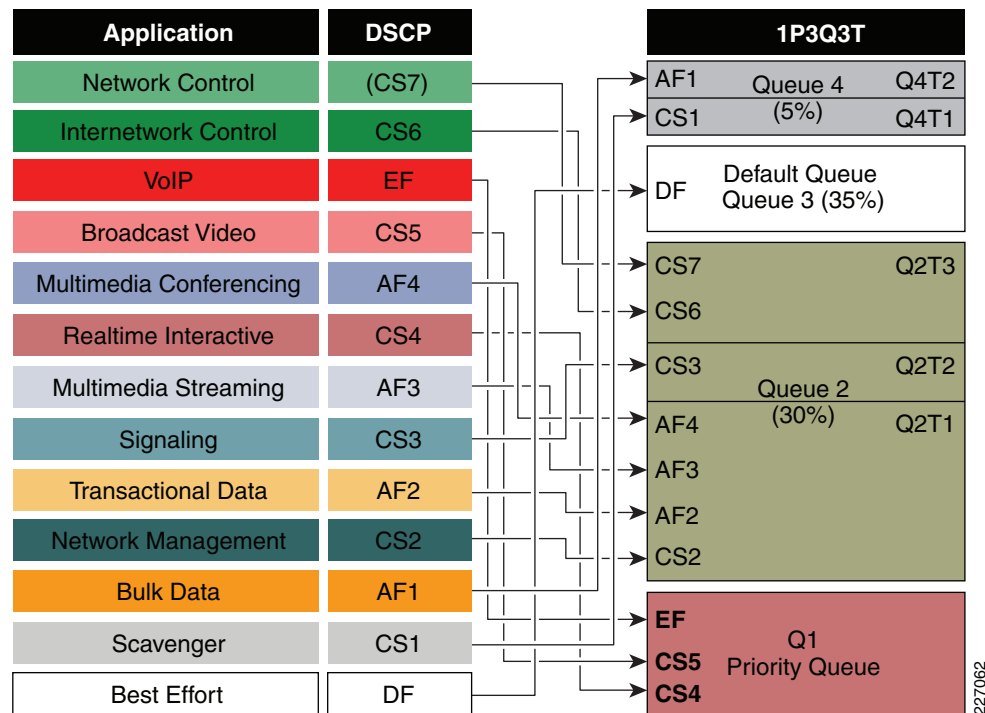
Two different egress queuing sets can be configured on the Catalyst 3750-E; however, to maintain consistent per-hop behaviors, it is generally recommended to use only one.

A unique feature of the Catalyst 3750-E is that it supports flexible buffer allocations to hardware queues, which may be dynamically loaned or borrowed against (as needed). Specifically, each queue can lend part of its buffering capacity, unless a specified minimum reserve threshold has been reached. Additionally, each queue may borrow up to four times its capacity from a common pool of buffers (which are not allocated to any specific queue) should these be available for use. The recommended buffer allocations for queues 1 through 4 are 20%, 30%, 35%, and 15%, respectively. Correspondingly, the recommended parameters for reserve thresholds and maximum (overload) thresholds for non-priority queues are 100% and 400%, respectively; for the priority queue, all thresholds should be set to 100%

Once the primary queuing set has been configured for 1P3Q3T egress queuing, WTD thresholds can be defined on Q2 and Q4 to provide intra-queue QoS. Specifically, Q2T1 can be explicitly set at 80% queue depth and Q2T2 can be explicitly set at 90% queue depth (while Q3 remains implicitly set at 100% queue depth). Also, Q4T1 can be explicitly set at 60% queue depth, while the other thresholds for Q4 remain at their default values (of 100% queue depth), with the exception of the maximum (overload) threshold, which can be set to 400%. This last setting allows for even Scavenger & Bulk traffic to benefit from the extended buffering capabilities of this platform, especially when considering that these are the least favored flows from a bandwidth perspective and thus will likely need the deepest queues.

With the queues and thresholds set, then VoIP (EF), broadcast video (CS5), and realtime interactive (CS4) traffic can be mapped to the strict priority egress queue (Q1). Network management (CS2), transactional data (AF2), multimedia streaming (AF3), and multimedia conferencing (AF4) traffic can be mapped to Q2T1. Signaling (CS3) traffic can be mapped to Q2T2. Network (CS7) and internetwork (CS6) traffic can be mapped to Q2T3. Default (DF) traffic can be mapped to Q3, the default queue. Scavenger (CS1) traffic can be mapped to Q4T1, while bulk data (AF1) is mapped to Q4T2. These 1P3Q3T egress queuing mappings for the Catalyst 3750-E are shown in [Figure 2-17](#).

Figure 2-17 Catalyst 3750-E 1P3Q3T Egress Queuing Model



The corresponding configuration for 1P3Q3T egress queuing on the Catalyst 3750-E is shown in [Example 2-24](#).

Example 2-24 1P3Q3T Egress Queuing Configuration Example on a Catalyst 3750-E

```
! This section configures buffers and thresholds on Q1 through Q4
C3750-E(config)#mls qos queue-set output 1 buffers 15 30 35 20
! Queue buffers are allocated
C3750-E(config)#mls qos queue-set output 1 threshold 1 100 100 100 100
! All Q1 (PQ) Thresholds are set to 100%
C3750-E(config)#mls qos queue-set output 1 threshold 2 80 90 100 400
! Q2T1 is set to 80%; Q2T2 is set to 90%;
! Q2 Reserve Threshold is set to 100%;
```



```

! Q2 Maximum (Overflow) Threshold is set to 400%
C3750-E(config)#mls qos queue-set output 1 threshold 3 100 100 100 400
! Q3T1 is set to 100%, as all packets are marked the same weight in Q3
! Q3 Reserve Threshold is set to 100%;
! Q3 Maximum (Overflow) Threshold is set to 400%
C3750-E(config)#mls qos queue-set output 1 threshold 4 60 100 100 400
! Q4T1 is set to 60%; Q4T2 is set to 100%
! Q4 Reserve Threshold is set to 100%;
! Q4 Maximum (Overflow) Threshold is set to 400%

! This section configures egress CoS-to-Queue mappings (if required)
C3750-E(config)#mls qos srr-queue output cos-map queue 1 threshold 3 4 5
! CoS 4 and 5 are mapped to egress Q1T3 (the tail of the PQ)
C3750-E(config)#mls qos srr-queue output cos-map queue 2 threshold 1 2
! CoS 2 is mapped to egress Q2T1
C3750-E(config)#mls qos srr-queue output cos-map queue 2 threshold 2 3
! CoS 3 is mapped to egress Q2T2
C3750-E(config)#mls qos srr-queue output cos-map queue 2 threshold 3 6 7
! CoS 6 and 7 are mapped to Q2T3
C3750-E(config)#mls qos srr-queue output cos-map queue 3 threshold 3 0
! CoS 0 is mapped to Q3T3 (the tail of the default queue)
C3750-E(config)#mls qos srr-queue output cos-map queue 4 threshold 3 1
! CoS 1 is mapped to Q4T3 (tail of the less-than-best-effort queue)

! This section configures egress DSCP-to-Queue mappings
C3750-E(config)# mls qos srr-queue output dscp-map queue 1 threshold 3 32 40 46
! DSCP CS4, CS5 and EF are mapped to egress Q1T3 (tail of the PQ)
C3750-E(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 16 18 20 22
! DSCP CS2 and AF2 are mapped to egress Q2T1
C3750-E(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 26 28 30 34 36 38
! DSCP AF3 and AF4 are mapped to egress Q2T1
C3750-E(config)#mls qos srr-queue output dscp-map queue 2 threshold 2 24
! DSCP CS3 is mapped to egress Q2T2
C3750-E(config)#mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
! DSCP CS6 and CS7 are mapped to egress Q2T3
C3750-E(config)#mls qos srr-queue output dscp-map queue 3 threshold 3 0
! DSCP DF is mapped to egress Q3T3 (tail of the best effort queue)
C3750-E(config)#mls qos srr-queue output dscp-map queue 4 threshold 1 8
! DSCP CS1 is mapped to egress Q4T1
C3750-E(config)# mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
! DSCP AF1 is mapped to Q4T3 (tail of the less-than-best-effort queue)

! This section configures interface egress queuing parameters
C3750-E(config)#interface range GigabitEthernet1/0/1-48
C3750-E(config-if-range)# queue-set 1
! The interface(s) is assigned to queue-set 1
C3750-E(config-if-range)# srr-queue bandwidth share 1 30 35 5
! The SRR sharing weights are set to allocate 30% BW to Q2
! 35% BW to Q3 and 5% BW to Q4
! Q1 SRR sharing weight is ignored, as it will be configured as a PQ
C3750-E(config-if-range)# priority-queue out
! Q1 is enabled as a strict priority queue

```

**Note**

CoS-to-queue mappings are only required if some switch ports are configured to trust-CoS on ingress. In which case, the CoS-to-DSCP map should also be modified to map CoS 5 to DSCP EF (as shown in [Example 2-3](#)). Additionally, it should be noted that due to the limited granularity of CoS-to-queue

mapping, it is not possible to assign multimedia-conferencing (AF4) and realtime interactive (CS4) traffic into separate queues (as both share the same CoS value of 4); nor is it possible to assign signaling (CS3) and multimedia streaming (AF3) traffic into separate queue thresholds (as both share the same CoS value of 3); nor is it possible to assign scavenger (CS1) and bulk data (AF1) traffic into separate queue thresholds (as both share the same CoS value of 1).

This configuration can be verified with the commands:

- **show mls qos queue-set** (shown in [Example 2-25](#))
- **show mls qos maps cos-output-q**
- **show mls qos maps dscp-output-q**
- **show mls qos interface interface x/y queueing** (shown in [Example 2-26](#))
- **show mls qos interface interface x/y statistics** (shown in [Example 2-27](#))

Example 2-25 Verifying Egress Queuing on a Catalyst 3750-E—show mls qos queue-set

```
C3750-E#show mls qos queue-set 1
Queueset: 1
Queue      :      1      2      3      4
-----
buffers    :      15     30     35     20
threshold1:     100     80    100     60
threshold2:     100     90    100    100
reserved   :     100    100    100    100
maximum    :     100    400    400    400
C3750-E#
```

[Example 2-26](#) shows that the queuing buffers, drop-thresholds, reserve-thresholds, and maximum (overload) thresholds have been configured correctly on a per-queue-set basis.

Example 2-26 Verifying Egress Queuing on a Catalyst 3750-E—show mls qos interface interface x/y queueing

```
C3750-E#show mls qos interface GigabitEthernet 1/0/1 queueing
GigabitEthernet1/0/1
Egress Priority Queue : enabled
Shaped queue weights (absolute) : 25 0 0 0
Shared queue weights   : 1 30 35 5
The port bandwidth limit : 100 (Operational Bandwidth:100.0)
The port is mapped to qset : 1

C3750-E#
```

[Example 2-26](#) shows that strict-priority queueing has been enabled on the interface, and that the queues Q2, Q3, and Q4 receive 30%, 35% and 5% of the remaining bandwidth, respectively.

Example 2-27 Verifying Egress Queuing on a Catalyst 3750-E—show mls qos interface interface x/y statistics

```
C3750-E#show mls qos interface GigabitEthernet 1/0/49 statistics
GigabitEthernet1/0/49 (All statistics are in packets)

dscp: incoming
-----

0 - 4 :      1729      0      0      0      0
5 - 9 :         0      0      0      0      0
```

```

10 - 14 :      0      0      0      0      0
15 - 19 :      0      0      0      0      0
20 - 24 :      0      0      0      0      0
25 - 29 :      0      0      0      0      0
30 - 34 :      0      0      0      0      0
35 - 39 :      0      0      0      0      0
40 - 44 :      0      0      0      0      0
45 - 49 :      0      127292      0      1263      0
50 - 54 :      0      0      0      0      0
55 - 59 :      0      0      0      0      0
60 - 64 :      0      0      0      0      0

```

dscp: outgoing

```

-----
 0 - 4 :      947678      0      0      0      0
 5 - 9 :          0      0      0      23842155      0
10 - 14 :     1190043      0      0      0      0
15 - 19 :          0      0      0      1061726      0
20 - 24 :          0      0      0      0      10372
25 - 29 :          0      0      0      0      0
30 - 34 :          0      0      0      0      8320623
35 - 39 :          0      0      0      0      0
40 - 44 :          0      0      0      0      0
45 - 49 :          0      127291      0      784      0
50 - 54 :          0      0      0      0      0
55 - 59 :          0      0      0      0      0
60 - 64 :          0      0      0      0      0

```

cos: incoming

```

-----
 0 - 4 :      130653      0      0      998      0
 5 - 7 :      127599      613      3156

```

cos: outgoing

```

-----
 0 - 4 :      947754      25032199      1061726      10372      8320623
 5 - 7 :      127291      784      3462

```

output queues enqueued:

queue: threshold1 threshold2 threshold3

```

-----
queue 0:          0      0      127291
queue 1:     9382416      10396      4246
queue 2:          0      0      947611
queue 3:     23842152      1190043      0

```

output queues dropped:

queue: threshold1 threshold2 threshold3

```

-----
queue 0:          0      0      0
queue 1:          0      0      0
queue 2:          0      0      0
queue 3:          892      0      0

```

Policer: Inprofile: 0 OutofProfile: 0

C3750-E#

[Example 2-27](#) shows a set of dynamically-updated packet statistic tables for an uplink port on an access layer Catalyst 3750-E switch that is primarily congested in the access-to-distribution direction. The first table shows the incoming DSCP values (from the distribution layer). DSCP values are broken into groups of 4. For example, incoming packets marked DSCP EF/46 are listed in the DSCP 45-49 row in the second column (in this case: 127,292 packets). The second table shows the outgoing packets (to the distribution layer) in a similar format. For example, DSCP CS1/8 is listed in the DSCP 5-9 row in the third column (23,842,155 packets). The third table shows incoming packets (from the distribution layer) by CoS values (again grouped in sets of 4); similarly the fourth table shows outgoing packets (to the distribution layer) by CoS values. The fifth and sixth tables are particularly interesting in terms of queuing statistics: the fifth table shows the number of packets assigned to each queue/threshold combination.

**Note**

The queue numbers are 1 lower than the numbers used in the configuration syntax (such that Q1 is shown here as Q0, Q2 is shown here as Q1, Q3 is shown here as Q2, and Q4 is shown here as Q3).

For example, from the fifth table, it can be seen that 127,291 packets were sent to the (tail of the) PQ (shown here as Q0); similarly, 23,842,155 packets were sent to the scavenger/bulk queue first threshold (shown here as Q3T1). Finally, the sixth table shows any drops that have occurred on a per-queue/per-threshold basis; from this table it can be seen that 892 drops occurred in the scavenger/bulk queue first threshold (scavenger class drops).

EtherChannel QoS Model

As discussed in [EtherChannel QoS](#), QoS policies applied to EtherChannel links on the Catalyst 2960-G/S, 2975-GS, 3560-G/E/X, and 3750-G/E/X family of switches are required to be identically-configured on each and every EtherChannel physical port-member interface; these include both ingress trust/classification/marketing/policing policies, as well as egress queuing policies (ingress queuing policies are globally-defined and as such are not bound by this requirement). If the policies are not identically-configured, even though they may appear in the configuration, these will not take effect. Also, it is recommended to load-balance across the EtherChannel by source-and-destination IP address. An example of an EtherChannel QoS Model for the Catalyst 3750-E family is shown in [Example 2-28](#).

**Note**

As the ingress queuing policies, as well as egress-queuing mappings, have not changed from the previous configuration examples, these are omitted from this EtherChannel QoS Model example.

Example 2-28 EtherChannel QoS Design on a Catalyst 3750-E

```
! This section configures EtherChannel source-and-destination load-balancing
C3750-E(config)# port-channel load-balance src-dst-ip
```

```
! This section configures the (logical) EtherChannel interface
C3750-E(config)# interface Port-channel1
C3750-E(config-if)# description ETHERCHANNEL-TRUNK-TO-DISTRIBUTION-LAYER
C3750-E(config-if)# switchport mode trunk
C3750-E(config-if)# switchport trunk encapsulation dot1q
C3750-E(config-if)# switchport trunk allowed vlan 10,110
```

```
! This section configures Trust-DSCP and (1P3Q3T) Egress Queuing across the
! (physical) EtherChannel member-ports
C3750-E(config)# interface range TenGigabitEthernet1/0/1-2
C3750-E(config-if-range)# description PORT-CHANNEL1-PHYSICAL-PORT-MEMBER
```

```

C3750-E(config-if-range)# switchport mode trunk
C3750-E(config-if-range)# switchport trunk encapsulation dot1q
C3750-E(config-if-range)# switchport trunk allowed vlan 10,110
C3750-E(config-if-range)# channel-group 1 mode auto
! Associates the physical ports with the logical EtherChannel bundle
C3750-E(config-if-range)# mls qos trust dscp
! The physical port-member interfaces are set to statically trust DSCP
C3750-E(config-if-range)# queue-set 1
! The interfaces are assigned to queue-set 1
C3750-E(config-if-range)# srr-queue bandwidth share 1 30 35 5
! The SRR sharing weights are set to allocate 30% BW to Q2
! 35% BW to Q3 and 5% BW to Q4
! Q1 SRR sharing weight is ignored, as it will be configured as a PQ
C3750-E(config-if-range)# priority-queue out
! Q1 is enabled as a strict priority queue

```

This configuration can be verified with the commands:

- **show mls qos interface**
- **show mls qos input-queue**
- **show mls qos maps cos-input-q**
- **show mls qos maps dscp-input-q**
- **show mls qos maps cos-output-q**
- **show mls qos maps dscp-output-q**
- **show mls qos interface *interface x/y* queueing**
- **show mls qos interface *interface x/y* statistics**

AutoQoS-SRND4 Models

As mentioned in [AutoQoS](#), as of August 2010, an updated version of AutoQoS was released for the Catalyst 2960-G/S, 2975-GS, 3560-G/E/X, and 3750-G/E/X family of switches with IOS release 12.2(55)SE. This release was directly based on the recommendations put forward in this design chapter to support medianet applications; in fact, the new keyword and name for this version of AutoQoS is AutoQoS-SRND4 (taken from Solution Reference Network Design guide version 4, which is the Cisco name for this design chapter). AutoQoS-SRND4 is the fastest and most accurate method to deploy the recommended QoS designs to support rich media applications across this family of switches.

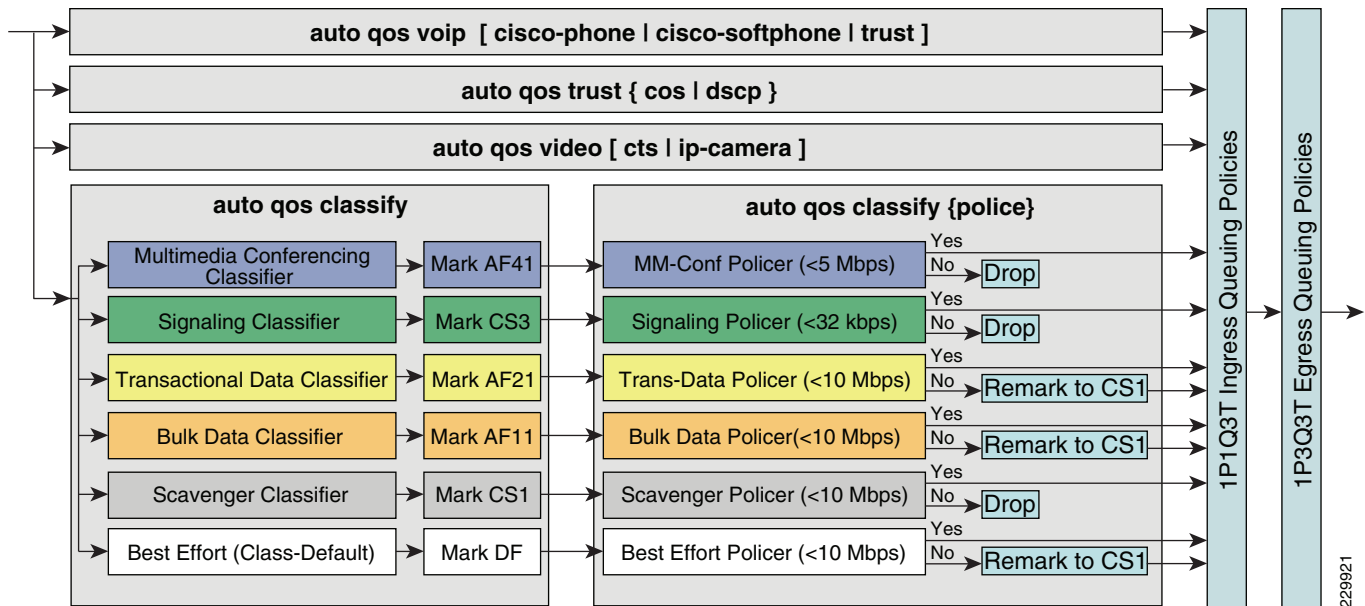
AutoQoS-SRND4—which can be shortened to “Auto QoS” for the sake of simplicity—presents the network administrator with four main ingress QoS policy options in interface-configuration mode:

- **auto qos voip [cisco-phone | cisco-softphone | trust]**—This option provides not only legacy support for Auto QoS VoIP IP Telephony deployments, but also expands on these models to include provisioning for additional classes of rich media applications and to include data-plane policing/scavenger-class QoS policy-elements to protect and secure these applications.
- **auto qos trust {cos | dscp}**—This option configures the port to statically trust either CoS or DSCP. If neither CoS nor DSCP are explicitly specified, then the **auto qos trust** command will configure—by default—CoS-trust on Layer 2 switch ports and DSCP-trust on Layer 3 routed interfaces.
- **auto qos video [cts | ip-camera]**—This new option provides automatic configuration support for both Cisco TelePresence Systems (via the **cts** keyword) as well as IP Video Surveillance cameras (via the **ip-camera** keyword).

- **auto qos classify {police}**—This option provides a generic template that can classify and mark up to six classes of medianet traffic, as well as optionally provision data-plane policing/scavenger-class QoS policy-elements for these traffic classes (via the optional **police** keyword).

Each ingress option is automatically complemented by a complete set of ingress and egress queuing configurations, complete with both CoS- and DSCP-to-Queue mappings, as shown in [Figure 2-18](#).

Figure 2-18 Auto QoS SRND4 Models



Note

Ingress queuing is not supported on the Catalyst 2960-S platform.

The complete configuration provisioned by each of these new **auto qos** options, along with the complete ingress and egress queuing configurations, will be detailed in the following sections. For the sake of logical development however, **auto qos voip** will be discussed last, as it combines several policy elements from other **auto qos ingress** model options.

Auto QoS Trust Models

The **auto qos trust** command configures static trust policies on the port(s) or interface(s) that it is configured on: if the port is operating as a Layer 2 switch port, then (by default) CoS-trust is configured (as shown in [Example 2-29](#)); whereas, if the port is operating as a Layer 3 routed interface then (by default) DSCP-trust is configured (as shown in [Example 2-33](#)).

All AutoQoS configurations assume **mls qos** to have been previously enabled.

Example 2-29 Auto QoS Trust Applied on a Layer 2 Switch Port

```
! This section configures auto qos trust on a L2 switch port
C3750-E(config)# interface GigabitEthernet1/0/1
C3750-E(config-if)# description L2-ACCESS-PORT
C3750-E(config-if)# switchport access vlan 10
C3750-E(config-if)# switchport voice vlan 110
C3750-E(config-if)# spanning-tree portfast
```

```
C3750-E(config-if)# auto qos trust
! Auto-configures static trust policy (and ingress and egress queuing policies)
```

The effect of this **auto qos trust** policy on a layer 2 switch port can be verified by the **show run interface** command, as shown in [Example 2-30](#).

Example 2-30 Auto QoS Trust Applied on a Layer 2 Switch Port Verification—show run interface

```
C3750-E# show run interface GigabitEthernet1/0/1
Building configuration...

Current configuration : 251 bytes
!
interface GigabitEthernet1/0/1
  description L2-ACCESS-PORT
  switchport access vlan 10
  switchport voice vlan 110
  srr-queue bandwidth share 1 30 35 5
  queue-set 2
  priority-queue out
  mls qos trust cos
  ! AutoQoS has configured the port to static CoS-trust
  auto qos trust
  spanning-tree portfast
end

C3750-E#
```

The **show run interface** command displays the effect of deploying **auto qos trust** on a Layer 2 switch port configuration: namely four additional lines of configuration (in bold) have been added automatically (as well as many more in the global configuration; these will all be discussed in detail in the following sections). From the configuration it can be seen that the interface trust state has been set to statically trust CoS.

This interface trust state and CoS-to-DSCP maps can be verified by the commands:

- **show mls qos interface** (as shown in [Example 2-31](#))
- **show mls qos maps cos-dscp** (as shown in [Example 2-32](#))

Example 2-31 Auto QoS Trust Verification on a Layer 2 Switch Port—show mls qos interface

```
C3750-E# show mls qos interface GigabitEthernet 1/0/1
GigabitEthernet1/0/1
trust state: trust cos
trust mode: trust cos
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
qos mode: port-based
```

[Example 2-31](#) verifies that—by default—**auto qos trust** has configured a static trust-CoS policy on this Layer 2 switchport.

Example 2-32 Auto QoS Trust Verification on a Layer 2 Switch Port—show mls qos maps cos-dscp

```
C3750-E# show mls qos maps cos-dscp
Cos-dscp map:
cos:    0   1   2   3   4   5   6   7
```

```
-----
dscp:    0  8 16 24 32 46 48 56
```

C3750-E#

[Example 2-32](#) shows that **auto qos trust** has also automatically modified the CoS-to-DSCP mapping table to ensure that CoS 5 is mapped to DSCP EF/46.

By way of contrast, [Example 2-33](#) through [Example 2-35](#) illustrate the effect of **auto qos trust** on a Layer 3 routed interface.

Example 2-33 Auto QoS Trust Applied on a Layer 3 Routed Interface

```
! This section configures auto qos trust on a L3 routed interface
C3750-E(config)# interface GigabitEthernet1/0/48
C3750-E(config-if)# description L3-ROUTED-INTERFACE
C3750-E(config-if)# no switchport
C3750-E(config-if)# ip address 10.0.1.103 255.255.255.0
C3750-E(config-if)# auto qos trust
! Auto-configures static trust policy (and ingress and egress queuing policies)
```

The effect of this **auto qos trust** policy on a Layer 3 routed interface can be verified by the **show run interface** command, as shown in [Example 2-34](#).

Example 2-34 Auto QoS Trust Applied on a Layer 3 Routed Interface Verification—show run interface

```
C3750-E# show run interface GigabitEthernet1/0/48
Building configuration...

Current configuration : 271 bytes
!
interface GigabitEthernet1/0/48
 description L3-ROUTED-INTERFACE
 no switchport
 ip address 10.0.1.103 255.255.255.0
 no ip route-cache
 no ip mroute-cache
 srr-queue bandwidth share 1 30 35 5
 queue-set 2
 priority-queue out
 mls qos trust dscp
 ! AutoQoS has configured the port to static DSCP-trust
 auto qos trust
end

C3750-E#
```

The **show run interface** command displays the effect of deploying **auto qos trust** on a Layer 3 routed interface configuration. From the configuration it can be seen that the interface trust state has been set to statically trust DSCP.

This interface trust state can be verified by the command:

- **show mls qos interface** (as shown in [Example 2-35](#))

Example 2-35 Auto QoS Trust Verification on a Layer 3 Routed Interface—show mls qos interface

```
C3750-E# show mls qos interface gigabitEthernet 1/0/48
```



```
GigabitEthernet1/0/48
trust state: trust dscp
trust mode: trust dscp
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
qos mode: port-based

C3750-E#
```

[Example 2-35](#) verifies that although the same **auto qos trust** policy is applied to a Layer 3 routed interface (in [Example 2-33](#)) as was applied to a Layer 2 switch port (in [Example 2-29](#)), this has resulted in a DSCP-trust policy on this interface, by default.

Since all interswitch links are recommended to be set to DSCP-trust, whether they are operating at Layer 2 or Layer 3, it is recommended to always use the **auto qos trust dscp** option on all interswitch links, as well as any access-edge ports connected to trusted devices (assuming that the trusted devices have the ability to mark/re-mark DSCP). The **auto qos trust dscp** option is shown in [Example 2-36](#).

Example 2-36 Auto QoS Trust DSCP Applied on a Layer 2 Switch Port

```
! This section configures auto qos trust dscp on a L2 switch port
C3750-E(config)# interface GigabitEthernet1/0/1
C3750-E(config-if)# description L2-ACCESS-PORT
C3750-E(config-if)# switchport access vlan 10
C3750-E(config-if)# switchport voice vlan 110
C3750-E(config-if)# spanning-tree portfast
C3750-E(config-if)# auto qos trust dscp
! Auto-configures static DSCP-trust policy (and ingress and egress queuing policies)
```

The effect of this **auto qos trust dscp** policy on a Layer 2 switch port can be verified by the **show run interface** command, as shown in [Example 2-37](#).

Example 2-37 Auto QoS Trust DSCP Applied on a Layer 2 Switch Port Verification—show run interface

```
C3750-E# show run interface GigabitEthernet1/0/1
Building configuration...

Current configuration : 256 bytes
!
interface GigabitEthernet1/0/1
  description L2-ACCESS-PORT
  switchport access vlan 10
  switchport voice vlan 110
  srr-queue bandwidth share 1 30 35 5
  queue-set 2
  priority-queue out
  mls qos trust dscp
  ! AutoQoS has configured the port to static DSCP-trust
  auto qos trust dscp
  spanning-tree portfast
end

C3750-E#
```

The **show run interface** command displays the effect of deploying **auto qos trust dscp** on a Layer 2 switch port interface configuration. From the configuration it can be seen that the interface trust state has been set to statically trust DSCP.

This interface trust state can be verified by the command:

- **show mls qos interface** (as shown in [Example 2-38](#))

Example 2-38 Auto QoS Trust DSCP Verification on a Layer 2 Switch Port—show mls qos interface

```
C3750-E# show mls qos interface GigabitEthernet 1/0/1
GigabitEthernet1/0/1
trust state: trust dscp
trust mode: trust dscp
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
qos mode: port-based

C3750-E#
```

[Example 2-38](#) verifies that although the interface is configured as a Layer 2 switch port, because of the use of the **dscp** keyword in conjunction with the **auto qos trust** interface command, DSCP-trust has been applied to the interface (rather than the default CoS-trust).

Auto QoS Video Models

Besides supporting IP Telephony devices such as Cisco IP Phones and Soft-Phones (via AutoQoS-VoIP), Auto QoS now also supports video devices, such as Cisco TelePresence Systems (CTS) and IP Video-Surveillance cameras, both of which support conditional trust via CDP-negotiation.

Cisco TelePresence Systems can mark their video flow and their audio flows with to CoS 4 and DSCP CS4. Additionally, any voice traffic originating from the Cisco 7975G IP Phone, that is an integral part of the CTS, is marked to CoS 5 and DSCP EF. Furthermore, any signaling traffic—whether for the CTS and/or the IP Phone—is marked CoS 3 and DSCP CS3. These CTS markings are illustrated in [Figure 2-7](#).

Similar to **auto qos trust** behavior, **auto qos video cts** will dynamically extend CoS-trust to CTS systems connecting to Layer 2 switch ports (by default) and will dynamically extend DSCP-trust to CTS systems connecting to Layer 3 routed interfaces (by default). Typically CTS systems are connected to Layer 2 switch ports, however, as shown in [Example 2-39](#).

Example 2-39 Auto QoS Video CTS Configuration on a Layer 2 Switch Port

```
C3750-E(config)# interface GigabitEthernet1/0/1
C3750-E(config-if)# description L2-ACCESS-PORT-TO-CTS
C3750-E(config-if)# switchport access vlan 10
C3750-E(config-if)# switchport voice vlan 110
C3750-E(config-if)# spanning-tree portfast
C3750-E(config-if)# auto qos video cts
! Auto-configures conditional-trust policy for CTS (+ ingress and egress queuing policies)
```

The effect of this **auto qos video cts** policy on a Layer 2 switch port can be verified by the **show run interface** command, as shown in [Example 2-40](#).

Example 2-40 Auto QoS Video CTS Applied on a Layer 2 Switch Port Verification—show run interface

```
C3750-E# show run interface GigabitEthernet1/0/1
Building configuration...
```

```

Current configuration : 288 bytes
!
interface GigabitEthernet1/0/1
 description L2-ACCESS-PORT-TO-CTS
 switchport access vlan 10
 switchport voice vlan 110
 srr-queue bandwidth share 1 30 35 5
 queue-set 2
 priority-queue out
 mls qos trust device cts
 ! AutoQoS has configured a conditional-trust policy for cts devices
 mls qos trust cos
 ! AutoQoS has configured CoS-trust to be dynamically extended
 auto qos video cts
 spanning-tree portfast
end

C3750-E#

```

The **show run interface** command displays the effect of deploying **auto qos video cts** on a Layer 2 switch port interface configuration. From the configuration it can be seen that the interface trust state has been set to conditionally-trust cts devices and to dynamically extend CoS-trust to these.

This interface trust state can be verified by the command:

- **show mls qos interface** (as shown in [Example 2-41](#))

Example 2-41 Auto QoS Video CTS Conditional Trust Verification on a Layer 2 Switch Port—show mls qos interface

```

C3750-E# show mls qos interface GigabitEthernet 1/0/1
GigabitEthernet1/0/1
trust state: trust cos
trust mode: trust cos
trust enabled flag: dis
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: cts
qos mode: port-based

C3750-E#

```

As shown by [Example 2-41](#), the **auto qos video cts** command has configured a conditional trust policy to dynamically extend CoS-trust to CTS systems. As a CTS device is currently attached to this port, the current trust state is trust-CoS.

Nonetheless, should an administrator choose to trust-DSCP instead of CoS, they can still do so while using the **auto qos video cts** command, simply by manually adding a **mls qos trust dscp** interface command to the configuration, as shown in [Example 2-42](#).

Example 2-42 Auto QoS Video CTS (DSCP-Trust) Configuration on a Layer 2 Switch Port

```

C3750-E(config)# interface GigabitEthernet1/0/1
C3750-E(config-if)# description L2-ACCESS-PORT-TO-CTS
C3750-E(config-if)# switchport access vlan 10
C3750-E(config-if)# switchport voice vlan 110
C3750-E(config-if)# spanning-tree portfast
C3750-E(config-if)# auto qos video cts
! Auto-configures conditional-trust policy for CTS (+ ingress and egress queuing
policies)

```

```
C3750-E(config-if)# mls qos trust dscp
! Manually configures DSCP-trust to be dynamically extended (rather than CoS-trust)
```

The effect of this **auto qos video cts** policy in conjunction with a manual override **mls qos trust dscp** policy on a Layer 2 switch port can be verified by the **show run interface** command, as shown in [Example 2-43](#).

Example 2-43 Auto QoS Video CTS and MLS QoS Trust DSCP Applied on a Layer 2 Switch Port Verification—show run interface

```
C3750-E# show run interface GigabitEthernet 1/0/1
Building configuration...

Current configuration : 289 bytes
!
interface GigabitEthernet1/0/1
 description L2-ACCESS-PORT-TO-CTS
 switchport access vlan 10
 switchport voice vlan 110
 srr-queue bandwidth share 1 30 35 5
 queue-set 2
 priority-queue out
 mls qos trust device cts
 ! AutoQoS has configured a conditional-trust policy for cts devices
 mls qos trust dscp
 ! The trust-state to be dynamically extended has been manually set to DSCP-trust
 auto qos video cts
 spanning-tree portfast
end

C3750-E#
```

The **show run interface** command displays the effect of deploying **auto qos video cts** on a Layer 2 switch port interface along with a **manual mls qos trust dscp** command. From the configuration it can be seen that the interface trust state has been set to conditionally-trust cts devices and to dynamically extend DSCP-trust to these.

This example demonstrates a simple, yet powerful point: AutoQoS configurations may be modified and tailored to specific administrative needs or preferences. In other words, deploying AutoQoS is not an “all-or-nothing” option, but rather one that may be viewed as a generic template on which custom-tailored designs may be overlaid. Even with a moderate amount of manual configuration, AutoQoS can still significantly expedite medianet QoS deployments and greatly reduce manual configuration errors in the process.

This interface trust state can be verified by the command:

- **show mls qos interface** (as shown in [Example 2-44](#))

Example 2-44 Auto QoS Video CTS (DSCP) Conditional Trust Verification on a Layer 2 Switch Port—show mls qos interface

```
C3750-E# show mls qos interface GigabitEthernet 1/0/1
GigabitEthernet1/0/1
trust state: trust dscp
trust mode: trust dscp
trust enabled flag: dis
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: cts
qos mode: port-based
```

C3750-E#

[Example 2-44](#) shows that the **auto qos video cts** interface command in conjunction with the **mls qos trust dscp** interface command has configured a conditional trust policy to dynamically extend DSCP-trust to CTS systems. As a CTS device is currently attached to this port, the current trust state is trust-DSCP.

Unlike CTS devices, IP Video Surveillance Cameras are only required to mark their video (and if supported, audio) flows at Layer 3 (typically to DSCP CS5/40). This allows for more flexible deployment models, as these cameras do not therefore have to be deployed in dedicated VLANs connecting to the access switch via an 802.1Q trunk. As such, the **auto qos video ip-camera** interface command dynamically extends DSCP-trust to such devices, once these have successfully identified themselves to the switch via CDP. DSCP-trust is dynamically extended whether the port is configured as a Layer 2 switch port or as a Layer 3 routed interface, as shown in [Example 2-45](#).

Example 2-45 Auto QoS Video IP-Camera Configuration on a Layer 2 Switch Port

```
C3750-E(config)# interface GigabitEthernet1/0/1
C3750-E(config-if)# description L2-ACCESS-PORT-TO-IPVS-CAMERA
C3750-E(config-if)# switchport access vlan 10
C3750-E(config-if)# switchport voice vlan 110
C3750-E(config-if)# spanning-tree portfast
C3750-E(config-if)# auto qos video ip-camera
! Auto-configures conditional-trust policy for IPVS (+ ingress and egress queuing
policies)
```

The effect of this **auto qos video ip-camera** policy on a Layer 2 switch port can be verified by the **show run interface** command, as shown in [Example 2-46](#).

Example 2-46 Auto QoS Video IP-Camera Applied on a Layer 2 Switch Port Verification—show run interface

```
C3750-E# show run interface GigabitEthernet 1/0/1
Building configuration...

Current configuration : 309 bytes
!
interface GigabitEthernet1/0/1
 description L2-ACCESS-PORT-TO-IPVS-CAMERA
 switchport access vlan 10
 switchport voice vlan 110
 srr-queue bandwidth share 1 30 35 5
 queue-set 2
 priority-queue out
 mls qos trust device ip-camera
 ! AutoQoS has configured a conditional-trust policy for ip-camera devices
 mls qos trust dscp
 ! AutoQoS has configured DSCP-trust to be dynamically extended
 auto qos video ip-camera
 spanning-tree portfast
end

C3750-E#
```

The **show run interface** command displays the effect of deploying **auto qos video ip-camera** on a Layer 2 switch port interface configuration. From the configuration it can be seen that the interface trust state has been set to conditionally-trust ip-camera devices and dynamically extend DSCP-trust to these.

This interface trust state can be verified by the command:

- **show mls qos interface** (as shown in [Example 2-47](#))

Example 2-47 Auto QoS Video IP-Camera Conditional Trust Verification on a Layer 2 Switch Port—show mls qos interface

```
C3750-E# show mls qos interface GigabitEthernet 1/0/1
GigabitEthernet1/0/1
trust state: trust dscp
trust mode: trust dscp
trust enabled flag: dis
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: ip-camera
qos mode: port-based

C3750-E#
```

[Example 2-47](#) confirms that the **auto qos video ip-camera** interface command has configured a conditional trust policy for IPVS cameras to dynamically extend DSCP-trust. As an IPVS camera is currently attached to this port, the current trust state is trust-DSCP.

In a similar vein to the CTS (DSCP-trust) example, should an administrator wish to extend CoS-trust instead of DSCP trust to IPVS cameras, they could add **mls qos trust cos** to the **auto qos video ip-camera** interface configuration.

Auto QoS Classify and Police Models

The AutoQoS Classify and Police models provide a generic template to support additional rich media and data applications, providing a classification (and optional policing) model for these. These models are most suitable for switch ports connecting to PC endpoint devices.

Six application classes (multimedia conferencing, signaling, transactional data, bulk data, scavenger, and best-effort) are automatically defined via class-maps. Each class-map references an associated extended IP access-list. These IP access lists define the TCP and UDP port numbers of the given class of applications based on the sample ports summarized in [Table 2-9](#). **However, it cannot be overemphasized that these are just generic application examples for these classes and the administrator can add/change/delete the access-list entries to match on their specific applications.**

[Example 2-48](#) shows the application of the **auto qos classify** command on a Layer 2 switch port.

Example 2-48 Auto QoS Classify Configuration on a Layer 2 Switch Port

```
C3750-E(config)#interface GigabitEthernet1/0/1
C3750-E(config-if)# description L2-ACCESS-PORT-TO-PC
C3750-E(config-if)# switchport access vlan 10
C3750-E(config-if)# spanning-tree portfast
C3750-E(config-if)# auto qos classify
! Auto-configures classify policy (+ ingress and egress queuing policies)
```

The effect of this **auto qos classify** policy on a Layer 2 switch port can be verified by the **show run** command, as shown in [Example 2-49](#).

Example 2-49 Auto QoS Classify Configuration on a Layer 2 Switch Port Verification—show run

```
C3750-E# show run
Building configuration...
<snip>
```

```

! This section defines the class-maps for AutoQoS-Classify
! Each Class-Map is associated with an Extended IP Access-List
class-map match-all AUTOQOS_MULTIMEDIACONF_CLASS
  match access-group name AUTOQOS-ACL-MULTIMEDIACONF

class-map match-all AUTOQOS_DEFAULT_CLASS
  match access-group name AUTOQOS-ACL-DEFAULT

class-map match-all AUTOQOS_TRANSACTION_CLASS
  match access-group name AUTOQOS-ACL-TRANSACTIONAL-DATA

class-map match-all AUTOQOS_SIGNALING_CLASS
  match access-group name AUTOQOS-ACL-SIGNALING

class-map match-all AUTOQOS_BULK_DATA_CLASS
  match access-group name AUTOQOS-ACL-BULK-DATA

class-map match-all AUTOQOS_SCAVANGER_CLASS
  match access-group name AUTOQOS-ACL-SCAVANGER
!

! This section defines the policy-map for AutoQoS-Classify
policy-map AUTOQOS-SRND4-CLASSIFY-POLICY
  class AUTOQOS_MULTIMEDIACONF_CLASS
    set dscp af41
    ! Marks Multi-media Conferencing traffic to AF41
  class AUTOQOS_BULK_DATA_CLASS
    set dscp af11
    ! Marks Bulk Data traffic to AF11
  class AUTOQOS_TRANSACTION_CLASS
    set dscp af21
    ! Marks Transactional Data traffic to AF21
  class AUTOQOS_SCAVANGER_CLASS
    set dscp cs1
    ! Marks Scavenger traffic to CS1
  class AUTOQOS_SIGNALING_CLASS
    set dscp cs3
    ! Marks Signaling traffic to CS3
  class AUTOQOS_DEFAULT_CLASS
    set dscp default
    ! An explicit default class marks Best Effort traffic to DF
!
<snip>

! This section applies the AutoQoS-Classify policy-map to the interface
interface GigabitEthernet1/0/1
  description L2-ACCESS-PORT-TO-PC
  switchport access vlan 10
  switchport voice vlan 110
  srr-queue bandwidth share 1 30 35 5
  queue-set 2
  priority-queue out
  auto qos classify
  spanning-tree portfast
  service-policy input AUTOQOS-SRND4-CLASSIFY-POLICY
  ! Attaches the AutoQoS-Classify service-policy to the interface
!
<snip>

```

```

! This section defines the Extended IP Access-Lists for AutoQoS-Classify
ip access-list extended AUTOQOS-ACL-BULK-DATA
  permit tcp any any eq 22
  permit tcp any any eq 465
  permit tcp any any eq 143
  permit tcp any any eq 993
  permit tcp any any eq 995
  permit tcp any any eq 1914
  permit tcp any any eq ftp
  permit tcp any any eq ftp-data
  permit tcp any any eq smtp
  permit tcp any any eq pop3

ip access-list extended AUTOQOS-ACL-DEFAULT
  permit ip any any

ip access-list extended AUTOQOS-ACL-MULTIENHANCED-CONF
  permit udp any any range 16384 32767

ip access-list extended AUTOQOS-ACL-SCAVANGER
  permit tcp any any range 2300 2400
  permit udp any any range 2300 2400
  permit tcp any any range 6881 6999
  permit tcp any any range 28800 29100
  permit tcp any any eq 1214
  permit udp any any eq 1214
  permit tcp any any eq 3689
  permit udp any any eq 3689
  permit tcp any any eq 11999

ip access-list extended AUTOQOS-ACL-SIGNALING
  permit tcp any any range 2000 2002
  permit tcp any any range 5060 5061
  permit udp any any range 5060 5061

ip access-list extended AUTOQOS-ACL-TRANSACTIONAL-DATA
  permit tcp any any eq 443
  permit tcp any any eq 1521
  permit udp any any eq 1521
  permit tcp any any eq 1526
  permit udp any any eq 1526
  permit tcp any any eq 1575
  permit udp any any eq 1575
  permit tcp any any eq 1630
  permit udp any any eq 1630
!
<snip>

```

As can be seen from the configuration output in [Example 2-49](#), the **auto qos classify** command generates class-maps, associated extended IP access-lists, and a policy map, which is attached to the interface (along with input and output queuing policies, which will be discussed in detail a following section). Again, it should be noted that the IP access-list entries are based on the sample ports summarized in [Table 2-9](#) and that these are just generic application examples for these classes and the administrator can add/change/delete the access-list entries to match on their specific applications.

This configuration can be verified with the commands:

- **show mls qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Additionally, should the administrator wish to enable data-plane policing/scavenger-class QoS policies on these application classes, they may do so by including the option keyword **police** in conjunction with the **auto qos classify** interface command, as shown in [Example 2-50](#).

Example 2-50 Auto QoS Classify and Police Configuration on a Layer 2 Switch Port

```
C3750-E(config)#interface GigabitEthernet1/0/1
C3750-E(config-if)# description L2-ACCESS-PORT-TO-PC
C3750-E(config-if)# switchport access vlan 10
C3750-E(config-if)# spanning-tree portfast
C3750-E(config-if)# auto qos classify police
! Auto-configures classify & police policy (+ ingress and egress queuing policies)
```

The effect of this **auto qos classify police** policy on a Layer 2 switch port can be verified by the **show run** command, as shown in [Example 2-51](#).



Note

For the sake of brevity and to minimize redundancy, the class-maps and extended IP access-lists, which are identical to those shown in [Example 2-50](#), are not repeated in [Example 2-51](#).

Example 2-51 Auto QoS Classify & Police Configuration on a Layer 2 Switch Port Verification—how run

```
C3750-E# show run
Building configuration...
<snip>

!
! This section configures the global policed-DSCP markdown map
mls qos map policed-dscp 0 10 18 to 8
! DSCP 0 (DF), 10 (AF11) and 18 (AF21) are marked down to 8 (CS1)
! if found to be in excess of their (respective) policing rates
!
<snip>

! This section defines the policy-map for AutoQoS-Classify-Police
policy-map AUTOQOS-SRND4-CLASSIFY-POLICE-POLICY
class AUTOQOS_MULTIHANCED_CONF_CLASS
  set dscp af41
  police 5000000 8000 exceed-action drop
  ! Multimedia-conferencing is marked AF41 and policed to drop at 5 Mbps
class AUTOQOS_BULK_DATA_CLASS
  set dscp af11
  police 10000000 8000 exceed-action policed-dscp-transmit
  ! Bulk-data is marked AF11 and policed to remark (to CS1) at 10 Mbps
class AUTOQOS_TRANSACTION_CLASS
  set dscp af21
  police 10000000 8000 exceed-action policed-dscp-transmit
  ! Trans-data is marked AF21 and policed to remark (to CS1) at 10 Mbps
class AUTOQOS_SCAVANGER_CLASS
  set dscp cs1
  police 10000000 8000 exceed-action drop
  ! Scavenger traffic is marked CS1 and policed to drop at 10 Mbps
class AUTOQOS_SIGNALING_CLASS
  set dscp cs3
  police 32000 8000 exceed-action drop
  ! Signaling is marked CS3 and policed to drop at 32 kbps
class AUTOQOS_DEFAULT_CLASS
  set dscp default
```

```

    police 10000000 8000 exceed-action policed-dscp-transmit
    ! An explicit default class marks all other IP traffic to DF
    ! and polices all other IP traffic to remark (to CS1) at 10 Mbps
    !
<snip>

! This section applies the AutoQoS-Classify-Police policy-map to the interface
interface GigabitEthernet1/0/1
description L2-ACCESS-PORT-TO-PC
switchport access vlan 10
switchport voice vlan 110
srr-queue bandwidth share 1 30 35 5
queue-set 2
priority-queue out
auto qos classify police
spanning-tree portfast
service-policy input AUTOQOS-SRND4-CLASSIFY-POLICE-POLICY
    ! Attaches the AutoQoS-Classify service-policy to the interface
    !
<snip>

```

As can be seen from the configuration output in [Example 2-51](#), the two principle changes in the configuration attributable to the **police** keyword used in conjunction with the **auto qos classify** command are:

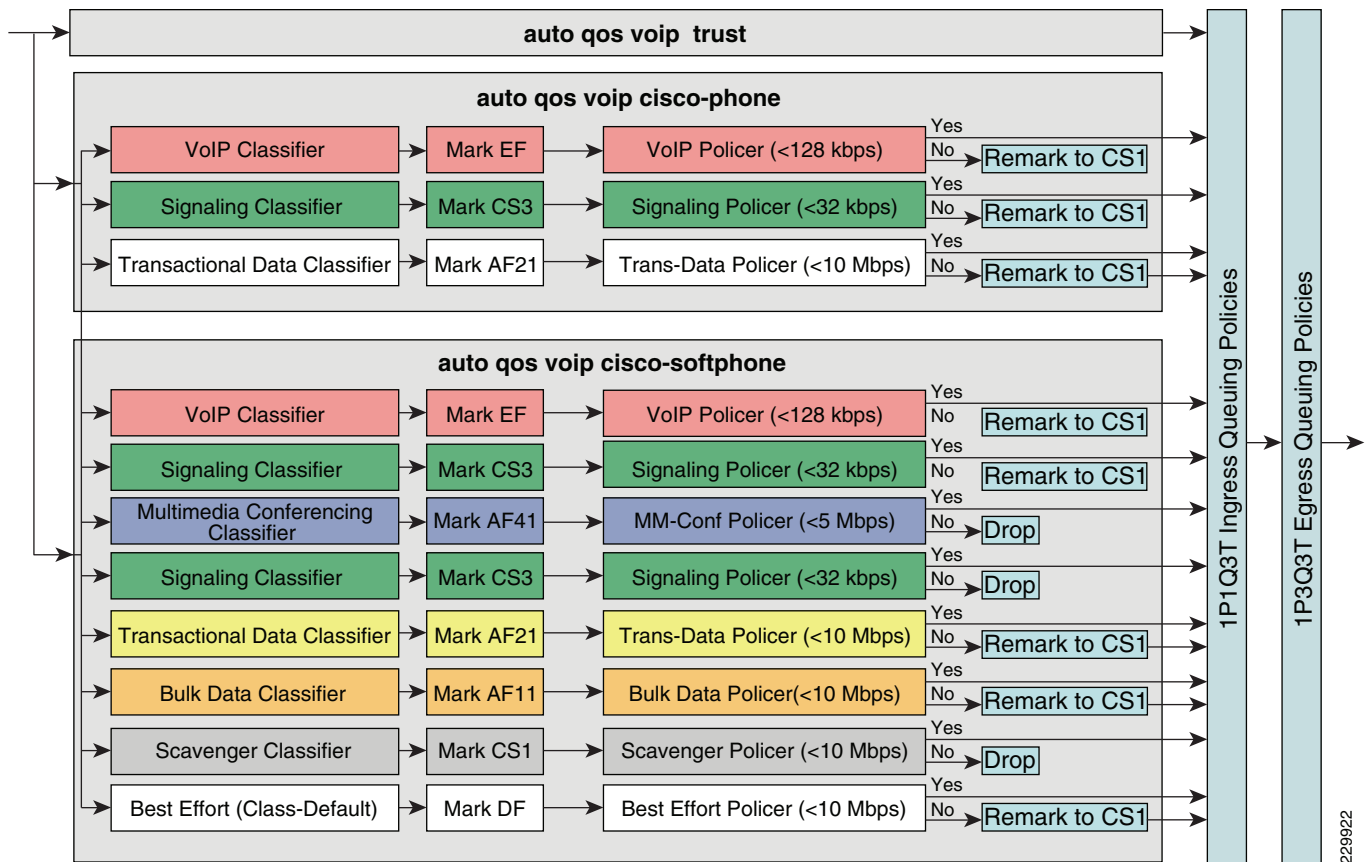
- A globally-defined **policed-dscp** map to mark down DF (0), AF11 (10), and AF21 (18) to CS1 (8)—if found to be exceeding their respective policing rates.
- An amended policy-map that polices multimedia conferencing traffic (to drop if exceeding 5 Mbps), bulk data (to remark if exceeding 10 Mbps), transactional data (to remark if exceeding 10 Mbps), scavenger (to drop if exceeding 10 Mbps), signaling (to drop if exceeding 32 Kbps), and best-effort traffic (to remark if exceeding 10 Mbps).

This configuration can be verified with the commands:

- **show mls qos maps policed-dscp**
- **show mls qos interface**
- **show mls qos interface interface x/y policers**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Auto QoS VoIP Models

As with legacy AutoQoS-VoIP, there are three deployment options for AutoQoS (SRND4) VoIP: **trust**, **cisco-phone**, and **cisco-softphone**. These updated **auto qos voip** deployment options—complete with ingress and egress queuing configurations—are illustrated in [Figure 2-19](#).

Figure 2-19 Auto QoS VoIP (SRND4) Models

Each of these **auto qos voip** deployment options will be detailed in turn.

The first point to be noted is that since the SRND4 versions of **auto qos voip** expands functionality beyond the original AutoQoS-VoIP feature, the administrator must indicate which version of this AutoQoS VoIP is desired. By default, simply entering **auto qos voip** interface-configuration commands will invoke legacy legacy AutoQoS-VoIP configurations; however, if the administrator first enters **auto qos srnd4** in the global configuration command *prior* to applying these **auto qos voip** interface-configuration commands, then the SRND4 versions of **auto qos voip** will be applied.

The **auto qos voip trust** option is a legacy deployment option, which has been largely relegated by the previously-discussed **auto qos trust** option. Like **auto qos trust**, **auto qos voip trust** will configure static CoS-trust on Layer 2 switch ports and static DSCP-trust on Layer 3 routed interfaces. However, unlike **auto qos trust** there is no additional **cos** or **dscp** keyword option to override these default trust-settings (but this may be manually overridden with an explicitly defined **mls qos trust [cos | dscp]** interface configuration command). [Example 2-52](#) shows **auto qos voip trust** being applied to a Layer 2 switch port.

Example 2-52 Auto QoS VoIP Trust Applied on a Layer 2 Switch Port

```
! This section specifies that SRND4 version of AutoQoS is to be enabled
C3750-E(config)# auto qos srnd4
! Globally defines the current version of AutoQoS to be SRND4

! This section configures auto qos trust on a L2 switch port
C3750-E(config)# interface GigabitEthernet1/0/1
```

```

C3750-E(config-if)# description L2-ACCESS-PORT
C3750-E(config-if)# switchport access vlan 10
C3750-E(config-if)# switchport voice vlan 110
C3750-E(config-if)# spanning-tree portfast
C3750-E(config-if)# auto qos voip trust
! Auto-configures static trust policy (and ingress and egress queuing policies)

```

The effect of this **auto qos voip trust** policy on a Layer 2 switch port can be verified by the **show run interface** command, as shown in [Example 2-53](#).

Example 2-53 Auto QoS VoIP Trust Applied on a Layer 2 Switch Port Verification—show run interface

```

C3750-E# show run interface GigabitEthernet1/0/1
Building configuration...

Current configuration : 251 bytes
!
interface GigabitEthernet1/0/1
  description L2-ACCESS-PORT
  switchport access vlan 10
  switchport voice vlan 110
  srr-queue bandwidth share 1 30 35 5
  queue-set 2
  priority-queue out
  mls qos trust cos
  ! AutoQoS has configured the port to static CoS-trust
  auto qos trust
  spanning-tree portfast
end

C3750-E#

```

The **show run interface** command shows that although **auto qos voip trust** was configured on the interface, this has been converted and replaced with **auto qos trust** (as this is functionally equivalent on this Layer 2 switch port interface).

This interface trust state and CoS-to-DSCP maps can be verified by the commands:

- **show mls qos interface**
- **show mls qos maps cos-dscp**

A second deployment option offered by the (SRND4) **auto qos voip** feature is to use the **cisco-phone** keyword. As previously mentioned, the administrator must first enter **auto qos srnd4** in the global configuration prior to entering **auto qos voip cisco-phone** on a specific interface(s). When **auto qos voip cisco-phone** is configured on a Layer 2 switch port, it will dynamically extend trust-CoS to Cisco IP Phones; when configured on Layer 3 routed interfaces, it will dynamically extend trust-DSCP to Cisco IP Phones. Additionally, this command will configure data-plane policing/scavenger-class QoS policies on voice, signaling and best-effort traffic, as shown in [Example 2-54](#) and [Example 2-55](#).

Example 2-54 Auto QoS VoIP Cisco-Phone (SRND4) Applied on a Layer 2 Switch Port

```

! This section specifies that SRND4 version of AutoQoS is to be enabled
C3750-E(config)# auto qos srnd4
! Globally defines the current version of AutoQoS to be SRND4

! This section applies AutoQoS (SRND4) to a layer 2 switch port
C3750-E(config)# interface GigabitEthernet1/0/1
C3750-E(config-if)# description L2-ACCESS-PORT
C3750-E(config-if)# switchport access vlan 10

```

```

C3750-E(config-if)# switchport voice vlan 110
C3750-E(config-if)# spanning-tree portfast
C3750-E(config-if)# auto qos voip cisco-phone
    ! Auto-configures conditional-trust and marking and policing policies for IP Phones
    ! As well as ingress and egress queuing policies

```

The effect of this **auto qos voip cisco-phone** policy on a Layer 2 switch port can be verified by the **show run** command, as shown in [Example 2-55](#).

Example 2-55 Auto QoS VoIP Cisco-Phone (SRND4) Applied on a Layer 2 Switch Port Verification—show run

```

C3750-E# show run
Building configuration...
<snip>

! This section confirms the AutoQoS version currently enabled
auto qos srnd4
!

! This section defines the AutoQoS-VoIP-Cisco-Phone (SRND4) Class-Maps
class-map match-all AUTOQOS_VOIP_DATA_CLASS
  match ip dscp ef
  ! Voice is matched on DSCP EF

class-map match-all AUTOQOS_DEFAULT_CLASS
  match access-group name AUTOQOS-ACL-DEFAULT
  ! An explicit default class matches all other traffic via IP ACL

class-map match-all AUTOQOS_VOIP_SIGNAL_CLASS
  match ip dscp cs3
  ! Signaling traffic is matched on CS3
!

! This section defines the AutoQoS-VoIP-Cisco-Phone (SRND4) Policy-Map
policy-map AUTOQOS-SRND4-CISCOPHONE-POLICY
  class AUTOQOS_VOIP_DATA_CLASS
    set dscp ef
    police 128000 8000 exceed-action policed-dscp-transmit
    ! Voice is marked to DSCP EF and policed (to remark) if exceeding 128 kbps
  class AUTOQOS_VOIP_SIGNAL_CLASS
    set dscp cs3
    police 32000 8000 exceed-action policed-dscp-transmit
    ! Signaling is marked to DSCP CS3 and policed (to remark) if exceeding 32 kbps
  class AUTOQOS_DEFAULT_CLASS
    set dscp default
    police 10000000 8000 exceed-action policed-dscp-transmit
    ! An explicit default class marks all other IP traffic to DF
    ! and polices all other IP traffic to remark (to CS1) at 10 Mbps
!

! This section attaches the AutoQoS-VoIP-Cisco-Phone (SRND4) Policy-Map to the interface
interface GigabitEthernet1/0/1
  description L2-ACCESS-PORT
  switchport access vlan 10
  switchport voice vlan 110
  srr-queue bandwidth share 1 30 35 5
  queue-set 2
  priority-queue out
  mls qos trust device cisco-phone

```

```

! AutoQoS has configured a conditional-trust policy for cisco-phone devices
mls qos trust cos
! AutoQoS has configured CoS-trust to be dynamically extended
auto qos voip cisco-phone
spanning-tree portfast
service-policy input AUTOQOS-SRND4-CISCOPHONE-POLICY
! Attaches the AutoQoS-VoIP-Cisco-Phone (SRND4) Policy-Map to the interface
!
<snip>

! This section defines the explicit-default extended IP ACL
ip access-list extended AUTOQOS-ACL-DEFAULT
permit ip any any
!

```

[Example 2-55](#) shows that the applied version of **auto qos voip** is **srnd4** and, as such, voice is policed to remark if exceeding 128 kbps, signaling is policed to remark if exceeding 32 kbps and best effort traffic is policed to remark to scavenger if exceeding 10 Mbps.

This configuration can be verified with the commands:

- **show mls qos maps policed-dscp**
- **show mls qos interface**
- **show mls qos interface interface x/y policers**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

A third deployment option offered by the (SRND4) **auto qos voip** feature is to use the **cisco-softphone** keyword. As previously mentioned, the administrator must first enter **auto qos srnd4** in the global configuration prior to entering **auto qos voip cisco-softphone** on specific interface(s).

In addition to the voice and signaling classes, six additional application classes (multimedia conferencing, signaling, transactional data, bulk data, scavenger and best-effort) are automatically defined via class-maps. Each class-map references an associated extended IP access-list. These IP access lists define the TCP and UDP port numbers of the given class of applications, based on the sample ports summarized in [Table 2-9](#). *However, it cannot be overemphasized that these are just generic application examples for these classes and the administrator can add/change/delete the access-list entries to match on their specific applications.*

[Example 2-56](#) shows the application of **auto qos voip cisco-softphone** on a Layer 2 switch port interface.

Example 2-56 Auto QoS VoIP Cisco-SoftPhone (SRND4) Applied on a Layer 2 Switch Port

```

! This section specifies that SRND4 version of AutoQoS is to be enabled
C3750-E(config)# auto qos srnd4
! Globally defines the current version of AutoQoS to be SRND4

! This section applies AutoQoS (SRND4) to a layer 2 switch port
C3750-E(config)# interface GigabitEthernet1/0/1
C3750-E(config-if)# description L2-ACCESS-PORT
C3750-E(config-if)# switchport access vlan 10
C3750-E(config-if)# switchport voice vlan 110
C3750-E(config-if)# spanning-tree portfast
C3750-E(config-if)# auto qos voip cisco-softphone
! Auto-configures conditional-trust and marking and policing policies for softphones

```

! As well as ingress and egress queuing policies

The effect of this **auto qos voip cisco-softphone** policy on a Layer 2 switch port can be verified by the **show run** command, as shown in [Example 2-57](#).



Note

For the sake of brevity and to minimize redundancy, the class-maps and extended IP access-lists, which are identical to those shown in [Example 2-51](#) and [Example 2-55](#), are not repeated here.

Example 2-57 Auto QoS VoIP Cisco-SoftPhone (SRND4) Applied on a Layer 2 Switch Port Verification—show run

```
C3750-E# show run
Building configuration...
<snip>

! This section confirms the AutoQoS version currently enabled
auto qos srnd4
!
<snip>

! This section defines the AutoQoS-VoIP-Cisco-SoftPhone (SRND4) Policy-Map
policy-map AUTOQOS-SRND4-SOFTPHONE-POLICY
class AUTOQOS_VOIP_DATA_CLASS
  set dscp ef
  police 128000 8000 exceed-action policed-dscp-transmit
  ! Voice is marked to DSCP EF and policed (to remark) if exceeding 128 kbps
class AUTOQOS_VOIP_SIGNAL_CLASS
  set dscp cs3
  police 32000 8000 exceed-action policed-dscp-transmit
  ! Signaling is marked to DSCP CS3 and policed (to remark) if exceeding 32 kbps
class AUTOQOS_MULTIMEDIA_CONF_CLASS
  set dscp af41
  police 5000000 8000 exceed-action drop
  ! MM-Conf is marked to DSCP AF41 and policed (to drop) if exceeding 5 Mbps
class AUTOQOS_BULK_DATA_CLASS
  set dscp af11
  police 10000000 8000 exceed-action policed-dscp-transmit
  ! Bulk Data is marked to DSCP AF11 and policed (to remark) if exceeding 10 Mbps
class AUTOQOS_TRANSACTION_CLASS
  set dscp af21
  police 10000000 8000 exceed-action policed-dscp-transmit
  ! Trans-Data is marked to DSCP AF21 and policed (to remark) if exceeding 10 Mbps
class AUTOQOS_SCAVANGER_CLASS
  set dscp cs1
  police 10000000 8000 exceed-action drop
  ! Scavenger is marked to DSCP CS1 and policed (to drop) if exceeding 10 Mbps
class AUTOQOS_SIGNALING_CLASS
  set dscp cs3
  police 32000 8000 exceed-action drop
  ! Signaling is marked to DSCP CS3 and policed (to drop) if exceeding 32 kbps
class AUTOQOS_DEFAULT_CLASS
  set dscp default
  ! An explicit default class marks all other IP traffic to DF
!
<snip>
```

```

! This section attaches the AutoQoS-VoIP-Cisco-SoftPhone (SRND4) Policy-Map to the
interface
interface GigabitEthernet1/0/1
description L2-ACCESS-PORT
switchport access vlan 10
switchport voice vlan 110
srr-queue bandwidth share 1 30 35 5
queue-set 2
priority-queue out
auto qos voip cisco-softphone
spanning-tree portfast
service-policy input AUTOQOS-SRND4-SOFTPHONE-POLICY
! Attaches the AutoQoS-VoIP-Cisco-SoftPhone (SRND4) Policy-Map to the interface
!

```

As shown in [Example 2-57](#), the **auto qos voip cisco-softphone** policy essentially combines the **auto qos voip cisco-phone** policy with the **auto qos classify police** policy.

This configuration can be verified with the commands:

- **show mls qos maps policed-dscp**
- **show mls qos interface**
- **show mls qos interface interface x/y policers**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Auto QoS 1P1Q3T Ingress Queuing Model

The AutoQoS SRND4 ingress queuing model is illustrated in [Figure 2-16](#). These ingress queuing policies are automatically configured along with any other AutoQoS SRND4 QoS model and are shown in [Example 2-58](#).

Example 2-58 Auto QoS (SRND4) 1P1Q3T Ingress Queuing Verification—show run

```

C3750-E# show run
Building configuration...
<snip>

! This section displays (non-default) input queue parameters
mls qos srr-queue input bandwidth 70 30
! Q1 is assigned 70% BW via SRR shared weights
! Q2 SRR shared weight is ignored (as it has been configured as a PQ)
mls qos srr-queue input threshold 1 80 90
! Q1 thresholds are configured at 80% (Q1T1) and 90% (Q1T2)
! Q1T3 is implicitly set at 100% (the tail of the queue)
! Q2 thresholds are all set (by default) to 100% (the tail of Q2)
mls qos srr-queue input priority-queue 2 bandwidth 30
! Q2 is enabled as a strict-priority ingress queue with 30% BW

! This section displays (non-default) ingress CoS-to-Queue mappings
mls qos srr-queue input cos-map queue 1 threshold 2 3
! CoS value 3 is mapped to ingress Q1T2
mls qos srr-queue input cos-map queue 1 threshold 3 6 7
! CoS values 6 and 7 are mapped to ingress Q1T3
mls qos srr-queue input cos-map queue 2 threshold 1 4
! CoS values 4 is mapped to ingress Q2 (the PQ)

```



```

! This section displays (non-default) ingress DSCP-to-Queue mappings
mls qos srr-queue input dscp-map queue 1 threshold 2 24
! DSCP CS3 is mapped to ingress Q1T2
mls qos srr-queue input dscp-map queue 1 threshold 3 48 49 50 51 52 53 54 55
! DSCP CS6 (48) and non-standard DSCPs 49-55 are mapped to Q1T3
mls qos srr-queue input dscp-map queue 1 threshold 3 56 57 58 59 60 61 62 63
! DSCP CS7 (56) and non-standard DSCPs 57-63 are mapped to Q1T3
mls qos srr-queue input dscp-map queue 2 threshold 3 32 33 40 41 42 43 44 45
! DSCP CS4 (32), CS5 (40) and non-standard DSCPs 33-45 are mapped to Q2T3
mls qos srr-queue input dscp-map queue 2 threshold 3 46 47
! DSCP EF (46) and non-standard DSCP 47 are mapped to Q2T3

```

**Note**

Ingress queuing is not supported on the Cisco Catalyst 2960-S.

Comparing the AutoQoS 1P3Q1T Ingress Queuing ([Example 2-58](#)) with the previously-defined manual 1P3Q1T Ingress Queuing Configuration ([Example 2-20](#)) will reveal virtually identical queuing models, with the only differences in configuration being that default settings and/or mappings are not shown in the AutoQoS example configuration; additionally, the AutoQoS example includes some mappings of non-standard DSCP values to queues (which-as previously noted-were omitted from previous examples for the sake of simplicity).

This configuration can be verified with the commands:

- **show mls qos input-queue**
- **show mls qos maps cos-input-q**
- **show mls qos maps dscp-input-q**

Auto QoS 1P3Q3T Egress Queuing Model

The AutoQoS SRND4 egress queuing model is illustrated in [Figure 2-17](#). These egress queuing policies are automatically configured along with any other AutoQoS SRND4 QoS model. The egress queuing policies automatically configured by AutoQoS SRND4 are shown in [Example 2-59](#).

Example 2-59 Auto QoS (SRND4) 1P3Q3T Egress Queuing Verification—show run

```

C3750-E# show run
Building configuration...
<snip>

! This section displays (non-default) egress CoS-to-Queue mappings
mls qos srr-queue output cos-map queue 1 threshold 3 4 5
! CoS 4 and 5 are mapped to egress Q1T3 (the tail of the PQ)
mls qos srr-queue output cos-map queue 2 threshold 1 2
! CoS 2 is mapped to egress Q2T1
mls qos srr-queue output cos-map queue 2 threshold 2 3
! CoS 3 is mapped to egress Q2T2
mls qos srr-queue output cos-map queue 2 threshold 3 6 7
! CoS 6 and 7 are mapped to Q2T3
mls qos srr-queue output cos-map queue 3 threshold 3 0
! CoS 0 is mapped to Q3T3 (the tail of the default queue)
mls qos srr-queue output cos-map queue 4 threshold 3 1
! CoS 1 is mapped to Q4T3 (tail of the less-than-best-effort queue)

! This section displays (non-default) egress DSCP-to-Queue mappings
mls qos srr-queue output dscp-map queue 1 threshold 3 32 33 40 41 42 43 44 45

```

```

! Maps CS4 (32) and non-standard DSCPs 33-45 to Q1T3 (the tail of the PQ)
mls qos srr-queue output dscp-map queue 1 threshold 3 46 47
! Maps EF (46) and non-standard DSCP 47 to Q1T3 (the tail of the PQ)
mls qos srr-queue output dscp-map queue 2 threshold 1 16 17 18 19 20 21 22 23
! Maps CS2 (16) and AF2 (18/20/22) and non-standard DSCPs between 17-23 to Q2T1
mls qos srr-queue output dscp-map queue 2 threshold 1 26 27 28 29 30 31 34 35
! Maps AF3 (26/28/30) and AF41 (34) and non-standard DSCPs between 27-35 to Q2T1
mls qos srr-queue output dscp-map queue 2 threshold 1 36 37 38 39
! Maps AF42 (36) and AF43 (38) and non-standard DSCPs between 37-39 to Q2T1
mls qos srr-queue output dscp-map queue 2 threshold 2 24
! Maps CS3 (24) to Q2T2
mls qos srr-queue output dscp-map queue 2 threshold 3 48 49 50 51 52 53 54 55
! Maps CS6 (48) and non-standard DSCPs between 49-55 to Q2T3
mls qos srr-queue output dscp-map queue 2 threshold 3 56 57 58 59 60 61 62 63
! Maps CS7 (56) and non-standard DSCPs between 57-63 to Q2T3
mls qos srr-queue output dscp-map queue 3 threshold 3 0 1 2 3 4 5 6 7
! Maps DF (0) and non-standard DSCPs between 1-7 to Q3T3 (tail of best-effort queue)
mls qos srr-queue output dscp-map queue 4 threshold 1 8 9 11 13 15
! Maps CS1 and non-standard DSCPs between 9-15 to Q4T1
mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
! Maps AF1 (10/12/14) to Q4T2

```

```

! This section displays (non-default) egress queue parameters
mls qos queue-set output 1 threshold 1 100 100 50 200
! Q1T1 is set to 100%; Q1T2 is set to 100%;
! Q1 (PQ) Reserve Threshold is set to 100%;
! Q1 (PQ) Maximum (Overflow) Threshold is set to 200%
mls qos queue-set output 1 threshold 2 125 125 100 400
! Q2T1 is set to 125%; Q2T2 is set to 125%;
! Q2 Reserve Threshold is set to 100%;
! Q2 Maximum (Overflow) Threshold is set to 400%
mls qos queue-set output 1 threshold 3 100 100 100 400
! Q3T1 is set to 100%, Q2T2 is set to 100%
! Q3 Reserve Threshold is set to 100%;
! Q3 Maximum (Overflow) Threshold is set to 400%
mls qos queue-set output 1 threshold 4 60 150 50 200
! Q4T1 is set to 60%; Q4T2 is set to 150%
! Q4 Reserve Threshold is set to 50%;
! Q4 Maximum (Overflow) Threshold is set to 200%
mls qos queue-set output 1 buffers 15 25 40 20
! Allocates 15% of buffers to Q1; 25% to Q2; 40% to Q3 and 20% to Q4
<snip>

```

```

! This section displays (non-default) interface egress queuing settings
interface GigabitEthernet1/0/1
description L2-ACCESS-PORT
switchport access vlan 10
switchport voice vlan 110
srr-queue bandwidth share 1 30 35 5
! The SRR sharing weights are set to allocate 30% BW to Q2
! 35% BW to Q3 and 5% BW to Q4
! Q1 SRR sharing weight is ignored, as it will be configured as a PQ
queue-set 2
! The interface(s) is assigned to queue-set 2
priority-queue out
! Q1 is enabled as a strict priority queue
mls qos trust cos
auto qos trust
spanning-tree portfast
!

```

Comparing the AutoQoS 1P3Q4T Egress Queuing ([Example 2-59](#)) with the previous-defined manual 1P3Q4T Egress Queuing Configuration ([Example 2-24](#)) will reveal virtually identical queuing models, with a few minor differences. These configuration differences are due to:

- That default settings and/or mappings are not shown in the AutoQoS example configuration.
- The AutoQoS example includes some mappings of non-standard DSCP values to queues (which—as previously noted—were omitted from previous examples for the sake of simplicity.
- Some minor administrative preferences on the part of the platform engineers versus the authors of this document (mainly relating to buffer and threshold fine-tuning).

This configuration can be verified with the commands:

- **show mls qos queue-set**
- **show mls qos maps cos-output-q**
- **show mls qos maps dscp-output-q**
- **show mls qos interface *interface x/y* queuing**
- **show mls qos interface *interface x/y* statistics**

Auto QoS and EtherChannels

At the time of writing, AutoQoS is not supported on EtherChannels. However, if AutoQoS SRND4 has been deployed on any other switch port or interface, it is a fairly simple matter to apply the same ingress and egress queuing policies generated by AutoQoS to EtherChannel port-member interfaces. Specifically, only four incremental interface-configuration commands are required to be applied to the EtherChannel physical port-member interfaces, as shown in [Example 2-60](#).

Example 2-60 EtherChannel AutoQoS SRND4 Design on a Catalyst 3750-E.

```
! This section configures Trust-DSCP and AutoQoS SRND4 Egress Queuing across the
! (physical) EtherChannel member-ports
! Also AutoQoS SRND4 1P3Q1T Ingress Queuing Policies will also automatically apply
C3750-E(config)# interface range TenGigabitEthernet1/0/1-2
C3750-E(config-if-range)# description PORT-CHANNEL1-PHYSICAL-PORT-MEMBER
C3750-E(config-if-range)# switchport mode trunk
C3750-E(config-if-range)# switchport trunk encapsulation dot1q
C3750-E(config-if-range)# switchport trunk allowed vlan 10,110
C3750-E(config-if-range)# channel-group 1 mode auto
! Associates the physical ports with the logical EtherChannel bundle
C3750-E(config-if-range)# mls qos trust dscp
! The physical port-member interfaces are set to statically trust DSCP
C3750-E(config-if-range)# queue-set 2
! The interfaces are assigned to queue-set 2
C3750-E(config-if-range)# srr-queue bandwidth share 1 30 35 5
! The SRR sharing weights are set to allocate 30% BW to Q2
! 35% BW to Q3 and 5% BW to Q4
! Q1 SRR sharing weight is ignored, as it will be configured as a PQ
C3750-E(config-if-range)# priority-queue out
! Q1 is enabled as a strict priority queue
```

This configuration can be verified with the commands:

- **show mls qos interface**
- **show mls qos input-queue**
- **show mls qos maps cos-input-q**

- `show mls qos maps dscp-input-q`
- `show mls qos maps cos-output-q`
- `show mls qos maps dscp-output-q`
- `show mls qos interface interface x/y queueing`
- `show mls qos interface interface x/y statistics`

Auto QoS Removal

Some administrators may be a bit surprised to see the amount of incremental QoS-related configuration generated by AutoQoS. This surprise may at times even turn into alarm if they attempt a `no auto qos...` command and still see a significant amount of QoS policies lingering around in their switch configuration. The reason this feature has been implemented in this manner is that if an administrator changes a switch port's QoS role, the switch does not change its entire queuing, mapping, buffer, and threshold configurations—which may potentially adversely effect traffic. Nonetheless, some administrators feel more comfortable with a script that can quickly revert all QoS configurations back to default. To this end, such a script is shown in [Example 2-61](#).



Note

This script can only be run when all `auto qos` commands have been removed from all interfaces.

Example 2-61 Command Script to Remove AutoQoS Configurations

```
! This section removes AutoQoS-SRND4 Policy-Maps (as applicable)
C3750-E(config)# no policy-map AUTOQOS-SRND4-CLASSIFY-POLICY
C3750-E(config)# no policy-map AUTOQOS-SRND4-CLASSIFY-POLICE-POLICY
C3750-E(config)# no policy-map AUTOQOS-SRND4-CISCOPHONE-POLICY
C3750-E(config)# no policy-map AUTOQOS-SRND4-SOFTPHONE-POLICY

! This section removes AutoQoS SRND4 Class-Maps (as applicable)
C3750-E(config)# no class-map AUTOQOS_VOIP_DATA_CLASS
C3750-E(config)# no class-map AUTOQOS_VOIP_SIGNAL_CLASS
C3750-E(config)# no class-map AUTOQOS_MULTITIENHANCED_CONF_CLASS
C3750-E(config)# no class-map AUTOQOS_SIGNALING_CLASS
C3750-E(config)# no class-map AUTOQOS_TRANSACTION_CLASS
C3750-E(config)# no class-map AUTOQOS_BULK_DATA_CLASS
C3750-E(config)# no class-map AUTOQOS_SCAVANGER_CLASS
C3750-E(config)# no class-map AUTOQOS_DEFAULT_CLASS

! This section removes AutoQoS SRND4 IP ACLs (as applicable)
C3750-E(config)# no ip access-list extended AUTOQOS-ACL-MULTITIENHANCED-CONF
C3750-E(config)# no ip access-list extended AUTOQOS-ACL-SIGNALING
C3750-E(config)# no ip access-list extended AUTOQOS-ACL-TRANSACTIONAL-DATA
C3750-E(config)# no ip access-list extended AUTOQOS-ACL-BULK-DATA
C3750-E(config)# no ip access-list extended AUTOQOS-ACL-SCAVANGER
C3750-E(config)# no ip access-list extended AUTOQOS-ACL-DEFAULT

! This section resets QoS maps to default values
C3750-E(config)# default mls qos map cos-dscp
C3750-E(config)# default mls qos map policed-dscp

! This section resets Ingress Queuing to default values
C3750-E(config)# default mls qos srr-queue input cos-map
C3750-E(config)# default mls qos srr-queue input dscp-map
C3750-E(config)# default mls qos srr-queue input buffers
C3750-E(config)# default mls qos srr-queue input bandwidth
C3750-E(config)# default mls qos srr-queue input threshold 1
```

```
C3750-E(config)# default mls qos srr-queue input priority-queue 2
```

```
! This section resets Egress Queuing to default values
```

```
C3750-E(config)# default mls qos srr-queue output cos-map
```

```
C3750-E(config)# default mls qos srr-queue output dscp-map
```

```
C3750-E(config)# default mls qos queue-set output 1 buffers
```

```
C3750-E(config)# default mls qos queue-set output 1 threshold
```

```
! This section removes AutoQoS version from the Global Config
```

```
C3750-E(config)# no auto qos srnd4
```

These changes can be verified with a simple:

- **show run**

Cisco Catalyst 4500/4900 and 4500-E/4900M QoS Design

The Catalyst 4500 family of switches provides Layer 2 through Layer 4 network services, including advanced high availability, security, and QoS services in addition to integrated PoE to support unified communications. The Catalyst 4500 and 4900 share common features and syntax and as such are grouped together and discussed as a single switch family, namely the Catalyst 4500 (which may also be designated as ,Classic Supervisors,). Similarly, the Catalyst 4500-E and 4900M share common features and syntax and as such are also grouped together and discussed as a single switch family, namely the Catalyst 4500-E (which may also be designated as Supervisor 6-E).

Catalyst 4500/4500-E switches come in various chassis, supervisor, and linecard combinations, as are discussed in turn.

The Catalyst 4500 switch family offers chassis that support 3, 6, 7, and 10 slots; these models include the Catalyst 4503, 4506, 4507R, and 4510R, respectively (the latter two models supporting a redundant supervisor option). Catalyst 4500 chassis provide 6 Gbps of bandwidth per linecard slot.

Similarly, the Catalyst 4500-E switch family offers chassis that support 3, 6, 7, and 10 slots; these models include the Catalyst 4503-E, 4506-E, 4507R-E, and 4510R-E, respectively (the latter two models supporting a redundant supervisor option). Catalyst 4500 chassis provide 24 Gbps of bandwidth per linecard slot (or 6 Gbps of bandwidth per non “E-series” linecards).

Multiple supervisor options exist for the Catalyst 4500/4500-E family of switches, including:

- Supervisor II-Plus-TS—Supports basic Layer 2-Layer 4 services at up to 64 Gbps (48-millions of packets per second [mpps]) switching; includes 12 ports of wire-speed 10/100/1000 802.3af Power over Ethernet (PoE) and eight wire-speed SFP ports directly on the supervisor engine.
- Supervisor II-Plus—Supports basic Layer 2-Layer 4 services at up to 64 Gbps (48 mpps) switching; includes two GigabitEthernet uplinks.
- Supervisor II-Plus-10GE—Supports basic Layer 2-Layer 4 services at up to 81 Gbps (108 mpps) switching; includes four GigabitEthernet and two Ten-GigabitEthernet uplinks.
- Supervisor IV—Supports advanced Layer 2-Layer 4 services at up to 64 Gbps (48 mpps) switching; includes two GigabitEthernet uplinks.
- Supervisor V—Supports advanced Layer 2-Layer 4 services at up to 96 Gbps (72 mpps) switching; includes two GigabitEthernet uplinks.
- Supervisor V-10GE—Supports advanced Layer 2-Layer 4 services at up to 136 Gbps (102 mpps) switching; includes four GigabitEthernet and two Ten-GigabitEthernet uplinks.

- Supervisor 6-E—Supports advanced Layer 2-Layer 4 services at up to 320 Gbps (250 mpps) switching; includes two Ten-GigabitEthernet uplinks.
- Supervisor 7-E—Supports advanced Layer 2-Layer 4 services at up to 848 Gbps (250 mpps) switching with nonblocking 48 Gbps per slot; includes four Gigabit or Ten-GigabitEthernet uplinks.

All of the supervisors above—with the exception of the Supervisor 6-E and 7-E—are referred to as Classic Supervisors. There is a major difference between QoS functionality and syntax on the Catalyst 4500 Classic Supervisors, as compared to the Supervisor 6-E and 7-E (which is discussed further in the following section).

The Catalyst 4500/4500-E linecards that meet the minimum requirements for medianet switch ports (including Gigabit Ethernet support, as well as supporting a strict priority hardware queue with at least three additional hardware queues), at the time of writing, are listed in [Table 2-3](#).

Table 2-3 Catalyst 4500/4500-E Linecards for Medianet Campus Networks

Line Card	Number of Ports	Port Speed	Port Type	Wire Rate	Cisco Catalyst 4500 Series Min/Max Ports		
					4503-E	4506-E/ 4507R-E/ 4507R+E	4510R-E/ 4510R+E
E-Series 10 Gigabit Ethernet Line Cards							
WS-X4712-SFP+E	12	10GBASE-R	SFP+ or SFP	2.5-to-1 with SFP+ 1:1 with SFP	12/28	12/64	12/100
WS-X4606-X2-E	6	10GBASE-X	X2 or SFP with TwinGig Converter Module	2.5-to-1 with X2 1:1 with SFP	6/14 12/26	6/34 12/68	6/34 12/68
E-Series 10/100/1000 Line Cards							
WS-X4748-RJ45V+E	48	10/100/1000	RJ-45	1:1	48/96	48/240	48/384
WS-X4648-RJ45V+E	48	10/100/1000	RJ-45	2-to-1	48/96	48/240	48/384
WS-X4648-RJ45V-E	48	10/100/1000	RJ-45	2-to-1	48/96 ⁵	48/240	48/384
WS-X4648-RJ45-E	48	10/100/1000	RJ-45	2-to-1	48/96 ⁵	48/240	48/384
E-Series Gigabit Ethernet SFP Line Cards							
WS-X4624-SFP-E	24	1000	Pluggables	1:1	24/48	24/120	24/168
WS-X4612-SFP-E	12	1000	Pluggables	1:1	12/28	12/64	12/100
Classic 10/100/1000 Line Cards							
WS-X4548-RJ45V+	48	10/100/1000	RJ-45 PoE IEEE 802.3af, Cisco prestandard and PoEP-ready	8-to-1	48/96	48/240	48/384
WS-X4548-GB-RJ45V	48	10/100/1000	RJ-45 PoE IEEE 802.3af and Cisco prestandard	8-to-1	48/96	48/240	48/384


Table 2-3 Catalyst 4500/4500-E Linecards for Medianet Campus Networks

WS-X4524-GB-RJ45V	24	10/100/1000	RJ-45 PoE IEEE 802.3af and Cisco prestandard	4-to-1	24/48	24/120	24/168
WS-X4548-GB-RJ45	48	10/100/1000	RJ-45	8-to-1	48/96	48/240	48/384
Classic Gigabit Ethernet Fiber (GBIC or SFP) Line Cards							
WS-X4306-GB	6	1000BASE-X	GBIC	Yes	6/12	6/30	6/42
WS-X4418-GB	18	1000BASE-X	GBIC	2 ports full 16 ports 4-to-1	18/36	18/90	18/126
WS-X4448-GB-LX	48	1000BASE-LX	48 SFPs (included)	8-to-1	48/96	48/240	48/384
WS-X4448-GB-SFP	48	1000BASE-X	SFP	8-to-1	48/96	48/240	48/384
WS-X4506-GB-T	6 + 6	10/100/1000	1000BASE-X (SFP) RJ-45 PoE IEEE 802.3af and Cisco prestandard	Yes	6/12	6/30	6/42

Platform-Specific QoS Considerations

As mentioned, there is a significant difference in how QoS is implemented on the Catalyst 4500 Classic Supervisor family versus the Catalyst 4500-E Supervisor 6-E family; the former implements a switch-specific QoS (called the switch QoS model), while the latter implements Cisco IOS MQC (called the MQC model) on the switch.

The Catalyst 4500 Classic Supervisor switch QoS model is shown in [Figure 2-20](#).

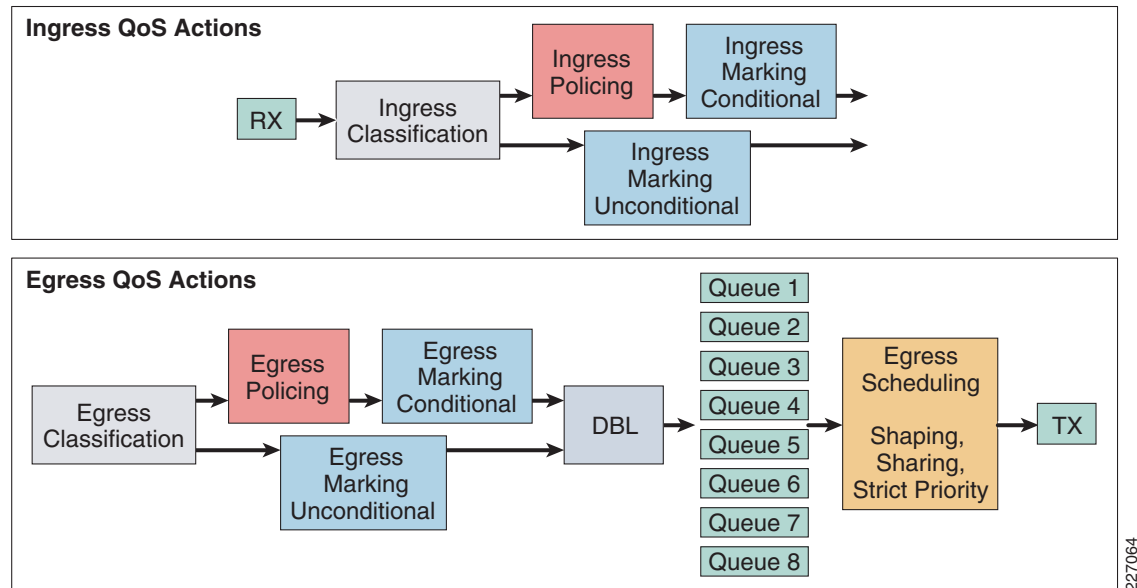
Figure 2-20 Catalyst 4500 (Classic Supervisor) Switch QoS Model

Traffic is classified on ingress, based on trust states, access lists, or class maps. Policers can be applied to flows, on either a per-port basis or a per-VLAN basis or even a per-port/per-VLAN basis. Similarly, marking policies can be applied on the same basis. Egress queuing is based on a 4Q1T or a 1P3Q1T model (the latter being preferred, as it supports the EF PHB), with a platform-specific proprietary congestion avoidance mechanism providing Active Queue Management (AQM), namely Dynamic Buffer Limiting (DBL). DBL tracks the queue length for each traffic flow in the switch. When the queue length of a flow exceeds its limit, DBL drop packets or sets the Explicit Congestion Notification (ECN) bits in the packet headers.

Catalyst 4500 Classic Supervisor syntax is essentially equivalent to the QoS syntax on other Catalyst platforms, with the exception that the **mls** command prefix is omitted on this platform. Thus **mls qos** is abbreviated to simply **qos** on the Catalyst 4500 Classic Supervisors.

In contrast, the Catalyst 4500 Series Switch using Supervisor Engine 6-E/7-E employs the MQC QoS model. In this model, QoS is applied via Modular QoS Command-Line Interface (MQC). As such, certain QoS features are implemented differently (and others are not supported; the following sections detail the differences). The Catalyst 4500 Supervisor 6-E/7-E MQC QoS model is shown in Figure 2-21.

Figure 2-21 Catalyst 4500-E (Supervisor 6-E) MQC Model



In the MQC packet, QoS policies are applied as follows:

- Step 1** The incoming packet is classified (based on different packet fields, receive port, or VLAN) to belong to a traffic class.
- Step 2** Depending on the traffic class, the packet is rate limited/policed and its priority is optionally marked (typically at the edge of the network) so that lower priority packets are dropped or marked with lower priority in the packet fields (DSCP and CoS).
- Step 3** After the packet has been marked, it is looked up for forwarding. This action obtains the transmit port and VLAN to transmit the packet.
- Step 4** The packet is classified in the output direction based on the transmit port or VLAN. The classification takes into account any marking of the packet by input QoS.
- Step 5** Depending on the output classification, the packet is policed, its priority is optionally (re-)marked, and the transmit queue for the packet is determined depending on the traffic class.
- Step 6** The transmit queue state is dynamically monitored via DBL and drop threshold configuration to determine whether the packet should be dropped or enqueued for transmission.
- Step 7** If eligible for transmission, the packet is enqueued to a transmit queue. The transmit queue is selected based on output QoS classification criteria. The selected queue provides the desired behavior in terms of latency and bandwidth.

**Note**

As the QoS feature set and syntax for the Supervisor 6-E and 7-E are identical, these are collectively referred to as the Supervisor 6-E for the remainder of this design chapter.

Enabling QoS (Classic Supervisors)

**Note**

QoS does not have to be explicitly enabled within the Catalyst 4500-E Supervisor 6-E MQC model, as it is enabled by default.

QoS must be enabled globally on the Catalyst 4500 Classic Supervisors. This is a critical first step to deploying QoS on these platforms. If this small, but important, step is overlooked, it can lead to frustration in troubleshooting QoS problems because the switch software accepts QoS commands and even displays these within the switch configuration, but none of the QoS commands are active until the **qos** global command is enabled, as shown in [Example 2-62](#).

Example 2-62 Enabling QoS on a Catalyst 4500 Classic Supervisor

```
C4500-CS(config)#qos
C4500-CS(config)#
```

This configuration can be verified with the command:

- **show qos** (as shown in [Example 2-63](#))

Example 2-63 Verifying Global QoS on a Catalyst 4500 Classic Supervisor—show qos

```
C4500-CS#show qos
QoS is enabled globally
IP header DSCP rewrite is enabled

C4500-CS#
```

Trust Models

Catalyst 4500 Classic Supervisor switch ports can be configured to statically trust CoS, DSCP, or to dynamically and conditionally trust Cisco IP phones. By default, with QoS enabled, all ports are set to an untrusted state.

In contrast, the Catalyst 4500-E Supervisor 6-E does not support trust CoS, as it considers all interfaces to be trusted (via DSCP-trust) by default; it does, however, support conditional trust.

Trust-CoS Model

A Catalyst 4500 Classic Supervisor switch port can be configured to trust CoS by configuring the interface with the **qos trust cos** command. However, if an interface is set to trust CoS, then it by default calculates a packet's internal DSCP to be the incoming packet's (CoS value * 8). While this may be suitable for most markings, this default mapping may not be suitable for VoIP, as VoIP is usually marked CoS 5,

which would map by default to DSCP 40 (and not 46, which is the EF PHB as defined by RFC 3246). Therefore, if an interface is set to trust CoS, then the default CoS-to-DSCP mapping table should be modified such that CoS 5 maps to DSCP 46, as shown in [Example 2-64](#).

**Note**

As previously mentioned, the Catalyst 4500-E Supervisor 6-E does not support CoS-trust, but considers all interfaces to be trusted—via DSCP-trust—by default.

Example 2-64 Configuring Trust CoS and CoS-to-DSCP Mapping Modification on a Catalyst 4500 Classic Supervisor

```
C4500-CS(config)#qos map cos 5 to dscp 46
! CoS 5 is mapped to DSCP 46 (EF)
C4500-CS(config)#interface GigabitEthernet 1/1
C4500-CS(config-if)#qos trust cos
! The interface is set to statically trust CoS
```

This configuration can be verified with the commands:

- **show qos map cos-dscp** (as shown in [Example 2-65](#))
- **show qos interface** (as shown in [Example 2-66](#))

Example 2-65 Verifying Global CoS-to-DSCP Mapping Modifications on a Catalyst 4500 Classic Supervisor—show qos map cos-dscp

```
C4500-CS#show qos map cos-dscp
CoS-DSCP Mapping Table
  CoS:   0  1  2  3  4  5  6  7
-----
  DSCP:   0  8 16 24 32 46 48 56

C4500-CS#
```

In [Example 2-65](#), the CoS-to-DSCP mapping value for CoS 5 has been modified from the default mapping of 40 (CoS 5 * 8) to 46 (to match the recommendation from RFC 3246 that realtime applications be marked DSCP 46/EF).

Example 2-66 Verifying Interface Trust Settings on a Catalyst 4500—show qos interface

```
C4500-CS#show qos interface GigabitEthernet 1/1
QoS is enabled globally
Port QoS is enabled
Administrative Port Trust State: 'cos'
Operational Port Trust State: 'cos'
Trust device: none
Default DSCP: 0 Default CoS: 0
Appliance trust: none
Tx-Queue   Bandwidth   ShapeRate   Priority   QueueSize
          (bps)         (bps)
1         250000000   disabled    N/A        1920
2         250000000   disabled    N/A        1920
3         250000000   disabled    normal     1920
4         250000000   disabled    N/A        1920

C4500-CS#
```

In [Example 2-66](#), the administrative port trust state is set to trust-cos and the current operation port trust state is also at trust-cos.

Trust-DSCP Model

Because of the additional granularity of DSCP versus QoS markings, it is generally recommended to trust DSCP rather than CoS (everything else being equal). A Catalyst 4500 Classic Supervisor switch port can be configured to trust DSCP with the **qos trust dscp** interface command, as shown in [Example 2-67](#).



Note

As previously mentioned, Catalyst 4500-E Supervisor 6-E considers all interfaces to be trusted—via DSCP-trust—by default

Example 2-67 Configuring Trust-DSCP on a Catalyst 4500 Classic Supervisor

```
C4500-CS(config)#interface GigabitEthernet 1/1
C4500-CS(config-if)#qos trust dscp
! The interface is set to statically trust DSCP
```

This configuration can be verified with the command:

- **show qos interface**

Conditional Trust Model

In addition to configuring switch ports to statically trust endpoints, the Catalyst 4500 Classic Supervisor and the Catalyst 4500-E Supervisor 6-E support dynamic, conditional trust with the **qos trust device** interface command, which can be configured with the **cisco-phone** keyword to extend trust to Cisco IP phones, after these have been verified via a CDP-negotiation. Additionally, the type of trust to be extended can be specified (either CoS or DSCP) on the Catalyst 4500 Classic Supervisor. When configuring conditional trust to Cisco IP Phones, it is recommended to dynamically extend CoS-Trust, as Cisco IP Phones can only remark PC QoS markings at Layer 2 (CoS) and not at Layer 3 (DSCP). For other endpoints that do not have this remarking limitation, it is recommended to dynamically extend DSCP-trust (over CoS-trust), because DSCP has greater marking granularity. An example of a dynamic, conditional trust policy that is set to extend CoS-trust to CDP-verified Cisco IP phones is shown in [Example 2-68](#).

Example 2-68 Configuring Conditional (CoS-mode) Trust on a Catalyst 4500 Classic Supervisor

```
C4500-CS(config)#qos map cos 5 to dscp 46
! CoS 5 is mapped to DSCP 46 (EF)
C4500-CS(config)#interface GigabitEthernet 1/1
C4500-CS(config-if)# switchport access vlan 10
C4500-CS(config-if)# switchport voice vlan 110
C4500-CS(config-if)# spanning-tree portfast
C4500-CS(config-if)# qos trust device cisco-phone
! The interface is set to conditionally-trust Cisco IP Phones
C4500-CS(config-if)# qos trust cos
! CoS-trust will be dynamically extended to Cisco IP Phones
```

This configuration can be verified with the command:

- **show qos interface** (as shown in [Example 2-69](#))

Example 2-69 Verifying Interface Trust Settings on a Catalyst 4500—show qos interface

```
C4500-CS#show qos interface GigabitEthernet 1/1
QoS is enabled globally
```

```

Port QoS is enabled
Administrative Port Trust State: 'cos'
Operational Port Trust State: 'cos'
Trust device: cisco-phone
Default DSCP: 0 Default CoS: 0
Appliance trust: none
Tx-Queue   Bandwidth   ShapeRate   Priority   QueueSize
          (bps)        (bps)
1          250000000    disabled    N/A        1920
2          250000000    disabled    N/A        1920
3          250000000    disabled    normal     1920
4          250000000    disabled    N/A        1920

```

C4500-CS#

In [Example 2-69](#), the trust device feature has been enabled, with the trusted device being specified as a **cisco-phone**. The administrative port trust state—that is, the mode of trust (CoS or DSCP) that is extended dynamically to the IP Phone—is set to trust CoS. Similarly, the current (dynamic) operational port trust state is shown as trusting CoS. This is because there is a Cisco IP phone currently connected to the switch port; if the IP phone is removed from this switch port, the operational port trust state toggles to “untrusted”.

Because the Catalyst 4500 Sup6-E and Sup7-E do not support trust CoS and the fact that Cisco IP phones can only remark CoS bits on PC-generated traffic, a workaround policy is required for switch ports on these Supervisors that may connect to Cisco IP phones. This workaround is a dynamic conditional trust policy applied to the port in conjunction with a simple MQC policy that explicitly matches CoS 5 (for voice) and CoS 3 (for signaling) and marks the DSCP values of these packets to EF and CS3, respectively (essentially performing a CoS-to-DSCP mapping). This workaround policy is shown in [Example 2-70](#).

Example 2-70 Configuring Conditional (CoS-mode) Trust on a Catalyst 4500 Sup-6E

```

! This section defines the class-maps to match Voice and Signaling
C4500-E(config-cmap)# class-map match-all VOICE
C4500-E(config-cmap)# match cos 5
C4500-E(config-cmap)# class-map match-all SIGNALING
C4500-E(config-cmap)# match cos 3

! This section defines the CoS-to-DSCP remarking policy-map
C4500-E(config-cmap)# policy-map CISCO-IPPHONE
C4500-E(config-pmap)# class VOICE
C4500-E(config-pmap-c)# set dscp ef
! Maps CoS 5 to DSCP EF
C4500-E(config-pmap-c)# class SIGNALING
C4500-E(config-pmap-c)# set dscp cs3
! Maps CoS 3 to DSCP CS3
C4500-E(config-pmap-c)# class class-default
C4500-E(config-pmap-c)# set dscp default
! All other traffic is set to DSCP DF

! This section applies conditional trust and the policy-map to the int(s)
C4500-E(config-pmap-c)# interface GigabitEthernet 3/1
C4500-E(config-if)# switchport access vlan 10
C4500-E(config-if)# switchport voice vlan 110
C4500-E(config-if)# spanning-tree portfast
C4500-E(config-if)# qos trust device cisco-phone
! Applies conditional-trust to the switch port
C4500-E(config-if)# service-policy input CISCO-IPPHONE
! Attaches the CoS-to-DSCP mapping policy-map

```

This configuration can be verified with the commands:

- **show qos interface**
- **show class-map**
- **show policy-map**

Marking Models

The Catalyst 4500 family of switches supports two main marking models:

- Per-port marking model
- Per-VLAN marking model

Each model is detailed in the following sections.

Per-Port Marking Model

The per-port marking model (based on [Figure 2-10](#)) matches VoIP and signaling traffic from the VVLAN by matching on DSCP EF and CS3, respectively. Multimedia conferencing traffic from the DVLAN is matched by UDP/RTP ports 16384-32767. Signaling traffic is matched on SCCP ports (TCP 2000-2002), as well as on SIP ports (TCP/UDP 5060-5061). Other transactional data traffic, bulk data, and scavenger traffic are matched on various ports (outlined in [Figure 2-9](#)). Unlike the Catalyst 3750-E examples, no explicit default class is required, as the implicit class default performs policy actions (such as marking or policing) on the Catalyst 4500. The service policy is applied to an interface range, along with (DSCP-mode) conditional trust, as shown in [Example 2-71](#).

Example 2-71 Per-Port Marking Configuration Example on a Catalyst 4500 Classic Supervisor or a Catalyst 4500-E Supervisor 6-E

```
! This first section configures IP access-lists to match applications
C4500-CS(config)#ip access-list extended MULTIMEDIA-CONFERENCING
C4500-CS(config-ext-nacl)# remark RTP
C4500-CS(config-ext-nacl)# permit udp any any range 16384 32767

C4500-CS(config)#ip access-list extended SIGNALING
C4500-CS(config-ext-nacl)# remark SCCP
C4500-CS(config-ext-nacl)# permit tcp any any range 2000 2002
C4500-CS(config-ext-nacl)# remark SIP
C4500-CS(config-ext-nacl)# permit tcp any any range 5060 5061
C4500-CS(config-ext-nacl)# permit udp any any range 5060 5061

C4500-CS(config)#ip access-list extended TRANSACTIONAL-DATA
C4500-CS(config-ext-nacl)# remark HTTPS
C4500-CS(config-ext-nacl)# permit tcp any any eq 443
C4500-CS(config-ext-nacl)# remark ORACLE-SQL*NET
C4500-CS(config-ext-nacl)# permit tcp any any eq 1521
C4500-CS(config-ext-nacl)# permit udp any any eq 1521
C4500-CS(config-ext-nacl)# remark ORACLE
C4500-CS(config-ext-nacl)# permit tcp any any eq 1526
C4500-CS(config-ext-nacl)# permit udp any any eq 1526
C4500-CS(config-ext-nacl)# permit tcp any any eq 1575
C4500-CS(config-ext-nacl)# permit udp any any eq 1575
C4500-CS(config-ext-nacl)# permit tcp any any eq 1630
C4500-CS(config-ext-nacl)# permit udp any any eq 1526

C4500-CS(config)#ip access-list extended BULK-DATA
```

```

C4500-CS(config-ext-nacl)# remark FTP
C4500-CS(config-ext-nacl)# permit tcp any any eq ftp
C4500-CS(config-ext-nacl)# permit tcp any any eq ftp-data
C4500-CS(config-ext-nacl)# remark SSH/SFTP
C4500-CS(config-ext-nacl)# permit tcp any any eq 22
C4500-CS(config-ext-nacl)# remark SMTP/SECURE SMTP
C4500-CS(config-ext-nacl)# permit tcp any any eq smtp
C4500-CS(config-ext-nacl)# permit tcp any any eq 465
C4500-CS(config-ext-nacl)# remark IMAP/SECURE IMAP
C4500-CS(config-ext-nacl)# permit tcp any any eq 143
C4500-CS(config-ext-nacl)# permit tcp any any eq 993
C4500-CS(config-ext-nacl)# remark POP3/SECURE POP3
C4500-CS(config-ext-nacl)# permit tcp any any eq pop3
C4500-CS(config-ext-nacl)# permit tcp any any eq 995
C4500-CS(config-ext-nacl)# remark CONNECTED PC BACKUP
C4500-CS(config-ext-nacl)# permit tcp any eq 1914 any

C4500-CS(config)#ip access-list extended SCAVENGER
C4500-CS(config-ext-nacl)# remark KAZAA
C4500-CS(config-ext-nacl)# permit tcp any any eq 1214
C4500-CS(config-ext-nacl)# permit udp any any eq 1214
C4500-CS(config-ext-nacl)# remark MICROSOFT DIRECT X GAMING
C4500-CS(config-ext-nacl)# permit tcp any any range 2300 2400
C4500-CS(config-ext-nacl)# permit udp any any range 2300 2400
C4500-CS(config-ext-nacl)# remark APPLE ITUNES MUSIC SHARING
C4500-CS(config-ext-nacl)# permit tcp any any eq 3689
C4500-CS(config-ext-nacl)# permit udp any any eq 3689
C4500-CS(config-ext-nacl)# remark BITTORRENT
C4500-CS(config-ext-nacl)# permit tcp any any range 6881 6999
C4500-CS(config-ext-nacl)# remark YAHOO GAMES
C4500-CS(config-ext-nacl)# permit tcp any any eq 11999
C4500-CS(config-ext-nacl)# remark MSN GAMING ZONE
C4500-CS(config-ext-nacl)# permit tcp any any range 28800 29100

! This section configures the class-maps
C4500-CS(config-cmap)#class-map match-all VVLAN-VOIP
C4500-CS(config-cmap)# match ip dscp ef
! VoIP is trusted (from the VVLAN)

C4500-CS(config-cmap)#class-map match-all VVLAN-SIGNALING
C4500-CS(config-cmap)# match ip dscp cs3
! Signaling is trusted (from the VVLAN)

C4500-CS(config-cmap)#class-map match-all MULTIMEDIA-CONFERENCING
C4500-CS(config-cmap)# match access-group name MULTIMEDIA-CONFERENCING
! Associates MULTIMEDIA-CONFERENCING access-list with class-map

C4500-CS(config-cmap)#class-map match-all SIGNALING
C4500-CS(config-cmap)# match access-group name SIGNALING
! Associates SIGNALING access-list with class-map

C4500-CS(config-cmap)#class-map match-all TRANSACTIONAL-DATA
C4500-CS(config-cmap)# match access-group name TRANSACTIONAL-DATA
! Associates TRANSACTIONAL-DATA access-list with class-map

C4500-CS(config-cmap)#class-map match-all BULK-DATA
C4500-CS(config-cmap)# match access-group name BULK-DATA
! Associates BULK-DATA access-list with class-map

C4500-CS(config-cmap)#class-map match-all SCAVENGER
C4500-CS(config-cmap)# match access-group name SCAVENGER
! Associates SCAVENGER access-list with class-map

! This section configures the Per-Port ingress marking policy-map

```

```

C4500-CS(config)#policy-map PER-PORT-MARKING
C4500-CS(config-pmap)# class VVLAN-VOIP
C4500-CS(config-pmap-c)# set dscp ef
    ! VoIP is marked EF
C4500-CS(config-pmap-c)# class VVLAN-SIGNALING
C4500-CS(config-pmap-c)# set dscp cs3
    ! Signaling (from the VVLAN) is marked CS3
C4500-CS(config-pmap-c)# class MULTIMEDIA-CONFERENCING
C4500-CS(config-pmap-c)# set dscp af41
    ! Multimedia-conferencing is marked AF41
C4500-CS(config-pmap-c)# class SIGNALING
C4500-CS(config-pmap-c)# set dscp cs3
    ! Signaling (from the DVLAN) is marked CS3
C4500-CS(config-pmap-c)# class TRANSACTIONAL-DATA
C4500-CS(config-pmap-c)# set dscp af21
    ! Transactional Data is marked AF21
C4500-CS(config-pmap-c)# class BULK-DATA
C4500-CS(config-pmap-c)# set dscp af11
    ! Bulk Data is marked AF11
C4500-CS(config-pmap-c)# class SCAVENGER
C4500-CS(config-pmap-c)# set dscp cs1
    ! Scavenger traffic is marked CS1
C4500-CS(config-pmap-c)# class class-default
C4500-CS(config-pmap-c)# set dscp default
    ! An implicit class-default marks all other traffic to DF

    ! This section attaches the service-policy to the interface(s)
C4500-CS(config)#interface range GigabitEthernet 2/1-48
C4500-CS(config-if-range)# switchport access vlan 10
C4500-CS(config-if-range)# switchport voice vlan 110
C4500-CS(config-if-range)# spanning-tree portfast
C4500-CS(config-if-range)# qos trust device cisco-phone
    ! The interface is set to conditionally-trust Cisco IP Phones
C4500-CS(config-if-range)# qos trust cos
    ! CoS-trust will be dynamically extended to Cisco IP Phones
C4500-CS(config-if-range)# service-policy input INGRESS-MARKING
    ! Attaches the Per-Port Marking policy to the interface(s)

```

**Note**

On the Catalyst 4500 Classic Supervisors, marking commands on an interface cannot be enabled until IP routing is enabled globally. If IP routing is disabled globally and you try to configure the service policy on an interface, the configuration is accepted but it does not take effect. You are prompted with the message: “Set command will not take effect since CEF is disabled. Please enable IP routing and CEF globally.” To enable IP routing globally, issue the **ip routing** and **ip cef** global configuration commands. After you do this, the marking commands take effect.

**Note**

The second interface trust command (**qos trust cos**) is not supported on the Catalyst 4500-E Supervisor 6-E.

This configuration can be verified with the commands:

- **show qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface** (as shown in [Example 2-72](#))

Example 2-72 Verifying Service Policies on a Catalyst 4500—show policy-map interface

```

C4500-CS#show policy-map interface GigabitEthernet 2/1
GigabitEthernet2/1

Service-policy input: PER-PORT-MARKING

Class-map: VVLAN-VOIP (match-all)
  65211 packets
  Match: ip dscp ef (46)
  QoS Set
    ip dscp ef

Class-map: VVLAN-SIGNALING (match-all)
  731 packets
  Match: ip dscp cs3 (24)
  QoS Set
    ip dscp cs3

Class-map: MULTIMEDIA-CONFERENCING (match-all)
  16290 packets
  Match: access-group name MULTIMEDIA-CONFERENCING
  QoS Set
    ip dscp af41

Class-map: SIGNALING (match-all)
  130 packets
  Match: access-group name SIGNALING
  QoS Set
    ip dscp cs3

Class-map: TRANSACTIONAL-DATA (match-all)
  13211 packets
  Match: access-group name TRANSACTIONAL-DATA
  QoS Set
    ip dscp af21

Class-map: BULK-DATA (match-all)
  16518 packets
  Match: access-group name BULK-DATA
  QoS Set
    ip dscp af11

Class-map: SCAVENGER (match-all)
  14238 packets
  Match: access-group name SCAVENGER
  QoS Set
    ip dscp cs1

Class-map: class-default (match-any)
  25881 packets
  Match: any
    25881 packets
  QoS Set
    ip dscp default
C4500-CS#

```

Example 2-72 shows that the **show policy-map interface** command on the Catalyst 4500 Classic Supervisors dynamically increments counters. However, it should be noted that these are slightly delayed and seem to increment only every 10-15 seconds.

Per-VLAN Marking Model (Classic Supervisors)

An alternative approach for deploying marking policies on the Catalyst 4500 Classic Supervisor platforms is to deploy these on a per-VLAN basis. In order to do so, the interfaces belonging to the VLANs need to be configured with the **qos vlan-based** interface command. Additionally, the policy map can be simplified/broken-apart, as applicable to each VLAN. Adapting the per-port marking example to a VLAN-based marking policing allows for the VVLAN-based policy map to be reduced to only two explicit classes, VoIP and signaling. Similarly, the DVLAN-based policy map is reduced to five explicit classes, multimedia conferencing, signaling, transactional data, bulk data, and scavenger. Both policy maps still retain an implicit default class for all other flows. A per-VLAN marking model is shown in [Example 2-73](#).



Note

As the access lists and class maps are identical to the previous example, these are omitted for brevity in this—and in following—examples for this switch platform family.

Example 2-73 Per-VLAN Marking Configuration Example on a Catalyst 4500 Classic Supervisor

```

! This section configures the ingress marking policy-map for the VVLAN
C4500-CS(config)#policy-map VVLAN-MARKING
C4500-CS(config-pmap)# class VVLAN-VOIP
C4500-CS(config-pmap-c)# set dscp ef
! VoIP is trusted (from the VVLAN)
C4500-CS(config-pmap-c)# class VVLAN-SIGNALING
C4500-CS(config-pmap-c)# set dscp cs3
! Signaling is trusted (from the VVLAN)
C4500-CS(config-pmap-c)# class class-default
C4500-CS(config-pmap-c)# set dscp default
! An implicit class-default marks all other VVLAN traffic to DF

! This section configures the ingress marking policy-map for the DVLAN
C4500-CS(config)#policy-map DVLAN-MARKING
C4500-CS(config-pmap)# class MULTIMEDIA-CONFERENCING
C4500-CS(config-pmap-c)# set dscp af41
! Multimedia-conferencing is marked AF41
C4500-CS(config-pmap-c)# class SIGNALING
C4500-CS(config-pmap-c)# set dscp cs3
! Signaling (from the DVLAN) is marked CS3
C4500-CS(config-pmap-c)# class TRANSACTIONAL-DATA
C4500-CS(config-pmap-c)# set dscp af21
! Transactional Data is marked AF21
C4500-CS(config-pmap-c)# class BULK-DATA
C4500-CS(config-pmap-c)# set dscp af11
! Bulk Data is marked AF11
C4500-CS(config-pmap-c)# class SCAVENGER
C4500-CS(config-pmap-c)# set dscp cs1
! Scavenger traffic is marked CS1
C4500-CS(config-pmap-c)# class class-default
C4500-CS(config-pmap-c)# set dscp default
! An implicit class-default marks all other DVLAN traffic to DF

! This section configures the interface(s) for conditional trust,
! with CoS-trust and enables VLAN-based QoS
C4500-CS(config)#interface range GigabitEthernet 2/1-48
C4500-CS(config-if-range)# switchport access vlan 10
C4500-CS(config-if-range)# switchport voice vlan 110
C4500-CS(config-if-range)# spanning-tree portfast
C4500-CS(config-if-range)# qos trust device cisco-phone

```

```

! The interface is set to conditionally-trust Cisco IP Phones
C4500-CS(config-if-range)# qos trust cos
! CoS-trust will be dynamically extended to Cisco IP Phones
C4500-CS(config-if-range)# qos vlan-based
! Enables VLAN-based QoS on the interface(s)

! This section attaches the service-policy to the DVLAN interface
C4500-CS(config)#interface Vlan 10
C4500-CS(config-if)# description DVLAN
C4500-CS(config-if)# ip route-cache cef
! Enables IP CEF on the VLAN interface (required for marking)
C4500-CS(config-if)# service-policy input DVLAN-MARKING
! Attaches the DVLAN Per-VLAN Marking policy to the DVLAN interface

! This section attaches the service-policy to the VVLAN interface
C4500-CS(config)#interface Vlan 110
C4500-CS(config-if)# description VVLAN
C4500-CS(config-if)# ip route-cache cef
! Enables IP CEF on the VLAN interface (required for marking)
C4500-CS(config-if)# service-policy input VVLAN-MARKING
! Attaches the VVLAN Per-VLAN Marking policy to the VVLAN interface

```

This configuration can be verified with the commands:

- **show qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Per-VLAN Marking Model (Supervisor 6-E/7-E)

The per-VLAN marking model is essentially the same for the Catalyst 4500-E Supervisor 6-E, except for the final set of interface and VLAN commands. The Catalyst 4500-E does not support the **qos vlan-based** interface command, neither is the **qos trust dscp** interface command required, and finally, service policies are attached to VLANs via a VLAN-configuration mode (instead of an interface configuration mode), as shown in [Example 2-74](#).

Example 2-74 Per-VLAN Marking Configuration Example on a Catalyst 4500-E Supervisor 6-E

```

! This section configures the ingress marking policy-map for the VVLAN
C4500-E(config)#policy-map VVLAN-MARKING
C4500-E(config-pmap)# class VVLAN-VOIP
C4500-E(config-pmap-c)# set dscp ef
! VoIP is trusted (from the VVLAN)
C4500-E(config-pmap-c)# class VVLAN-SIGNALING
C4500-E(config-pmap-c)# set dscp cs3
! Signaling is trusted (from the VVLAN)
C4500-E(config-pmap-c)# class class-default
C4500-E(config-pmap-c)# set dscp default
! An implicit class-default marks all other VVLAN traffic to DF

! This section configures the ingress marking policy-map for the DVLAN
C4500-E(config)#policy-map DVLAN-MARKING
C4500-E(config-pmap)# class MULTIMEDIA-CONFERENCING
C4500-E(config-pmap-c)# set dscp af41
! Multimedia-conferencing is marked AF41

```

```

C4500-E(config-pmap-c)# class SIGNALING
C4500-E(config-pmap-c)# set dscp cs3
    ! Signaling (from the DVLAN) is marked CS3
C4500-E(config-pmap-c)# class TRANSACTIONAL-DATA
C4500-E(config-pmap-c)# set dscp af21
    ! Transactional Data is marked AF21
C4500-E(config-pmap-c)# class BULK-DATA
C4500-E(config-pmap-c)# set dscp af11
    ! Bulk Data is marked AF11
C4500-E(config-pmap-c)# class SCAVENGER
C4500-E(config-pmap-c)# set dscp cs1
    ! Scavenger traffic is marked CS1
C4500-E(config-pmap-c)# class class-default
C4500-E(config-pmap-c)# set dscp default
    ! An implicit class-default marks all other DVLAN traffic to DF

    ! This section configures the interface(s) for conditional trust,
C4500-E(config)#interface range GigabitEthernet 2/1-48
C4500-E(config-if-range)# switchport access vlan 10
C4500-E(config-if-range)# switchport voice vlan 110
C4500-E(config-if-range)# spanning-tree portfast
C4500-E(config-if-range)# qos trust device cisco-phone

    ! This section attaches the marking policy to the DVLAN
C4500-E(config)#vlan config 10
C4500-E(config-vlan-config)# service-policy input DVLAN-MARKING

    ! This section attaches the marking policy to the VVLAN
C4500-E(config)#vlan config 110
C4500-E(config-vlan-config)# service-policy input VVLAN-MARKING

```

This configuration can be verified with the commands:

- **show qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map vlan *vlan-number*** (this command is virtually identical to **show policy-map interface**, except that it references a VLAN directly, rather than a VLAN interface)

Policing Models

The Catalyst 4500 Classic Supervisors support these policing options:

- Only single rate (two color marker) policers are supported
- Support for 1024 policers on ingress and on egress
 - The Supervisor Engine V-10GE supports 8192 policers on ingress and on egress
 - Additionally the Supervisor V-10GE supports 512 flow-based policers (on ingress Layer 3 interfaces only) which can police over 100,000 microflows
- Policing accuracy of +/- 1.5% of configured rate

In contrast, the Catalyst 4500 Supervisor 6-E supports:

- Single rate policer (two color marker)
 - 16,000 single rate policers are supported
- Single rate three color marker (srTCM) (RFC 2697)
 - 8000 single rate three color markers are supported
- Two rate three color marker (trTCM) (RFC 2698)
 - 8000 two rate three color markers are supported
- Policing accuracy of 0.75% of configured policer rate.

The Catalyst 4500 family of switches supports these ingress policing models:

- Per-port policing model—This model attaches policers to physical switch port interfaces.
- Per-VLAN policing model—This model attaches policers to logical VLAN interfaces; however, there is an inherent limitation with this policing model. It only supports a single aggregate policer per VLAN and—since the number of ports associated with a VLAN is dynamic and variable—thus is quite restricted in overall policing effectiveness. Therefore, it is generally recommended to use the per-port/per-VLAN policing model instead, as it offers more discrete policing options.
- Per-port/per-VLAN policing model—This model attaches policers to discrete VLANs traversing a single switch trunk interface.
- User-Based Rate Limiting (UBRL) models—This model (supported on the Supervisor V-10GE only) applies flow-based policers to Layer 3 interfaces to police microflows on a per-source or per-destination basis; UBRL may be applied on a per-port or per-port/per-VLAN basis.

The per-port and per-port/per-VLAN policing models and the UBRL models for the Catalyst 4500 family of switches are detailed in the following sections.

Per-Port Policing Model (Classic Supervisors)

The per-port policing model is quite similar to the per-port marking model, except that the policy action includes a policing function—in some cases to drop, in others to remark. As shown in [Figure 2-10](#), the VoIP and signaling traffic from the VVLAN can be policed to drop at 128 kbps and 32 kbps, respectively (as any excessive traffic matching this criteria would be indicative of network abuse). Similarly, the multimedia conferencing, signaling, and scavenger traffic from the DVLAN can be policed to drop. On the other hand, data plane policing policies can be applied to transactional, bulk, and best effort data traffic, such that these flows are subject to being remarked (but not dropped at the ingress edge) when severely out-of-profile. Remarketing is performed by configuring a policed-DSCP map with the global configuration command **qos map dscp policed**, which specifies which DSCP values are subject to remarketing if out-of-profile and what value these should be remarked as (which in the case of data plane policing/scavenger class QoS policies, this value is CS1/DSCP 8). A per-port policing for the Catalyst 4500 Classic Supervisor is shown in [Example 2-75](#).

Example 2-75 Per-Port Policing Configuration Example on a Catalyst 4500 Classic Supervisor

```
! This section configures the global policed-DSCP markdown map
C4500-CS(config)#qos map dscp policed 0 10 18 to dscp 8
! DSCP 0 (DF), 10 (AF11) and 18 (AF21) are marked down to 8 (CS1)
! if found to be in excess of their (respective) policing rates

! This section configures the Per-Port policing policy-map
C4500-CS(config)#policy-map PER-PORT-POLICING
C4500-CS(config-pmap)# class VVLAN-VOIP
C4500-CS(config-pmap-c)# set dscp ef
```

```

C4500-CS(config-pmap-c)# police 128k 8000 exceed-action drop
! VoIP is marked EF and policed to drop at 128 kbps
C4500-CS(config-pmap-c)# class VVLAN-SIGNALING
C4500-CS(config-pmap-c)# set dscp cs3
C4500-CS(config-pmap-c)# police 32k 8000 exceed-action drop
! (VVLAN) Signaling is marked CS3 and policed to drop at 32 kbps
C4500-CS(config-pmap-c)# class MULTIMEDIA-CONFERENCING
C4500-CS(config-pmap-c)# set dscp af41
C4500-CS(config-pmap-c)# police 5m 8000 exceed-action drop
! Multimedia-conferencing is marked AF41 and policed to drop at 5 Mbps
C4500-CS(config-pmap-c)# class SIGNALING
C4500-CS(config-pmap-c)# set dscp cs3
C4500-CS(config-pmap-c)# police 32k 8000 exceed-action drop
! (DVLAN) Signaling is marked CS3 and policed to drop at 32 kbps
C4500-CS(config-pmap-c)# class TRANSACTIONAL-DATA
C4500-CS(config-pmap-c)# set dscp af21
C4500-CS(config-pmap-c)# police 10m 8000 exceed-action policed-dscp-transmit
! Trans-data is marked AF21 and policed to remark (to CS1) at 10 Mbps
C4500-CS(config-pmap-c)# class BULK-DATA
C4500-CS(config-pmap-c)# set dscp af11
C4500-CS(config-pmap-c)# police 10m 8000 exceed-action policed-dscp-transmit
! Bulk-data is marked AF11 and policed to remark (to CS1) at 10 Mbps
C4500-CS(config-pmap-c)# class SCAVENGER
C4500-CS(config-pmap-c)# set dscp cs1
C4500-CS(config-pmap-c)# police 10m 8000 exceed-action drop
! Scavenger traffic is marked CS1 and policed to drop at 10 Mbps
C4500-CS(config-pmap-c)# class class-default
C4500-CS(config-pmap-c)# set dscp default
C4500-CS(config-pmap-c)# police 10m 8000 exceed-action policed-dscp-transmit
! The implicit default class marks all other traffic to DF
! and polices all other traffic to remark (to CS1) at 10 Mbps

! This section attaches the service-policy to the interface(s)
C4500-CS(config)#interface range GigabitEthernet 2/1-48
C4500-CS(config-if-range)# switchport access vlan 10
C4500-CS(config-if-range)# switchport voice vlan 110
C4500-CS(config-if-range)# spanning-tree portfast
C4500-CS(config-if-range)# qos trust device cisco-phone
! The interface is set to conditionally-trust Cisco IP Phones
C4500-CS(config-if-range)# qos trust cos
! CoS-trust will be dynamically extended to Cisco IP Phones
C4500-CS(config-if-range)# service-policy input PER-PORT-POLICING
! Attaches the Per-Port Policing policy to the interface(s)

```

**Note**

Catalyst 4500 software allows for policing rates to be entered using the postfixes **k** (for kilobits), **m** (for megabits), and **g** (for gigabits), as shown in [Example 2-16](#). Additionally, decimal points are allowed in conjunction with these postfixes; for example, a rate of 10.5 Mbps could be entered with the policy map command **police 10.5m**. While these policing rates are converted to their full bps values within the configuration, it makes the entering of these rate more user-friendly and less error prone (as could easily be the case when having to enter up to 10 zeros to define the policing rate).

This configuration can be verified with the commands:

- **show qos maps dscp policed** (as shown in [Example 2-76](#))
- **show qos interface**
- **show class-map**
- **show policy-map**

- show policy-map interface

Example 2-76 Verifying Global Policing Markdown Mappings on a Catalyst 4500 Classic Supervisor—show qos maps dscp policed

```
C4500-CS#show qos maps dscp policed
Policed DSCP Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    08 01 02 03 04 05 06 07 08 09
1 :    08 11 12 13 14 15 16 17 08 19
2 :    20 21 22 23 24 25 26 27 28 29
3 :    30 31 32 33 34 35 36 37 38 39
4 :    40 41 42 43 44 45 46 47 48 49
5 :    50 51 52 53 54 55 56 57 58 59
6 :    60 61 62 63
```

C4500-CS#

In [Example 2-76](#), the policing DSCP-markdown mapping is shown. The first digit of the DSCP value of a packet offered to a policer is shown along the Y-axis of the table; the second digit of the DSCP value of a packet offered to a policer is shown along the X-axis of the table. For example, the DSCP value for the transactional data application class (AF21/18) is found in the row d1=1 and column d2=8. And, as shown, packets with this offered DSCP value (along with DF/0 and AF11/10) are remarked to CS1 (08) if found to be in excess of the policing rate.

Per-Port Policing Model (Supervisor 6-E)

The per-port policing model is essentially the same for the Catalyst 4500-E Supervisor 6-E, except that it does not require a global policed-DSCP map and thus the policing commands are slightly different, also no trust-DSCP statement is required on the interface(s), as shown in [Example 2-77](#).

Example 2-77 Per-Port Policing Configuration Example on a Catalyst 4500-E Supervisor 6-E

```
! This section configures the Per-Port policing policy-map
C4500-E(config)#policy-map PER-PORT-POLICING
C4500-E(config-pmap)# class VVLAN-VOIP
C4500-E(config-pmap-c)# set dscp ef
C4500-E(config-pmap-c)# police 128k bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! VoIP is marked EF and policed to drop at 128 kbps
C4500-E(config-pmap)# class VVLAN-SIGNALING
C4500-E(config-pmap-c)# set dscp cs3
C4500-E(config-pmap-c)# police 32k bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! (VVLAN) Signaling is marked CS3 and policed to drop at 32 kbps
C4500-E(config-pmap)# class MULTIMEDIA-CONFERENCING
C4500-E(config-pmap-c)# set dscp af41
C4500-E(config-pmap-c)# police 5m bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! Multimedia-conferencing is marked AF41 and policed to drop at 5 Mbps
C4500-E(config-pmap)# class SIGNALING
C4500-E(config-pmap-c)# set dscp cs3
C4500-E(config-pmap-c)# police 32k bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
```

```

! (DVLAN) Signaling is marked CS3 and policed to drop at 32 kbps
C4500-E(config-pmap)# class TRANSACTIONAL-DATA
C4500-E(config-pmap-c)# set dscp af21
C4500-E(config-pmap-c)# police 10m bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action set-dscp-transmit cs1
! Trans-data is marked AF21 and policed to remark (to CS1) at 10 Mbps
C4500-E(config-pmap)# class BULK-DATA
C4500-E(config-pmap-c)# set dscp af11
C4500-E(config-pmap-c)# police 10m bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action set-dscp-transmit cs1
! Bulk-data is marked AF11 and policed to remark (to CS1) at 10 Mbps
C4500-E(config-pmap)# class SCAVENGER
C4500-E(config-pmap-c)# set dscp cs1
C4500-E(config-pmap-c)# police 10m bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! Scavenger traffic is marked CS1 and policed to drop at 10 Mbps
C4500-E(config-pmap)# class class-default
C4500-E(config-pmap-c)# set dscp default
C4500-E(config-pmap-c)# police 10m bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action set-dscp-transmit cs1
! The implicit default class marks all other traffic to DF
! and polices all other traffic to remark (to CS1) at 10 Mbps

! This section attaches the service-policy to the interface(s)
C4500-E(config)#interface range GigabitEthernet 2/1-48
C4500-E(config-if-range)# switchport access vlan 10
C4500-E(config-if-range)# switchport voice vlan 110
C4500-E(config-if-range)# spanning-tree portfast
C4500-E(config-if-range)# qos trust device cisco-phone
! The interface is set to conditionally-trust Cisco IP Phones
C4500-E(config-if-range)# service-policy input PER-PORT-POLICING
! Attaches the Per-Port Policing policy to the interface(s)

```

**Note**

Advanced network administrators can leverage the Catalyst 4500-E Supervisor 6-E's support of three color markers—either the RFC 2697 single rate three color marker (srTCM) or the RFC 2698 two rate three color marker (trTCM)—such that the exceeding policing action for the transactional data and bulk data policers would be to remark to AF22 and AF12 (respectively), while the violating policing action for these classes would be to remark to CS1.

This configuration can be verified with the commands:

- **show qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Per-Port/Per-VLAN Policing Model (Classic Supervisors)

An alternative—and more discrete—approach for deploying policing policies on the Catalyst 4500 platforms is to deploy these on a per-port/per-VLAN basis. The Catalyst 4500 has a very elegant syntax for deploying per-port/per-VLAN policies (as compared to the Catalyst 3750-E syntax, for example), where policies are applied within a VLAN mode within a switch port's interface configuration mode, as shown in [Example 2-78](#).

In [Example 2-78](#), three policers are applied to the VVLAN of a given access edge trunk port: one to police VoIP to 128 kbps, another to police signaling to 32 kbps, and a third to police everything else to 32 kbps. On the other hand, six policers are applied to the DVLAN of a given access edge trunk port: one to police multimedia conferencing traffic to drop at 5 Mbps, a second to police signaling to drop at 32 kbps, a third to police transactional data to remark (to CS1) at 10 Mbps, a fourth to police bulk data to remark (to CS1) at 10 Mbps, a fifth to police scavenger to drop at 10 Mbps, and a sixth to police everything else to remark (to CS1) at 10 Mbps.

As in the previous examples, remarking is performed by configuring a policed-DSCP map with the global configuration command **qos map dscp policed**, which specifies which DSCP values are subject to remarking (if out-of-profile) and what these values should be remarked to (which in the case of scavenger-class QoS policies, the remarking value is CS1/DSCP 8).

Example 2-78 Per-Port/Per-VLAN Policing Configuration Example on a Catalyst 4500 Classic Supervisor

```
! This section configures the global policed-DSCP markdown map
C4500-CS(config)#qos map dscp policed 0 10 18 to dscp 8
! DSCP 0 (DF), 10 (AF11) and 18 (AF21) are marked down to 8 (CS1)
! if found to be in excess of their (respective) policing rates

! This section configures the policy-map for the VVLAN policers
C4500-CS(config)#policy-map VVLAN-POLICERS
C4500-CS(config-pmap)# class VVLAN-VOIP
C4500-CS(config-pmap-c)# set dscp ef
C4500-CS(config-pmap-c)# police 128k 8000 exceed-action drop
! VoIP is marked EF and policed to drop at 128 kbps
C4500-CS(config-pmap-c)# class VVLAN-SIGNALING
C4500-CS(config-pmap-c)# set dscp cs3
C4500-CS(config-pmap-c)# police 32k 8000 exceed-action drop
! (VVLAN) Signaling is marked CS3 and policed to drop at 32 kbps
C4500-CS(config-pmap-c)# class class-default
C4500-CS(config-pmap-c)# set dscp default
C4500-CS(config-pmap-c)# police 32k 8000 exceed-action drop
! The implicit default class marks all other VVLAN traffic to DF
! and polices to drop at 32 kbps

! This section configures the policy-map for the DVLAN policers
C4500-CS(config)#policy-map DVLAN-POLICERS
C4500-CS(config-pmap)# class MULTIMEDIA-CONFERENCING
C4500-CS(config-pmap-c)# set dscp af41
C4500-CS(config-pmap-c)# police 5m 8000 exceed-action drop
! Multimedia-conferencing is marked AF41 and policed to drop at 5 Mbps
C4500-CS(config-pmap-c)# class SIGNALING
C4500-CS(config-pmap-c)# set dscp cs3
C4500-CS(config-pmap-c)# police 32k 8000 exceed-action drop
! (DVLAN) Signaling is marked CS3 and policed to drop at 32 kbps
C4500-CS(config-pmap-c)# class TRANSACTIONAL-DATA
C4500-CS(config-pmap-c)# set dscp af21
C4500-CS(config-pmap-c)# police 10m 8000 exceed-action policed-dscp-transmit
```



```

! Trans-data is marked AF21 and policed to remark (to CS1) at 10 Mbps
C4500-CS(config-pmap-c)# class BULK-DATA
C4500-CS(config-pmap-c)# set dscp af11
C4500-CS(config-pmap-c)# police 10m 8000 exceed-action policed-dscp-transmit
! Bulk-data is marked AF11 and policed to remark (to CS1) at 10 Mbps
C4500-CS(config-pmap-c)# class SCAVENGER
C4500-CS(config-pmap-c)# set dscp cs1
C4500-CS(config-pmap-c)# police 10m 8000 exceed-action drop
! Scavenger traffic is marked CS1 and policed to drop at 10 Mbps
C4500-CS(config-pmap-c)# class class-default
C4500-CS(config-pmap-c)# set dscp default
C4500-CS(config-pmap-c)# police 10m 8000 exceed-action policed-dscp-transmit
! The implicit default class marks all other traffic to DF
! and polices all other traffic to remark (to CS1) at 10 Mbps

! This section attaches the policy to the VLANs on a Per-Port basis
C4500-CS(config)#interface range GigabitEthernet 2/1-48
C4500-CS(config-if-range)# switchport access vlan 10
C4500-CS(config-if-range)# switchport voice vlan 110
C4500-CS(config-if-range)# spanning-tree portfast
C4500-CS(config-if-range)# qos trust device cisco-phone
! The interface is set to conditionally-trust Cisco IP Phones
C4500-CS(config-if-range)# qos trust cos
! CoS-trust will be dynamically extended to Cisco IP Phones
C4500-CS(config-if-range)# vlan 10
C4500-CS(config-if-vlan-range)# service-policy input DVLAN-POLICERS
! Attaches the Per-Port/Per-VLAN DVLAN Policing policy to the
! DVLAN of the trunked interface(s)
C4500-CS(config-if-range)# vlan 110
C4500-CS(config-if-vlan-range)# service-policy input VVLAN-POLICERS
! Attaches the Per-Port/Per-VLAN VVLAN Policing policy to the
! VVLAN of the trunked interface(s)

```

This configuration can be verified with the commands:

- **show qos maps dscp policed**
- **show qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**
- **show policy-map interface** *interface x/y vlan vlan-number*

Per-Port/Per-VLAN Policing Model (Supervisor 6-E/7-E)

The per-port/per-VLAN policing model is essentially the same for the Catalyst 4500-E Supervisor 6-E, except that it does not require a global policed-DSCP map and thus the policing commands are slightly different; also no trust-DSCP statement is required on the interface(s), as shown in [Example 2-79](#).

Example 2-79 Per-Port/Per-VLAN Policing Configuration Example on a Catalyst 4500-E Supervisor 6-E

```

! This section configures the policy-map for the VVLAN policers
C4500-E(config)#policy-map VVLAN-POLICERS
C4500-E(config-pmap)# class VVLAN-VOIP
C4500-E(config-pmap-c)# set dscp ef
C4500-E(config-pmap-c)# police 128k bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop

```

```

! VoIP is marked EF and policed to drop at 128 kbps
C4500-E(config-pmap)# class VVLAN-SIGNALING
C4500-E(config-pmap-c)# set dscp cs3
C4500-E(config-pmap-c)# police 32k bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! (VVLAN) Signaling is marked CS3 and policed to drop at 32 kbps
C4500-E(config-pmap)# class class-default
C4500-E(config-pmap-c)# set dscp default
C4500-E(config-pmap-c)# police 32k bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! The implicit default class marks all other VVLAN traffic to DF
! and polices to drop at 32 kbps

! This section configures the policy-map for the DVLAN policers
C4500-E(config)#policy-map DVLAN-POLICERS
C4500-E(config-pmap)# class MULTIMEDIA-CONFERENCING
C4500-E(config-pmap-c)# set dscp af41
C4500-E(config-pmap-c)# police 5m bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! Multimedia-conferencing is marked AF41 and policed to drop at 5 Mbps
C4500-E(config-pmap)# class SIGNALING
C4500-E(config-pmap-c)# set dscp cs3
C4500-E(config-pmap-c)# police 32k bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! (DVLAN) Signaling is marked CS3 and policed to drop at 32 kbps
C4500-E(config-pmap)# class TRANSACTIONAL-DATA
C4500-E(config-pmap-c)# set dscp af21
C4500-E(config-pmap-c)# police 10m bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action set-dscp-transmit cs1
! Trans-data is marked AF21 and policed to remark (to CS1) at 10 Mbps
C4500-E(config-pmap)# class BULK-DATA
C4500-E(config-pmap-c)# set dscp af11
C4500-E(config-pmap-c)# police 10m bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action set-dscp-transmit cs1
! Bulk-data is marked AF11 and policed to remark (to CS1) at 10 Mbps
C4500-E(config-pmap)# class SCAVENGER
C4500-E(config-pmap-c)# set dscp cs1
C4500-E(config-pmap-c)# police 10m bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! Scavenger traffic is marked CS1 and policed to drop at 10 Mbps
C4500-E(config-pmap)# class class-default
C4500-E(config-pmap-c)# set dscp default
C4500-E(config-pmap-c)# police 10m bc 8000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action set-dscp-transmit cs1
! The implicit default class marks all other traffic to DF
! and polices all other traffic to remark (to CS1) at 10 Mbps

! This section attaches the policy to the VLANs on a Per-Port basis
C4500-E(config)#interface range GigabitEthernet 2/1-48
C4500-E(config-if-range)# switchport access vlan 10
C4500-E(config-if-range)# switchport voice vlan 110
C4500-E(config-if-range)# spanning-tree portfast
C4500-E(config-if-range)# qos trust device cisco-phone
! The interface is set to conditionally-trust Cisco IP Phones

```

```

C4500-E(config-if-range)# vlan 10
C4500-E(config-if-vlan-range)# service-policy input DVLAN-POLICERS
! Attaches the Per-Port/Per-VLAN DVLAN Policing policy to the
! DVLAN of the trunked interface(s)
C4500-E(config-if-range)# vlan 110
C4500-E(config-if-vlan-range)# service-policy input VVLAN-POLICERS
! Attaches the Per-Port/Per-VLAN VVLAN Policing policy to the
! VVLAN of the trunked interface(s)

```

**Note**

Advanced network administrators can leverage the Catalyst 4500-E Supervisor 6-E's support of three color markers—either the RFC 2697 single rate three color marker (srTCM) or the RFC 2698 two rate three color marker (trTCM)—such that the exceeding policing action for the transactional data and bulk data policers would be to remark to AF22 and AF12 (respectively), while the violating policing action for these classes would be to remark to CS1.

This configuration can be verified with the commands:

- **show qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**
- **show policy-map interface** *interface x/y* **vlan** *vlan-number*

Per-Port User-Based Rate Limiting (Supervisor V-10GE)

UBRL adopts microflow policing capability to dynamically learn traffic flows and rate limit each unique flow to an individual rate and, as such, is a highly effective and efficient policing tool, particularly at the distribution layer in a medianet campus network.

UBRL is available on Supervisor Engine V-10GE with NetFlow support. UBRL can be applied to ingress traffic on routed interfaces and is typically used in environments where a per-user, granular rate limiting mechanism is required, such as at the distribution layer, to provide a second line of policing defense in the campus. Like other policers, UBRL can be used to drop or remark exceeding flows.

A flow is defined by five-tuples (IP source address, IP destination address, IP protocol field, Layer 4 protocol source, and destination ports), which are the same for each packet in the flow. Flow-based policers apply a single policy to discrete flows without having to specify the virtually-infinite tuple-combinations. UBRL can also be applied with source or destination flow masks; these masks apply an aggregate microflow policing policy to multiple flows sharing the same source or IP destination addresses.

In the per-port UBRL Model, a class map matches on a microflow basis and aggregates these by source addresses. Then a policer applies an aggregate limit to all microflows sharing a common source IP address, remarking traffic in excess of the policing rate.

Remarking is performed by configuring a policed-DSCP map with the global configuration command **qos map dscp policed**, which specifies which DSCP values are subject to remarking (if out-of-profile) and what these values should be remarked to (which in the case of scavenger class QoS policies, the remarking value is CS1/DSCP 8).

UBRL is supported on Layer 3 interfaces and can be applied on either a per-port or per-port/per-VLAN-basis, as shown in [Example 2-80](#) and [Example 2-81](#), respectively.

In [Example 2-80](#), the campus distribution block is using a routed access design and, as such, has Layer 3 interfaces (TenGigabitEthernet 1/1 and 1/2) connecting it to the access layer switches. UBRL is applied to all flows to ensure that any endpoint transmitting at more than 5% capacity (an example value) of the access edge 10/100/1000 switch ports are subject to data plane policing/scavenger class QoS.

Example 2-80 Per-Port UBRL Configuration Example on a Catalyst 4500 Supervisor V-10GE

```
! This section configures the global policed-DSCP markdown map
C4500-CS(config)#qos map dscp policed 0 10 18 24 34 46 to dscp 8
! DSCP 0 (DF), 10 (AF11) and 18 (AF21), 24 (CS3), 34 (AF41) & 46 (EF)
! are marked down to 8 (CS1) if found to be in excess of their
! (respective) policing rates

! This section defines the sourced-based microflow class-map
C4500-SupV-10GE(config)#class-map match-all ENDPOINTS
C4500-SupV-10GE(config-cmap)# match flow ip source-address
! All flows sharing a unique source IP address will be matched

! This section defines the aggregate per-source-IP UBRL policer
C4500-SupV-10GE(config)#policy-map UBRL
C4500-SupV-10GE(config-pmap)# class ENDPOINTS
C4500-SupV-10GE(config-pmap-c)# police 50m 8000 byte conform-action transmit exceed-action
policed-dscp-transmit
! Any flows from a single source IP address
! will be remarked to CS1 if exceeding 50 Mbps

! This section attaches the UBRL policy to a Layer 3 interface
C4500-SupV-10GE(config)#interface range TenGigabitEthernet1/1-2
C4500-SupV-10GE(config-if-range)# description L3-Dwnlnk to Access-Layer
C4500-SupV-10GE(config-if-range)# no switchport
C4500-SupV-10GE(config-if-range)# qos trust dscp
! Sets the interface(s) trust state to statically trust-DSCP
! As this is a Distribution-Layer downlink
C4500-SupV-10GE(config-if-range)# service-policy input UBRL
! Attaches the UBRL policy to the interface(s)
```

This configuration can be verified with the commands:

- **show qos maps dscp policed**
- **show qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Per-Port/Per-VLAN User-Based Rate Limiting (Supervisor V-10GE)

In contrast with the previous example, if the campus distribution block is using a Layer 2/Layer 3 design, and as such has Layer 2 trunked interfaces (TenGigabitEthernet 1/1 and 1/2) connecting it to the access layer switches, then UBRL can be applied on a per-port/per-VLAN basis. In this case, separate UBRL policies can be applied to each VLAN traversing the trunked interfaces—via per-port/per-VLAN UBRL policies—as each VLAN is routed through the switch.

To highlight policy flexibility, additional levels of classification are included in this second UBRL example (which incidentally can also be applied to the per-port UBRL model). Instead of applying a blanket UBRL policy to all endpoints, separate UBRL policies can be applied to different types of endpoints or application-and-endpoint-combinations. For example, VoIP from Cisco IP phones in the VVLAN can be rate limited to 128 Kbps, while signaling traffic from these endpoints can be limited to 32 kbps. Similarly, TelePresence endpoints in the VVLAN (which mark their media flows to CS4) can be limited to 25 Mbps. All other endpoint-generated traffic in the VVLAN can be limited to 32 kbps per endpoint.

Similar policy granularity can be applied to the DVLAN policer, if desired. However in this example, a simplified DVLAN policer is applied to all flows to ensure that any DVLAN endpoint transmitting at more than 5% capacity (an example value) of the access edge 10/100/1000 switch ports are subject to data plane policing/scavenger class QoS.

Static DSCP-trust is configured on the physical ports and the per-port/per-VLAN UBRL policies are applied to their respective VLANs within the trunked interface, as shown in [Example 2-81](#).

Example 2-81 Per-Port/Per-VLAN UBRL Configuration Example on a Catalyst 4500 Supervisor V-10GE

```
! This section configures the global policed-DSCP markdown map
C4500-CS(config)#qos map dscp policed 0 10 18 34 to dscp 8
! DSCP 0 (DF), 10 (AF11) and 18 (AF21) and 34 (AF41)
! are marked down to 8 (CS1) if found to be in excess of their
! (respective) policing rates

! This section defines the sourced-based microflow class-maps
4507-E(config)#class-map match-all VOIP-ENDPOINTS
4507-E(config-cmap)# match ip dscp ef
4507-E(config-cmap)# match flow ip source-address
! All flows marked EF from a single source IP will be matched

4507-E(config)#class-map match-all TELEPRESENCE-ENDPOINTS
4507-E(config-cmap)# match ip dscp cs4
4507-E(config-cmap)# match flow ip source-address
! All flows marked CS4 from a single source IP will be matched

4507-E(config)#class-map match-all SIGNALING-ENDPOINTS
4507-E(config-cmap)# match ip dscp cs3
4507-E(config-cmap)# match flow ip source-address
! All flows marked CS3 from a single source IP will be matched

C4500-SupV-10GE(config)#class-map match-all ENDPOINTS
C4500-SupV-10GE(config-cmap)# match flow ip source-address
! All flows sharing a unique source IP address will be matched

! This section defines the aggregate per-source-IP VVLAN UBRL policer
4507-E(config)#policy-map VVLAN-UBRL
4507-E(config-pmap)# class VOIP-ENDPOINTS
4507-E(config-pmap-c)# police 128k 8000 byte conform-action transmit exceed-action drop
! All flows marked EF from a single IP are policed to drop at 128 kbps
4507-E(config-pmap)# class TELEPRESENCE-ENDPOINTS
4507-E(config-pmap-c)# police 25m 256000 byte conform-action transmit exceed-action drop
! All flows marked CS4 from a single IP are policed to drop at 25 Mbps
4507-E(config-pmap)# class SIGNALING-ENDPOINTS
4507-E(config-pmap-c)# police 32k 8000 byte conform-action transmit exceed-action drop
! All flows marked CS3 from a single IP are policed to drop at 32 kbps
4507-E(config-pmap)# class ENDPOINTS
4507-E(config-pmap-c)# police 32k 8000 byte conform-action transmit exceed-action drop
! All other flows from a single VVLAN IP are policed to 32 kbps
```

```

! This section defines the aggregate per-source-IP DVLAN UBRL policer
C4500-SupV-10GE(config)#policy-map DVLAN-UBRL
C4500-SupV-10GE(config-pmap)# class ENDPOINTS
C4500-SupV-10GE(config-pmap-c)# police 50m 8000 byte conform-action transmit exceed-action
policed-dscp-transmit
! Any flows from a single source IP address within the DVLAN
! will be remarked to CS1 if exceeding 50 Mbps

! This section configures static DSCP trust on the trunked interfaces
! And attaches the UBRL policies to their respective VLANs
C4500-SupV-10GE(config)#interface range TenGigabitEthernet1/1-2
C4500-SupV-10GE(config-if-range)# description L2-Dwnlnk to Access-Layer
C4500-SupV-10GE(config-if)# switchport trunk encapsulation dot1q
C4500-SupV-10GE(config-if)# switchport trunk allowed vlan 10,110
C4500-SupV-10GE(config-if)# switchport mode trunk
C4500-SupV-10GE(config-if)# qos trust dscp
! Sets the interface(s) trust state to statically trust-DSCP
! As this is a Distribution-Layer (trunked) downlink
C4500-SupV-10GE(config)#int vlan 10
C4500-SupV-10GE(config-if)# service-policy input DVLAN-UBRL
! Attaches the Per-Port/Per-VLAN DVLAN UBRL policy to the
! DVLAN of the trunked interface(s)
C4500-SupV-10GE(config)#int vlan 110
C4500-SupV-10GE(config-if)# service-policy input VVLAN-UBRL
! Attaches the Per-Port/Per-VLAN VVLAN UBRL policy to the
! DVLAN of the trunked interface(s)

```

This configuration can be verified with the commands:

- **show qos maps dscp policed**
- **show qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**
- **show policy-map interface** *interface x/y* **vlan** *vlan-number*

Queuing Models

The Catalyst 4500 switch family only supports egress queuing models, which can be configured on the Classic Supervisor branch to operate in either a 4Q1T mode or a 1P3Q1T mode (the latter of which is recommended for medianet campus networks, as it supports the EF PHB) and on the Supervisor 6-E branch can be configured (via MQC) to provide a flexible queuing structure, to a maximum of 1P7Q1T.

Additionally, the Catalyst 4500 family uses a platform-specific congestion avoidance algorithm to provide Active Queue Management (AQM), namely Dynamic Buffer Limiting (DBL). DBL tracks the queue length for each traffic flow in the switch. When the queue length of a flow exceeds its limit, DBL drop packets or set the Explicit Congestion Notification (ECN) bits in the packet headers.

Furthermore, the Catalyst 4500 supports DSCP-to-queue mapping on both branches.

The Catalyst 4500 Classic Supervisor 1P3Q1T+DBL model and the Supervisor 6-E 1P7Q1T+DBL models are detailed in the following sections.

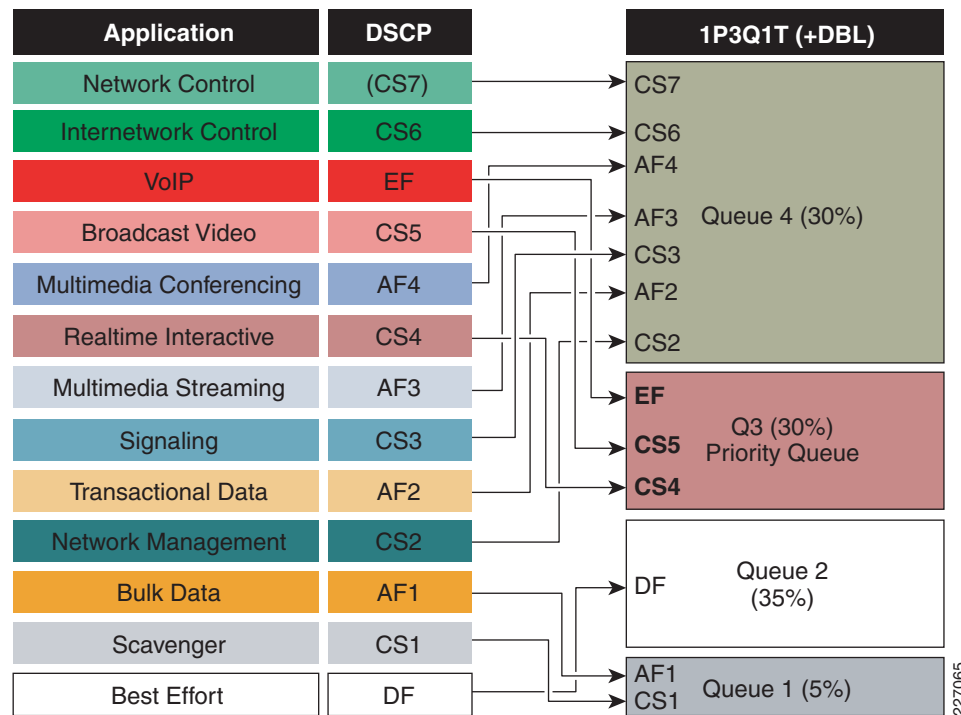
Egress Queuing 1P3Q1T+DBL (Classic Supervisors) Model

On the Catalyst 4500 Classic Supervisors, queue 3 can be enabled as a strict-priority queue. Once enabled, Q4 can be used as a guaranteed bandwidth queue, Q2 can be used as the default best effort queue, and Q1 can be used as a less than best effort (scavenger) queue. Bandwidth can be assigned as: 5%, 35%, 30%, and 30% for queues 1 through 4, respectively.

DBL can be enabled on all DSCP values, with the exception of DSCP values that are mapped to the PQ (specifically, CS4/32, CS5/40, and EF/46), as this may cause drops to occur on these real time flows. Additionally, DBL can be enabled to mark the IP Explicit Congestion Notification (IP ECN) bits within the IP ToS Byte in the event of congestion. A service policy can be configured to enable DBL on all flows (except those already noted) and applied to each interface on which queuing is enabled.

Once these queues have been configured, then VoIP (EF), broadcast video (CS5), and realtime interactive (CS4) traffic can be mapped to the strict priority queue (Q3). Network control (CS7), internetwork control (CS6), signaling (CS3), and management (CS2) traffic can be mapped to Q4, along with multimedia conferencing (AF4), multimedia streaming (AF3), and transactional data (AF2). Best effort traffic is sent to the default queue (Q2), while bulk data (AF1) and scavenger (CS1) traffic are mapped to the deferential queue (Q1). These 1P3Q1T+DBL egress queuing mappings for the Catalyst 4500 Classic Supervisor are shown in Figure 2-22.

Figure 2-22 Catalyst 4500 Classic Supervisor 1P3Q1T+DBL Egress Queuing Model



The corresponding configuration for 1P3Q1T+DBL egress queuing on a Catalyst 4500 Classic Supervisor is shown in Example 2-82.

Example 2-82 1P3Q1T+DBL Egress Queuing Configuration Example on a Catalyst 4500 Classic Supervisor

```
! This section enables and configures DBL
C4500-CS(config)#qos db1
! DBL is globally enabled
```

```

C4500-CS(config)#no qos dbl dscp-based 32
C4500-CS(config)#no qos dbl dscp-based 40
C4500-CS(config)#no qos dbl dscp-based 46
! DBL is explicitly disabled on DSCP CS4, CS5 and EF
! as these DSCP values are assigned to the PQ
! and as such should never experience congestion avoidance drops
C4500-CS(config)#qos dbl exceed-action ecn
! DBL will mark IP ECN bits in the event of congestion

! This section configures the DBL policy-map
C4500-CS(config)#policy-map DBL
C4500-CS(config-pmap)# class class-default
C4500-CS(config-pmap-c)#  dbl
! DBL is enabled on all flows
! (with the exception of DSCP CS4, CS5 and EF)

! This section configures the DSCP-to-Queue mappings
C4500-CS(config)#qos map dscp 8 10 12 14 to tx-queue 1
! DSCP CS1 and AF1 are mapped to Q1 (the less than best effort queue)
C4500-CS(config)#qos map dscp 0 to tx-queue 2
! DSCP DF is mapped to Q2 (the best effort/default queue)
C4500-CS(config)#qos map dscp 32 40 46 to tx-queue 3
! DSCP CS4, CS5 and EF are mapped to Q3 (the PQ)
C4500-CS(config)#qos map dscp 16 18 20 22 to tx-queue 4
! DSCP CS2 and AF2 are mapped to Q4 (guaranteed BW queue)
C4500-CS(config)#qos map dscp 24 26 28 30 to tx-queue 4
! DSCP CS3 and AF3 are mapped to Q4 (guaranteed BW queue)
C4500-CS(config)#qos map dscp 34 36 38 to tx-queue 4
! DSCP AF4 is mapped to Q4 (guaranteed BW queue)
C4500-CS(config)#qos map dscp 48 56 to tx-queue 4
! DSCP CS6 and CS7 are mapped to Q4 (guaranteed BW queue)

! This section configures the interface(s) for egress queuing
C4500-CS(config)#interface range GigabitEthernet 1/1-2
C4500-CS(config-if-range)# tx-queue 1
C4500-CS(config-if-tx-queue)#  bandwidth percent 5
! Q1 (less than best effort queue) is assigned 5% BW
C4500-CS(config-if-tx-queue)# tx-queue 2
C4500-CS(config-if-tx-queue)#  bandwidth percent 35
! Q2 (default/best effort queue) is assigned 35% BW
C4500-CS(config-if-tx-queue)# tx-queue 3
C4500-CS(config-if-tx-queue)#  priority high
C4500-CS(config-if-tx-queue)#  bandwidth percent 30
C4500-CS(config-if-tx-queue)#  shape percent 30
! Q3 is enabled as a PQ and assigned 30% BW
! Additionally Q3 is shaped (limited) to 30%
C4500-CS(config-if-tx-queue)# tx-queue 4
C4500-CS(config-if-tx-queue)#  bandwidth percent 30
! Q4 (guaranteed BW queue) is assigned 30% BW
C4500-CS(config-if-range)# service-policy output DBL
! DBL policy-map is attached to the interface(s)

```

This configuration can be verified with the commands:

- **show qos dbl** (as shown in [Example 2-83](#))
- **show qos maps dscp tx-queue** (as shown in [Example 2-84](#))
- **show qos interface** (as shown in [Example 2-85](#))

- **show class-map**
- **show policy-map**
- **show policy-map interface**

Example 2-83 Verifying DBL on a Catalyst 4500 Classic Supervisor—show qos dbl

```
C4500-CS#show qos dbl
QoS is enabled globally
DBL is enabled globally on DSCP values:
    0-31,33-39,41-45,47-63
DBL flow includes vlan
DBL flow includes layer4-ports
DBL uses ecn to indicate congestion
DBL exceed-action probability: 15%
DBL max credits: 15
DBL aggressive credit limit: 10
DBL aggressive buffer limit: 2 packets

C4500-CS#
```

Example 2-83 shows that DBL has been globally enabled and is active on all DSCP values with the exception of CS4/32, CS5/40, and EF/46. Also that DBL uses IP ECN to indicate congestion.

Example 2-84 Verifying DSCP-to-Queue Mappings on a Catalyst 4500 Classic Supervisor—show qos maps dscp tx-queue

```
C4500-CS#show qos maps dscp tx-queue
DSCP-TxQueue Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    02 01 01 01 01 01 01 01 01 01
1 :    01 01 01 01 01 01 04 02 04 02
2 :    04 02 04 02 04 02 04 02 04 02
3 :    04 02 03 03 04 03 04 03 04 03
4 :    03 03 03 03 03 03 03 03 04 04
5 :    04 04 04 04 04 04 04 04 04 04
6 :    04 04 04 04

C4500-CS#
```

Example 2-84 shows the ingress DSCP-to-queue mappings. The first digit of the DSCP value of a packet is shown along the Y-axis of the table; the second digit of the DSCP value of a packet is shown along the X-axis of the table. The mapping table corresponds to Figure 2-22. It can be noted that CS4 (DSCP 32), CS5 (DSCP 40), and EF (DSCP 46) are all mapped to Q3 (the PQ). It should also be noted that internal DSCP values 32 through 47 are mapped to Q2 by default, which is why the table shows additional values being mapped to this queue.

Example 2-85 Verifying Queuing Settings on a Catalyst 4500 Classic Supervisor—show qos interface

```
C4500-CS#show qos interface GigabitEthernet 1/1
QoS is enabled globally
Port QoS is enabled
Administrative Port Trust State: 'dscp'
Operational Port Trust State: 'dscp'
Trust device: none
Default DSCP: 0 Default CoS: 0
Appliance trust: none
Tx-Queue   Bandwidth   ShapeRate   Priority   QueueSize
          (bps)         (bps)
-----
```

1	<u>50000000</u>	disabled	N/A	1920
2	<u>350000000</u>	disabled	N/A	1920
3	<u>300000000</u>	disabled	high	1920
4	<u>300000000</u>	disabled	N/A	1920

C4500-CS#

[Example 2-85](#) shows that interface GigabitEthernet 1/1 has been configured such that Q1 through Q4 receive 5%, 35%, 30%, and 30% (respectively) of the interface bandwidth (1 Gbps) and that Q3 has been enabled as a high priority/strict priority queue.

Egress Queuing 1P7Q1T+DBL (Supervisor 6-E/7-E) Model

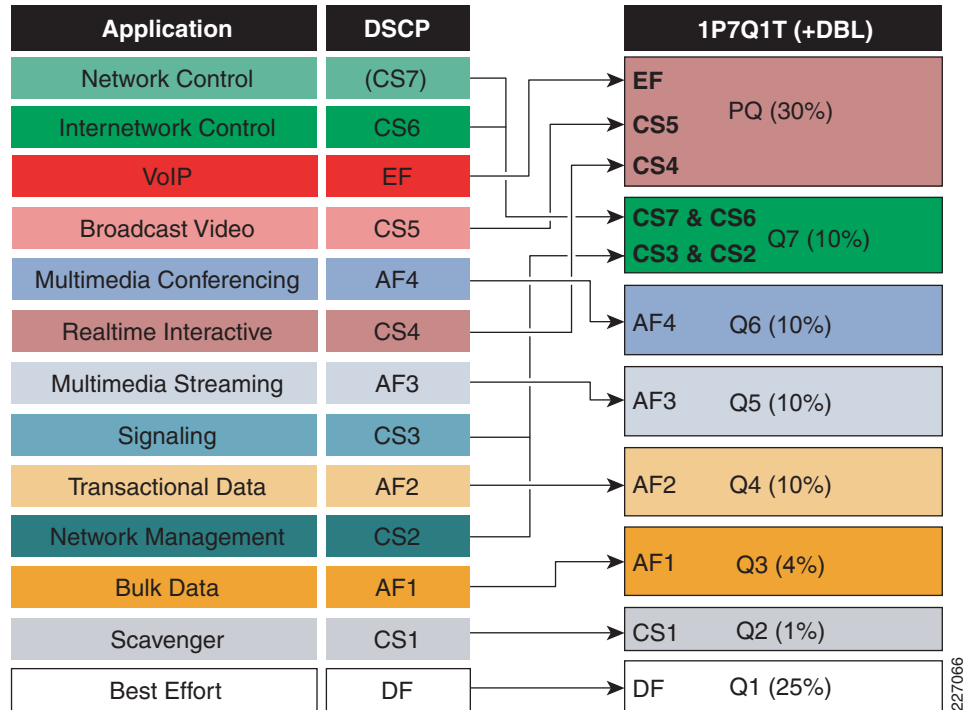
The Catalyst 4500-E Supervisor 6-E hardware supports (up to) eight transmit queues per port. Queues are assigned when an output policy is attached to a port with one or more queuing related actions for one or more classes of traffic. Because there are only eight queues per port, there can be at most eight classes of traffic (including the reserved class, class default) with queuing actions defined.

On the Catalyst 4500-E Supervisor 6-E only one transmit queue on a port can be configured as strict priority queue (which, in effect, constitutes a hardware low latency queue) with the **priority** policy-map class action command. The priority queue is serviced first until it is empty or until it is under its limited rate. Only one traffic stream can be destined for the priority queue per class level policy (in other words, multiple hardware LLQs are not supported on the Supervisor 6-E). The Supervisor 6-E supports an unconditional (explicit) policer to rate limit packets enqueued to the strict priority queue. When the priority queue is configured on one class of a policy map without a policer, only **bandwidth remaining percent** is accepted on other classes (guaranteeing a minimum bandwidth for other classes from the remaining bandwidth of what is left after using the priority queue); however, when the priority queue is configured with a policer, then either **bandwidth percent** or **bandwidth remaining percent** is accepted on the other queuing classes.

However, if queuing policies are to be applied to EtherChannel interfaces, then it is recommended not to police the priority queue. This is because two policy maps would be needed in this case: one policy-map would be needed to police the priority queue (which would have to be applied to the logical EtherChannel interface in the egress direction) and a second policy-map would be needed to define the queuing policy (using bandwidth remaining percent), which would be applied to all EtherChannel physical port-member interfaces in the egress direction. Thus to simplify the queuing policy and to increase its portability and modularity, the priority queue is not policed in [Example 2-86](#).

Additionally, as with the Classic Supervisors, DBL can be enabled on a per-class basis, but is most effective when applied against TCP-based traffic flows (as opposed to UDP-based traffic flows).

Thus the Catalyst 4500-E Supervisor 6-E can be configured to operate in a 1P7Q1T+DBL mode. VoIP (EF), broadcast video (CS5), and realtime interactive (CS4) flows can be assigned to the strict priority queue. Network and internetwork control (CS6 and CS7, respectively), along with signaling and management (CS3 and CS2, respectively), can all share a control/management queue. This allows for dedicated queues to be provisioned for multimedia conferencing (AF4), multimedia streaming (AF3), transactional data (AF2), and bulk data (AF1). Also, scavenger (CS1) traffic can share a bandwidth-constrained “less than best effort” queue, while all other traffic is assigned to the default/best effort queue. The recommended 1P7Q1T Sup 6-E egress queuing configuration for the C4500-E Supervisor 6-E is illustrated in [Figure 2-23](#).

Figure 2-23 Catalyst 4500-E Supervisor 6-E 1P7Q1T+DBL Egress Queuing Model

The corresponding configuration for 1P7Q1T+DBL egress queuing on a Catalyst 4500-E Supervisor 6-E is shown in [Example 2-86](#).

Example 2-86 1P7Q1T+DBL Egress Queuing Configuration Example on a Catalyst 4500-E Supervisor-6E

```

! This section configures the class-maps for the egress queuing policy
! Note: these class-maps require unique names from any ingress
!     policy class-maps; otherwise classification errors may occur
!     due to overlapping classification logic
C4500-E(config)#class-map match-any PRIORITY-QUEUE
C4500-E(config-cmap)# match dscp ef
C4500-E(config-cmap)# match dscp cs5
C4500-E(config-cmap)# match dscp cs4
! VoIP (EF), Broadcast Video (CS5) and Realtime Interactive (CS4)
! are all mapped to the PQ
C4500-E(config)#class-map match-any CONTROL-MGMT-QUEUE
C4500-E(config-cmap)# match dscp cs7
C4500-E(config-cmap)# match dscp cs6
C4500-E(config-cmap)# match dscp cs3
C4500-E(config-cmap)# match dscp cs2
! Network Control (CS7), Internetwork Control (CS6),
! Signaling (CS3) and Management (CS2) are mapped
! to a Control/Management Queue
C4500-E(config)#class-map match-all MULTIMEDIA-CONFERENCING-QUEUE
C4500-E(config-cmap)# match dscp af41 af42 af43
! Multimedia Conferencing (AF4) is assigned a dedicated queue
C4500-E(config)#class-map match-all MULTIMEDIA-STREAMING-QUEUE
C4500-E(config-cmap)# match dscp af31 af32 af33
! Multimedia Streaming (AF3) is assigned a dedicated queue
C4500-E(config)#class-map match-all TRANSACTIONAL-DATA-QUEUE
C4500-E(config-cmap)# match dscp af21 af22 af23
! Transactional Data (AF2) is assigned a dedicated queue

```

```

C4500-E(config)#class-map match-all BULK-DATA-QUEUE
C4500-E(config-cmap)# match dscp af11 af12 af13
! Bulk Data (AF1) is assigned a dedicated queue
C4500-E(config)#class-map match-all SCAVENGER-QUEUE
C4500-E(config-cmap)# match dscp cs1
! Scavenger (CS1) is assigned a dedicated queue

! This section configures the 1P7Q1T+DBL egress queuing policy-map
C4500-E(config)#policy-map 1P7Q1T
C4500-E(config-pmap-c)# class PRIORITY-QUEUE
C4500-E(config-pmap-c)# priority
! Defines a priority queue
C4500-E(config-pmap-c)# class CONTROL-MGMT-QUEUE
C4500-E(config-pmap-c)# bandwidth remaining percent 10
! Defines a control/management queue with 10% BW remaining
C4500-E(config-pmap-c)# class MULTIMEDIA-CONFERENCING-QUEUE
C4500-E(config-pmap-c)# bandwidth remaining percent 10
! Defines a multimedia conferencing queue with 10% BW remaining
C4500-E(config-pmap-c)# class MULTIMEDIA-STREAMING-QUEUE
C4500-E(config-pmap-c)# bandwidth remaining percent 10
! Defines a multimedia streaming queue with 10% BW remaining
C4500-E(config-pmap-c)# class TRANSACTIONAL-DATA-QUEUE
C4500-E(config-pmap-c)# bandwidth remaining percent 10
C4500-E(config-pmap-c)# db1
! Defines a transactional data queue with 10% BW remaining + DBL
C4500-E(config-pmap-c)# class BULK-DATA-QUEUE
C4500-E(config-pmap-c)# bandwidth remaining percent 4
C4500-E(config-pmap-c)# db1
! Defines a bulk data queue with 10% BW remaining + DBL
C4500-E(config-pmap-c)# class SCAVENGER-QUEUE
C4500-E(config-pmap-c)# bandwidth remaining percent 1
! Defines a (minimal) scavenger queue with 1% BW remaining/limit
C4500-E(config-pmap-c)# class class-default
C4500-E(config-pmap-c)# bandwidth remaining percent 25
C4500-E(config-pmap-c)# db1
! Provisions the default/Best Effort queue with 25% BW remaining + DBL

! This section attaches the egress queuing policy to the interface(s)
C4500-E(config)#interface range TenGigabitEthernet 1/1-2
C4500-E(config-if-range)# service-policy output 1P7Q1T

```

**Note**

As noted within the comments in [Example 2-86](#), unique class map names must be used for these egress queuing policies, so that logical incompatibilities—resulting in classification errors—are not introduced.

This configuration can be verified with the commands:

- **show class-map**
- **show policy-map**
- **show policy-map interface** (as shown in [Example 2-87](#))

Example 2-87 Verifying Queuing Policies on a Catalyst 4500-E Supervisor-6E—show policy-map interface

```

C4500-E#show policy-map interface TenGigabitEthernet 1/1
TenGigabitEthernet1/1

Service-policy output: 1P7Q1T

```

```
Class-map: PRIORITY-QUEUE (match-any)
  102598 packets
  Match: dscp ef (46)
    102598 packets
  Match: dscp cs5 (40)
    0 packets
  Match: dscp cs4 (32)
    0 packets
  priority queue:
    Transmit: 22782306 Bytes, Queue Full Drops: 0 Packets

Class-map: CONTROL-MGMT-QUEUE (match-any)
  24847 packets
  Match: dscp cs7 (56)
    0 packets
  Match: dscp cs6 (48)
    0 packets
  Match: dscp cs3 (24)
    24847 packets
  Match: dscp cs2 (16)
    0 packets
  bandwidth remaining 10 (%)
    Transmit: 24909844 Bytes, Queue Full Drops: 0 Packets

Class-map: MULTIMEDIA-CONFERENCING-QUEUE (match-all)
  22280511 packets
  Match: dscp af41 (34) af42 (36) af43 (38)
  bandwidth remaining 10 (%)
    Transmit: 4002626800 Bytes, Queue Full Drops: 0 Packets

Class-map: MULTIMEDIA-STREAMING-QUEUE (match-all)
  0 packets
  Match: dscp af31 (26) af32 (28) af33 (30)
  bandwidth remaining 10 (%)
    Transmit: 0 Bytes, Queue Full Drops: 0 Packets

Class-map: TRANSACTIONAL-DATA-QUEUE (match-all)
  235852 packets
  Match: dscp af21 (18) af22 (20) af23 (22)
  bandwidth remaining 10 (%)
    Transmit: 247591260 Bytes, Queue Full Drops: 0 Packets
  db1
    Probabilistic Drops: 0 Packets
    Belligerent Flow Drops: 0 Packets

Class-map: BULK-DATA-QUEUE (match-all)
  2359020 packets
  Match: dscp af11 (10) af12 (12) af13 (14)
  bandwidth remaining 4 (%)
    Transmit: 2476460700 Bytes, Queue Full Drops: 0 Packets
  db1
    Probabilistic Drops: 0 Packets
    Belligerent Flow Drops: 0 Packets

Class-map: SCAVENGER-QUEUE (match-all)
  78607323 packets
  Match: dscp cs1 (8)
  bandwidth remaining 1 (%)
    Transmit: 98144078642 Bytes, Queue Full Drops: 26268 Packets

Class-map: class-default (match-any)
  12388183 packets
  Match: any
```

```

12388183 packets
bandwidth remaining 25 (%)
  Transmit: 13001465825 Bytes, Queue Full Drops: 0 Packets
dbl
  Probabilistic Drops: 0 Packets
  Belligerent Flow Drops: 0 Packets
C4500-E#

```

[Example 2-87](#) shows various queuing classes and their associated packet and byte counts, including 26,268 queuing drops noted on the scavenger-queue.

EtherChannel QoS Models

Similar to the trust and queuing policies on the Catalyst 4500/4500-E family, there are two sets of EtherChannel QoS models: one for the Classic Supervisors and another for the Supervisor 6-E/7-E series. Each of these will be presented in turn.

Classic Supervisors EtherChannel QoS Model

For EtherChannel interfaces configured on Catalyst 4500 Classic Supervisors, the ingress QoS policies (including trust, classification, marking and/or policing policies) are configured on the logical Port-Channel interface, whereas the egress queuing and DBL policies are applied directly to the physical port-member interfaces, as shown in [Example 2-88](#).

Example 2-88 EtherChannel QoS Design on a Catalyst 4500 Classic Supervisor

```

! This section configures the logical Port-Channel Interface and sets DSCP-trust
C4500-CS(config)# interface Port-channel1
C4500-CS(config-if)# description ETHERCHANNEL-TRUNK-TO-DISTRIBUTION-LAYER
C4500-CS(config-if)# switchport mode trunk
C4500-CS(config-if)# switchport trunk encapsulation dot1q
C4500-CS(config-if)# switchport trunk allowed vlan 10,110
C4500-CS(config-if)# qos trust dscp
! The logical EtherChannel interface is set to statically trust DSCP

! This section configures 1P3Q1T+DBL Queuing on Physical Port-Member Interfaces
C4500-CS(config)# interface range TenGigabitEthernet4/1-2
C4500-CS(config-if-range)# description PORT-CHANNEL1-PHYSICAL-PORT-MEMBER
C4500-CS(config-if-range)# switchport mode trunk
C4500-CS(config-if-range)# switchport trunk encapsulation dot1q
C4500-CS(config-if-range)# switchport trunk allowed vlan 10,110
C4500-CS(config-if-range)# channel-group 1 mode auto
C4500-CS(config-if-range)# tx-queue 1
C4500-CS(config-if-tx-queue)# bandwidth percent 5
! Q1 (less than best effort queue) is assigned 5% BW
C4500-CS(config-if-tx-queue)# tx-queue 2
C4500-CS(config-if-tx-queue)# bandwidth percent 35
! Q2 (default/best effort queue) is assigned 35% BW
C4500-CS(config-if-tx-queue)# tx-queue 3
C4500-CS(config-if-tx-queue)# priority high
C4500-CS(config-if-tx-queue)# bandwidth percent 30
C4500-CS(config-if-tx-queue)# shape percent 30
! Q3 is enabled as a PQ and assigned 30% BW
! Additionally Q3 is shaped (limited) to 30%
C4500-CS(config-if-tx-queue)# tx-queue 4
C4500-CS(config-if-tx-queue)# bandwidth percent 30
! Q4 (guaranteed BW queue) is assigned 30% BW

```

```
C4500-CS(config-if-range)# service-policy output DBL
! DBL policy-map is attached to the physical port-member interfaces
```

This configuration can be verified with the commands:

- **show qos dbl**
- **show qos maps dscp tx-queue**
- **show qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Supervisor 6-E/7-E EtherChannel QoS Model

For EtherChannel interfaces configured on Catalyst 4500-E Supervisor 6-E/7-E switches, the ingress QoS policies (including classification, marking, and/or policing policies) are applied via MQC service-policy statements (in the *ingress* direction using the **input** keyword) configured on the logical Port-Channel interface. Trust statements are not required, as these supervisors implicitly trust by default. Additionally, these supervisors support egress marking and/or policing policies to be similarly applied via MQC service-policy statements (in the *egress* direction using the **output** keyword) on the logical Port-Channel interface.

Egress queueing policies, however, are applied via MQC service-policy statements (in the *egress* direction using the **output** keyword) on the physical port-member interfaces, as shown in [Example 2-89](#).

Example 2-89 EtherChannel QoS Design on a Catalyst 4500 Supervisor 6-E/7-E

```
! This section configures the logical Port-Channel Interface and sets DSCP-trust
C4500-E(config)# interface Port-channel1
C4500-E(config-if)# description ETHERCHANNEL-TRUNK-TO-DISTRIBUTION-LAYER
C4500-E(config-if)# switchport mode trunk
C4500-E(config-if)# switchport trunk encapsulation dot1q
C4500-E(config-if)# switchport trunk allowed vlan 10,110

! This section configures 1P3Q1T+DBL Queuing on Physical Port-Member Interfaces
C4500-E(config)# interface range TenGigabitEthernet1/1-2
C4500-E(config-if-range)# description PORT-CHANNEL1-PHYSICAL-PORT-MEMBER
C4500-E(config-if-range)# switchport mode trunk
C4500-E(config-if-range)# switchport trunk encapsulation dot1q
C4500-E(config-if-range)# switchport trunk allowed vlan 10,110
C4500-E(config-if-range)# channel-group 1 mode auto
C4500-E(config-if-range)# service-policy output 1P7Q1T
! Applies 1P7Q1T+DBL (MQC) queuing policy to physical port-member interfaces
```

This configuration can be verified with the commands:

- **show class-map**
- **show policy-map**
- **show policy-map interface**

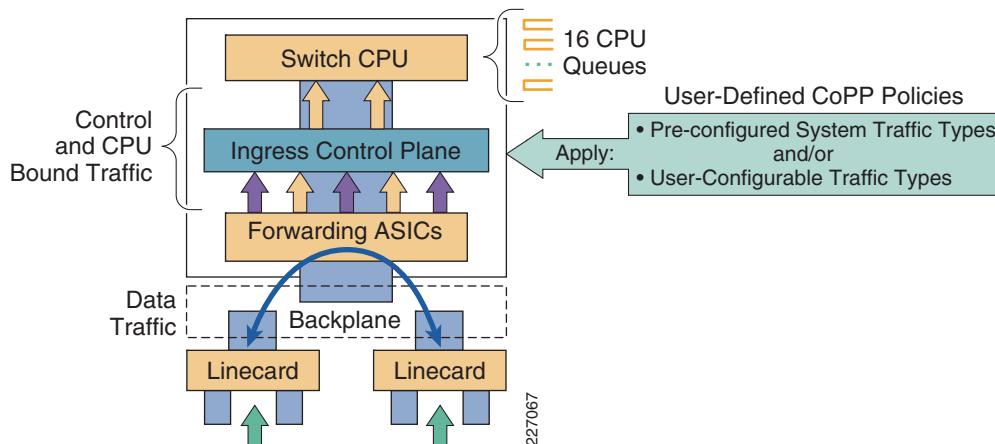
**Note**

As noted in the previous section, the queuing policies will only attach to EtherChannel port-member physical interfaces if the priority queue is not explicitly policed. If policing the priority queue is desired, then a separate policy map needs to be constructed to do so and attached to the logical EtherChannel interface in the egress direction.

Control Plane Policing

The Catalyst 4500 Series switches support CoPP on all supervisor engines compatible with Cisco IOS release 12.2(31)SG. In this platform CoPP is implemented in hardware in a centralized, non-distributed fashion. CoPP policies are centrally configured under the control plane configuration mode and then enforced in hardware by the classification TCAM and QoS policers of the supervisor engine. This CoPP model is shown in Figure 2-24.

Figure 2-24 Catalyst 4500 Control Plane Policing Implementation



CoPP Configuration

The Catalyst 4500 implementation of CoPP uses the modular QoS command line interface (MQC) to define traffic classification criteria and to specify the configurable policy actions for the classified traffic. MQC uses class maps to define packets for a particular traffic class. After you have classified the traffic, you can create policy maps to enforce policy actions for the identified traffic. The **control-plane** global configuration command allows the CoPP service policy to be directly attached to the control plane.

Additionally, Catalyst 4500 CoPP supports the definition of non-IP traffic classes in addition to IP traffic classes. With this, instead of using the default class for handling all non-IP traffic, you can define separate policies for non-IP traffic. This results in better and more granular control over non-IP protocols, such as ARP, IPX, BPDUs, CDP, and SSTP.

One particular characteristic of Catalyst 4500 CoPP is that the CoPP policy must be named **system-cpp policy**. In fact, **system-cpp-policy** is the only policy map that can be attached to the control-plane. To facilitate the configuration of the system-cpp-policy, Catalyst 4500's CoPP provides a global macro function (called **system-cpp**) that automatically generates and applies CoPP policies to the control plane. The resulting configuration uses a collection of system-defined class maps for common Layer 3

and Layer 2 control plane traffic. The names of all system-defined CoPP class maps and their matching ACLs contain the prefix **system-cpp-**. By default, no action is specified on any of the system predefined traffic classes. Table 2-4 lists the predefined system ACLs.

Table 2-4 Catalyst 4500 System Pre-Defined CoPP ACLs

Pre-defined Named ACL	Description
system-cpp-dot1x	MAC DA = 0180.C200.0003
system-cpp-lldp	MAC DA=0180.c200.000E
system-cpp-mcast-cfm	MAC DA=0100.0ccc.ccc0 - 0100.0ccc.ccc7
system-cpp-ucast-cfm	MAC DA=0100.0ccc.ccc0
system-cpp-bpdu-range	MAC DA = 0180.C200.0000 - 0180.C200.000F
system-cpp-cdp	MAC DA = 0100.0CCC.CCCC (UDLD/DTP/VTP/Pagp)
system-cpp-sstp	MAC DA = 0100.0CCC.CCCD
system-cpp-cgmp	MAC DA = 01-00-0C-DD-DD-DD
system-cpp-ospf	IP Protocol = OSPF, IPDA matches 224.0.0.0/24
system-cpp-igmp	IP Protocol = IGMP, IPDA matches 224.0.0.0/3
system-cpp-pim	IP Protocol = PIM, IPDA matches 224.0.0.0/24
system-cpp-all-systems-on-subnet	IPDA = 224.0.0.1
system-cpp-all-routers-on-subnet	IPDA = 224.0.0.2
system-cpp-ripv2	IPDA = 224.0.0.9
system-cpp-ip-mcast-linklocal	IP DA = 224.0.0.0/24
system-cpp-dhcp-cs	IP Protocol = UDP, L4SrcPort = 68, L4DstPort = 67
system-cpp-dhcp-sc	IP Protocol = UDP, L4SrcPort = 67, L4DstPort = 68
system-cpp-dhcp-ss	IP Protocol = UDP, L4SrcPort = 67, L4DstPort = 67

In addition to the predefined classes, you can configure your own class maps matching other control plane traffic. In order to take effect, these user-defined class maps need to be added to the **system-cpp-policy** policy-map.

To summarize, CoPP is enabled on Catalyst 4500 Series switches by performing these steps:

-
- Step 1** Enable QoS with the **qos** global configuration command.
 - Step 2** Run the **macro global apply system-cpp** global macro to create the system-cpp-policy policy-map and attach it to the control-plane.
 - Step 3** Optionally, define the necessary ACLs to be used to match your own traffic classes.
 - Step 4** Next, classify the control plane traffic using the **class-map** command.
 - Step 5** After the traffic is classified, configure a **policy-map** with a **police** policy action to each class, indicating whether to permit all packets, to drop all packets, or to drop packets crossing a specified rate limit for that particular class.
-

**Note**

For more information refer to the Configuring Control Plane Policing documentation for the Catalyst 4500 at:

http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/50sg/configuration/guide/cntl_pln.html.

CoPP Considerations and Restrictions

The following are important considerations and known restrictions that should be taken into account prior configuring CoPP on the Catalyst 4500:

- CoPP is not enabled unless the global QoS is enabled and police action is specified.
- Only ingress CoPP is supported, so only the **input** keyword is supported in control plane-related CLIs.
- Use the system-defined class maps for policing control plane traffic.
- ARP support is limited to gratuitous ARPs (destination MAC in the 0180.C200.0020 - 0180.C200.002F range). Broadcast ARPs are not currently supported by CoPP.
- Control plane traffic can be policed only using CoPP. Traffic cannot be policed at the input interface or VLAN even though a policy map containing the control plane traffic is accepted when the policy map is attached to an interface or VLAN.
- System-defined class maps cannot be used in policy maps for regular QoS.
- Use ACLs and class maps to identify data plane and management plane traffic that are handled by CPU. User-defined class maps should be added to the **system-cpp-policy** policy map for CoPP.
- The policy map named **system-cpp-policy** is dedicated for CoPP. When attached to the control plane, it cannot be detached.
- The default **system-cpp-policy** policy map does not define actions for the system-defined class maps, which means no policing.
- The only action supported in **system-cpp-policy** policy map is police.
- Do not use the log keyword in the CoPP policy ACLs.
- Both MAC and IP ACLs can be used to define data plane and management plane traffic classes.
- The exceeding action **policed-dscp-transmit** is not supported for CoPP.

CoPP Model

CoPP can be deployed on the Catalyst 4500 in one of two main ways:

- The global macro **macro global apply system-cpp** can be used to pre-configure CoPP access lists, class maps, and a **system-cpp-policy** policy map (with no class actions configured); this template can then be modified and tuned by the administrator to suit specific environments (this is the recommended approach for Catalyst 4500 Classic Supervisors).
- The CoPP policy can be generated manually.

In [Example 2-90](#), CoPP has been deployed manually (to keep the policy as consistent as possible between the Catalyst 4500 and 6500 examples), inline with the recommendations for CoPP class definitions and deployment models presented earlier in this chapter.

Example 2-90 Control Plane Policing Model on a Catalyst 4500-E Supervisor 6-E

```

! This section defines the access-lists for the CoPP traffic classes
C4500-E(config)# ip access-list extended COPP-BGP
C4500-E(config-ext-nacl)# remark BGP
C4500-E(config-ext-nacl)# permit tcp host 192.168.1.1 host 10.1.1.1 eq bgp
C4500-E(config-ext-nacl)# permit tcp host 192.168.1.1 eq bgp host 10.1.1.1

C4500-E(config)# ip access-list extended COPP-IGP
C4500-E(config-ext-nacl)# remark IGP (OSPF)
C4500-E(config-ext-nacl)# permit ospf any host 224.0.0.5
C4500-E(config-ext-nacl)# permit ospf any host 224.0.0.6
C4500-E(config-ext-nacl)# permit ospf any any

C4500-E(config)# ip access-list extended COPP-INTERACTIVE-MANAGEMENT
C4500-E(config-ext-nacl)# remark TACACS (return traffic)
C4500-E(config-ext-nacl)# permit tcp host 10.2.1.1 host 10.1.1.1 established
C4500-E(config-ext-nacl)# remark SSH
C4500-E(config-ext-nacl)# permit tcp 10.2.1.0 0.0.0.255 host 10.1.1.1 eq 22
C4500-E(config-ext-nacl)# remark SNMP
C4500-E(config-ext-nacl)# permit udp host 10.2.2.2 host 10.1.1.1 eq snmp
C4500-E(config-ext-nacl)# remark NTP
C4500-E(config-ext-nacl)# permit udp host 10.2.2.3 host 10.1.1.1 eq ntp

C4500-E(config)# ip access-list extended COPP-FILE-MANAGEMENT
C4500-E(config-ext-nacl)# remark (initiated) FTP (active and passive)
C4500-E(config-ext-nacl)# permit tcp 10.2.1.0 0.0.0.255 eq 21 host 10.1.1.1 gt 1023
established
C4500-E(config-ext-nacl)# permit tcp 10.2.1.0 0.0.0.255 eq 20 host 10.1.1.1 gt 1023
C4500-E(config-ext-nacl)# permit tcp 10.2.1.0 0.0.0.255 gt 1023 host 10.1.1.1 gt 1023
established
C4500-E(config-ext-nacl)# remark (initiated) TFTP
C4500-E(config-ext-nacl)# permit udp 10.2.1.0 0.0.0.255 gt 1023 host 10.1.1.1 gt 1023

C4500-E(config)# ip access-list extended COPP-MONITORING
C4500-E(config-ext-nacl)# remark PING-ECHO
C4500-E(config-ext-nacl)# permit icmp any any echo
C4500-E(config-ext-nacl)# remark PING-ECHO-REPLY
C4500-E(config-ext-nacl)# permit icmp any any echo-reply
C4500-E(config-ext-nacl)# remark TRACEROUTE
C4500-E(config-ext-nacl)# permit icmp any any ttl-exceeded
C4500-E(config-ext-nacl)# permit icmp any any port-unreachable

C4500-E(config)# ip access-list extended COPP-CRITICAL-APPLICATIONS
C4500-E(config-ext-nacl)# remark HSRP
C4500-E(config-ext-nacl)# permit ip any host 224.0.0.2
C4500-E(config-ext-nacl)# remark DHCP
C4500-E(config-ext-nacl)# permit udp host 0.0.0.0 host 255.255.255.255 eq bootps
C4500-E(config-ext-nacl)# permit udp host 10.2.2.8 eq bootps any eq bootps

C4500-E(config)# ip access-list extended COPP-UNDESIRABLE
C4500-E(config-ext-nacl)# remark UNDESIRABLE TRAFFIC
C4500-E(config-ext-nacl)# permit udp any any eq 1434

! This section defines the CoPP Policy Class-Maps
C4500-E(config)# class-map match-all COPP-BGP
C4500-E(config-cmap)# match access-group name COPP-BGP
! Associates COPP-BGP ACL with class-map

C4500-E(config)# class-map match-all COPP-IGP
C4500-E(config-cmap)# match access-group name COPP-IGP

```

```

! Associates COPP-IGP ACL with class-map

C4500-E(config)# class-map match-all COPP-INTERACTIVE-MANAGEMENT
C4500-E(config-cmap)# match access-group name COPP-INTERACTIVE-MANAGEMENT
! Associates COPP-INTERACTIVE-MANAGEMENT ACL with class-map

C4500-E(config)# class-map match-all COPP-FILE-MANAGEMENT
C4500-E(config-cmap)# match access-group name COPP-FILE-MANAGEMENT
! Associates COPP-FILE-MANAGEMENT with class-map

C4500-E(config)# class-map match-all COPP-MONITORING
C4500-E(config-cmap)# match access-group name COPP-MONITORING
! Associates COPP-MONITORING ACL with class-map

C4500-E(config)# class-map match-all COPP-CRITICAL-APPLICATIONS
C4500-E(config-cmap)# match access-group name COPP-CRITICAL-APPLICATIONS
! Associates COPP-CRITICAL-APPLICATIONS ACL with class-map

C4500-E(config)# class-map match-all COPP-UNDESIRABLE
C4500-E(config-cmap)# match access-group name COPP-UNDESIRABLE
! Associates COPP-UNDESIRABLE ACL with class-map

! This section defines the CoPP Policy
C4500-E(config-cmap)#policy-map system-cpp-policy
C4500-E(config-pmap)# class COPP-BGP
C4500-E(config-pmap-c)# police cir 4000000 bc 400000 be 400000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! Polices BGP to 4 Mbps
C4500-E(config-pmap)# class COPP-IGP
C4500-E(config-pmap-c)# police cir 300000 bc 3000 be 3000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! Polices IGP to 300 kbps
C4500-E(config-pmap)# class COPP-INTERACTIVE-MANAGEMENT
C4500-E(config-pmap-c)# police cir 500000 bc 5000 be 5000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! Polices Interactive Management to 500 kbps
C4500-E(config-pmap)# class COPP-FILE-MANAGEMENT
C4500-E(config-pmap-c)# police cir 6000000 bc 60000 be 60000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! Polices File Management to 6 Mbps
C4500-E(config-pmap)# class COPP-MONITORING
C4500-E(config-pmap-c)# police cir 900000 bc 9000 be 9000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! Polices Monitoring to 900 kbps
C4500-E(config-pmap)# class COPP-CRITICAL-APPLICATIONS
C4500-E(config-pmap-c)# police cir 900000 bc 9000 be 9000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop
! Polices Critical Applications to 900 Kbps
C4500-E(config-pmap)# class COPP-UNDESIRABLE
C4500-E(config-pmap-c)# police cir 32000 bc 3000 be 3000
C4500-E(config-pmap-c-police)# conform-action drop
C4500-E(config-pmap-c-police)# exceed-action drop
! Polices all Undesirable traffic (conform-action is drop)
C4500-E(config-pmap)# class class-default
C4500-E(config-pmap-c)# police cir 500000 bc 5000 be 5000
C4500-E(config-pmap-c-police)# conform-action transmit
C4500-E(config-pmap-c-police)# exceed-action drop

```

```
! Polices all other Control Plane traffic to 500 kbps
```

```
! This section attaches the CoPP policy to the Control Plane
C4500-E(config)#control-plane
C4500-E(config-cp)# service-policy input system-cpp-policy
! Attaches CoPP policy to control plane
```

**Note**

As previously mentioned, to apply this policy to the control plane of a Catalyst 4500 Classic Supervisor, the global macro **macro global apply system-cpp** should be added to the configuration above (prior to the definition of the system-cpp-policy policy map). Additionally, as the Catalyst 4500 Classic Supervisors only support single rate policers, the policing commands need to be adapted to a single rate syntax, as has been shown in the per-port and per-port/per-VLAN policing model examples for Classic Supervisors (see [Example 2-75](#) and [Example 2-78](#), respectively).

This configuration can be verified with the commands:

- **show class-map**
- **show policy-map**
- **show policy-map control-plane** (as shown in [Example 2-91](#))

Example 2-91 Verifying Control Plane Policing on a Catalyst 4500—show policy-map control-plane

```
C4500-E#show policy-map control-plane
Control Plane

Service-policy input: system-cpp-policy

Class-map: COPP-ACL-BGP (match-all)
  23277 packets
  Match: access-group name COPP-ACL-BGP
  police:
    cir 4000000 bps, bc 400000 bytes, be 400000 bytes
    conformed Packet count - n/a, 16854098 bytes; actions:
      transmit
    exceeded Packet count - n/a, 0 bytes; actions:
      drop
    violated Packet count - n/a, 0 bytes; actions:
      drop
    conformed 34000 bps, exceed 0 bps

Class-map: COPP-ACL-IGP (match-all)
  1135 packets
  Match: access-group name COPP-ACL-IGP
  police:
    cir 300000 bps, bc 3000 bytes, be 3000 bytes
    conformed Packet count - n/a, 87438 bytes; actions:
      transmit
    exceeded Packet count - n/a, 0 bytes; actions:
      drop
    violated Packet count - n/a, 0 bytes; actions:
      drop
    conformed 0 bps, exceed 0 bps

Class-map: COPP-ACL-INTERACTIVE-MANAGEMENT (match-all)
  1251 packets
  Match: access-group name COPP-ACL-INTERACTIVE-MANAGEMENT
  police:
```

```

        cir 500000 bps, bc 5000 bytes, be 5000 bytes
        conformed Packet count - n/a, 84014 bytes; actions:
            transmit
        exceeded Packet count - n/a, 0 bytes; actions:
            drop
        violated Packet count - n/a, 0 bytes; actions:
            drop
        conformed 0 bps, exceed 0 bps

Class-map: COPP-ACL-FILE-MANAGEMENT (match-all)
43254 packets
Match: access-group name COPP-ACL-FILE-MANAGEMENT
police:
    cir 6000000 bps, bc 60000 bytes, be 60000 bytes
    conformed Packet count - n/a, 24475620 bytes; actions:
        transmit
    exceeded Packet count - n/a, 1124 bytes; actions:
        drop
    violated Packet count - n/a, 0 bytes; actions:
        drop
    conformed 328000 bps, exceed 0 bps

Class-map: COPP-ACL-MONITORING (match-all)
283 packets
Match: access-group name COPP-ACL-MONITORING
police:
    cir 900000 bps, bc 9000 bytes, be 9000 bytes
    conformed Packet count - n/a, 21528 bytes; actions:
        transmit
    exceeded Packet count - n/a, 0 bytes; actions:
        drop
    violated Packet count - n/a, 0 bytes; actions:
        drop
    conformed 0 bps, exceed 0 bps

Class-map: COPP-ACL-CRITICAL-APPLICATIONS (match-all)
10 packets
Match: access-group name COPP-ACL-CRITICAL-APPLICATIONS
police:
    cir 900000 bps, bc 9000 bytes, be 9000 bytes
    conformed Packet count - n/a, 3973 bytes; actions:
        transmit
    exceeded Packet count - n/a, 0 bytes; actions:
        drop
    violated Packet count - n/a, 0 bytes; actions:
        drop
    conformed 0 bps, exceed 0 bps

Class-map: COPP-ACL-UNDESIRABLE (match-all)
0 packets
Match: access-group name COPP-ACL-UNDESIRABLE
police:
    cir 32000 bps, bc 3000 bytes, be 3000 bytes
    conformed Packet count - n/a, 0 bytes; actions:
        drop
    exceeded Packet count - n/a, 0 bytes; actions:
        drop
    violated Packet count - n/a, 0 bytes; actions:
        drop
    conformed 0 bps, exceed 0 bps

Class-map: class-default (match-any)
61193 packets
Match: any

```

```

61193 packets
police:
  cir 500000 bps, bc 5000 bytes, be 5000 bytes
  conformed Packet count - n/a, 646069 bytes; actions:
    transmit
  exceeded Packet count - n/a, 562 bytes; actions:
    drop
  violated Packet count - n/a, 46646 bytes; actions:
    drop
  conformed 0 bps, exceed 0 bps
C4500-E#

```

[Example 2-91](#) shows sample traffic being matched across various control plane traffic classes.

**Note**

To clear the counters on the control plane, enter the **clear control-plane *** command.

**Note**

As previously mentioned, to apply this policy to the control plane of a Catalyst 4500 Classic Supervisor, the global macro **macro global apply system-cpp** should be added to the configuration above (prior to the definition of the system-cpp-policy policy map). Additionally, as the Catalyst 4500 Classic Supervisors only support single rate policers, the policing commands need to be adapted to a single rate syntax, as has been shown in the per-port and per-port/per-VLAN policing model examples for Classic Supervisors (see [Example 2-75](#) and [Example 2-78](#), respectively).

Cisco Catalyst 6500 and 6500-E QoS Design

The Cisco Catalyst 6500/6500-E series switches represent the flagship of Cisco's switching portfolio, delivering innovative, secure, converged services throughout the campus, from the access edge wiring closet to the distribution to the core to the data center to the WAN/VPN edge. The Catalyst 6500/6500-E provides leading-edge Layer 2-Layer 7 services, including rich high availability, manageability, virtualization, security, and QoS feature sets, as well as integrated Power-over-Ethernet (PoE), allowing for maximum flexibility in virtually any role within the campus.

Catalyst 6500/6500-E switches come in various chassis, supervisor, feature-cards, and linecard combinations, as are discussed in turn.

Catalyst 6500 (regular) chassis are available in 3, 6, 9, or 13 slot combinations; namely, the 6503, 6506, 6509, and 6513, respectively. Additionally, the 6509 is available in a Network Equipment Building Standards (NEBS) design, where the network modules are presented vertically (as opposed to the standard horizontal design), as the 6509-NEB-A. Catalyst 6500 chassis provide up to 32 Gbps of bandwidth per linecard slot.

Also, Catalyst 6500 chassis are available in Enhanced models, as designated by a -E suffix (such as 6503-E) in 3, 4, 6, and 9 slot combinations; namely, the 6503-E, 6504-E, 6506-E, and the 6509-E. Additionally, the 6509-E is also available in an Enhanced Vertical (NEBS compliant) design, as the 6509-V-E. Catalyst 6500-E chassis provide up to 80 Gbps of bandwidth per linecard slot.

At the time of writing, three supervisor engine options are available for the Catalyst 6500 series switches (not including Virtual Switch Supervisor Engines):

- **Supervisor Engine 720**—The Supervisor Engine 720 is part of the Catalyst 6500's third generation suite of supervisor modules and increases slot efficiency by integrating a high performance 720 Gbps switch fabric backplane with a new routing and forwarding engine, including a third generation policy feature card (PFC3). With such an architecture, the Supervisor 720 delivers

scalable performance, achieving centralized forwarding (CEF) at 48Mpps/720Gbps, accelerated CEF at 400Mpps (peak) /720Gbps, and distributed forwarding (dCEF) 400Mpps (sustained)/720Gbps.

- **Supervisor Engine 32**—The Supervisor 32 is designed for enterprise campus LAN access switches and extends Supervisor Engine 720 level of advanced services into the access layer through the Policy Feature Card 3B (PFC3B).
- **Supervisor Engine 32-10GE (with PISA)**—The Supervisor Engine 32-10GE with Programmable Intelligent Services Accelerator (PISA) embeds multi-gigabit deep packet inspection capability into Cisco's flagship switching platform. This enables hardware acceleration of services that offers new levels of application intelligence, integrated security, and operational manageability for enterprise campus access and WAN routing. In addition to PISA technology, this product incorporates all the functionality of the Supervisor Engine 32, while supporting hardware-accelerated Stateful Application Intelligence (SAI) proactively optimized network traffic for application delivery based on more than 100 different protocols.

These supervisors, in turn, can leverage various feature cards, including the Multilayer Switch Feature Card (MSFC), which serves as the routing engine, the Policy Feature Card (PFC), which serves as the primary QoS engine, as well as various Distributed Feature Cards (DFCs), which serve to scale policies and processing. Specifically relating to QoS, the PFC sends a copy of the QoS policies to the DFC to provide local support for the QoS policies, which enables the DFCs to support the same QoS features that the PFC supports.

The QoS features supported on currently shipping PFCs and DFCs are summarized in [Table 2-5](#).

Table 2-5 QoS Features Supported on Catalyst 6500 PFCs and DFCs

Feature	PFC3A and DFC3A	PFC3B and DFC3B	PFC3BXL and DFC3BXL	PFC3C and DFC3C	PFC3CXL and DFC3CXL
Support for DFCs	Yes	Yes	Yes	Yes	Yes
Flow granularity	Source Destination	Source Destination	Source Destination	Source Destination	Source Destination
QoS ACLs	IP, MAC	IP, MAC	IP, MAC	IP, MAC	IP, MAC
DSCP transparency	Optional	Optional	Optional	Optional	Optional
Egress ToS rewrite	Optional	Optional	Optional	Optional	Optional
Policing:					
Ingress aggregate policers	Yes	Yes	Yes	Yes	Yes
Egress aggregate policers	Yes	Yes	Yes	Yes	Yes
Number of aggregate policers	1022	1022	1022	1022	1022
Microflow policers	64 rates	64 rates	64 rates	64 rates	64 rates
Number of flows per Microflow policer	64,000	110,000	240,000	110,000	240,000
Unit of measure for policer statistics	Bytes	Bytes	Bytes	Bytes	Bytes
Basis of policer operation	Layer 2 length	Layer 2 length	Layer 2 length	Layer 2 length	Layer 2 length

**Note**

When PFCs are mixed with DFCs, the switch will adopt the lowest common denominator for QoS features. For additional details, refer to:

https://www.cisco.com/en/US/products/hw/switches/ps708/products_qanda_item09186a00809a7673.shtml.

Additionally, the Catalyst 6500 linecards that meet the minimum requirements for medianet switch ports (including Gigabit Ethernet support, as well as supporting a strict priority hardware queue with at least three additional hardware queues), at the time of writing, are listed in [Table 2-6](#) (10 Gigabit Ethernet Modules), [Table 2-7](#) (Gigabit Ethernet Modules), and [Table 2-8](#) (10/100/1000 Ethernet Modules), respectively.

Table 2-6 Catalyst 6500 10 Gigabit Ethernet Modules

Part Number	Product Description
WS-X6704-10GE	4-Port 10 Gigabit Ethernet
WS-X6708-10G-3C	8-port 10 Gigabit Ethernet
WS-X6716-10T-3C	Cisco Catalyst 6500 16-Port 10 Gigabit Ethernet Copper Module with DFC3C
WS-X6716-10T-3CXL	Cisco Catalyst 6500 16-Port 10 Gigabit Ethernet Copper Module with DFC3CXL
WS-X6716-10G-3C	Cisco Catalyst 6500 16-Port 10 Gigabit Ethernet Module with DFC3C, requires X2
WS-X6716-10G-3CXL	Cisco Catalyst 6500 16-Port 10 Gigabit Ethernet Module with DFC3CXL, requires X2
WS-X6708-10G-3C	Cisco Catalyst 6500 8-Port 10 Gigabit Ethernet Module with DFC3C, requires X2
WS-X6708-10G-3CXL	Cisco Catalyst 6500 8-Port 10 Gigabit Ethernet Module with DFC3CXL, requires X2
WS-X6704-10GE	Cisco Catalyst 6500 4-Port 10 Gigabit Ethernet Module, requires XENPAK

Table 2-7 Catalyst 6500 Gigabit Ethernet Modules

Part Number	Product Description
WS-X6748-SFP	Fabric-Enabled 48-Port Small Form-Factor Pluggable (SFP)-Based Gigabit Ethernet Module
WS-X6724-SFP	Fabric-Enabled 24-Port SFP-Based Gigabit Ethernet Module

Table 2-8 Catalyst 6500 10/100/1000 Ethernet Modules

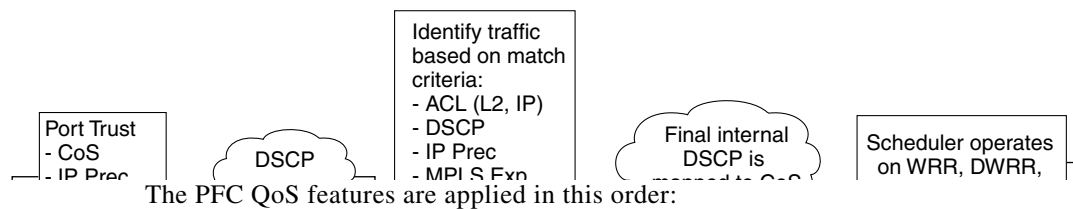
Part Number	Product Description
WS-X6748-GE-TX	Cisco Catalyst 6500 Series 48-Port 10/100/1000 RJ-45 Cisco Express Forwarding 720 Interface Module; field-upgradable to support distributed forwarding with the addition of the distributed forwarding daughter card (part number WS-F6700-DFC3A=)
WS-X6548-GE-TXWS	Cisco Catalyst 6500 Series 48-Port 10/100/1000 RJ-45 Cisco Express Forwarding 256 Interface Module; field-upgradable to support Cisco Prestandard PoE daughter card (part number WS-F6K-VPWR-GE=) or 802.3af PoE daughter card (part number WS-F6K-GE48-AF=)
WS-X6548-GE-45AF	Cisco Catalyst 6500 Series 48-Port 10/100/1000 RJ-45 Cisco Express Forwarding 256 Interface Module with 802.3af PoE daughter card (that is, includes daughter card [part number WS-F6K-GE48-AF=])
WS-X6548V-GE-TX	Cisco Catalyst 6500 Series 48-Port 10/100/1000 RJ-45 Cisco Express Forwarding 256 Interface Module with Cisco Prestandard PoE Daughter Card (that is, includes daughter card [part number WS-F6K-VPWR-GE=])

**Note**

Even though, at the time of writing, the WS-X6516A-GBIC, the WS-X6408A-GBIC, the WS-X6548-GE-TX, the WS-X6548-GE-45AF, and the WS-X6816-GBIC are currently shipping Gigabit Ethernet or 10/100/1000 Ethernet modules, these modules do not support the minimum recommended queuing structure (of 1P3QyT) for medianet campus networks (as discussed at the beginning of this chapter) and, as such, these linecards are not included in this chapter.

Platform-Specific QoS Considerations

Figure 2-25 shows the Catalyst 6500 PFC QoS model.

Figure 2-25 Catalyst 6500 PFC QoS Model

The PFC QoS features are applied in this order:

1. Ingress port PFC QoS features:
 - Port trust state—Trust CoS, IP precedence, or DSCP.
 - Layer 2 CoS remarking—PFC QoS applies Layer 2 CoS remarking, which marks the incoming frame with the port CoS value, in these situations:
 - If the traffic is not in an ISL, 802.1Q, or 802.1p frame.
 - If a port is configured as untrusted.
 - Ingress queuing and congestion avoidance—If you configure an Ethernet LAN port to trust CoS or DSCP, QoS classifies the traffic on the basis of its Layer 2 CoS value or its Layer 3 DSCP value and assigns it to an ingress queue to provide congestion avoidance.



Note Layer 3 DSCP-based queue mapping is available only on WS-X6708-10GE, WS-X6716-10GE, and Supervisor Engine 720-10GE ports.

2. PFC and DFC QoS features:

- Internal DSCP—On the PFC and DFCs, QoS associates an internal DSCP value with all traffic to classify it for processing through the system. There is an initial internal DSCP based on the traffic trust state and a final internal DSCP. The final internal DSCP can be the same as the initial value or an MQC policy map can set it to a different value.
 - MQC policy maps—MQC policy maps can do one or more of these operations:
 - Change the trust state of the traffic (bases the internal DSCP value on a different QoS label)
 - Set the initial internal DSCP value (only for traffic from untrusted ports)
 - Mark the traffic
 - Police the traffic
3. Egress Ethernet LAN port QoS features:
- Layer 3 DSCP marking with the final internal DSCP (optional)
 - Layer 2 CoS marking mapped from the final internal DSCP
 - Layer 2 CoS-based and Layer 3 DSCP-based queuing and congestion avoidance.



Note Layer 3 DSCP-based queue mapping is available only on WS-X6708-10GE, WS-X6716-10GE, and Supervisor Engine 720-10GE ports.

The buffering, ingress, and egress queuing structure details for Catalyst 6500/6500-E Supervisor Engines, Gigabit and 10/100/1000 modules, and 10 Gigabit Ethernet modules that meet the minimum queuing requirements for medianet campus networks are summarized in [Table 2-9](#), [Table 2-10](#), and [Table 2-11](#), respectively.

Table 2-9 Catalyst 6500 Supervisor Engine Module Queue Structures

Supervisor Engines	Ingress Queue and Drop Thresholds	Ingress Queue Scheduler	Egress Queue and Drop Thresholds	Egress Queue Scheduler	Total Buffer Size	Ingress Buffer Size	Egress Buffer Size
WS-SUP720	1P1Q4T	—	1P2Q2T	Weighted Round Robin (WRR)	512 KB	73 KB	439 KB
WS-SUP720-3B							
WS-SUP720-3BXL							
WS-SUP32-10GE	2Q8T	WRR	1P3Q8T	Distributed WRR (DWRR) Or Shaped Round Robin (SRR)			
10-Gigabit Ethernet ports					193 MB	105 MB	88 MB
Gigabit Ethernet port					17.7 MB	9.6 MB	8.1 MB
WS-SUP32-GE					17.7 MB	9.6 MB	8.1 MB



Note

As the ports on the Supervisor 720 only support a 1P2Q2T queuing structure, and as the minimum recommended queuing structure for medianet campus networks is 1P3QyT, it is recommended to use alternate ports, whenever possible.


Note

To disable the Supervisor Engine 720-10GE Gigabit Ethernet ports, enter **shutdown** interface configuration mode commands for the Supervisor Engine 720-10GE Gigabit Ethernet ports and then enter the **mls qos 10g-only** global configuration command, which disables the Gigabit Ethernet ports on the Supervisor Engine 720-10GE.

Table 2-10 Catalyst 6500 Gigabit Ethernet and 10/100/1000 Ethernet Module Queue Structures

Modules	Ingress Queue and Drop Thresholds	Ingress Queue Scheduler	Egress Queue and Drop Thresholds	Egress Queue Scheduler	Total Buffer Size	Ingress Buffer Size	Egress Buffer Size
WS-X6748-GE-TX with DFC3	2Q8T	WRR	1P3Q8T	DWRR	1.3 MB	166 KB	1.2 MB
WS-X6748-GE-TX with CFC	1Q8T	—					
WS-X6748-SFP with DFC3	2Q8T	WRR					
WS-X6748-SFP with CFC	1Q8T	—					
WS-X6724-SFP with DFC3	2Q8T	WRR					
WS-X6724-SFP with CFC	1Q8T	—					
WS-X6148A-GE-TX	1Q2T	WRR	1P3Q8T	WRR	5.5 MB	120 KB	5.4 MB
WS-X6148A-GE-45AF							

Table 2-11 10 Catalyst 6500 10 Gigabit Ethernet Modules Queue Structures

Modules	Ingress Queue and Drop Thresholds	Ingress Queue Scheduler	Egress Queue and Drop Thresholds	Egress Queue Scheduler	Total Buffer Size	Ingress Buffer Size	Egress Buffer Size
WS-X6716-10 GE							
Performance mode	8Q4T	DWRR	1P7Q4T	DWRR SRR	198 MB	108 MB per port	90 MB per port
Oversubscription mode	1P7Q2T				91 MB	90 MB per port	1 MB per port group
WS-X6708-10 GE	8Q4T	DWRR	1P7Q4T	DWRR SRR	198 MB	108 MB	90 MB
WS-X6704-10 GE with DFC3	8Q8T	WRR	1P7Q8T	DWRR	16 MB	2 MB	14 MB
WS-X6704-10 GE with CFC	1Q8T	—					

**Note**

At the time of writing, only the WS-X6708-10GE, WS-X6716-10GE, and Supervisor Engine 720-10GE ports support DSCP-to-queue mapping. All other Supervisor and Ethernet switch module ports for the Catalyst 6500/6500-E family support CoS-to-queue mapping (only).

Enabling QoS

QoS must be enabled globally on the Catalyst 6500-E series switches. This is a critical first step to deploying QoS on these platforms. If this small—but important—step is overlooked, it can lead to frustration in troubleshooting QoS problems because the switch software accepts QoS commands and even displays these within the switch configuration, but none of the QoS commands are active until the **mls qos** global command is enabled, as shown in [Example 2-92](#).

**Note**

To reduce wordiness, the Catalyst 6500 and 6500-E series switches are collectively referred to as Catalyst 6500-E in this chapter, unless otherwise noted.

Example 2-92 Enabling QoS on a Catalyst 6500-E

```
C6500-E(config)#mls qos
C6500-E(config)#
```

This configuration can be verified with the command:

- **show mls qos** (as shown in [Example 2-93](#))

Example 2-93 Verifying Global QoS on a Catalyst 6500-E—show mls qos

```
C6500-E#show mls qos
QoS is enabled globally
Policy marking depends on port_trust
QoS ip packet dscp rewrite enabled globally
Input mode for GRE Tunnel is Pipe mode
Input mode for MPLS is Pipe mode
Vlan or Portchannel(Multi-Earl) policies supported: Yes
Egress policies supported: Yes

----- Module [1] -----
QoS global counters:
  Total packets: 2743
  IP shortcut packets: 1117
  Packets dropped by policing: 0
  IP packets with TOS changed by policing: 106
  IP packets with COS changed by policing: 2
  Non-IP packets with COS changed by policing: 0
  MPLS packets with EXP changed by policing: 0
C6500-E#
```

Trust Models

The Catalyst 6500-E switch ports can be configured to statically trust CoS, DSCP, and IP Precedence (although this is considered to be relegated by DSCP-trust) or to dynamically and conditionally trust Cisco IP phones. By default, with QoS enabled, all ports are set to an untrusted state.

Trust-CoS Model

A Catalyst 6500-E switch port can be configured to trust CoS by configuring the interface with the **mls qos trust cos** command. However, if an interface is set to trust CoS, then it by default calculates a packet's internal DSCP to be the incoming packet's (CoS value * 8). While this may be suitable for most markings, this default mapping may not be suitable for VoIP, as VoIP is usually marked CoS 5, which would map by default to DSCP 40 (and not 46, which is the EF PHB as defined by RFC 3246). Therefore, if an interface is set to trust CoS, then the default CoS-to-DSCP mapping table should be modified such that CoS 5 maps to DSCP 46, as shown in [Example 2-94](#).

Example 2-94 Configuring Trust CoS and CoS-to-DSCP Mapping Modification on a Catalyst 6500-E

```
C6500-E(config)#mls qos map cos-dscp 0 8 16 24 32 46 48 56
! CoS 5 (the sixth CoS value, starting from 0) is mapped to 46
C6500-E(config)#interface GigabitEthernet 2/1
C6500-E(config-if)#mls qos trust cos
! The interface is set to statically trust CoS
```

This configuration can be verified with the commands:

- **show mls qos**
- **show mls qos map cos-dscp** (as shown in [Example 2-95](#))
- **show mls qos module** (as shown in [Example 2-96](#))

Example 2-95 Verifying Global CoS-to-DSCP Mapping Modifications on a Catalyst 6500-E—show mls qos map cos-dscp

```
C6500-E#show mls qos map cos-dscp
Cos-dscp map:
    cos:   0   1   2   3   4   5   6   7
-----
    dscp:  0  8 16 24 32 46 48 56

C6500-E#
```

In [Example 2-95](#), the CoS-to-DSCP mapping value for CoS 5 has been modified from the default mapping of 40 (CoS 5 * 8) to 46 (to match the recommendation from RFC 3246 that realtime applications be marked DSCP 46/EF).

Example 2-96 Verifying Interface Trust Settings on a Catalyst 6500-E—show mls qos module

```
C6500-E#show mls qos module 2
QoS is enabled globally
Policy marking depends on port_trust
QoS ip packet dscp rewrite enabled globally
Input mode for GRE Tunnel is Pipe mode
Input mode for MPLS is Pipe mode
QoS Trust state is CoS on the following interface:
Gi2/1
Vlan or Portchannel(Multi-Earl) policies supported: Yes
Egress policies supported: Yes

No forwarding engine in module [2]
C6500-E#
```

Unlike the previously discussed switch platforms, the Catalyst 6500-E does not support the **show mls qos interface** verification command, but rather uses a **show mls qos module** command, as shown in [Example 2-96](#), wherein the port trust mode for interface Gigabit 2/1 is verified to be set to trust CoS.

Trust-DSCP Model

Because of the additional granularity of DSCP versus QoS markings, it is generally recommended to trust DSCP rather than CoS (everything else being equal). A Catalyst 6500-E switch port can be configured to trust DSCP with the **mls qos trust dscp** interface command, as shown in [Example 2-97](#).

Example 2-97 Configuring Trust-DSCP on a Catalyst 6500-E

```
C6500-E(config)#interface GigabitEthernet 2/1
C6500-E(config-if)#mls qos trust dscp
! The interface is set to statically trust DSCP
```

This configuration can be verified with the commands:

- **show mls qos**
- **show mls qos module**

Conditional-Trust Model

Beginning with IOS Release 12.2(33)SX11, the Catalyst 6500-E family supports dynamic, conditional trust with the **mls qos trust device** interface command, which can be configured with the **cisco-phone** keyword to extend trust to Cisco IP phones, after these have been verified via a CDP-negotiation. Additionally, the type of trust to be extended must be specified (either CoS or DSCP). When configuring conditional trust to Cisco IP Phones, it is recommended to dynamically extend CoS-Trust, as Cisco IP Phones can only remark PC QoS markings at Layer 2 (CoS) and not at Layer 3 (DSCP). For other endpoints that do not have this remarking limitation, it is recommended to dynamically extend DSCP-trust (over CoS-trust). An example of a dynamic, conditional trust policy that is set to extend CoS-trust to CDP-verified Cisco IP phones is shown in [Example 2-98](#).

Example 2-98 Configuring (CoS-mode) Conditional Trust on a Catalyst 6500-E

```
C6500-E(config)#mls qos map cos-dscp 0 8 16 24 32 46 48 56
! CoS 5 (the sixth CoS value, starting from 0) is mapped to 46
C6500-E(config)#interface GigabitEthernet2/1
C6500-E(config-if)# switchport
C6500-E(config-if)# switchport access vlan 10
C6500-E(config-if)# switchport voice vlan 110
C6500-E(config-if)# spanning-tree portfast edge
C6500-E(config-if)# mls qos trust device cisco-phone
! The interface is set to conditionally-trust Cisco IP Phones
C6500-E(config-if)# mls qos trust cos
! CoS-trust will be dynamically extended to Cisco IP Phones
```

This configuration can be verified with the commands:

- **show mls qos**
- **show mls qos module**
- **show queueing interface** (as shown in [Example 2-99](#))

Example 2-99 Verifying Interface Trust Settings on a Catalyst 6500-E—show queueing interface

```

C6500-E#show queueing interface GigabitEthernet 2/1
Interface GigabitEthernet2/1 queueing strategy:  Weighted Round-Robin
  Port QoS is enabled
Trust boundary enabled

  Trust state: trust CoS
  Extend trust state: not trusted [COS = 0]
  Default COS is 0
...
[Additional output omitted for brevity.]

```

In [Example 2-99](#), the trust boundary/conditional trust feature has been enabled and the current (dynamic) trust state is shown as trust CoS. This is because there is a Cisco IP phone currently connected to the switch port; if the IP phone is removed from this switch port, the trust state toggles to “Port is untrusted”.

Marking Models

The Catalyst 6500 family of switches supports two main marking models:

- Per-port marking model
- Per-VLAN marking model

Additionally, classification for each policy model may be performed by using access lists or (on the Supervisor Engine 32 with PISA) with Network Based Application Recognition. Each marking model is detailed in the following sections, with the per-port marking model showing both classification options.

Per-Port Marking Model (Access-List Based Classification)

The access list-based per-port marking model (based on [Figure 2-10](#)) matches VoIP and signaling traffic from the VVLAN by matching on DSCP EF and CS3, respectively. Multimedia conferencing traffic from the DVLAN is matched by UDP/RTP ports 16384-32767. Signaling traffic is matched on SCCP ports (TCP 2000-2002), as well as on SIP ports (TCP/UDP 5060-5061). Other transactional data traffic, bulk data, and scavenger traffic are matched on various ports (outlined in [Figure 2-9](#)). The service policy is applied to an interface range, along with (DSCP-mode) conditional trust, as shown in [Example 2-100](#).

Example 2-100 Example 2-62 (ACL-Based) Per-Port Marking Configuration Example on a Catalyst 6500-E

```

! This first section configures IP access-lists to match applications
C6500-E(config)#ip access-list extended MULTIMEDIA-CONFERENCING
C6500-E(config-ext-nacl)# remark RTP
C6500-E(config-ext-nacl)# permit udp any any range 16384 32767

C6500-E(config)#ip access-list extended SIGNALING
C6500-E(config-ext-nacl)# remark SCCP
C6500-E(config-ext-nacl)# permit tcp any any range 2000 2002
C6500-E(config-ext-nacl)# remark SIP
C6500-E(config-ext-nacl)# permit tcp any any range 5060 5061
C6500-E(config-ext-nacl)# permit udp any any range 5060 5061

C6500-E(config)#ip access-list extended TRANSACTIONAL-DATA
C6500-E(config-ext-nacl)# remark HTTPS
C6500-E(config-ext-nacl)# permit tcp any any eq 443
C6500-E(config-ext-nacl)# remark ORACLE-SQL*NET
C6500-E(config-ext-nacl)# permit tcp any any eq 1521

```

```

C6500-E(config-ext-nacl)# permit udp any any eq 1521
C6500-E(config-ext-nacl)# remark ORACLE
C6500-E(config-ext-nacl)# permit tcp any any eq 1526
C6500-E(config-ext-nacl)# permit udp any any eq 1526
C6500-E(config-ext-nacl)# permit tcp any any eq 1575
C6500-E(config-ext-nacl)# permit udp any any eq 1575
C6500-E(config-ext-nacl)# permit tcp any any eq 1630
C6500-E(config-ext-nacl)# permit udp any any eq 1526

C6500-E(config)#ip access-list extended BULK-DATA
C6500-E(config-ext-nacl)# remark FTP
C6500-E(config-ext-nacl)# permit tcp any any eq ftp
C6500-E(config-ext-nacl)# permit tcp any any eq ftp-data
C6500-E(config-ext-nacl)# remark SSH/SFTP
C6500-E(config-ext-nacl)# permit tcp any any eq 22
C6500-E(config-ext-nacl)# remark SMTP/SECURE SMTP
C6500-E(config-ext-nacl)# permit tcp any any eq smtp
C6500-E(config-ext-nacl)# permit tcp any any eq 465
C6500-E(config-ext-nacl)# remark IMAP/SECURE IMAP
C6500-E(config-ext-nacl)# permit tcp any any eq 143
C6500-E(config-ext-nacl)# permit tcp any any eq 993
C6500-E(config-ext-nacl)# remark POP3/SECURE POP3
C6500-E(config-ext-nacl)# permit tcp any any eq pop3
C6500-E(config-ext-nacl)# permit tcp any any eq 995
C6500-E(config-ext-nacl)# remark CONNECTED PC BACKUP
C6500-E(config-ext-nacl)# permit tcp any eq 1914 any

C6500-E(config)#ip access-list extended SCAVENGER
C6500-E(config-ext-nacl)# remark KAZAA
C6500-E(config-ext-nacl)# permit tcp any any eq 1214
C6500-E(config-ext-nacl)# permit udp any any eq 1214
C6500-E(config-ext-nacl)# remark MICROSOFT DIRECT X GAMING
C6500-E(config-ext-nacl)# permit tcp any any range 2300 2400
C6500-E(config-ext-nacl)# permit udp any any range 2300 2400
C6500-E(config-ext-nacl)# remark APPLE ITUNES MUSIC SHARING
C6500-E(config-ext-nacl)# permit tcp any any eq 3689
C6500-E(config-ext-nacl)# permit udp any any eq 3689
C6500-E(config-ext-nacl)# remark BITTORRENT
C6500-E(config-ext-nacl)# permit tcp any any range 6881 6999
C6500-E(config-ext-nacl)# remark YAHOO GAMES
C6500-E(config-ext-nacl)# permit tcp any any eq 11999
C6500-E(config-ext-nacl)# remark MSN GAMING ZONE
C6500-E(config-ext-nacl)# permit tcp any any range 28800 29100

! This section configures the class-maps
C6500-E(config-cmap)# class-map match-all VVLAN-VOIP
C6500-E(config-cmap)# match dscp ef
! VoIP is trusted (from the VVLAN)

C6500-E(config-cmap)# class-map match-all VVLAN-SIGNALING
C6500-E(config-cmap)# match dscp cs3
! Signaling is trusted (from the VVLAN)

C6500-E(config-cmap)# class-map match-all MULTIMEDIA-CONFERENCING
C6500-E(config-cmap)# match access-group name MULTIMEDIA-CONFERENCING
! Associates MULTIMEDIA-CONFERENCING access-list with class-map

C6500-E(config-cmap)# class-map match-all SIGNALING
C6500-E(config-cmap)# match access-group name SIGNALING
! Associates SIGNALING access-list with class-map

C6500-E(config-cmap)# class-map match-all TRANSACTIONAL-DATA
C6500-E(config-cmap)# match access-group name TRANSACTIONAL-DATA

```

```

! Associates TRANSACTIONAL-DATA access-list with class-map

C6500-E(config-cmap)# class-map match-all BULK-DATA
C6500-E(config-cmap)# match access-group name BULK-DATA
! Associates BULK-DATA access-list with class-map

C6500-E(config-cmap)# class-map match-all SCAVENGER
C6500-E(config-cmap)# match access-group name SCAVENGER
! Associates SCAVENGER access-list with class-map

! This section configures the Per-Port ingress marking policy-map
C6500-E(config)# policy-map PER-PORT-MARKING
C6500-E(config-pmap)# class VVLAN-VOIP
C6500-E(config-pmap-c)# set dscp ef
! VoIP is marked EF
C6500-E(config-pmap-c)# class VVLAN-SIGNALING
C6500-E(config-pmap-c)# set dscp cs3
! Signaling (from the VVLAN) is marked CS3
C6500-E(config-pmap-c)# class MULTIMEDIA-CONFERENCING
C6500-E(config-pmap-c)# set dscp af41
! Multimedia-conferencing is marked AF41
C6500-E(config-pmap-c)# class SIGNALING
C6500-E(config-pmap-c)# set dscp cs3
! Signaling (from the DVLAN) is marked CS3
C6500-E(config-pmap-c)# class TRANSACTIONAL-DATA
C6500-E(config-pmap-c)# set dscp af21
! Transactional Data is marked AF21
C6500-E(config-pmap-c)# class BULK-DATA
C6500-E(config-pmap-c)# set dscp af11
! Bulk Data is marked AF11
C6500-E(config-pmap-c)# class SCAVENGER
C6500-E(config-pmap-c)# set dscp cs1
! Scavenger traffic is marked CS1
C6500-E(config-pmap-c)# class class-default
C6500-E(config-pmap-c)# set dscp default
! An implicit class-default marks all other traffic to DF

! This section attaches the service-policy to the interface(s)
C6500-E(config)#interface range GigabitEthernet 2/1-48
C6500-E(config-if-range)# switchport
C6500-E(config-if-range)# switchport access vlan 10
C6500-E(config-if-range)# switchport voice vlan 110
C6500-E(config-if-range)# spanning-tree portfast edge
C6500-E(config-if-range)# service-policy input PER-PORT-MARKING
! Attaches the Per-Port Marking policy to the interface(s)

```

**Note**

Access-lists—along with other policy elements—consume Ternary Content Addressable Memory (TCAM) resources on the Catalyst 6500 platform. The **show platform hardware capacity forwarding** monitoring command can be used to ensure that TCAM resources are being managed effectively.

**Note**

The **mls qos trust** interface commands are not functionally compatible in conjunction with a **service-policy** interface command on the Catalyst 6500-E and thus should not be used in conjunction with them.

This configuration can be verified with the commands:

- **show mls qos**
- **show mls qos module**
- **show queueing interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface** (as shown in [Example 2-101](#))

Example 2-101 Verifying Service Policies on a Catalyst 6500-E—show policy-map interface

```
C6500-E#show policy-map interface GigabitEthernet 2/1
```

```
GigabitEthernet2/1
```

```
Service-policy input: PER-PORT-MARKING
```

```
class-map: VVLAN-VOIP (match-all)
  Match: dscp ef (46)
  set dscp 46:
  Earl in slot 1 :
    9029996 bytes
    5 minute offered rate 82704 bps
    aggregate-forwarded 9029996 bytes
```

```
class-map: VVLAN-SIGNALING (match-all)
  Match: dscp cs3 (24)
  set dscp 24:
  Earl in slot 1 :
    40940 bytes
    5 minute offered rate 360 bps
    aggregate-forwarded 40940 bytes
```

```
class-map: MULTIMEDIA-CONFERENCING (match-all)
  Match: access-group name MULTIMEDIA-CONFERENCING
  set dscp 34:
  Earl in slot 1 :
    53286740 bytes
    5 minute offered rate 792024 bps
    aggregate-forwarded 53286740 bytes
```

```
class-map: SIGNALING (match-all)
  Match: access-group name SIGNALING
  set dscp 24:
  Earl in slot 1 :
    5315494 bytes
    5 minute offered rate 79456 bps
    aggregate-forwarded 5315494 bytes
```

```
class-map: TRANSACTIONAL-DATA (match-all)
  Match: access-group name TRANSACTIONAL-DATA
  set dscp 18:
  Earl in slot 1 :
    105965882 bytes
    5 minute offered rate 1579600 bps
    aggregate-forwarded 105965882 bytes
```

```
class-map: BULK-DATA (match-all)
  Match: access-group name BULK-DATA
  set dscp 10:
  Earl in slot 1 :
```

```

528425740 bytes
5 minute offered rate 7886976 bps
aggregate-forwarded 528425740 bytes

class-map: SCAVENGER (match-all)
  Match: access-group name SCAVENGER
  set dscp 8:
  Earl in slot 1 :
    526573950 bytes
    5 minute offered rate 7873888 bps
    aggregate-forwarded 526573950 bytes

class-map: class-default (match-any)
  Match: any
  set dscp 0:
  Earl in slot 1 :
    745010286 bytes
    5 minute offered rate 11148752 bps
    aggregate-forwarded 745010286 bytes
C6500-E#

```

[Example 2-101](#) shows that the **show policy-map interface** command on the Catalyst 6500-E dynamically increments counters. However, it should be noted that these are slightly delayed and seem to increment only every 10-15 seconds.

Per-Port Marking Model (NBAR-based Classification)

The NBAR-based per-port marking model matches traffic based on NBAR Packet Description Language Module (PDL) keywords used in conjunction with the **match protocol** class map command.

Specifically, VoIP (from Cisco IP phones or Cisco IP Communicators) can be matched with the **cisco-phone** keyword. Multimedia conferencing can be matched with the **rtp** keyword. Signaling traffic can be matched by the **h323**, **sip**, and **skinny** keywords for H.323, Session Initiation Protocol (SIP), and Skinny Call Control Protocol (SCCP/Skinny) protocols, respectively. Transactional data can be matched (among other options) by the **notes**, **ora-srv**, **sap**, **secure-http**, and **sqlnet** keywords for Lotus Notes, Oracle, SAP, and SQL*NET, respectively. Bulk data can be matched (among other options) by the **exchange**, **ftp**, **secure-ftp**, **imap**, **secure-imap**, **pop3**, **secure-pop3**, and **smtp** keywords for Microsoft Exchange, FTP/S-FTP, IMAP/S-IMAP, POP3/S-POP3, and SMTP, respectively. Finally, scavenger traffic can be matched (among other options) by the **bittorrent**, **blizwow**, **doom**, **fasttrack**, **gnutella**, **kazaa2**, **streamwork**, and **youtube** keywords for BitTorrent, World of Warcraft Gaming Protocol, Doom, FastTrack traffic (KaZaA, Morpheus, Grokster, etc.), Gnutella Version2 traffic (BearShare, Shareaza, Morpheus, etc.), Kazaa Version 2 traffic, StreamWorks player traffic, and YouTube traffic, respectively.

It should be noted that NBAR PDL matching can be used in conjunction with any other type of matching criteria, including DSCP values and access lists. Additionally, it is good to keep in mind that the **match-any** operator keyword should be used when defining a class map with multiple (mutually exclusive) match statements (such as multiple NBAR protocols), otherwise the classification logic fails.

An NBAR-based per-port marking model is shown in [Example 2-102](#).

Example 2-102 (NBAR-Based) Per-Port Marking Configuration Example on a Catalyst 6500-E Supervisor Engine 32 with PISA

```

! This section configures the NBAR-based class-maps
C6500-E-SUP32-PISA(config)#class-map match-any VOIP
C6500-E-SUP32-PISA(config-cmap)# match protocol cisco-phone
! Matches Cisco IP Phones and PC-based Unified Communicators

```

```

C6500-E-SUP32-PISA(config-cmap)#class-map match-any MULTIMEDIA-CONFERENCING
C6500-E-SUP32-PISA(config-cmap)# match protocol rtp
! Matches Real Time Protocol

C6500-E-SUP32-PISA(config-cmap)#class-map match-any SIGNALING
C6500-E-SUP32-PISA(config-cmap)# match protocol h323
! Matches H323 Protocol
C6500-E-SUP32-PISA(config-cmap)# match protocol skinny
! Matches Skinny Call Control Protocol
C6500-E-SUP32-PISA(config-cmap)# match protocol sip
! Matches Session Initiation Protocol

C6500-E-SUP32-PISA(config-cmap)#class-map match-any TRANSACTIONAL-DATA
C6500-E-SUP32-PISA(config-cmap)# match protocol notes
! Matches Lotus Notes
C6500-E-SUP32-PISA(config-cmap)# match protocol ora-srv
! Matches Oracle TCP/IP Listener
C6500-E-SUP32-PISA(config-cmap)# match protocol sap
! Matches SAP Systems Applications Product in Data processing
C6500-E-SUP32-PISA(config-cmap)# match protocol secure-http
! Matches Secured HTTP
C6500-E-SUP32-PISA(config-cmap)# match protocol sqlnet
! Matches SQL*NET for Oracle

C6500-E-SUP32-PISA(config-cmap)#class-map match-any BULK-DATA
C6500-E-SUP32-PISA(config-cmap)# match protocol exchange
! Matches MS-RPC for Exchange
C6500-E-SUP32-PISA(config-cmap)# match protocol ftp
! Matches File Transfer Protocol
C6500-E-SUP32-PISA(config-cmap)# match protocol imap
! Matches Internet Message Access Protocol
C6500-E-SUP32-PISA(config-cmap)# match protocol pop3
! Matches Post Office Protocol
C6500-E-SUP32-PISA(config-cmap)# match protocol secure-ftp
! Matches FTP over TLS/SSL
C6500-E-SUP32-PISA(config-cmap)# match protocol secure-imap
! Matches Internet Message Access Protocol over TLS/SSL
C6500-E-SUP32-PISA(config-cmap)# match protocol secure-pop3
! Matches Post Office Protocol over TLS/SSL
C6500-E-SUP32-PISA(config-cmap)# match protocol smtp
! Matches Simple Mail Transfer Protocol

C6500-E-SUP32-PISA(config-cmap)#class-map match-any SCAVENGER
C6500-E-SUP32-PISA(config-cmap)# match protocol bittorrent
! Matches BitTorrent
C6500-E-SUP32-PISA(config-cmap)# match protocol blizwow
! Matches World of Warcraft Gaming Protocol
C6500-E-SUP32-PISA(config-cmap)# match protocol doom
! Matches Doom Id Software
C6500-E-SUP32-PISA(config-cmap)# match protocol fasttrack
! Matches FastTrack Traffic - KaZaA, Morpheus, Grokster
C6500-E-SUP32-PISA(config-cmap)# match protocol gnutella
! Matches Gnutella Version2 Traffic - BearShare, Shareaza, Morpheus
C6500-E-SUP32-PISA(config-cmap)# match protocol kazaa2
! Matches Kazaa Version 2
C6500-E-SUP32-PISA(config-cmap)# match protocol streamwork
! Matches Xing Technology StreamWorks player
C6500-E-SUP32-PISA(config-cmap)# match protocol youtube
! Matches Youtube streams

! This section configures the NBAR-based Per-Port Marking policy-map
C6500-E-SUP32-PISA(config-pmap)#policy-map PER-PORT-NBAR-MARKING
C6500-E-SUP32-PISA(config-pmap-c)# class VOIP

```

```

C6500-E-SUP32-PISA(config-pmap-c)# set dscp ef
! VoIP is marked EF
C6500-E-SUP32-PISA(config-pmap-c)# class MULTIMEDIA-CONFERENCING
C6500-E-SUP32-PISA(config-pmap-c)# set dscp af41
! Multimedia-conferencing is marked AF41
C6500-E-SUP32-PISA(config-pmap-c)# class SIGNALING
C6500-E-SUP32-PISA(config-pmap-c)# set dscp cs3
! Signaling (from either VLAN) is marked CS3
C6500-E-SUP32-PISA(config-pmap-c)# class TRANSACTIONAL-DATA
C6500-E-SUP32-PISA(config-pmap-c)# set dscp af21
! Transactional Data is marked AF21
C6500-E-SUP32-PISA(config-pmap-c)# class BULK-DATA
C6500-E-SUP32-PISA(config-pmap-c)# set dscp af11
! Bulk Data is marked AF11
C6500-E-SUP32-PISA(config-pmap-c)# class SCAVENGER
C6500-E-SUP32-PISA(config-pmap-c)# set dscp cs1
! Scavenger traffic is marked CS1
C6500-E-SUP32-PISA(config-pmap-c)# class class-default
C6500-E-SUP32-PISA(config-pmap-c)# set dscp default
! An implicit class-default marks all other traffic to DF

! This section attaches the service-policy to the interface(s)
C6500-E-SUP32-PISA(config)#interface range GigabitEthernet 2/1-48
C6500-E-SUP32-PISA(config-if-range)# switchport
C6500-E-SUP32-PISA(config-if-range)# switchport access vlan 10
C6500-E-SUP32-PISA(config-if-range)# switchport voice vlan 110
C6500-E-SUP32-PISA(config-if-range)# spanning-tree portfast
C6500-E-SUP32-PISA(config-if-range)# access-group mode prefer port
C6500-E-SUP32-PISA(config-if-range)# no mls qos trust
! Port is set to an untrusted state
C6500-E-SUP32-PISA(config-if-range)# service-policy input PER-PORT-NBAR-MARKING
! Attaches the NBAR-based Per-Port Marking policy to the interface(s)

```

**Note**

The **mls qos trust** interface commands are not functionally compatible in conjunction with a **service-policy** interface command on the Catalyst 6500-E, and thus should not be used in conjunction with them.

**Note**

The number of filters in any given class map is limited to eight on the Catalyst 6500 Supervisor Engine 32. Therefore, no more than eight PDLMs can be used to match an application class. However, access lists can be used in conjunction with PDLMs, as applicable, to increase the number of applications matched by a given class map.

This configuration can be verified with the commands:

- **show mls qos**
- **show mls qos module**
- **show queueing interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Per-VLAN Marking Model

An alternative approach for deploying marking policies on the Catalyst 6500-E is to deploy these on a per-VLAN basis. In order to do so, the interfaces belonging to the VLANs need to be configured with the **mls qos vlan-based** interface command. Additionally, the policy map can be simplified/broken-apart, as applicable to each VLAN. Adapting the previous example to a VLAN-based marking policing allows for the VVLAN-based policy map to be reduced to only three classes, VoIP, signaling, and a default class. Similarly, the DVLAN-based policy map is reduced to six classes, multimedia conferencing, signaling, transactional data, bulk data, scavenger, and a default class. A per-VLAN marking model is shown in [Example 2-103](#).



Note

As the access lists and class maps are identical to the previous examples, these are omitted for brevity in this—and in following—examples for this switch platform family.

Example 2-103 Per-VLAN Marking Configuration Example on a Catalyst 6500-E

```

! This section configures the ingress marking policy-map for the VVLAN
C6500-E(config)#policy-map VVLAN-MARKING
C6500-E(config-pmap-c)# class VVLAN-VOIP
C6500-E(config-pmap-c)# set dscp ef
! VoIP is trusted (from the VVLAN)
C6500-E(config-pmap-c)# class VVLAN-SIGNALING
C6500-E(config-pmap-c)# set dscp cs3
! Signaling is trusted (from the VVLAN)
C6500-E(config-pmap-c)# class class-default
C6500-E(config-pmap-c)# set dscp default
! The implicit default class marks all other VVLAN IP traffic to DF

! This section configures the ingress marking policy-map for the DVLAN
C6500-E(config)#policy-map DVLAN-MARKING
C6500-E(config-pmap-c)# class MULTIMEDIA-CONFERENCING
C6500-E(config-pmap-c)# set dscp af41
! Multimedia-conferencing is marked AF41
C6500-E(config-pmap-c)# class SIGNALING
C6500-E(config-pmap-c)# set dscp cs3
! Signaling (from the DVLAN) is marked CS3
C6500-E(config-pmap-c)# class TRANSACTIONAL-DATA
C6500-E(config-pmap-c)# set dscp af21
! Transactional Data is marked AF21
C6500-E(config-pmap-c)# class BULK-DATA
C6500-E(config-pmap-c)# set dscp af11
! Bulk Data is marked AF11
C6500-E(config-pmap-c)# class SCAVENGER
C6500-E(config-pmap-c)# set dscp cs1
! Scavenger traffic is marked CS1
C6500-E(config-pmap-c)# class class-default
C6500-E(config-pmap-c)# set dscp default
! The implicit default class marks all other DVLAN IP traffic to DF

! This section configures the interface(s) for VLAN-based QoS
C6500-E(config)#interface range GigabitEthernet 2/1-48
C6500-E(config-if-range)# switchport
C6500-E(config-if-range)# switchport access vlan 10
C6500-E(config-if-range)# switchport voice vlan 110
C6500-E(config-if-range)# spanning-tree portfast edge
C6500-E(config-if-range)# mls qos vlan-based
! Enables VLAN-based QoS on the interface(s)

```



```

! This section attaches the DVLAN policy to the DVLAN interface
C6500-E(config)#interface Vlan 10
C6500-E(config-if)# description DVLAN
C6500-E(config-if)# service-policy input DVLAN-MARKING
! Attaches the DVLAN Per-VLAN Marking policy to the DVLAN interface

! This section attaches the VVLAN policy to the VVLAN interface
C6500-E(config)#interface Vlan 110
C6500-E(config-if)# description VVLAN
C6500-E(config-if)# service-policy input VVLAN-MARKING
! Attaches the VVLAN Per-VLAN Marking policy to the VVLAN interface

```

This configuration can be verified with the commands:

- **show mls qos**
- **show mls qos module**
- **show queueing interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Policing Models

Table 2-5 summarizes policing feature specifics for the Catalyst 6500-E series switches. The Catalyst 6500-E supports these ingress policing models:

- Per-port policing model—This model attaches policers to physical switch port interfaces.
- Per-VLAN policing model—This model attaches policers to logical VLAN interfaces; however, there is an inherent limitation with this policing model; it only supports a single-aggregate policer per VLAN and—since the number of ports associated with a VLAN is dynamic and variable—thus is quite restricted in overall policing effectiveness; therefore, it is generally recommended to use the microflow policing model instead, as it offers more discrete policing options.
- Microflow policing models—This model applies flow-based policers to Layer 3 interfaces to police microflows on a per-source or per-destination basis; microflow policing may be applied on a per-port or per-VLAN basis.



Note

Unlike the previously discussed Catalyst switch platforms, the Catalyst 6500-E does not support per-port/per-VLAN policing.

The per-port policing model and the microflow policing models for the Catalyst 6500-E family of switches are detailed in the following sections.

Per-Port Policing Model

The per-port policing model is quite similar to the per-port marking model, except that the policy action includes a policing function—in some cases to drop, in others to remark. As shown in Figure 2-10, the VoIP and signaling traffic from the VVLAN can be policed to drop at 128 kbps and 32 kbps, respectively (as any excessive traffic matching this criteria would be indicative of network abuse). Similarly, the

multimedia conferencing, signaling, and scavenger traffic from the DVLAN can be policed to drop. On the other hand, data plane policing policies can be applied to transactional, bulk, and best effort data traffic, such that these flows are subject to being remarked (but not dropped at the ingress edge) when severely out-of-profile.

Remarking is performed by configuring policed-DSCP maps with the global configuration commands **mls qos map policed-dscp normal-burst** (which specifies the *exceeding* remarking action) and **mls qos map policed-dscp max-burst** (which specifies the *violating* remarking action, in the case of a dual-rate policer). These commands specify which DSCP values are subject to remarking if out-of-profile and what value these should be remarked as (which in the case of data plane policing/scavenger class QoS policies, this value is CS1/DSCP 8). Even if single rate policers are used, it is recommended to configure the **mls qos map dscp policed max-burst** markdown map, as the *maximum_burst_bytes* parameter for the policer is set to equal to the *normal_burst_bytes* parameter, unless explicitly specified otherwise. In other words, the PIR is set to equal the CIR, unless explicitly specified otherwise, and thus the **exceed-action policed-dscp-transmit** keywords causes PFC QoS to mark traffic down DSCP values as defined by the **policed-dscp max-burst** markdown map (and not the **policed-dscp normal-burst** markdown map).

A per-port policing for the Catalyst 6500-E is shown in [Example 2-104](#).

Example 2-104 Per-Port Policing Configuration Example on a Catalyst 6500-E

```
! This section configures the global policed-DSCP markdown map
C6500-E(config)#mls qos map policed-dscp normal-burst 0 10 18 to 8
! DSCP 0 (DF), 10 (AF11) and 18 (AF21) are marked down to 8 (CS1)
! if found to be exceeding their (respective) policing rates
C6500-E(config)#mls qos map policed-dscp max-burst 0 10 18 to 8
! DSCP 0 (DF), 10 (AF11) and 18 (AF21) are marked down to 8 (CS1)
! if found to be violating their (respective) policing rates

! This section configures the Per-Port policing policy-map
C6500-E(config)# policy-map PER-PORT-POLICING
C6500-E(config-pmap-c)# class VVLAN-VOIP
C6500-E(config-pmap-c)# police 128k 8000
C6500-E(config-pmap-c-police)# conform-action set-dscp-transmit ef
C6500-E(config-pmap-c-police)# exceed-action drop
! Conforming VoIP is marked EF and policed to drop at 128 kbps
C6500-E(config-pmap-c)# class VVLAN-SIGNALING
C6500-E(config-pmap-c)# police 32k 8000
C6500-E(config-pmap-c-police)# conform-action set-dscp-transmit cs3
C6500-E(config-pmap-c-police)# exceed-action drop
! Conforming (VVLAN) Sig is marked CS3 and policed to drop at 32 kbps
C6500-E(config-pmap-c)# class MULTIMEDIA-CONFERENCING
C6500-E(config-pmap-c)# police 5m 8000
C6500-E(config-pmap-c-police)# conform-action set-dscp-transmit af41
C6500-E(config-pmap-c-police)# exceed-action drop
! Conforming MM-Conf is marked AF41 and policed to drop at 5 Mbps
C6500-E(config-pmap-c)# class SIGNALING
C6500-E(config-pmap-c)# police 32k 8000
C6500-E(config-pmap-c-police)# conform-action set-dscp-transmit cs3
C6500-E(config-pmap-c-police)# exceed-action drop
! Conforming (DVLAN) Sig is marked CS3 and policed to drop at 32 kbps
C6500-E(config-pmap-c)# class TRANSACTIONAL-DATA
C6500-E(config-pmap-c)# police 10m 8000
C6500-E(config-pmap-c-police)# conform-action set-dscp-transmit af21
C6500-E(config-pmap-c-police)# exceed-action policed-dscp-transmit
! Conforming Transactional Data is marked AF21 and
! is policed to remark (to CS1) at 10 Mbps
C6500-E(config-pmap-c)# class BULK-DATA
C6500-E(config-pmap-c)# police 10m 8000
```

```

C6500-E(config-pmap-c-police)# conform-action set-dscp-transmit af11
C6500-E(config-pmap-c-police)# exceed-action policed-dscp-transmit
! Conforming Bulk Data is marked AF11 and
! is policed to remark (to CS1) at 10 Mbps
C6500-E(config-pmap-c)# class SCAVENGER
C6500-E(config-pmap-c)# police 10m 8000
C6500-E(config-pmap-c-police)# conform-action set-dscp-transmit cs1
C6500-E(config-pmap-c-police)# exceed-action drop
! Conforming Scavenger traffic is marked CS1 and
! is policed to drop at 10 Mbps

C6500-E(config-pmap-c)# class class-default
C6500-E(config-pmap-c)# police 10m 8000
C6500-E(config-pmap-c-police)# conform-action set-dscp-transmit default
C6500-E(config-pmap-c-police)# exceed-action policed-dscp-transmit
! The implicit default class marks all other traffic to DF and
! polices all other traffic to remark (to CS1) at 10 Mbps

! This section attaches the service-policy to the interface(s)
C6500-E(config)#interface range GigabitEthernet 2/1-48
C6500-E(config-if-range)# switchport
C6500-E(config-if-range)# switchport access vlan 10
C6500-E(config-if-range)# switchport voice vlan 110
C6500-E(config-if-range)# spanning-tree portfast edge
C6500-E(config-if-range)# mls qos trust device cisco-phone
! The interface(s) is set to conditionally-trust Cisco IP Phones
C6500-E(config-if-range)# service-policy input PER-PORT-POLICING
! Attaches the Per-Port Marking policy to the interface(s)

```

**Note**

Catalyst 6500-E software allows for policing rates to be entered using the postfixes **k** (for kilobits), **m** (for megabits), and **g** (for gigabits), as shown in [Example 2-104](#). Additionally, decimal points are allowed in conjunction with these postfixes; for example, a rate of 10.5 Mbps could be entered with the policy map command **police 10.5m**. While these policing rates are converted to their full bps values within the configuration, it makes the entering of these rate more user-friendly and less error prone (as could easily be the case when having to enter up to 10 zeros to define the policing rate).

**Note**

Advanced network administrators can leverage the Catalyst 6500-E support of dual-rate policers—corresponding to the RFC 2698 two rate three color marker (trTCM)—such that the exceeding policing-action for the transactional data and bulk data policers would be to remark to AF22 and AF12 (respectively), while the violating policing action for these classes would be to remark to CS1.

This configuration can be verified with the commands:

- **show mls qos**
- **show mls qos module**
- **show queueing interface**
- **show mls qos maps policed-dscp** (as shown in [Example 2-105](#))
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Example 2-105 Verifying Global Policing Markdown Mappings on a Catalyst 6500-E—show mls qos maps policed-dscp

```

C6500-E#show mls qos maps policed-dscp
Normal Burst Policed-dscp map: (dscp= d1d2)
d1 : d2 0  1  2  3  4  5  6  7  8  9
-----
0 :    08 01 02 03 04 05 06 07 08 09
1 :    08 11 12 13 14 15 16 17 08 19
2 :    20 21 22 23 24 25 26 27 28 29
3 :    30 31 32 33 34 35 36 37 38 39
4 :    40 41 42 43 44 45 46 47 48 49
5 :    50 51 52 53 54 55 56 57 58 59
6 :    60 61 62 63

Maximum Burst Policed-dscp map: (dscp= d1d2)
d1 : d2 0  1  2  3  4  5  6  7  8  9
-----
0 :    08 01 02 03 04 05 06 07 08 09
1 :    08 11 12 13 14 15 16 17 08 19
2 :    20 21 22 23 24 25 26 27 28 29
3 :    30 31 32 33 34 35 36 37 38 39
4 :    40 41 42 43 44 45 46 47 48 49
5 :    50 51 52 53 54 55 56 57 58 59
6 :    60 61 62 63

```

C6500-E#

In [Example 2-105](#), the policing DSCP-markdown mappings are shown in two tables:

- The first table (the normal burst policed-DSCP map) defines the remarking action for packets exceeding the CIR.
- The second table (the maximum burst policed-DSCP map) defines the remarking action for packets exceeding the PIR (which, as previously mentioned, is set to equal the CIR, unless explicitly specified otherwise).

The first digit of the DSCP value of a packet offered to a policer is shown along the Y-axis of the table; the second digit of the DSCP value of a packet offered to a policer is shown along the X-axis of the table. For example, the DSCP value for the transactional data application class (AF21/18) is found in both tables in the row d1=1 and column d2=8. And, as shown, packets with this offered DSCP value (along with DF/0 and AF11/10) are remarked to CS1 (08) if found to be in excess of the policing rate or in violation of the policing rate.

Per-Port Microflow Policing Model

Microflow policing dynamically learns traffic flows and rate limits each unique flow to an individual rate and as such, is a highly effective and efficient policing tool—particularly at the distribution layer in a medianet campus network.

Microflow policing can be applied to ingress traffic on routed interfaces and is typically used in environments where a per-user, granular rate limiting mechanism is required—such as at the distribution layer—to provide a second-line of policing defense in the campus. Like other policers, microflow policing can be used to drop or remark exceeding flows.

Microflow policers are enabled with the **police flow** policy-map class-action command. A flow is defined by five-tuples (IP source address, IP destination address, IP protocol field, Layer 4 protocol source, and destination ports), which are the same for each packet in the flow. Microflow policers apply a single policy to discrete traffic flows, without having to specify the virtually-infinite tuple-combinations. Microflow policing can also be applied with source or destination flow masks (with

the **mask src-only** and **mask dest-only** optional keywords, respectively); these masks apply an aggregate microflow policing policy to multiple flows sharing the same source or IP destination addresses.

In the per-port microflow policing model, a flow-based policer is applied with a **mask src-only** option and applies an aggregate limit to all microflows sharing a common source IP address, remarking traffic in excess of the policing rate.

Remarking is performed by configuring policed-DSCP maps with the global configuration commands **mls qos map policed-dscp normal-burst** (which specifies the *exceeding* remarking action) and **mls qos map policed-dscp max-burst** (which specifies the *violating* remarking action, in the case of a dual rate policer). These commands specify which DSCP values are subject to remarking if out-of-profile and what value these should be remarked as (which in the case of data plane policing/scavenger class QoS policies, this value is CS1/DSCP 8). Even if single rate policers are used, it is recommended to configure the **mls qos map dscp policed max-burst** markdown map, as the *maximum_burst_bytes* parameter for the policer is set to equal to the *normal_burst_bytes* parameter, unless explicitly specified otherwise. In other words, the PIR is set to equal the CIR, unless explicitly specified otherwise, and thus the **exceed-action policed-dscp-transmit** keywords causes PFC QoS to mark traffic down DSCP values as defined by the **policed-dscp max-burst** markdown map (and not the **policed-dscp normal-burst** markdown map).

In [Example 2-106](#), the campus distribution block is using a routed access design and, as such, has Layer 3 interfaces (TenGigabitEthernet 3/1 and 3/2) connecting it to the access layer switches. Microflow policing is applied to all flows to ensure that any endpoint transmitting at more than 5% capacity (an example value) of the access edge 10/100/1000 switch ports are subject to data plane policing/scavenger class QoS.

Example 2-106 Per-Port Microflow Policing Configuration Example on a Catalyst 6500

```
! This section configures the global policed-DSCP markdown map
C6500-E(config)# mls qos map policed-dscp normal-burst 0 10 18 24 34 46 to 8
! DSCP 0 (DF), 10 (AF11), 18 (AF21), 24 (CS3), 34 (AF41) or 46 (EF)
! are marked down to 8 (CS1) if found to be exceeding the aggregate
! per-source microflow policing rate
C6500-E(config)# mls qos map policed-dscp max-burst 0 10 18 24 34 46 to 8
! DSCP 0 (DF), 10 (AF11), 18 (AF21), 24 (CS3), 34 (AF41) or 46 (EF)
! are marked down to 8 (CS1) if found to be violating the aggregate
! per-source microflow policing rate

C6500-E(config)#policy-map MICROFLOW-POLICING
C6500-E(config-pmap)# class class-default
C6500-E(config-pmap-c)# police flow mask src-only 50m 8000 conform-action transmit
exceed-action policed-dscp-transmit
! Any flows from a single source IP address
! will be remarked to CS1 if exceeding 50 Mbps

! This section attaches the microflow policer to the L3 interface(s)
C6500-E(config-if)#interface range TenGigabitEthernet 3/1-2
C6500-E(config-if-range)# description L3-Dwnlnk to Access-Layer
C6500-E(config-if-range)# no switchport
C6500-E(config-if-range)# ip flow ingress
! Enables ingress Netflow on L3 interface (required for microflow)
C6500-E(config-if-range)# service-policy input MICROFLOW-POLICING
! Attaches the microflow policer to the L3 interface(s)
```

This configuration can be verified with the commands:

- **show mls qos**

- **show mls qos module**
- **show mls qos maps policed-dscp**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Per-VLAN Microflow Policing Model

In contrast with the previous example, if the campus distribution block is using a Layer 2/Layer 3 design, and as such has Layer 2 trunked interfaces (TenGigabitEthernet 3/1 and 3/2) connecting it to the access layer switches, then microflow policing can be applied on a per-VLAN basis. In this case, separate microflow policing policies can be applied to each VLAN.

To highlight policy flexibility, additional levels of classification are included in this second microflow policing example (which incidentally can also be applied to the per-port microflow policing model). Instead of applying a blanket microflow policer to all endpoints, separate microflow policers can be applied to different types of endpoints or application-and-endpoint-combinations. For example, VoIP from Cisco IP phones in the VVLAN can be policed to 128 Kbps, while signaling traffic from these endpoints can be policed to 32 kbps. Similarly, TelePresence endpoints in the VVLAN (which mark their media flows to CS4) can be policed to 25 Mbps. All other endpoint-generated traffic in the VVLAN can be policed to 32 kbps per endpoint.

Similar policy granularity can be applied to the DVLAN policer, if desired. However in this example, a simplified DVLAN policer is applied to all flows to ensure that any DVLAN endpoint transmitting at more than 5% capacity (an example value) of the access edge 10/100/1000 switch ports are subject to data plane policing/scavenger class QoS.

An example per-VLAN microflow policing model is shown in [Example 2-107](#).

Example 2-107 Per-VLAN Microflow Policing Configuration Example on a Catalyst 6500

```
! This section configures the global policed-DSCP markdown map
C6500-E(config)#mls qos map policed-dscp normal-burst 0 10 18 34 to 8
! DSCP 0 (DF), 10 (AF11), 18 (AF21) and 34 (AF41) are
! marked down to 8 (CS1) if found to be exceeding the aggregate
! per-source microflow policing rate
C6500-E(config)#mls qos map policed-dscp max-burst 0 10 18 34 to 8
! DSCP 0 (DF), 10 (AF11), 18 (AF21) and 34 (AF41) are
! marked down to 8 (CS1) if found to be violating the aggregate
! per-source microflow policing rate

! This section configures the class-maps
C6500-E(config)#class-map match-all VOIP-ENDPOINTS
C6500-E(config-cmap)# match dscp ef
! Matches VoIP (EF)
C6500-E(config)#class-map match-all TELEPRESENCE-ENDPOINTS
C6500-E(config-cmap)# match dscp cs4
! Matches TelePresence (CS4)
C6500-E(config)#class-map match-all SIGNALING-ENDPOINTS
C6500-E(config-cmap)# match dscp cs3
! Matches Signaling (CS3)

! This section configures the VVLAN microflow policing policy-map
C6500-E(config)#policy-map VVLAN-MICROFLOW-POLICING
C6500-E(config-pmap)# class VOIP-ENDPOINTS
```

```

C6500-E(config-pmap-c)# police flow mask src-only 128k 8000 conform-action transmit
exceed-action drop
    ! All EF flows from a single VVLAN IP are policed to drop at 128 kbps
C6500-E(config-pmap)# class TELEPRESENCE-ENDPOINTS
C6500-E(config-pmap-c)# police flow mask src-only 25m 256000 conform-action transmit
exceed-action drop
    ! All CS4 flows from a single VVLAN IP are policed to drop at 25 Mbps
C6500-E(config-pmap-c)# class SIGNALING-ENDPOINTS
C6500-E(config-pmap-c)# police flow mask src-only 32k 8000 conform-action transmit
exceed-action drop
    ! All CS3 flows from a single VVLAN IP are policed to drop at 32 kbps
C6500-E(config-pmap)# class class-default
C6500-E(config-pmap-c)# police flow mask src-only 32k 8000 conform-action transmit
exceed-action drop
    ! All other flows from a single VVLAN IP are policed to drop at 32kbps

    ! This section configures the DVLAN microflow policing policy-map
C6500-E(config)#policy-map DVLAN-MICROFLOW-POLICING
C6500-E(config-pmap)# class class-default
C6500-E(config-pmap-c)# police flow mask src-only 50m 8000 conform-action transmit
exceed-action policed-dscp-transmit
    ! Any flows from a single source IP address within the DVLAN
    ! will be remarked to CS1 if exceeding 50 Mbps

    ! This section applies the VVLAN microflow policing policy to the VVLAN
C6500-E(config)#interface vlan 110
C6500-E(config-if)# description VVLAN
C6500-E(config-if)# ip flow ingress
    ! Enables ingress Netflow on L3 VLAN interface
C6500-E(config-if)# service-policy input VVLAN-MICROFLOW-POLICING
    ! Attaches the VVLAN microflow policing policy to the VVLAN interface

    ! This section applies the DVLAN microflow policing policy to the VVLAN
C6500-E(config)#interface vlan 10
C6500-E(config-if)# description DVLAN
C6500-E(config-if)# ip flow ingress
    ! Enables ingress Netflow on L3 VLAN interface
C6500-E(config-if)# service-policy input DVLAN-MICROFLOW-POLICING
    ! Attaches the DVLAN microflow policing policy to the DVLAN interface

    ! This section enables VLAN-based QoS on the L2 (trunked) interface(s)
C6500-E(config)#interface range TenGigabitEthernet3/1-2
C6500-E(config-if-range)# description L2-Dwnlnk to Access-Layer
C6500-E(config-if-range)# switchport
C6500-E(config-if-range)# switchport trunk encapsulation dot1q
C6500-E(config-if-range)# switchport trunk allowed vlan 10,110
C6500-E(config-if-range)# switchport mode trunk
C6500-E(config-if-range)# mls qos vlan-based
    ! Enabled VLAN-Based QoS for the L2 interface(s)

```

This configuration can be verified with the commands:

- **show mls qos**
- **show mls qos module**
- **show mls qos maps policed-dscp**
- **show class-map**
- **show policy-map**

- **show policy-map interface**

Queuing Models

As summarized from the information presented in [Table 2-9](#), [Table 2-10](#), and [Table 2-11](#), medianet campus Catalyst 6500/6500-E switch modules can be grouped by egress queuing structures, as shown in [Table 2-12](#).

Table 2-12 Catalyst 6500 Switch Modules by Egress Queuing Structures

1P3Q8T (CoS-to-Queue)	1P7Q8T (CoS-to-Queue)	1P7Q4T (DSCP-to-Queue)
WS-SUP32-GE	WS-X6704-10GE	WS-X6708-10GE
WS-SUP32-10GE		WS-X6716-10GE
WS-X6148A-GE-TX		
WS-X6148A-GE-45AF		
WS-X6724-SFP		
WS-X6748-SFP		
WS-X6748-GE-TX		

Each of these Catalyst 6500/6500-E egress queuing models (1P3Q8T, 1P7Q8T, and 1P7Q4T) is covered in subsequent sections, but first, consideration has to be given to ingress queuing models

There are two main considerations relevant to ingress queuing design on the Catalyst 6500/6500-E:

- The degree of oversubscription (if any) of the linecard
- Whether the linecard requires trust-CoS to be enabled to engage ingress queuing

To the first consideration, some linecards may be designed to support a degree of oversubscription, meaning that theoretically more traffic may be offered to the linecard via the sum of all GE/10GE switch ports than can collectively access the backplane at once. Since such a scenario is extremely unlikely, it is often more cost-effective to utilize linecards that have a degree of oversubscription within the campus network. However, if this design choice has been made, it is important for administrators to recognize the potential for drops due to oversubscribed linecard architectures. To manage application-class service levels during such extreme scenarios, ingress queuing models may be enabled.

While the presence of oversubscribed linecard architectures may be viewed as the sole consideration as to enabling ingress queuing or not, a second important consideration should also be kept in mind, namely that many Catalyst 6500/6500-E linecards (at the time of writing) only support CoS-based ingress queuing models (and thus require trust-CoS to be enabled on these switch ports). Enabling trust-CoS reduces classification and marking granularity—limiting the administrator to an 8-class 802.1Q/p model. However, as previously discussed, RFC 4594-based medianet models may require up to 12 classes of service. Once CoS is trusted, DSCP values are overwritten (via the CoS-to-DSCP mapping table) and application classes sharing the same CoS values are longer distinguishable from one another. Therefore, given this classification and marking limitation and the fact that the value of enabling ingress queuing is only achieved in extremely rare scenarios, it is not recommended to enable CoS-based ingress queuing on the Catalyst 6500/6500-E; rather, limit such linecards to the access layer of a medianet campus network and deploy either non-oversubscribed linecards and/or linecards supporting DSCP-based queuing at the distribution and core layers of the campus network.

Table 2-13 helps summarize these considerations by listing the medianet switch models (presented in Table 2-12) and including their oversubscription ratios and whether the ingress queuing models are CoS or DSCP-based.

Table 2-13 Catalyst 6500 Switch Module Ingress Queuing Architectures

Switch Module	Maximum Input	Maximum Output (to Backplane)	Oversubscription Ratio	Ingress Queuing Structure	CoS/DSCP Based	Ingress Queuing Recommendations
WS-SUP32-GE	8 Gbps (8 x GE)	32 Gbps	-	2Q2T	CoS-Based	Not required
WS-SUP32-10GE	20 Gbps (2 x 10GE)	32 Gbps	-	2Q2T	CoS-Based	Not required
WS-X6148A-GE-TX	48 Gbps (48 x GE)	32 Gbps	6:5	1Q2T	CoS-Based	Not recommended (use linecard at access-layer only)
WS-X6148A-GE-45 AF	48 Gbps (48 x GE)	32 Gbps	6:5	1Q2T	CoS-Based	Not recommended (use linecard at access-layer only)
WS-X6724-SFP	24 Gbps (24 x GE)	40 Gbps (2 x 20 Gbps)	-	2Q8T/1Q8T	CoS-Based	Not required
WS-X6748-SFP	48 Gbps (48 x GE)	40 Gbps (2 x 20 Gbps)	6:5	2Q8T/1Q8T	CoS-Based	Not recommended (use linecard at access-layer only)
WS-X6748-GE-TX	48 Gbps (48 x GE)	40 Gbps (2 x 20 Gbps)	6:5	2Q8T/1Q8T	CoS-Based	Not recommended (use linecard at access-layer only)
WS-X6704-10GE	40 Gbps (4 x 10GE)	40 Gbps (2 x 20 Gbps)	-	8Q8T	CoS or DSCP-based	Not required
WS-X6708-10GE	80 Gbps (8 x 10GE)	40 Gbps (2 x 20 Gbps)	2:1	8Q4T	CoS or DSCP-based	Use DSCP-based 8Q4T ingress queuing
WS-X6716-10GE	160 Gbps (16 x 10GE)	40 Gbps (2 x 20 Gbps)	4:1	8Q4T/1P7Q2T*	CoS or DSCP-based	Use DSCP-based 1P7Q2T ingress queuing



Note

The Catalyst WS-X6716-10GE can be configured to operate in Performance Mode (with an 8Q4T ingress queuing structure) or in Oversubscription Mode (with a 1P7Q2T ingress queuing structure). In Performance mode, only one port in every group of four is operational (while the rest are administratively shut down), which eliminates any oversubscription on this linecard and as such ingress queuing is not required (as only 4 x 10GE ports are active in this mode and the backplane access rate is also at 40 Gbps). In Oversubscription Mode (the default mode), all ports are operational and the maximum oversubscription ratio is 4:1. Therefore it is recommended to enable 1P7Q2T DSCP-based ingress queuing on this linecard in Oversubscription Mode.

Additional details on these WS-X6716-10GE operational modes can be found at:

http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/qa_cisco_catalyst_6500_series_16port_10gigabit_ethernet_module.html

Therefore, if 6708 and 6716 linecards (with the latter operating in oversubscription mode) are used in the distribution and core layers of the medianet campus network, then 8Q4T DSCP-based ingress queuing and 1P7Q2T DSCP-based ingress queuing (respectively) are recommended to be enabled. These queuing models are detailed in the following sections.

8Q4T (DSCP-Based) Ingress Queuing Model

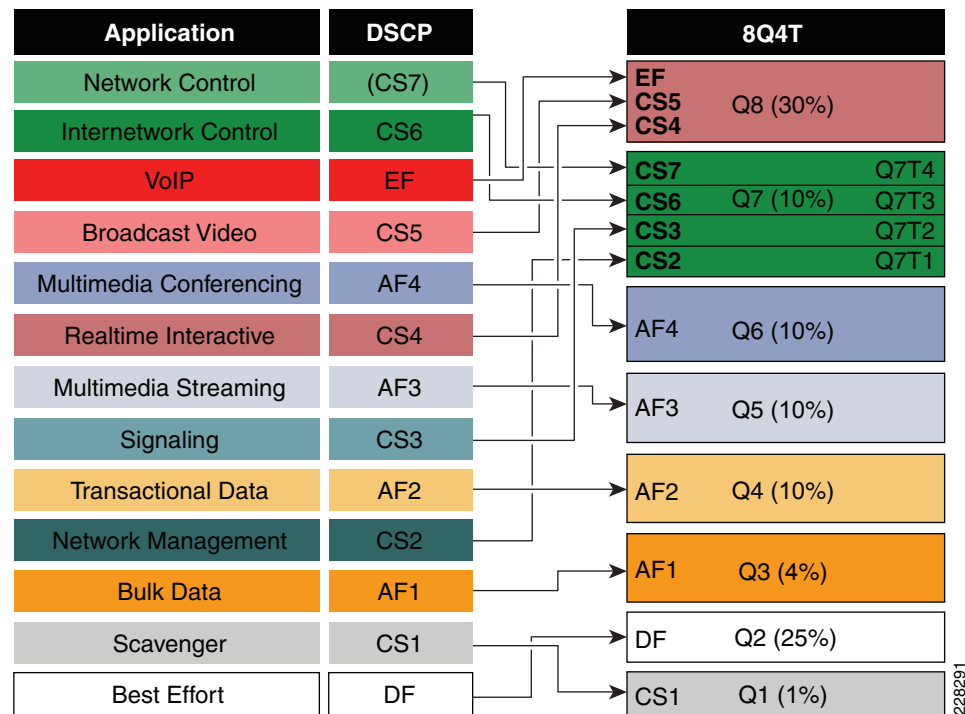
In the 8Q4T (DSCP-Based) ingress queuing model, 30% of the link bandwidth can be allocated for Q8, 10% for Q7, 10% for Q6, 10% for Q5, 10% for Q4, 4% for Q3, 25% for Q2 (the best effort queue), and 1% for Q1 (the scavenger queue). In turn, 15% of the buffers can be allocated for Q8, 10% each for Q3-Q7, 25% for Q2 (the best effort queue), and 5% for Q1 (the scavenger queue).

Additionally, WRED can be enabled on queues 1 through 7. Only basic WRED functionality is required for queues 1 and 2 (as only a single DSCP value is assigned to each); therefore the first minimum WRED thresholds for these queues can be set to 80% and the first maximum WRED thresholds for these queues can be set to 100%. As queues 3 through 6 have AF PHBs assigned to them, the WRED thresholds can be set to correspond to the three drop-precedence levels per AF class. Thus, the first three minimum WRED thresholds for these queues can be set to 70%, 80%, and 90%, respectively; and the first three maximum WRED thresholds for these queues can be set to 80%, 90%, and 100%, respectively.

Additionally, since Q7 has 4 separate DSCP values assigned to it, intra-queue QoS can be achieved by mapping these to different WRED thresholds. Thus, the minimum WRED thresholds for Q7T1, Q7T2, Q7T3, and Q7T4 can be set to 60%, 70%, 80%, and 90%, respectively; and the minimum WRED thresholds for Q7T1, Q7T2, Q7T3, and Q7T4 can be set to 70%, 80%, 90%, and 100%, respectively.

DSCP EF (VoIP), CS5 (broadcast video) and CS4 (realtime interactive) can be mapped to Q8. CS7 (network control) can be mapped to Q7T4; CS6 (internetwork control) can be mapped to Q7T3; CS3 (signaling) can be mapped to Q7T2; and CS2 (network management) can be mapped to Q7T1. AF4 (multimedia conferencing) can be mapped to Q6. AF3 (multimedia streaming) can be mapped to Q5. AF2 (transactional data) can be mapped to Q4. AF1 (bulk data) can be mapped to Q3. DF (best effort) can be mapped to Q2. CS1 can be mapped to Q1.

These 8Q4T DSCP-to-queue mappings are illustrated in [Figure 2-26](#).

Figure 2-26 Catalyst 6500-E 8Q4T (DSCP-to-Queue) Ingress Queuing Model

The corresponding configuration for 8Q8T (DSCP-to-Queue) ingress queuing on a Catalyst 6500-E is shown in [Example 2-108](#).

Example 2-108 8Q8T (DSCP-to-Queue) Ingress Queuing Configuration Example on a Catalyst 6500-E

```
! This section configures the port for DSCP-based Ingress Queuing
C6500-E(config)#interface range TenGigabitEthernet 2/1-8
C6500-E(config-if-range)# mls qos queue-mode mode-dscp
! Enables DSCP-to-Queue mapping
C6500-E(config-if-range)# mls qos trust dscp
! Enables DSCP-trust for ingress DSCP-based queuing

! This section configures the receive queues BW and limits
C6500-E(config-if-range)# rcv-queue queue-limit 10 25 10 10 10 10 10 15
! Allocates 10% to Q1, 25% to Q2, 10% to Q3, 10% to Q4,
! Allocates 10% to Q5, 10% to Q6, 10% to Q7 and 15% to Q8
C6500-E(config-if-range)# rcv-queue bandwidth 1 25 4 10 10 10 10 30
! Allocates 1% BW to Q1, 25% BW to Q2, 4% BW to Q3, 10% BW to Q4,
! Allocates 10% BW to Q5, 10% BW to Q6, 10% BW to Q7 & 30% BW to Q8

! This section enables WRED on all queues except Q8
C6500-E(config-if-range)# rcv-queue random-detect 1
! Enables WRED on Q1
C6500-E(config-if-range)# rcv-queue random-detect 2
! Enables WRED on Q2
C6500-E(config-if-range)# rcv-queue random-detect 3
! Enables WRED on Q3
C6500-E(config-if-range)# rcv-queue random-detect 4
! Enables WRED on Q4
C6500-E(config-if-range)# rcv-queue random-detect 5
! Enables WRED on Q5
C6500-E(config-if-range)# rcv-queue random-detect 6
```

```

! Enables WRED on Q6
C6500-E(config-if-range)# rcv-queue random-detect 7
! Enables WRED on Q7
C6500-E(config-if-range)# no rcv-queue random-detect 8
! Disables WRED on Q8

! This section configures WRED thresholds for Queues 1 through 7
C6500-E(config-if-range)# rcv-queue random-detect max-threshold 1 100 100 100 100
! Sets all WRED max thresholds on Q1 to 100%
C6500-E(config-if-range)# rcv-queue random-detect min-threshold 1 80 100 100 100
! Sets Q1T1 min WRED threshold to 80%
C6500-E(config-if-range)# rcv-queue random-detect min-threshold 2 80 100 100 100
! Sets Q2T1 min WRED threshold to 80%
C6500-E(config-if-range)# rcv-queue random-detect max-threshold 2 100 100 100 100
! Sets all WRED max thresholds on Q2 to 100%
C6500-E(config-if-range)# rcv-queue random-detect min-threshold 3 70 80 90 100
! Sets WRED min thresholds for Q3T1, Q3T2, Q3T3 to 70 %, 80% and 90%
C6500-E(config-if-range)# rcv-queue random-detect max-threshold 3 80 90 100 100
! Sets WRED max thresholds for Q3T1, Q3T2, Q3T3 to 80%, 90% and 100%
C6500-E(config-if-range)# rcv-queue random-detect min-threshold 4 70 80 90 100
! Sets WRED min thresholds for Q4T1, Q4T2, Q4T3 to 70 %, 80% and 90%
C6500-E(config-if-range)# rcv-queue random-detect max-threshold 4 80 90 100 100
! Sets WRED max thresholds for Q4T1, Q4T2, Q4T3 to 80%, 90% and 100%
C6500-E(config-if-range)# rcv-queue random-detect min-threshold 5 70 80 90 100
! Sets WRED min thresholds for Q5T1, Q5T2, Q5T3 to 70 %, 80% and 90%
C6500-E(config-if-range)# rcv-queue random-detect max-threshold 5 80 90 100 100
! Sets WRED max thresholds for Q5T1, Q5T2, Q5T3 to 80%, 90% and 100%
C6500-E(config-if-range)# rcv-queue random-detect min-threshold 6 70 80 90 100
! Sets WRED min thresholds for Q6T1, Q6T2, Q6T3 to 70 %, 80% and 90%
C6500-E(config-if-range)# rcv-queue random-detect max-threshold 6 80 90 100 100
! Sets WRED max thresholds for Q6T1, Q6T2, Q6T3 to 80%, 90% and 100%
C6500-E(config-if-range)# rcv-queue random-detect min-threshold 7 60 70 80 90
! Sets WRED min thresholds for Q7T1, Q7T2, Q7T3 and Q7T4
! to 60%, 70%, 80% and 90%, respectively
C6500-E(config-if-range)# rcv-queue random-detect max-threshold 7 70 80 90 100
! Sets WRED max thresholds for Q7T1, Q7T2, Q7T3 and Q7T4
! to 70%, 80%, 90% and 100%, respectively

! This section configures the DSCP-to-Receive-Queue mappings
C6500-E(config-if-range)# rcv-queue dscp-map 1 1 8
! Maps CS1 (Scavenger) to Q1T1
C6500-E(config-if-range)# rcv-queue dscp-map 2 1 0
! Maps DF (Best Effort) to Q2T1
C6500-E(config-if-range)# rcv-queue dscp-map 3 1 14
! Maps AF13 (Bulk Data-Drop Precedence 3) to Q3T1
C6500-E(config-if-range)# rcv-queue dscp-map 3 2 12
! Maps AF12 (Bulk Data-Drop Precedence 2) to Q3T2
C6500-E(config-if-range)# rcv-queue dscp-map 3 3 10
! Maps AF11 (Bulk Data-Drop Precedence 1) to Q3T3
C6500-E(config-if-range)# rcv-queue dscp-map 4 1 22
! Maps AF23 (Transactional Data-Drop Precedence 3) to Q4T1
C6500-E(config-if-range)# rcv-queue dscp-map 4 2 20
! Maps AF22 (Transactional Data-Drop Precedence 2) to Q4T2
C6500-E(config-if-range)# rcv-queue dscp-map 4 3 18
! Maps AF21 (Transactional Data-Drop Precedence 1) to Q4T3
C6500-E(config-if-range)# rcv-queue dscp-map 5 1 30
! Maps AF33 (Multimedia Streaming-Drop Precedence 3) to Q5T1
C6500-E(config-if-range)# rcv-queue dscp-map 5 2 28
! Maps AF32 (Multimedia Streaming-Drop Precedence 2) to Q5T2
C6500-E(config-if-range)# rcv-queue dscp-map 5 3 26
! Maps AF31 (Multimedia Streaming-Drop Precedence 1) to Q5T3

```

```

C6500-E(config-if-range)# rcv-queue dscp-map 6 1 38
! Maps AF43 (Multimedia Conferencing-Drop Precedence 3) to Q6T1
C6500-E(config-if-range)# rcv-queue dscp-map 6 2 36
! Maps AF42 (Multimedia Conferencing-Drop Precedence 2) to Q6T2
C6500-E(config-if-range)# rcv-queue dscp-map 6 3 34
! Maps AF41 (Multimedia Conferencing-Drop Precedence 1) to Q6T3
C6500-E(config-if-range)# rcv-queue dscp-map 7 1 16
! Maps CS2 (Network Management) to Q7T1
C6500-E(config-if-range)# rcv-queue dscp-map 7 2 24
! Maps CS3 (Signaling) to Q7T2
C6500-E(config-if-range)# rcv-queue dscp-map 7 3 48
! Maps CS6 (Internetwork Control) to Q7T3
C6500-E(config-if-range)# rcv-queue dscp-map 7 4 56
! Maps CS7 (Network Control) to Q7T4

C6500-E(config-if-range)# rcv-queue dscp-map 8 4 32 40 46
! Maps CS4 (Realtime Interactive), CS5 (Broadcast Video),
! and EF (VoIP) to Q8

```

This configuration can be verified with the command:

- **show queueing interface | begin Rx** (as shown in [Example 2-109](#))

Example 2-109 Verifying Ingress Queuing on a Catalyst 6500-E-show queueing interface | begin Rx

```

C6500-E#show queueing interface TenGigabitEthernet 1/8 | begin Rx
Queueing Mode In Rx direction: mode-dscp
Receive queues [type = 8q4t]:
Queue Id      Scheduling  Num of thresholds
-----
      01         WRR              04
      02         WRR              04
      03         WRR              04
      04         WRR              04
      05         WRR              04
      06         WRR              04
      07         WRR              04
      08         WRR              04

WRR bandwidth ratios:      1[queue 1] 25[queue 2]  4[queue 3] 10[queue 4]
10[queue 5] 10[queue 6] 10[queue 7] 30[queue 8]
queue-limit ratios:      10[queue 1] 25[queue 2] 10[queue 3] 10[queue 4]
10[queue 5] 10[queue 6] 10[queue 7] 15[queue 8]

queue tail-drop-thresholds
-----
1      70[1] 80[2] 90[3] 100[4]
2      100[1] 100[2] 100[3] 100[4]
3      100[1] 100[2] 100[3] 100[4]
4      100[1] 100[2] 100[3] 100[4]
5      100[1] 100[2] 100[3] 100[4]
6      100[1] 100[2] 100[3] 100[4]
7      100[1] 100[2] 100[3] 100[4]
8      100[1] 100[2] 100[3] 100[4]

queue random-detect-min-thresholds
-----
1      80[1] 100[2] 100[3] 100[4]
2      80[1] 100[2] 100[3] 100[4]
3      70[1] 80[2] 90[3] 100[4]
4      70[1] 80[2] 90[3] 100[4]
5      70[1] 80[2] 90[3] 100[4]

```

```
6      70[1] 80[2] 90[3] 100[4]
7      60[1] 70[2] 80[3] 90[4]
8      100[1] 100[2] 100[3] 100[4]
```

```
queue random-detect-max-thresholds
-----
1      100[1] 100[2] 100[3] 100[4]
2      100[1] 100[2] 100[3] 100[4]
3      80[1] 90[2] 100[3] 100[4]
4      80[1] 90[2] 100[3] 100[4]
5      80[1] 90[2] 100[3] 100[4]
6      80[1] 90[2] 100[3] 100[4]
7      70[1] 80[2] 90[3] 100[4]
8      100[1] 100[2] 100[3] 100[4]
```

WRED disabled queues: 8

```
queue thresh cos-map
-----
1      1      0 1
1      2      2 3
1      3      4
1      4      6 7
2      1
2      2
2      3
2      4
3      1
3      2
3      3
3      4
4      1
4      2
4      3
4      4
5      1
5      2
5      3
5      4
6      1
6      2
6      3
6      4
7      1
7      2
7      3
7      4
8      1      5
8      2
8      3
8      4
```

```
queue thresh dscp-map
-----
1      1      0 1 2 3 4 5 6 7 8 9 11 13 15 16 17 19 21 23 25 27 29 31 33 39 4
1 42 43 44 45 47
1      2
1      3
1      4
2      1      14
2      2      12
2      3      10
2      4
```

```

3      1      22
3      2      20
3      3      18
3      4
4      1      24 30
4      2      28
4      3      26
4      4
5      1      32 34 35 36 37 38
5      2
5      3
5      4
6      1      48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63
6      2
6      3
6      4
7      1
7      2
7      3
7      4
8      1      40 46
8      2
8      3
8      4

Packets dropped on Transmit:
  BPDU packets:  0

queue          dropped  [dscp-map]
-----
1              0  [0 1 2 3 4 5 6 7 8 9 11 13 15 16 17 19 21 23 25
27 29 31 33 39 41 42 43 44 45 47 ]
2              0  [14 12 10 ]
3              0  [22 20 18 ]
4              0  [24 30 28 26 ]
5              0  [32 34 35 36 37 38 ]
6              0  [48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63
]
8              0  [40 46 ]

Packets dropped on Receive:
  BPDU packets:  0

queue          dropped  [dscp-map]
-----
1              0  [0 1 2 3 4 5 6 7 8 9 11 13 15 16 17 19 21 23 25
27 29 31 33 39 41 42 43 44 45 47 ]
2              0  [14 12 10 ]
3              0  [22 20 18 ]
4              0  [24 30 28 26 ]
5              0  [32 34 35 36 37 38 ]
6              0  [48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63
]
8              0  [40 46 ]
C6500-E#

```

Example 2-109 verifies that 8Q8T (DSCP-based) ingress queuing has been enabled on the interface with the queue limits, bandwidth allocations, WRED thresholds, and DSCP-to-queue mappings as described at the beginning of this section.

1P7Q2T (DSCP-Based) Ingress Queuing Model

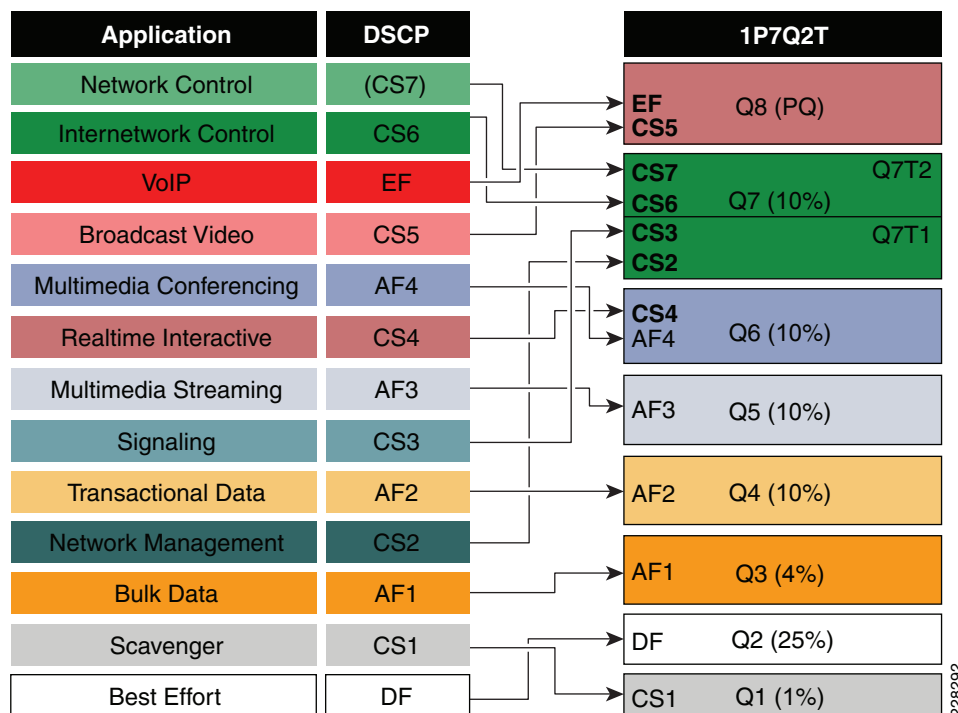
In the 1P7Q2T (DSCP-Based) ingress queuing model, 10% of the link bandwidth can be allocated for Q7, 10% for Q6, 10% for Q5, 10% for Q4, 4% for Q3, 25% for Q2 (the best effort queue), and 1% for Q1 (the scavenger queue); the bandwidth allocated for the strict-priority queue (Q8) is not configurable. In turn, 10% of the buffers can be allocated (each) for Q3-Q7, 25% for Q2 (the best effort queue), and 10% for Q1 (the scavenger queue).

Additionally, the 1P7Q2T structure supports—not WRED—but two tail-drop thresholds (one is configurable and the other is simply the tail of the queue). As such, this functionality would not be needed on queues 1 and 2 (as only a single DSCP value is mapped to each and there is no point tail-dropping flows sharing the same DSCP-value earlier than necessary). However, this functionality can be leveraged on queues 3 to 6 to loosely mimic the AF PHB (although only two levels of dropping would be supported, rather than the three specified in RFC 2597); specifically, AFx2 and AFx3 can be mapped to the first tail-drop threshold (set at 80%) and AFx1 can be mapped to the second drop threshold (the tail at 100%). Similarly, the first tail-drop threshold on Q7 can also be set to 80%, with the second remaining at 100%.

The 1P7Q2T model does not support explicit DSCP-mapping to the strict priority queue (at the time of writing); by default, DSCP EF (VoIP) and CS5 (broadcast video) are mapped to the strict priority queue (Q8). Therefore, CS4 needs to be mapped to another queue, which in this case can be Q6 (as the bandwidth allocated to it has been increased accordingly). Additionally, CS7 (network control) and CS6 (internetwork control) can be mapped to Q7T2; CS3 (signaling) and CS2 (network management) can be mapped to Q7T1. AF4 (multimedia conferencing) can be mapped to Q6. AF3 (multimedia streaming) can be mapped to Q5. AF2 (transactional data) can be mapped to Q4. AF1 (bulk data) can be mapped to Q3. DF (best effort) can be mapped to Q2. CS1 can be mapped to Q1.

These 1P7Q2T DSCP-to-queue mappings are illustrated in [Figure 2-27](#).

Figure 2-27 Catalyst 6500-E 1P7Q2T (DSCP-to-Queue) Ingress Queuing Model



The corresponding configuration for 1P7Q2T (DSCP-to-Queue) ingress queuing on a Catalyst 6500-E is shown in [Example 2-110](#).

Example 2-110 1P7Q2T (DSCP-to-Queue) Ingress Queuing Configuration Example on a Catalyst 6500-E

```
! This section configures the port for DSCP-based Ingress Queueing
C6500-E(config)#interface range TenGigabitEthernet 2/1-16
C6500-E(config-if-range)# mls qos queue-mode mode-dscp
! Enables DSCP-to-Queue mapping
C6500-E(config-if-range)# mls qos trust dscp
! Enables DSCP-trust for ingress DSCP-based queuing

! This section configures the receive queues BW and limits
C6500-E(config-if-range)# rcv-queue bandwidth 1 25 4 10 10 10 10
! Allocates 1% BW to Q1, 25% BW to Q2, 4% BW to Q3, 10% BW to Q4,
! Allocates 10% BW to Q5, 10% BW to Q6 and 10% BW to Q7
C6500-E(config-if-range)# rcv-queue queue-limit 10 25 10 10 10 10 10
! Allocates 10% to Q1, 25% to Q2, 10% to Q3, 10% to Q4,
! Allocates 10% to Q5, 10% to Q6 and 10% to Q7

! This section configures tail-dropping thresholds for Q1-Q7
C6500-E(config-if-range)# rcv-queue threshold 1 100 100
! No early tail-dropping threshold is set for Q1
C6500-E(config-if-range)# rcv-queue threshold 2 100 100
! No early tail-dropping threshold is set for Q2
C6500-E(config-if-range)# rcv-queue threshold 3 80 100
! An early tail-dropping threshold is set for Q3 at 80%
C6500-E(config-if-range)# rcv-queue threshold 4 80 100
! An early tail-dropping threshold is set for Q4 at 80%
C6500-E(config-if-range)# rcv-queue threshold 5 80 100
! An early tail-dropping threshold is set for Q5 at 80%
C6500-E(config-if-range)# rcv-queue threshold 6 80 100
! An early tail-dropping threshold is set for Q6 at 80%
C6500-E(config-if-range)# rcv-queue threshold 7 80 100
! An early tail-dropping threshold is set for Q7 at 80%

! This section configures the DSCP-to-Receive-Queue mappings
C6500-E(config-if-range)# rcv-queue dscp-map 1 2 8
! Maps CS1 (Scavenger) to Q1T2
C6500-E(config-if-range)# rcv-queue dscp-map 2 2 0
! Maps DF (Best Effort) to Q2T1
C6500-E(config-if-range)# rcv-queue dscp-map 3 1 12 14
! Maps AF12 and AF13 (Bulk Data) to Q3T1
C6500-E(config-if-range)# rcv-queue dscp-map 3 2 10
! Maps AF11 (Bulk Data) to Q3T2
C6500-E(config-if-range)# rcv-queue dscp-map 4 1 20 22
! Maps AF22 and AF23 (Transactional Data) to Q4T1
C6500-E(config-if-range)# rcv-queue dscp-map 4 2 18
! Maps AF21 (Transactional Data) to Q4T2
C6500-E(config-if-range)# rcv-queue dscp-map 5 1 28 30
! Maps AF32 and AF33 (Multimedia Streaming) to Q5T1
C6500-E(config-if-range)# rcv-queue dscp-map 5 2 26
! Maps AF31 (Multimedia Streaming) to Q5T2
C6500-E(config-if-range)# rcv-queue dscp-map 6 1 36 38
! Maps AF42 and AF43 (Multimedia Conferencing) to Q6T1
C6500-E(config-if-range)# rcv-queue dscp-map 6 2 32 34
! Maps CS4 (Realtime Interactive) and AF41 (Multimedia Conferencing) to Q6T2
C6500-E(config-if-range)# rcv-queue dscp-map 7 1 16 24
! Maps CS2 (Network Management) and CS3 (Signaling) to Q7T1
C6500-E(config-if-range)# rcv-queue dscp-map 7 2 48 56
```

```
! Maps CS6 (Internetwork Control) and CS7 (Network Control) to Q7T4
! DSCP EF (VoIP) and CS5 (Broadcast Video) are mapped by default to Q8/PQ
```

This configuration can be verified with the command:

- **show queueing interface | begin Rx**

1P3Q8T (CoS-Based) Egress Queuing Model

In the 1P3Q8T (CoS-Based) egress queuing model, 30% of the link bandwidth can be allocated for the priority queue (Q4), 40% for the non-realtime queue (Q3), 25% for the best effort queue (Q2), and 5% for the scavenger/bulk queue (Q1). In turn, 15% of the buffers can be allocated for the PQ (Q4), 40% for Q3, 25% for Q2, and 20% for Q1.

Additionally, WRED can be enabled on Q1, Q2, and Q3. Q1 and Q2 need only basic WRED functionality, as only a single CoS value is assigned to each; therefore the first minimum WRED thresholds can be set to 80% for these queues and the first maximum WRED thresholds can be set to 100% for these queues. Since Q3 has 4 separate CoS values assigned to it, intra-queue QoS can be achieved by mapping these to different WRED thresholds. Thus, the minimum WRED thresholds for Q3T1, Q3T2, Q3T3, and Q3T4 can be set to 60%, 70%, 80%, and 90%, respectively; and the maximum WRED thresholds for Q3T1, Q3T2, Q3T3, and Q3T4 can be set to 70%, 80%, 90%, and 100%, respectively.

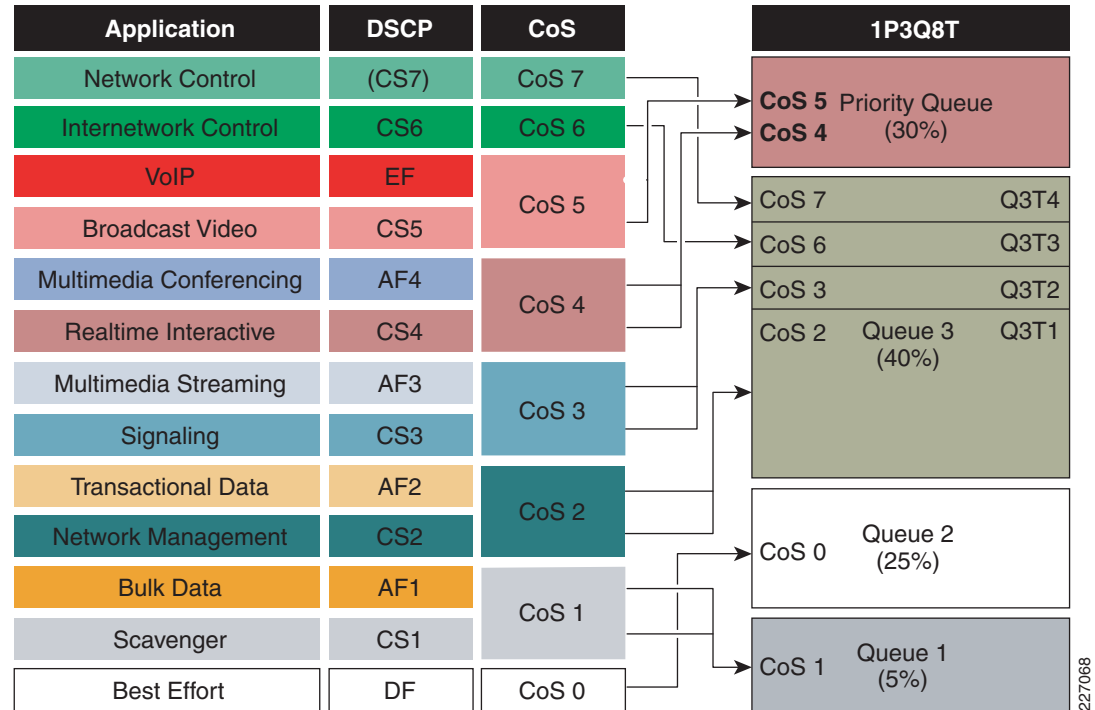
Following this, CoS values 5 (VoIP and broadcast video) and 4 (realtime interactive and multimedia conferencing) can be mapped to the priority queue. CoS 7 (network control) can be mapped to Q3T4, CoS 6 (internetwork control) can be mapped to Q3T3, CoS 3 (signaling and multimedia streaming) can be mapped to Q3T2, and CoS 2 (network management and transactional data) can be mapped to Q3T1. CoS 0 (Best Effort) can be mapped to Q2 and CoS 1 (scavenger and bulk) can be mapped to Q1.



Note

In the 1P3Q8T CoS-to-queue model, certain application classes that are normally mapped to differing queue/threshold-combinations must be mapped to the same queue/threshold because of the limited CoS-to-queue mapping level-of-granularity. These include realtime interactive and multimedia conferencing (both sharing CoS 4), signaling and multimedia streaming (both sharing CoS 3), network management and transactional data (both sharing CoS 2), and scavenger and bulk data (both sharing CoS 1).

These 1P3Q8T CoS-to-queue mappings are illustrated in [Figure 2-28](#).

Figure 2-28 Catalyst 6500-E 1P3Q8T (CoS-to-Queue) Egress Queuing Model

The corresponding configuration for 1P3Q8T (CoS-to-Queue) egress queuing on a Catalyst 6500-E is shown in [Example 2-111](#).

Example 2-111 1P3Q8T (CoS-to-Queue) Egress Queuing Configuration Example on a Catalyst 6500-E

```

! This section configures 1P3Q8T (CoS-Based) Egress Queuing
C6500-E(config)#interface range GigabitEthernet 2/1-48
C6500-E(config-if-range)# wrr-queue queue-limit 20 25 40
! Allocates 20% of the buffers to Q1, 25% to Q2 and 40% to Q3
C6500-E(config-if-range)# priority-queue queue-limit 15
! Allocates 15% of the buffers to the PQ
C6500-E(config-if-range)# wrr-queue bandwidth 5 25 40
! Allocates 5% BW to Q1, 25% BW to Q2 and 40% BW to Q3

! This section enables WRED on Queues 1 through 3
C6500-E(config-if-range)# wrr-queue random-detect 1
! Enables WRED on Q1
C6500-E(config-if-range)# wrr-queue random-detect 2
! Enables WRED on Q2
C6500-E(config-if-range)# wrr-queue random-detect 3
! Enables WRED on Q3

! This section configures WRED thresholds for Queues 1 through 3
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 1 100 100 100 100 100 100
100 100
! Sets all WRED max thresholds on Q1 to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 1 80 100 100 100 100 100
100 100
! Sets Q1T1 min WRED threshold to 80%; all others set to 100%

```

```

C6500-E(config-if-range)# wrr-queue random-detect max-threshold 2 100 100 100 100 100 100
100 100
! Sets all WRED max thresholds on Q2 to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 2 80 100 100 100 100 100
100 100
! Sets Q2T1 min WRED threshold to 80%; all others set to 100%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 3 70 80 90 100 100 100 100
100
! Sets Q3T1 max WRED threshold to 70%; Q3T2 max WRED threshold to 80%;
! Sets Q3T3 max WRED threshold to 90%; Q3T4 max WRED threshold to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 3 60 70 80 90 100 100 100
100
! Sets Q3T1 min WRED threshold to 60%; Q3T2 min WRED threshold to 70%;
! Sets Q3T3 min WRED threshold to 80%; Q3T4 min WRED threshold to 90%

! This section configures the CoS-to-Queue/Threshold mappings
C6500-E(config-if-range)# wrr-queue cos-map 1 1 1
! Maps CoS 1 (Scavenger and Bulk Data) to Q1T1
C6500-E(config-if-range)# wrr-queue cos-map 2 1 0
! Maps CoS 0 (Best Effort) to Q2T1
C6500-E(config-if-range)# wrr-queue cos-map 3 1 2
! Maps CoS 2 (Network Management and Transactional Data) to Q3T1
C6500-E(config-if-range)# wrr-queue cos-map 3 2 3
! Maps CoS 3 (Signaling and Multimedia Streaming) to Q3T2
C6500-E(config-if-range)# wrr-queue cos-map 3 3 6
! Maps CoS 6 (Internetwork Control) to Q3T3
C6500-E(config-if-range)# wrr-queue cos-map 3 4 7
! Maps CoS 7 (Network Control) to Q3T4
C6500-E(config-if-range)# priority-queue cos-map 1 4 5
! Maps CoS 4 (Realtime Interactive and Multimedia Conferencing) to PQ
! Maps CoS 5 (VoIP and Broadcast Video) to the PQ

```

This configuration can be verified with the command:

- **show queueing interface** (as shown in [Example 2-112](#))

Example 2-112 Verifying Queuing on a Catalyst 6500-E—show queueing interface

```

C6500-E#show queueing interface GigabitEthernet 2/1
Interface GigabitEthernet2/1 queueing strategy: Weighted Round-Robin
  Port QoS is enabled
  Trust boundary enabled

  Port is untrusted
  Extend trust state: not trusted [COS = 0]
  Default COS is 0
  Queueing Mode In Tx direction: mode-cos
  Transmit queues [type = 1p3q8t]:
  Queue Id      Scheduling  Num of thresholds
  -----
      01          WRR              08
      02          WRR              08
      03          WRR              08
      04          Priority          01

  WRR bandwidth ratios: 5[queue 1] 35[queue 2] 30[queue 3]
  queue-limit ratios: 20[queue 1] 25[queue 2] 40[queue 3] 15[Pri Queue]

  queue tail-drop-thresholds
  -----
      1      70[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
      2      70[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]

```

```
3      100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
```

```
queue random-detect-min-thresholds
```

```
-----
1      80[11] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
2      80[11] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
3      60[11] 70[2] 80[3] 90[4] 100[5] 100[6] 100[7] 100[8]
```

```
queue random-detect-max-thresholds
```

```
-----
1      100[11] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
2      100[11] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
3      70[11] 80[2] 90[3] 100[4] 100[5] 100[6] 100[7] 100[8]
```

WRED disabled queues:

```
queue thresh cos-map
```

```
-----
1      1      1
1      2
1      3
1      4
1      5
1      6
1      7
1      8
2      1      0
2      2
2      3
2      4
2      5
2      6
2      7
2      8
3      1      2
3      2      3
3      3      6
3      4      7
3      5
3      6
3      7
3      8
4      1      4 5
```

Queueing Mode In Rx direction: mode-cos

Receive queues [type = 1q8t]:

Queue Id Scheduling Num of thresholds

```
-----
01              WRR                      08
```

WRR bandwidth ratios: 100[queue 1]

queue-limit ratios: 100[queue 1]

```
queue tail-drop-thresholds
```

```
-----
1      100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
```

```
queue thresh cos-map
```

```
-----
1      1      0 1 2 3 4 5 6 7
1      2
1      3
1      4
1      5
```

```

1      6
1      7
1      8

Packets dropped on Transmit:
  BPDUs: 0

queue          dropped  [cos-map]
-----
1              8771    [1 ]
2              0       [0 ]
3              0       [2 3 6 7 ]
4              0       [4 5 ]

Packets dropped on Receive:
  BPDUs: 0

queue          dropped  [cos-map]
-----
1              0       [0 1 2 3 4 5 6 7 ]

```

C6500-E#

[Example 2-112](#) verifies that 1P3Q8T (CoS-based) egress queuing has been enabled on the interface, with the queue limits, bandwidth allocations, WRED thresholds, and CoS-to-queue mappings as described at the beginning of this section. Additionally, the “Packets Dropped on Transmit” table shows that 8771 packets were dropped from Q1 (the scavenger/bulk queue).

1P7Q8T (CoS-Based) Egress Queuing Model

In the 1P7Q8T (CoS-Based) egress queuing model, 15% of the queuing buffers and link bandwidth can be allocated for the priority queue (Q8), 15% for Q7, 5% for Q6, 5% for Q5, 15% for Q4, 15% for Q3, 25% for Q2 (the best effort queue), and 5% for Q1 (the scavenger/bulk queue). In this model, queue limits can be set to match the bandwidth allocations.

Additionally, WRED can be enabled on queues 1 through 7. Only basic WRED functionality is required for these queues (as only a single CoS value is assigned to each); therefore the first minimum WRED thresholds can be set to 80% for these queues and the first maximum WRED thresholds can be set to 100% for these queues. Since Q7 only has UDP-based flows assigned to it, the first minimum WRED threshold can also be set to 100% (to effectively disable WRED for this queue).

As eight queues exist for this queuing model, each CoS value can be assigned to a dedicated queue. CoS 5 (VoIP and broadcast video) can be mapped to the priority queue. CoS 4 (realtime interactive and multimedia conferencing) can be mapped to Q7. CoS 7 (network control) can be mapped to Q6. CoS 6 (internetwork control) can be mapped to Q5. CoS 3 (signaling and multimedia streaming) can be mapped to Q4. CoS 2 (network management and transactional data) can be mapped to Q3. CoS 0 (best effort) can be mapped to Q2 and CoS 1 (scavenger and bulk) can be mapped to Q1.

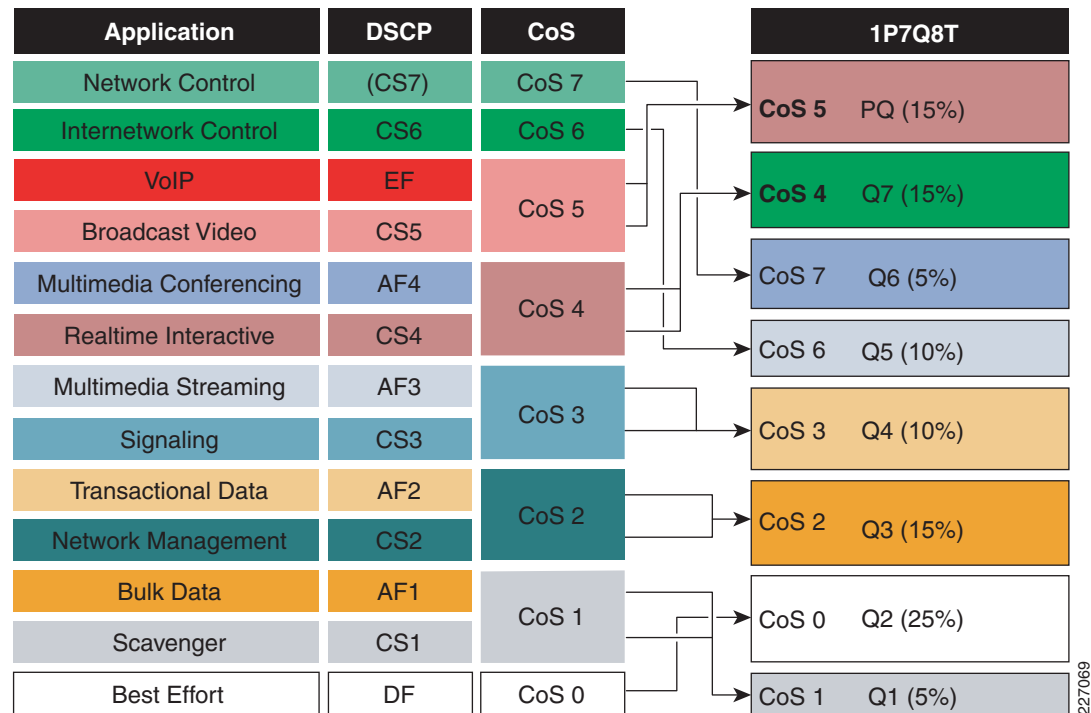


Note

As with the 1P3Q8T CoS-to-queue model, certain application classes that are normally mapped to differing queue/threshold-combinations must be mapped to the same queue/threshold in the 1P7Q8T model, because of the limited CoS-to-queue mapping level-of-granularity. These include realtime interactive and multimedia conferencing (both sharing CoS 4), signaling and multimedia streaming (both sharing CoS 3), network management and transactional data (both sharing CoS 2), and scavenger and bulk data (both sharing CoS 1).

These 1P7Q8T CoS-to-queue mappings are illustrated in [Figure 2-29](#).

Figure 2-29 Catalyst 6500-E 1P7Q8T (CoS-to-Queue) Egress Queuing Model



The corresponding configuration for 1P7Q8T (CoS-to-queue) egress queuing on a Catalyst 6500-E is shown in [Example 2-113](#).

Example 2-113 1P7Q8T (CoS-to-Queue) Egress Queuing Configuration Example on a Catalyst 6500-E

```
! This section configures 1P7Q8T (CoS-Based) Egress Queuing
C6500-E(config)#interface range TenGigabitEthernet 3/1-4
C6500-E(config-if-range)# wrr-queue queue-limit 5 25 15 15 5 5 15
! Allocates 5% to Q1, 25% to Q2, 15% to Q3, 15% to Q4,
! Allocates 5% to Q5, 5% to Q6 and 15% to Q7
C6500-E(config-if-range)# wrr-queue bandwidth 5 25 15 15 5 5 15
! Allocates 5% BW to Q1, 25% BW to Q2, 15% BW to Q3, 15% BW to Q4,
! Allocates 5% BW to Q5, 5% BW to Q6 and 15% BW to Q7
C6500-E(config-if-range)# priority-queue queue-limit 15
! Allocates 15% to the PQ
```

```
! This section enables WRED on Queues 1 through 7
C6500-E(config-if-range)# wrr-queue random-detect 1
! Enables WRED on Q1
C6500-E(config-if-range)# wrr-queue random-detect 2
! Enables WRED on Q2
C6500-E(config-if-range)# wrr-queue random-detect 3
! Enables WRED on Q3
C6500-E(config-if-range)# wrr-queue random-detect 4
! Enables WRED on Q4
C6500-E(config-if-range)# wrr-queue random-detect 5
! Enables WRED on Q5
C6500-E(config-if-range)# wrr-queue random-detect 6
! Enables WRED on Q6
```

```

C6500-E(config-if-range)# wrr-queue random-detect 7
! Enables WRED on Q7

! This section configures WRED thresholds for Queues 1 through 7
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 1 100 100 100 100 100 100
100 100
! Sets all WRED max thresholds on Q1 to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 1 80 100 100 100 100 100
100 100
! Sets Q1T1 min WRED threshold to 80%; all others set to 100%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 2 100 100 100 100 100 100
100 100
! Sets all WRED max thresholds on Q2 to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 2 80 100 100 100 100 100
100 100
! Sets Q2T1 min WRED threshold to 80%; all others set to 100%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 3 100 100 100 100 100 100
100 100
! Sets all WRED max thresholds on Q3 to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 3 80 100 100 100 100 100
100 100
! Sets Q3T1 min WRED threshold to 80%; all others set to 100%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 4 100 100 100 100 100 100
100 100
! Sets all WRED max thresholds on Q4 to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 4 80 100 100 100 100 100
100 100
! Sets Q4T1 min WRED threshold to 80%; all others set to 100%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 5 100 100 100 100 100 100
100 100
! Sets all WRED max thresholds on Q5 to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 5 80 100 100 100 100 100
100 100
! Sets Q5T1 min WRED threshold to 80%; all others set to 100%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 6 100 100 100 100 100 100
100 100
! Sets all WRED max thresholds on Q6 to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 6 80 100 100 100 100 100
100 100
! Sets Q6T1 min WRED threshold to 80%; all others set to 100%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 7 100 100 100 100 100 100
100 100
! Sets all WRED max thresholds on Q7 to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 7 100 100 100 100 100 100
100 100
! Sets all WRED max thresholds on Q7 to 100% (disabling WRED)

! This section configures the CoS-to-Queue/Threshold mappings
C6500-E(config-if-range)# wrr-queue cos-map 1 1 1
! Maps CoS 1 (Scavenger and Bulk Data) to Q1T1
C6500-E(config-if-range)# wrr-queue cos-map 2 1 0
! Maps CoS 0 (Best Effort) to Q2T1
C6500-E(config-if-range)# wrr-queue cos-map 3 1 2
! Maps CoS 2 (Network Management and Transactional Data) to Q3T1
C6500-E(config-if-range)# wrr-queue cos-map 4 1 3
! Maps CoS 3 (Signaling and Multimedia Streaming) to Q4T1
C6500-E(config-if-range)# wrr-queue cos-map 5 1 6
! Maps CoS 6 (Internetwork Control) to Q5T1
C6500-E(config-if-range)# wrr-queue cos-map 6 1 7
! Maps CoS 7 (Network Control) to Q6T1

```



```
C6500-E(config-if-range)# wrr-queue cos-map 7 1 4
! Maps CoS 4 (Realtime Interactive & Multimedia Conferencing) to Q7T1
C6500-E(config-if-range)# priority-queue cos-map 1 5
! Maps CoS 5 (VoIP and Broadcast Video) to the PQ
```

This configuration can be verified with the command:

- **show queueing interface**

1P7Q4T (DSCP-Based) Egress Queuing Model

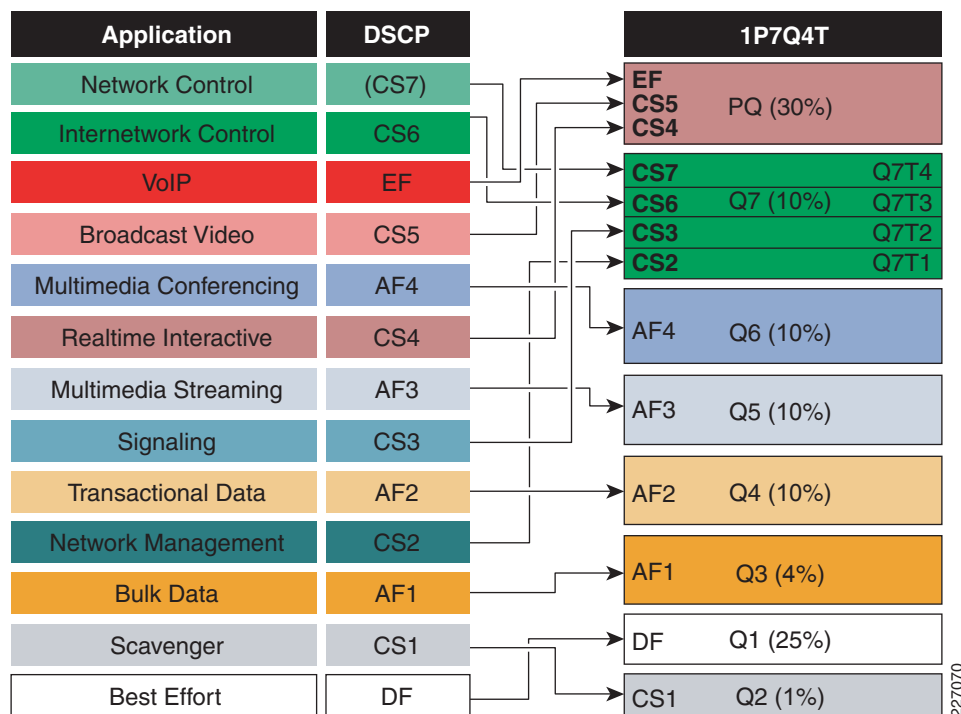
In the 1P7Q4T (DSCP-Based) egress queuing model, 30% of the link bandwidth can be allocated for the priority queue (Q8), 10% for Q7, 10% for Q6, 10% for Q5, 10% for Q4, 4% for Q3, 25% for Q2 (the best effort queue), and 1% for Q1 (the scavenger queue). In turn, 15% of the buffers can be allocated to the PQ (Q8), 10% of the buffers can be allocated (each) for Q3-Q7, 25% for Q2 (the best effort queue), and 10% for Q1 (the scavenger queue).

Additionally, WRED can be enabled on queues 1 through 7. Only basic WRED functionality is required for queues 1 and 2 (as only a single DSCP value is assigned to each); therefore the first minimum WRED thresholds for these queues can be set to 80% and the first maximum WRED thresholds for these queues can be set to 100%. As queues 3 through 6 have AF PHBs assigned to them, the WRED thresholds can be set to correspond to the three drop-precedence levels per AF class. Thus, the first three minimum WRED thresholds for these queues can be set to 70%, 80%, and 90%, respectively; and the first three maximum WRED thresholds for these queues can be set to 80%, 90%, and 100%, respectively.

Additionally, since Q7 has 4 separate DSCP values assigned to it, intra-queue QoS can be achieved by mapping these to different WRED thresholds. Thus, the minimum WRED thresholds for Q7T1, Q7T2, Q7T3, and Q7T4 can be set to 60%, 70%, 80%, and 90%, respectively; and the minimum WRED thresholds for Q7T1, Q7T2, Q7T3, and Q7T4 can be set to 70%, 80%, 90%, and 100%, respectively.

DSCP EF (VoIP), CS5 (broadcast video), and CS4 (realtime interactive) can be mapped to the priority queue. CS7 (network control) can be mapped to Q7T4; CS6 (internetwork control) can be mapped to Q7T3; CS3 (signaling) can be mapped to Q7T2; and CS2 (network management) can be mapped to Q7T1. AF4 (multimedia conferencing) can be mapped to Q6. AF3 (multimedia streaming) can be mapped to Q5. AF2 (transactional data) can be mapped to Q4. AF1 (bulk data) can be mapped to Q3. DF (best effort) can be mapped to Q2. And CS1 can be mapped to Q1.

These 1P7Q4T DSCP-to-queue mappings are illustrated in [Figure 2-30](#).

Figure 2-30 Catalyst 6500-E 1P7Q4T (DSCP-to-Queue) Egress Queuing Model

The corresponding configuration for 1P7Q4T (DSCP-to-queue) egress queuing on a Catalyst 6500-E is shown in [Example 2-114](#).

Example 2-114 1P7Q4T (DSCP-to-Queue) Egress Queuing Configuration Example on a Catalyst 6500-E

```

! This section configures 1P7Q4T (DSCP-Based) Egress Queuing
C6500-E(config)#interface range TenGigabitEthernet 4/1-8
C6500-E(config-if-range)# wrr-queue queue-limit 10 25 10 10 10 10 10
! Allocates 10% of the buffers to Q1, 25% to Q2, 10% to Q3, 10% to Q4,
! Allocates 10% to Q5, 10% to Q6 and 10% to Q7
C6500-E(config-if-range)# wrr-queue bandwidth 1 25 4 10 10 10 10
! Allocates 1% BW to Q1, 25% BW to Q2, 4% BW to Q3, 10% BW to Q4,
! Allocates 10% BW to Q5, 10% BW to Q6 and 10% BW to Q7
C6500-E(config-if-range)# priority-queue queue-limit 15
! Allocates 15% of the buffers to the PQ

! This section enables WRED on Queues 1 through 7
C6500-E(config-if-range)# wrr-queue random-detect 1
! Enables WRED on Q1
C6500-E(config-if-range)# wrr-queue random-detect 2
! Enables WRED on Q2
C6500-E(config-if-range)# wrr-queue random-detect 3
! Enables WRED on Q3
C6500-E(config-if-range)# wrr-queue random-detect 4
! Enables WRED on Q4
C6500-E(config-if-range)# wrr-queue random-detect 5
! Enables WRED on Q5
C6500-E(config-if-range)# wrr-queue random-detect 6
! Enables WRED on Q6
C6500-E(config-if-range)# wrr-queue random-detect 7
! Enables WRED on Q7

```

```

! This section configures WRED thresholds for Queues 1 through 7
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 1 100 100 100 100
! Sets all WRED max thresholds on Q1 to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 1 80 100 100 100
! Sets Q1T1 min WRED threshold to 80%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 2 100 100 100 100
! Sets all WRED max thresholds on Q2 to 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 2 80 100 100 100
! Sets Q2T1 min WRED threshold to 80%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 3 80 90 100 100
! Sets WRED max thresholds for Q3T1, Q3T2, Q3T3 to 80%, 90% and 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 3 70 80 90 100
! Sets WRED min thresholds for Q3T1, Q3T2, Q3T3 to 70 %, 80% and 90%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 4 70 80 90 100
! Sets WRED min thresholds for Q4T1, Q4T2, Q4T3 to 70 %, 80% and 90%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 4 80 90 100 100
! Sets WRED max thresholds for Q4T1, Q4T2, Q4T3 to 80%, 90% and 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 5 70 80 90 100
! Sets WRED min thresholds for Q5T1, Q5T2, Q5T3 to 70 %, 80% and 90%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 5 80 90 100 100
! Sets WRED max thresholds for Q5T1, Q5T2, Q5T3 to 80%, 90% and 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 6 70 80 90 100
! Sets WRED min thresholds for Q6T1, Q6T2, Q6T3 to 70 %, 80% and 90%
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 6 80 90 100 100
! Sets WRED max thresholds for Q6T1, Q6T2, Q6T3 to 80%, 90% and 100%
C6500-E(config-if-range)# wrr-queue random-detect min-threshold 7 60 70 80 90
! Sets WRED min thresholds for Q7T1, Q7T2, Q7T3 and Q7T4
! to 60%, 70%, 80% and 90%, respectively
C6500-E(config-if-range)# wrr-queue random-detect max-threshold 7 70 80 90 100
! Sets WRED max thresholds for Q7T1, Q7T2, Q7T3 and Q7T4
! to 70%, 80%, 90% and 100%, respectively

```

```

! This section enables DSCP-to-Queue/Threshold mappings
! and configures the DSCP-to-Queue/Threshold mappings
C6500-E(config-if-range)# mls qos queue-mode mode-dscp
! Enables DSCP-to-Queue mapping
C6500-E(config-if-range)# wrr-queue dscp-map 1 1 8
! Maps CS1 (Scavenger) to Q1T1
C6500-E(config-if-range)# wrr-queue dscp-map 2 1 0
! Maps DF (Best Effort) to Q2T1
C6500-E(config-if-range)# wrr-queue dscp-map 3 1 14
! Maps AF13 (Bulk Data-Drop Precedence 3) to Q3T1
C6500-E(config-if-range)# wrr-queue dscp-map 3 2 12
! Maps AF12 (Bulk Data-Drop Precedence 2) to Q3T2
C6500-E(config-if-range)# wrr-queue dscp-map 3 3 10
! Maps AF11 (Bulk Data-Drop Precedence 1) to Q3T3
C6500-E(config-if-range)# wrr-queue dscp-map 4 1 22
! Maps AF23 (Transactional Data-Drop Precedence 3) to Q4T1
C6500-E(config-if-range)# wrr-queue dscp-map 4 2 20
! Maps AF22 (Transactional Data-Drop Precedence 2) to Q4T2
C6500-E(config-if-range)# wrr-queue dscp-map 4 3 18
! Maps AF21 (Transactional Data-Drop Precedence 1) to Q4T3
C6500-E(config-if-range)# wrr-queue dscp-map 5 1 30
! Maps AF33 (Multimedia Streaming-Drop Precedence 3) to Q5T1
C6500-E(config-if-range)# wrr-queue dscp-map 5 2 28
! Maps AF32 (Multimedia Streaming-Drop Precedence 2) to Q5T2
C6500-E(config-if-range)# wrr-queue dscp-map 5 3 26
! Maps AF31 (Multimedia Streaming-Drop Precedence 1) to Q5T3
C6500-E(config-if-range)#
C6500-E(config-if-range)# wrr-queue dscp-map 6 1 38

```

```

! Maps AF43 (Multimedia Conferencing-Drop Precedence 3) to Q6T1
C6500-E(config-if-range)# wrr-queue dscp-map 6 2 36
! Maps AF42 (Multimedia Conferencing-Drop Precedence 2) to Q6T2
C6500-E(config-if-range)# wrr-queue dscp-map 6 3 34
! Maps AF41 (Multimedia Conferencing-Drop Precedence 1) to Q6T3
C6500-E(config-if-range)# wrr-queue dscp-map 7 1 16
! Maps CS2 (Network Management) to Q7T1
C6500-E(config-if-range)# wrr-queue dscp-map 7 2 24
! Maps CS3 (Signaling) to Q7T2
C6500-E(config-if-range)# wrr-queue dscp-map 7 3 48
! Maps CS6 (Internetwork Control) to Q7T3
C6500-E(config-if-range)# wrr-queue dscp-map 7 4 56
! Maps CS7 (Network Control) to Q7T4
C6500-E(config-if-range)# priority-queue dscp-map 1 32 40 46
! Maps CS4 (Realtime Interactive), CS5 (Broadcast Video),
! and EF (VoIP) to the PQ

```

**Note**

Due to the default WRED threshold settings, at times the maximum threshold needs to be configured before the minimum (as is the case on queues 1 through 3 in the example above); at other times, the minimum threshold needs to be configured before the maximum (as is the case on queues 4 through 7 in the example above).

This configuration can be verified with the command:

- **show queueing interface**

EtherChannel QoS Model

As discussed in [EtherChannel QoS](#), QoS policies on the Catalyst 6500/6500-E need to be separated, such that ingress trust, classification, marking, and/or policing policies are applied to the logical Port-Channel interface, whereas queuing policies (both ingress—if required—and egress) are applied directly on the physical port-member interfaces, as shown in [Example 2-115](#).

Example 2-115 EtherChannel QoS Design on a Catalyst 6500/6500-E

```

! This section configures the (logical) EtherChannel interface
C6500-E(config)# interface Port-channel1
C6500-E(config-if)# description ETHERCHANNEL-TRUNK-TO-DISTRIBUTION-LAYER
C6500-E(config-if)# switchport mode trunk
C6500-E(config-if)# switchport trunk encapsulation dot1q
C6500-E(config-if)# switchport trunk allowed vlan 10,110
C6500-E(config-if)# mls qos trust dscp
! The logical EtherChannel interface is set to statically trust DSCP

! This section configures (Optional Ingress and) Egress Queuing across the
! (physical) EtherChannel member-ports
C6500-E(config)# interface range TenGigabitEthernet6/1-2
C6500-E(config-if-range)# description PORT-CHANNEL1-PHYSICAL-PORT-MEMBER
C6500-E(config-if-range)# switchport mode trunk
C6500-E(config-if-range)# switchport trunk encapsulation dot1q
C6500-E(config-if-range)# switchport trunk allowed vlan 10,110
C6500-E(config-if-range)# channel-group 1 mode auto
C6500-E(config-if-range)# <Optional Linecard-specific Ingress Queuing Configuration>
! Optional: Either 8Q4T Ingress Queuing Config (if supported)
C6500-E(config-if-range)# <Linecard-specific Egress Queuing Configuration>
! 1P3Q8T (CoS-based) or 1P7Q8T (CoS-based) or 1P7Q4T (DSCP-based) Egress Queuing

```

This configuration can be verified with the command:

- **show queueing interface**

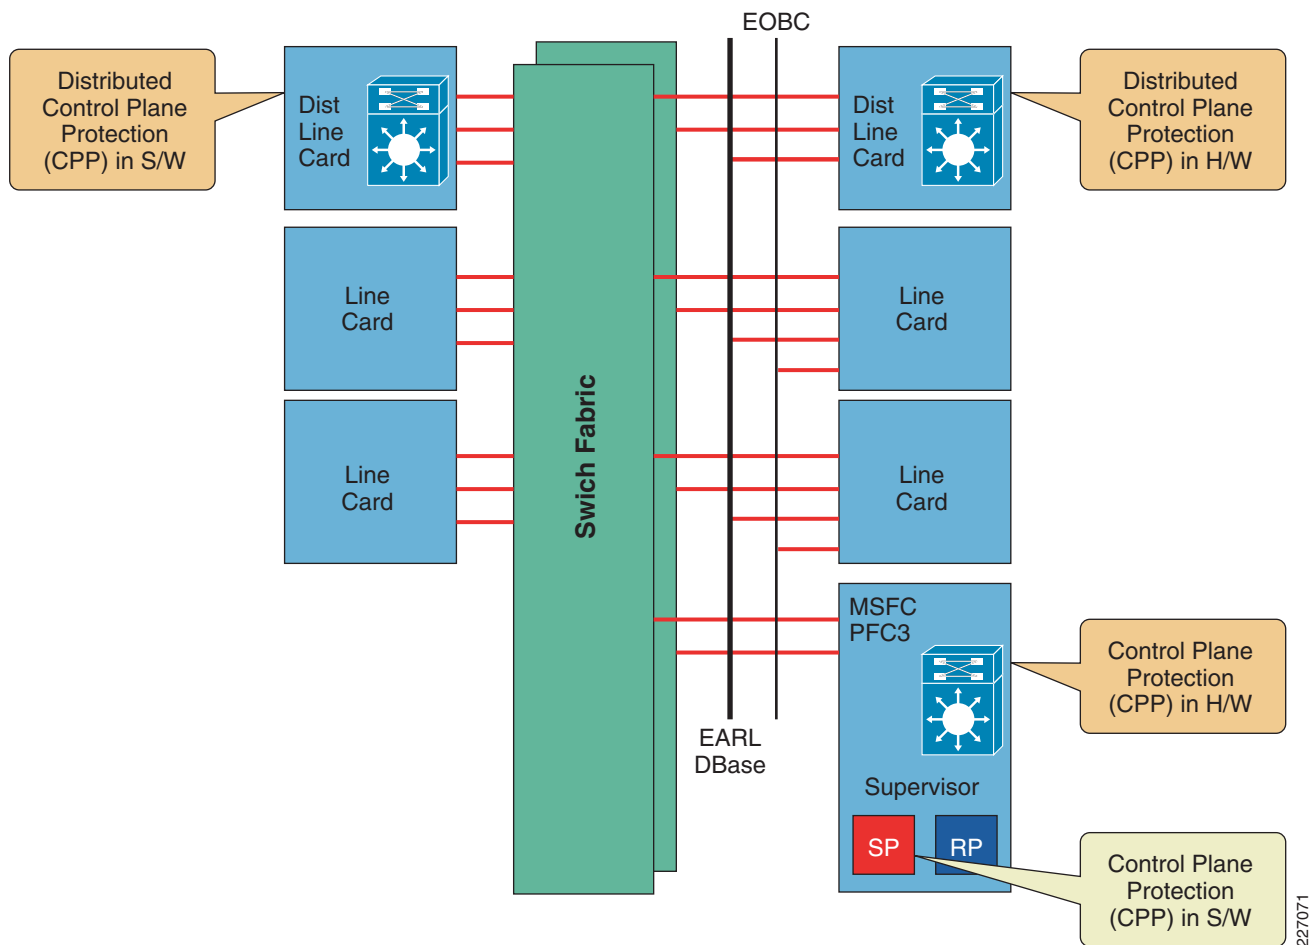
Control Plane Policing

As previously stated, the Catalyst 4500 and Catalyst 6500 Series switches implement CoPP similarly; however, CoPP has been enhanced on both platforms to leverage the benefits of their hardware architectures and as a result each platform provides unique features.

This section describes the implementation details of CoPP on Supervisors 720 and 32.

In the Catalyst 6500 Series switches, CoPP takes advantage of the processing power present on line-cards by implementing a distributed CoPP model. In this platform, the class QoS policies are centrally configured under the control plane configuration mode. When configured, these policies are first applied at the route processor (MSFC) level and then they get automatically pushed to the Policy Feature Card (PFC) and each Distributed Forwarding Card (DFC). This CoPP model is illustrated in Figure 2-31.

Figure 2-31 Catalyst 6500 Supervisor 720/Supervisor 32 Control Plane Policing Implementation



227071

CoPP at the RP is performed in software, while on the PFC and DFCs it is processed in hardware, with no performance degradation or increased latency. In this way, CoPP on Supervisors 32 and 720 provides two layers of protection: first, at wire speed on the PFC and DFCs, and second, at the Route Processor (RP) level. This helps to ensure that only the amount of traffic specified by the user actually reaches the control plane.

**Note**

The PFC3 and DFC3 provide hardware support for CoPP. However, CoPP is not enforced in hardware unless MLS QoS is globally enabled using the **mls qos** global command.

The Cisco Catalyst 6500 supports CoPP on the Supervisor 720 and Supervisor 32 in hardware starting with Cisco IOS release 12.2(18)SXD1. CoPP supports IPv4 in hardware, while multicast and broadcast traffic are only supported in software. Support for IPv6 traffic has been introduced in IOS release 12.2(18)SXE.

Another important characteristic of CoPP in Supervisors 720 and 32 is that it does not support the definition of non-IP traffic classes, with the exception of the class default. Class-default is a default class for all remaining traffic destined to the RP that does not match any other class. This default class allows you to specify how to treat traffic that is not explicitly associated with any other user-defined classes. The class-default is the only class in CoPP capable of handling both IP and non-IP traffic. User-defined classes can only handle IP traffic.

CoPP helps protect the RP of Catalyst 6500 Series switches in multiple ways. From a policing perspective, by filtering traffic sent to the RP, CoPP ensures that only the expected protocols are allowed. This effectively shields the control plane from unwanted and potentially malicious traffic. On the other hand, by rate limiting the traffic sent to the RP, CoPP provides protection against large volumes of packets that might be part of a DoS attack, which helps maintain network stability even during an attack.

CoPP Configuration

To configure CoPP on Supervisors 720 and 32 (Catalyst 6500):

-
- Step 1** Ensure QoS is enabled on Supervisors 32 and 720 with the **mls qos** global configuration command.
 - Step 2** Optionally, define the necessary ACLs to be used to match traffic classes.
 - Step 3** Classify the control plane traffic using the **class-map** command.
 - Step 4** After the traffic is classified, you apply a **policy-map** with a **police** action to each class, indicating whether to permit all packets, to drop all packets, or to drop packets crossing a specified rate limit for that particular class.
 - Step 5** Apply the defined CoPP policy to the control plane by using the **service-policy** command from control plane configuration mode.
-

**Note**

For more information refer to the Configuring Control Plane Policing documentation for the Catalyst 6500 at:
<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/copp.html>.

CoPP Considerations and Restrictions

The following are important considerations and known restrictions that should be taken into account prior to configuring CoPP on the Catalyst 6500:

- Because CoPP relies on the QoS implementation, CoPP policies are downloaded to the PFC and DFCs only if QoS is enabled. For this reason, ensure that the **mls qos** command is enabled at the global configuration mode for the PFC and each DFC where CoPP is required.
- CoPP does not support the definition of non-IP traffic classes except for the class-default. ACLs can be used instead of non-IP classes to drop non-IP traffic. At the same time, class-default can be used to limit non-IP traffic that reaches the RP CPU.
- On Supervisors 32 and 720, ARP policing is done with a QoS rate limiter rather than CoPP. Even though there is a match protocol arp for CoPP on these supervisors, this type of traffic is processed in software. Therefore, ARP policing should be configured with the hardware-based QoS rate limiter using the **mls qos protocol arp police bps** command.
- Prior to Cisco IOS software Release 12.2(18)SXE, only one match criteria was allowed for each traffic class. When using one of these earlier releases, to define multiple match rules with a match-any criteria, split the match access-group statements among multiple class maps instead of grouping them together.
- Prior to Cisco IOS software Release 12.2(18)SXE, the MQC class-default was not supported on Supervisor 720. This is a minor limitation because the class-default could be emulated with a normal class configured with an ip permit any rule.
- Omitting the policy parameters in a class causes the class to be handled by software-based CoPP. Use the **police** command and set the policy parameters to ensure the class is handled by hardware-based CoPP.
- Currently, multicast packets are handled only by the software-based CoPP at the RP level. However, there are CPU rate limiters available that can rate limit multicast packets to the CPU in hardware. These CPU rate limiters include the Multicast FIB-miss rate limiter and the Multicast Partial-SC rate limiter. These CPU rate limiters can be used in combination with ACLs and software CoPP to provide protection against multicast and DoS attacks.
- CoPP is not supported in hardware for broadcast packets. The combination of ACLs, traffic storm control, and CoPP software protection provides protection against broadcast DoS attacks.
- With PFC3A, egress QoS and CoPP cannot be configured at the same time. In this situation, CoPP is performed in software and a warning message is generated.
- In the rare situation where a large QoS configuration is being used, it is possible that the system could run out of TCAM space. When this scenario occurs, CoPP can be performed in software. Use the **show platform hardware capacity** command to monitor TCAM space.
- You must ensure that the CoPP policy does not filter critical traffic such as routing protocols or interactive access to the switches. Filtering this traffic could prevent remote access to the switch, requiring a console connection.
- Supervisor Engines 32 and 720 support built-in special-case rate limiters, which are useful for situations where an ACL cannot be used (for example, TTL, MTU, and IP options). When you enable the special-case rate limiters, you should be aware that the special-case rate limiters override the CoPP policy for packets matching the rate-limiter criteria.
- CoPP does not support ACEs with the **log** keyword.
- CoPP uses hardware QoS TCAM resources. Use the **show platform hardware capacity** and **show tcam utilization** commands to verify the TCAM use.

- ACE hit counters in hardware are only for ACL logic. You can rely on software ACE hit counters and the **show access-list**, **show policy-map control-plane**, and **show mls ip qos** commands to troubleshoot evaluate CPU traffic.

CoPP Model

In [Example 2-116](#), CoPP has been deployed on the Catalyst 6500 inline with the recommendations for CoPP class definitions and deployment models presented earlier in this chapter.

Example 2-116 Control Plane Policing Model on a Catalyst 6500

```
!This section defines the CoPP Access-Lists
C6500-E(config)#ip access-list extended COPP-ACL-BGP
C6500-E(config-ext-nacl)# remark BGP
C6500-E(config-ext-nacl)# permit tcp host 192.168.1.1 host 10.1.1.1 eq bgp
C6500-E(config-ext-nacl)# permit tcp host 192.168.1.1 eq bgp host 10.1.1.1

C6500-E(config)#ip access-list extended COPP-ACL-IGP
C6500-E(config-ext-nacl)# remark IGP (OSPF)
C6500-E(config-ext-nacl)# permit ospf any host 224.0.0.5
C6500-E(config-ext-nacl)# permit ospf any host 224.0.0.6
C6500-E(config-ext-nacl)# permit ospf any any

C6500-E(config)#ip access-list extended COPP-ACL-INTERACTIVE-MANAGEMENT
C6500-E(config-ext-nacl)# remark TACACS (return traffic)
C6500-E(config-ext-nacl)# permit tcp host 10.2.1.1 host 10.1.1.1 established
C6500-E(config-ext-nacl)# remark SSH
C6500-E(config-ext-nacl)# permit tcp 10.2.1.0 0.0.0.255 host 10.1.1.1 eq 22
C6500-E(config-ext-nacl)# remark SNMP
C6500-E(config-ext-nacl)# permit udp host 10.2.2.2 host 10.1.1.1 eq snmp
C6500-E(config-ext-nacl)# remark NTP
C6500-E(config-ext-nacl)# permit udp host 10.2.2.3 host 10.1.1.1 eq ntp

C6500-E(config)#ip access-list extended COPP-ACL-FILE-MANAGEMENT
C6500-E(config-ext-nacl)# remark (initiated) FTP (active and passive)
C6500-E(config-ext-nacl)# permit tcp 10.2.1.0 0.0.0.255 eq 21 host 10.1.1.1 gt 1023
established
C6500-E(config-ext-nacl)# permit tcp 10.2.1.0 0.0.0.255 eq 20 host 10.1.1.1 gt 1023
C6500-E(config-ext-nacl)# permit tcp 10.2.1.0 0.0.0.255 gt 1023 host 10.1.1.1 gt 1023
established
C6500-E(config-ext-nacl)# remark (initiated) TFTP
C6500-E(config-ext-nacl)# permit udp 10.2.1.0 0.0.0.255 gt 1023 host 10.1.1.1 gt 1023

C6500-E(config)#ip access-list extended COPP-ACL-MONITORING
C6500-E(config-ext-nacl)# remark PING-ECHO
C6500-E(config-ext-nacl)# permit icmp any any echo
C6500-E(config-ext-nacl)# remark PING-ECHO-REPLY
C6500-E(config-ext-nacl)# permit icmp any any echo-reply
C6500-E(config-ext-nacl)# remark TRACEROUTE
C6500-E(config-ext-nacl)# permit icmp any any ttl-exceeded
C6500-E(config-ext-nacl)# permit icmp any any port-unreachable

C6500-E(config)#ip access-list extended COPP-ACL-CRITICAL-APPLICATIONS
C6500-E(config-ext-nacl)# remark HSRP
C6500-E(config-ext-nacl)# permit ip any host 224.0.0.2
C6500-E(config-ext-nacl)# remark DHCP
C6500-E(config-ext-nacl)# permit udp host 0.0.0.0 host 255.255.255.255 eq bootps
C6500-E(config-ext-nacl)# permit udp host 10.2.2.8 eq bootps any eq bootps

C6500-E(config)#ip access-list extended COPP-ACL-UNDESIRABLE
C6500-E(config-ext-nacl)# remark UNDESIRABLE
```



```

C6500-E(config-ext-nacl)# permit udp any any eq 1434

! This section defines the CoPP Policy Class-Maps
C6500-E(config)# class-map match-all COPP-ACL-BGP
C6500-E(config-cmap)# match access-group name COPP-ACL-BGP
! Associates COPP-BGP ACL with class-map

C6500-E(config)#class-map match-all COPP-ACL-IGP
C6500-E(config-cmap)# match access-group name COPP-ACL-IGP
! Associates COPP-IGP ACL with class-map

C6500-E(config)#class-map match-all COPP-ACL-INTERACTIVE-MANAGEMENT
C6500-E(config-cmap)# match access-group name COPP-ACL-INTERACTIVE-MANAGEMENT
! Associates COPP-INTERACTIVE-MANAGEMENT ACL with class-map

C6500-E(config)#class-map match-all COPP-ACL-FILE-MANAGEMENT
C6500-E(config-cmap)# match access-group name COPP-ACL-FILE-MANAGEMENT
! Associates COPP-FILE-MANAGEMENT with class-map

C6500-E(config)#class-map match-all COPP-ACL-MONITORING
C6500-E(config-cmap)# match access-group name COPP-ACL-MONITORING
! Associates COPP-MONITORING ACL with class-map

C6500-E(config)#class-map match-all COPP-ACL-CRITICAL-APPLICATIONS
C6500-E(config-cmap)# match access-group name COPP-ACL-CRITICAL-APPLICATIONS
! Associates COPP-CRITICAL-APPLICATIONS ACL with class-map

C6500-E(config)#class-map match-all COPP-ACL-UNDESIRABLE
C6500-E(config-cmap)# match access-group name COPP-ACL-UNDESIRABLE
! Associates COPP-UNDESIRABLE ACL with class-map

! This section defines the CoPP Policy
C6500-E(config)# policy-map COPP-POLICY
C6500-E(config-pmap)# class COPP-ACL-BGP
C6500-E(config-pmap-c)# police cir 4000000 bc 400000 be 400000
C6500-E(config-pmap-c-police)# conform-action transmit
C6500-E(config-pmap-c-police)# exceed-action drop
! Polices BGP to 4 Mbps
C6500-E(config-pmap)# class COPP-ACL-IGP
C6500-E(config-pmap-c)# police cir 300000 bc 3000 be 3000
C6500-E(config-pmap-c-police)# conform-action transmit
C6500-E(config-pmap-c-police)# exceed-action drop
! Polices IGP to 300 kbps
C6500-E(config-pmap)# class COPP-ACL-INTERACTIVE-MANAGEMENT
C6500-E(config-pmap-c)# police cir 500000 bc 5000 be 5000
C6500-E(config-pmap-c-police)# conform-action transmit
C6500-E(config-pmap-c-police)# exceed-action drop
! Polices Interactive Management to 500 kbps
C6500-E(config-pmap)# class COPP-ACL-FILE-MANAGEMENT
C6500-E(config-pmap-c)# police cir 6000000 bc 60000 be 60000
C6500-E(config-pmap-c-police)# conform-action transmit
C6500-E(config-pmap-c-police)# exceed-action drop
! Polices File Management to 6 Mbps
C6500-E(config-pmap)# class COPP-ACL-MONITORING
C6500-E(config-pmap-c)# police cir 900000 bc 9000 be 9000
C6500-E(config-pmap-c-police)# conform-action transmit
C6500-E(config-pmap-c-police)# exceed-action drop
! Polices Monitoring to 900 kbps
C6500-E(config-pmap)# class COPP-ACL-CRITICAL-APPLICATIONS
C6500-E(config-pmap-c)# police cir 900000 bc 9000 be 9000

```

```

C6500-E(config-pmap-c-police)# conform-action transmit
C6500-E(config-pmap-c-police)# exceed-action drop
! Polices Critical Applications to 900 Kbps
C6500-E(config-pmap-)# class COPP-ACL-UNDESIRABLE
C6500-E(config-pmap-c)# police cir 32000 bc 3000 be 3000
C6500-E(config-pmap-c-police)# conform-action drop
C6500-E(config-pmap-c-police)# exceed-action drop
! Polices all Undesirable traffic (conform-action is drop)
C6500-E(config-pmap)# class class-default
C6500-E(config-pmap-c)# police cir 500000 bc 5000 be 5000
C6500-E(config-pmap-c-police)# conform-action transmit
C6500-E(config-pmap-c-police)# exceed-action drop
! Polices all other Control Plane traffic to 500 kbps

! This section attaches the CoPP policy to the Control Plane
C6500-E(config)#control-plane
C6500-E(config-cp)# service-policy input COPP-POLICY
! Attaches CoPP policy to control plane

```

This configuration can be verified with the commands:

- **show class-map**
- **show policy-map**
- **show policy-map control-plane** (as shown in [Example 2-117](#))

Example 2-117 Verifying Control Plane Policing on a Catalyst 6500—show policy-map control-plane

```
C6500-E#show policy-map control-plane
```

Control Plane Interface

Service-policy input: COPP-POLICY

Hardware Counters:

```

class-map: COPP-ACL-BGP (match-all)
  Match: access-group name COPP-ACL-BGP
  police :
    4000000 bps 400000 limit 400000 extended limit
  Earl in slot 1 :
    307094 bytes
    5 minute offered rate 6072 bps
    aggregate-forwarded 307094 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 10544 bps exceed 0 bps
  Earl in slot 4 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 0 bps exceed 0 bps

```

Software Counters:

```

Class-map: COPP-ACL-BGP (match-all)
  1375 packets, 1011718 bytes
  5 minute offered rate 11000 bps, drop rate 0000 bps
  Match: access-group name COPP-ACL-BGP
  police:
    cir 4000000 bps, bc 400000 bytes, be 400000 bytes
    conformed 1381 packets, 1016146 bytes; actions:

```

```

        transmit
    exceeded 0 packets, 0 bytes; actions:
        drop
    violated 0 packets, 0 bytes; actions:
        drop
    conformed 11000 bps, exceed 0000 bps, violate 0000 bps

```

Hardware Counters:

```

class-map: COPP-ACL-IGP (match-all)
  Match: access-group name COPP-ACL-IGP
  police :
    296000 bps 3000 limit 3000 extended limit
  Earl in slot 1 :
    243718 bytes
    5 minute offered rate 4880 bps
    aggregate-forwarded 243718 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 8368 bps exceed 0 bps
  Earl in slot 4 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 0 bps exceed 0 bps

```

Software Counters:

```

Class-map: COPP-ACL-IGP (match-all)
  625 packets, 646994 bytes
  5 minute offered rate 8000 bps, drop rate 0000 bps
  Match: access-group name COPP-ACL-IGP
  police:
    cir 300000 bps, bc 3000 bytes, be 3000 bytes
    conformed 628 packets, 650120 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    violated 0 packets, 0 bytes; actions:
      drop
    conformed 8000 bps, exceed 0000 bps, violate 0000 bps

```

Hardware Counters:

```

class-map: COPP-ACL-INTERACTIVE-MANAGEMENT (match-all)
  Match: access-group name COPP-ACL-INTERACTIVE-MANAGEMENT
  police :
    496000 bps 5000 limit 5000 extended limit
  Earl in slot 1 :
    8586 bytes
    5 minute offered rate 64 bps
    aggregate-forwarded 8586 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 0 bps exceed 0 bps
  Earl in slot 4 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 0 bps exceed 0 bps

```

Software Counters:

```

Class-map: COPP-ACL-INTERACTIVE-MANAGEMENT (match-all)

```

```

1953 packets, 130426 bytes
5 minute offered rate 0000 bps, drop rate 0000 bps
Match: access-group name COPP-ACL-INTERACTIVE-MANAGEMENT
police:
    cir 500000 bps, bc 5000 bytes, be 5000 bytes
    conformed 1890 packets, 126646 bytes; actions:
        transmit
    exceeded 76 packets, 4560 bytes; actions:
        drop
    violated 0 packets, 0 bytes; actions:
        drop
    conformed 0000 bps, exceed 0000 bps, violate 0000 bps

```

Hardware Counters:

```

class-map: COPP-ACL-FILE-MANAGEMENT (match-all)
Match: access-group name COPP-ACL-FILE-MANAGEMENT
police :
    6000000 bps 60000 limit 60000 extended limit
Earl in slot 1 :
    2622292 bytes
    5 minute offered rate 49808 bps
    aggregate-forwarded 2622292 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 75080 bps exceed 0 bps
Earl in slot 4 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 0 bps exceed 0 bps

```

Software Counters:

```

Class-map: COPP-ACL-FILE-MANAGEMENT (match-all)
14035 packets, 7729999 bytes
5 minute offered rate 84000 bps, drop rate 0000 bps
Match: access-group name COPP-ACL-FILE-MANAGEMENT
police:
    cir 6000000 bps, bc 60000 bytes, be 60000 bytes
    conformed 14080 packets, 7754177 bytes; actions:
        transmit
    exceeded 0 packets, 0 bytes; actions:
        drop
    violated 0 packets, 0 bytes; actions:
        drop
    conformed 84000 bps, exceed 0000 bps, violate 0000 bps

```

Hardware Counters:

```

class-map: COPP-ACL-MONITORING (match-all)
Match: access-group name COPP-ACL-MONITORING
police :
    896000 bps 9000 limit 9000 extended limit
Earl in slot 1 :
    359982 bytes
    5 minute offered rate 7040 bps
    aggregate-forwarded 359982 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 12336 bps exceed 0 bps
Earl in slot 4 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: transmit

```

```

exceeded 0 bytes action: drop
aggregate-forward 0 bps exceed 0 bps

```

Software Counters:

```

Class-map: COPP-ACL-MONITORING (match-all)
  812 packets, 627354 bytes
  5 minute offered rate 12000 bps, drop rate 0000 bps
Match: access-group name COPP-ACL-MONITORING
police:
  cir 900000 bps, bc 9000 bytes, be 9000 bytes
  conformed 818 packets, 631956 bytes; actions:
    transmit
  exceeded 0 packets, 0 bytes; actions:
    drop
  violated 0 packets, 0 bytes; actions:
    drop
  conformed 12000 bps, exceed 0000 bps, violate 0000 bps

```

Hardware Counters:

```

class-map: COPP-ACL-CRITICAL-APPLICATIONS (match-all)
  Match: access-group name COPP-ACL-CRITICAL-APPLICATIONS
  police :
    896000 bps 9000 limit 9000 extended limit
  Earl in slot 1 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 0 bps exceed 0 bps
  Earl in slot 4 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 0 bps exceed 0 bps

```

Software Counters:

```

Class-map: COPP-ACL-CRITICAL-APPLICATIONS (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
Match: access-group name COPP-ACL-CRITICAL-APPLICATIONS
police:
  cir 900000 bps, bc 9000 bytes, be 9000 bytes
  conformed 0 packets, 0 bytes; actions:
    transmit
  exceeded 0 packets, 0 bytes; actions:
    drop
  violated 0 packets, 0 bytes; actions:
    drop
  conformed 0000 bps, exceed 0000 bps, violate 0000 bps

```

Hardware Counters:

```

class-map: COPP-ACL-UNDESIRABLE (match-all)
  Match: access-group name COPP-ACL-UNDESIRABLE
  police :
    32000 bps 3000 limit 3000 extended limit
  Earl in slot 1 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: drop

```

```

exceeded 0 bytes action: drop
aggregate-forward 0 bps exceed 0 bps
Earl in slot 4 :
0 bytes
5 minute offered rate 0 bps
aggregate-forwarded 0 bytes action: drop
exceeded 0 bytes action: drop
aggregate-forward 0 bps exceed 0 bps

```

Software Counters:

```

Class-map: COPP-ACL-UNDESIRABLE (match-all)
0 packets, 0 bytes
5 minute offered rate 0000 bps, drop rate 0000 bps
Match: access-group name COPP-ACL-UNDESIRABLE
police:
  cir 32000 bps, bc 3000 bytes, be 3000 bytes
  conformed 0 packets, 0 bytes; actions:
    drop
  exceeded 0 packets, 0 bytes; actions:
    drop
  violated 0 packets, 0 bytes; actions:
    drop
  conformed 0000 bps, exceed 0000 bps, violate 0000 bps

```

Hardware Counters:

```

class-map: class-default (match-any)
Match: any
police :
  496000 bps 5000 limit 5000 extended limit
Earl in slot 1 :
40150 bytes
5 minute offered rate 968 bps
aggregate-forwarded 40150 bytes action: transmit
exceeded 0 bytes action: drop
aggregate-forward 1376 bps exceed 0 bps
Earl in slot 4 :
0 bytes
5 minute offered rate 0 bps
aggregate-forwarded 0 bytes action: transmit
exceeded 0 bytes action: drop
aggregate-forward 0 bps exceed 0 bps

```

Software Counters:

```

Class-map: class-default (match-any)
6778 packets, 553042 bytes
5 minute offered rate 0000 bps, drop rate 0000 bps
Match: any
6778 packets, 553042 bytes
5 minute rate 0 bps
police:
  cir 500000 bps, bc 5000 bytes, be 5000 bytes
  conformed 6690 packets, 547731 bytes; actions:
    transmit
  exceeded 92 packets, 5563 bytes; actions:
    drop
  violated 0 packets, 0 bytes; actions:
    drop
  conformed 0000 bps, exceed 0000 bps, violate 0000 bps

```

C6500-E#

[Example 2-117](#) shows sample traffic being matched across various control plane traffic classes, including some traffic in class-default that is being dropped (92 packets).

Summary

This chapter had four main sections. The first section discussed QoS design considerations that relate to enterprise medianet campus networks, and the second, third, and fourth sections, in turn, presented detailed design recommendations for the Catalyst desktop switch family (Catalyst 2960 & 2975, 3560G & 3750G, and 3560-E & 3750-E), Catalyst 4500/4500-E switch family, and the Catalyst 6500/6500-E switch family, respectively.

QoS design considerations discussed included Internal DSCP, trust states and operations, port-based or VLAN-based QoS, campus QoS models and port roles, and control plane policing.

The internal DSCP was shown to be the primary mechanism for QoS processing in most Cisco Catalyst switches and is determined by the trust state of the interface on which the packet enters the switch. These trust states include trust CoS, trust DSCP, and conditional trust. Trust CoS and trust DSCP are static port trust states that accept the Layer 2 or Layer 3 QoS markings of a packet, respectively. Conditional trust performs a CDP-based negotiation between the access switch and the endpoint, which—if successful and permitted by policy—results in a dynamic extension of either trust CoS or trust DSCP to the endpoint.

Beyond discussing basic ingress QoS policies, like trust, more complex ingress QoS policies were also presented, including applying QoS policies to physical interfaces (port-based QoS), logical interfaces (VLAN-based QoS), or to a combination of physical and logical interfaces (per-port/per-VLAN based QoS), as in the case of trunked switch ports. Per-port/per-VLAN based QoS was shown to provide the highest levels of policy granularity, particularly for policing policies.

Next, the four steps for deploying campus QoS were outlined, including:

1. Enable QoS.
2. Apply an ingress QoS model to assign trust or to explicitly classify and mark flows, to (optionally) police flows, and to enable ingress queuing (if required).
3. Apply an egress QoS model to assign flows to transmit queues, enable dropping policies, and egress policing (if supported and required).
4. Enable control plane policing (on platforms that support this feature).

Ingress QoS models were detailed at length, providing flexible template policies that applied to most access edge scenarios. Similarly, best practice egress QoS models were presented, showing that at a medianet campus Gigabit/Ten-Gigabit Ethernet interfaces should support a minimum a 1P3QyT model, including a:

- Realtime queue (to support a RFC 3246 EF PHB service), which should not exceed 33% of the link's bandwidth.
- Guaranteed-bandwidth queue (to support RFC 2597 AF PHB services).
- Default queue (to support a RFC 2474 DF service), which should be at least 25% of the link's bandwidth.
- Bandwidth-constrained queue (to support a RFC 3662 scavenger service), which should not exceed 5% of the link's bandwidth.

A flexible queuing model was also presented that would form as an egress queuing policy template to provide consistent per-node queuing behavior across discrete and disparate Catalyst queuing structures.

Following this, various medianet campus switch port QoS roles were defined, including:

- Switch ports connecting to untrusted endpoints

- Switch ports connecting to trusted endpoints
- Switch ports connecting to conditionally-trusted endpoints
- Switch ports connecting to switch ports (or router interfaces)

Control plane policing was discussed next and general best practice guidelines were presented for deploying CoPP within the medianet campus on both the Catalyst 4500 and 6500 switch platforms.

AutoQoS and SmartPorts were briefly overviewed as to their respective merits and caveats relating to medianet campus QoS deployments.

The second main section then applied these considerations to platform-specific designs for the Cisco Catalyst desktop/stackable switch family (specifically the Catalyst 2960 & 2975, 3560G & 3750G, and 3560-E & 3750-E). Trust models and per-port and per-VLAN marking models were presented for this switch family, as were per-port policing and per-port/per-VLAN policing (via hierarchical QoS). Additionally, both the 1P1Q3T ingress queueing model and the 1P3Q3T egress queueing model for this switch family were detailed.

The third section detailed designs for the Catalyst 4500 switch family (both the Classic Supervisors and the Supervisor 6-E). Trust models and per-port and per-VLAN marking models were presented for this switch family, as were per-port policing, per-port/per-VLAN policing, and UBRL. Additionally, both the 1P3Q1T+DBL and the 1P7Q1T+DBL egress queueing models for this switch family were detailed. Also control plane policing policy recommendations were specified for this switch family.

And the fourth main section presented design recommendations for the Catalyst 6500 switch family (both the Catalyst 6500 and 6500-E switches, for the Supervisor Engine 720, the Supervisor Engine 32, and the Supervisor Engine 32-10GE [with PISA]). Trust models and per-port and per-VLAN marking models—for both ACL-based or NBAR-based classification—were presented for this switch family, as were per-port policing and microflow policing. Additionally, the 1P3Q8T (CoS-based) queueing model, the 1P7Q4T (CoS-based) queueing model, and the 1P7Q4T (DSCP-based) queueing model for this switch family were detailed. Finally, control plane policing policy recommendations were specified for this switch family.

References

IETF RFCs:

- RFC 2474 Definition of the Differentiated Services Field
<http://www.ietf.org/rfc/rfc2474>
- RFC 2597 Assured Forwarding PHB Group
<http://www.ietf.org/rfc/rfc2597>
- RFC 3246 An Expedited Forwarding PHB
<http://www.ietf.org/rfc/rfc3246>
- RFC 3662 A Lower Effort Per-Domain Behavior for Differentiated Services
<http://www.ietf.org/rfc/rfc3662>
- RFC 4594 Configuration Guidelines for DiffServ Service Classes
<http://www.ietf.org/rfc/rfc4594>

Cisco Documentation:

- Cisco Catalyst 2960 QoS Configuration Guide
http://www.cisco.com/en/US/docs/switches/lan/catalyst2960/software/release/12.2_50_se/configuration/guide/swqos.html

- Cisco Catalyst 2975 QoS Configuration Guide
http://www.cisco.com/en/US/docs/switches/lan/catalyst2975/software/release/12.2_46_ex/configuration/guide/swqos.html
- Cisco Catalyst 3560/3750 QoS Configuration Guide
http://www.cisco.com/en/US/docs/switches/lan/catalyst3560/software/release/12.2_50_se/configuration/guide/swqos.html
- Cisco Catalyst 3560-E/3750-E QoS Configuration Guide
http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_50_se/configuration/guide/scg.html
- Cisco Catalyst 4500 Classic Supervisors QoS Configuration Guide
<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/50sg/configuration/guide/qos.html>
- Cisco Catalyst 4500 Supervisor 6-E QoS Configuration Guide
<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/50sg/configuration/guide/qos.html#wp1474085>
- Cisco Catalyst 4500 Control Plane Policing Configuration Guide
http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/50sg/configuration/guide/cntl_pln.html
- Cisco Catalyst 6500 PFC QoS Configuration Guide
<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/qos.html>
- Cisco Catalyst 6500 Supervisor Engine 32 PISA PFC QoS Configuration Guide
<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2ZY/configuration/guide/qos.html>
- Cisco Catalyst 6500 Control Plane Policing Configuration Guide
<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/copp.html>

White Papers:

- Overview of a Medianet Architecture
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/vrn.html>
- Enterprise Medianet Quality of Service Design 4.0-Overview
http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSIntro_40.html
- Infrastructure Protection on Cisco Catalyst 6500 and 4500 Series Switches
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns171/c649/ccmigration_09186a0080825564.pdf

