



CHAPTER 5

Campus QoS Design for TelePresence

Overview

The campus is the primary Place-in-the-Network (PIN) where TelePresence endpoints connect to the network infrastructure. Specifically, the 10/100/1000 NIC on the TelePresence primary codec connects—typically via an Intermediate Distribution Frame (IDF)—to the campus access layer edge switch port. It is at this switch port that the initial QoS polices required to support TelePresence are enabled. Additional QoS policies are also required on all campus inter-switch links, such as uplinks/downlinks between the access and distribution layers, as well as uplinks/downlinks from the distribution-to-core layers, and all core-layer links. Let's consider each of these port-specific QoS requirements.

Access Edge Switch Port QoS Considerations

The first QoS operation that needs to be performed is to define the trust boundary. The trust boundary is the point in the network at which 802.1Q/p CoS markings and/or IP DSCP markings are accepted or overridden by the network.

At the access-layer, the network administrator can enable the infrastructure to:

- Trust the endpoints (CoS and/or DSCP)
- Not trust the endpoints and manually re-mark TelePresence traffic using administratively-defined policies within the access-edge switch
- Conditionally trust the endpoints (trust is extended only after a successful CDP negotiation)

In Phase 1 deployments of Cisco TelePresence (Intra-Enterprise Point-to-Point Deployment Models, as discussed in [Chapter 3, “TelePresence Network Deployment Models”](#)) it is recommended to have a dedicated Communications Manager (CUCM) to support TelePresence. By default, CUCM marks any and all video traffic (including TelePresence) to AF41. It is recommended that this parameter be modified to mark video (i.e., TelePresence only, in this dedicated CUCM context) to CS4.



Note

The reason behind this recommendation is that CUCM does not (yet) have the ability to distinguish between different types of video. Therefore CUCM by default marks both generic Videoconferencing/Video Telephony (from applications like Cisco Unified Video Advantage, for example) as well as TelePresence to AF41.

If a dedicated CUCM is being used for managing TelePresence endpoints, and it has been configured to mark video (i.e., TelePresence) traffic to DSCP CS4, then the TelePresence primary codec marks all TelePresence call traffic (both video and audio) to CS4, but Call-Signaling traffic to CS3. The Cisco

7975G IP Phone similarly marks DSCP values correctly, marking VoIP traffic to EF and Call-Signaling to CS3. Therefore, the switch port connecting to the TelePresence primary codec can be configured to trust DSCP.

Alternatively, the access switch ports can be set to trust CoS, as both the Cisco 7975G IP Phone and the TelePresence primary codec are assigned to the Voice VLAN (VVLAN) and tag their traffic with 802.1Q/p CoS values. The 7975G IP Phone marks VoIP traffic to CoS 5 and Call-Signaling traffic to CoS 3. The Cisco TelePresence codec marks TelePresence traffic (both video and audio) to CoS 4 and Call-Signaling traffic to CoS 3.

However, if the switch port is configured to trust CoS, then it generates an internal DSCP value for all traffic flows via the CoS-to-DSCP map. Only one change is recommended to be made to the default CoS-to-DSCP map, which is to map CoS 5 to EF (46) instead of leaving the default mapping of CoS 5 to CS5 (40). The recommended CoS-to-DSCP map for access-switches connecting to Cisco TelePresence primary codecs is illustrated in [Table 5-1](#).

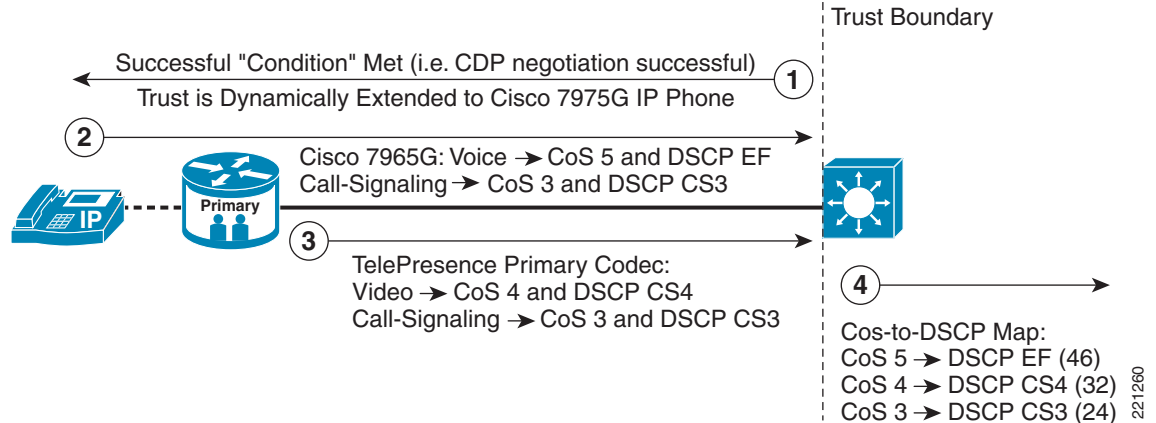
Table 5-1 Recommended Global CoS-to-DSCP Mapping for TelePresence Campus Switches

CoS Value	DSCP Value	PHB	Application
7	56	-	Network Control
6	48	CS6	Internetwork Control
5	46	EF	Voice
4	32	CS4	TelePresence
3	24	CS3	Call-Signaling
2	16	CS2	Management
1	8	CS1	Scavenger
0	0	DF	Default Forwarding/Best Effort

Finally, the access switch may be set to conditionally trust the TelePresence endpoint. This is because Cisco IP Telephony devices, including the Cisco Unified 7979G IP phone that is an intrinsic part of the TelePresence endpoint system, have the ability to identify themselves, via Cisco Discovery Protocol (CDP) to the network infrastructure. Upon a successful CDP negotiation/identification, the network infrastructure dynamically extends trust to the endpoints, which include both the Cisco Unified 7975G IP phone and the TelePresence primary codec. The primary functionality that conditional trust brings is to allow for user-mobility within the IP Telephony-enabled enterprise (users can add/move/change where their IP Phones are connected and the network automatically adapts without requiring an administrator to manually change switch port trust policies). This user-mobility is not a crucial functionality to support TelePresence, since TelePresence units are rarely moved around (due to sheer size). Nonetheless, this conditional trust functionality is supported by TelePresence codecs and adds a minor element of security in the event that the TelePresence codec is physically disconnected from the wall network jack by an unknowing and/or disgruntled individual, who then connects some other device (such as their laptop) to this trusted switch port. In this case, by using a conditional trust policy, the abuser's traffic would no longer be trusted.

The operation of conditional trust policies, as well as endpoint CoS markings and the CoS-to-DSCP mappings of the access-edge switch for TelePresence scenarios, is illustrated in [Figure 5-1](#).

Figure 5-1 Conditional Trust, CoS Markings, and Mappings for TelePresence



Note that if trust CoS is used (as opposed to conditional trust), steps 2, 3, and 4 still apply. The only difference is that the switch would skip step 1; the port would always be trusted regardless of CDP.

An optional recommendation for the access-edge switch port connecting to a TelePresence primary codec is to configure a policer to prevent network abuse in case of a compromise of this trusted port. Similar to the example previously given, this recommendation is to prevent an unknowing and/or disgruntled individual that gains physical access to the TelePresence switch port and decides to send rogue traffic over the network that can hijack voice or video queues and easily ruin call or video quality. Therefore, the administrator may choose to limit the scope of damage that such network abuse may present by configuring access-edge policers on TelePresence switch ports to drop (or remark to Scavenger - CS1) out-of-profile traffic originating on these ports. This is not only a Cisco recommended best practice, but is also reflected in RFC 4594 which recommends edge policing the Real-Time Interactive service class via a single-rate policer.

If such a policer is configured, it is recommended to use Per-Port/Per-VLAN policers, whenever supported. In this manner, a set of policers may be applied to the Voice VLAN to ensure that voice, video, and call signaling traffic are performing within normal levels and a separate, more stringent, policer can be applied to the data VLAN.



Note

When configuring policers for TelePresence, make sure you allow for the appropriate burst intervals, as defined in [Burst Requirements](#) in Chapter 4, “Quality of Service Design for TelePresence.”

Finally, to ensure guaranteed levels of service, queuing needs to be configured on all nodes where the potential for congestion exists, regardless of how infrequently it may occur.

In Catalyst switches, queuing (along with all other QoS operations) is performed in hardware. Therefore, there are a fixed number of hardware queues that vary by platform, as well as by linecards. The nomenclature for Catalyst queuing is 1PxQyT, where:

- 1P represents a strict priority (Expedite Forwarding) queue
- xQ represents a number of non-priority queues
- yT represents a number of drop-thresholds per queue

**Note**

As discussed, due to the strict service levels required by Cisco TelePresence, it is recommended to assign TelePresence flows to a strict priority queue, whether this is implemented in Cisco Catalyst hardware or in Cisco IOS software. However, some older Catalyst platforms and linecards do not support a strict priority queue. For example, some Catalyst 6500 linecards support only a 2Q2T egress queuing model and as such would not be recommended within a Cisco TelePresence campus network design.

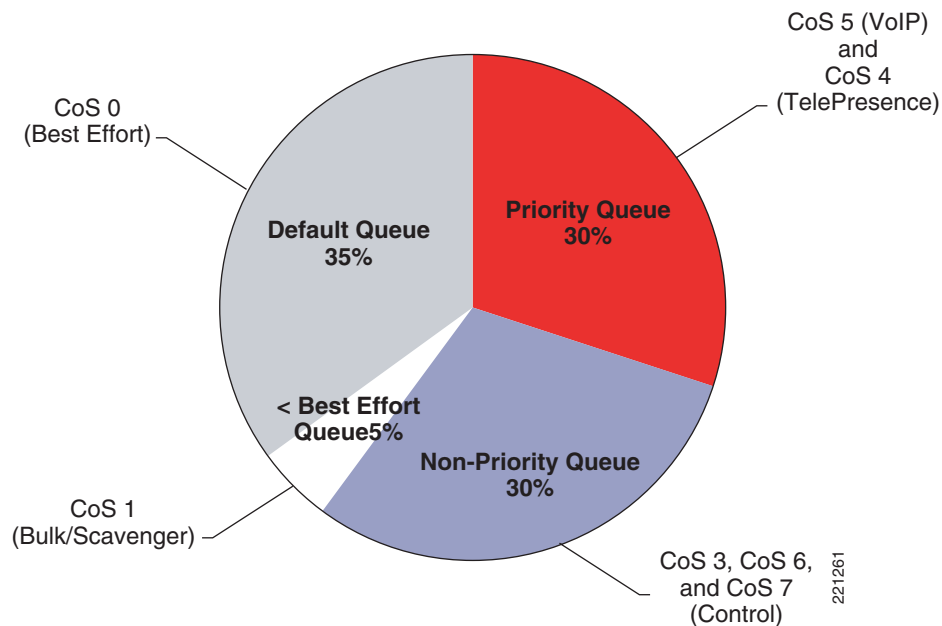
It is highly recommended that all Catalyst switches and linecards within a Cisco TelePresence campus design support a 1PxQyT queuing model.

For example, a Catalyst 6500 48-port 10/100/1000 RJ-45 Module (WS-X6748-GE-TX) has a 1P3Q8T, meaning 1 strict priority queue (which, incidentally, on this linecard is Queue 4) and 3 additional non-priority queues each with 8 configurable Weighted Random Early Detect (WRED) drop thresholds per queue.

Cisco Enterprise Systems Engineering (ESE) testing has shown that the optimal and most consistent service levels for TelePresence are achieved when TelePresence is provisioned with strict-priority hardware queuing (typically in conjunction with VoIP Telephony traffic), provided that the total bandwidth assigned to these realtime applications is less than 33% of the link—but this is virtually always the case on high-speed campus links in the range of 100 Mbps to 10 Gbps Ethernet. For example, consider a Catalyst 6513 provisioned with 11 x 48-port linecards, with each port configured to support G.711 VoIP (128 kbps max per port). Such a configuration would only require 67.584 Mbps or 6.8% of a GigE uplink. Even if a CTS-3000 system were connected to each of the 11 linecards, the total realtime bandwidth would be $[(11 \times 15 \text{ Mbps}) + (11 \times 48 \times 128 \text{ kbps})]$ 232.584 Mbps or 23.3% of a GigE uplink (which is still within the 33% LLQ Rule allowance).

As a generic campus queuing guide, it would be recommended to assign CoS values 4 (TelePresence) and 5 (VoIP) to the strict priority queue, CoS 3 (Call-Signaling) to a (non-default) non-priority queue, and CoS 0 (Best Effort) to the default queue. Finally, CoS 1 (Bulk and/or Scavenger traffic) should be assigned to a (minimally provisioned) less than Best Effort non-priority queue. These guidelines are illustrated in [Figure 5-2](#).

Figure 5-2 Generic Campus Queuing Provisioning and Mapping Guidelines



Campus Inter-Switch Link QoS Considerations

Once the trust boundary has been established and optimal access-edge policers have been enabled, then the DSCP values on all other inter-switch links and campus-to-WAN hand off links can be trusted. Therefore, it is recommended to trust DSCP (not CoS) on all inter-switch links, whether these are uplinks/downlinks to/from the distribution layer, uplinks/downlinks to/from the core layer, intra-core links, or links to WAN Aggregation routers.

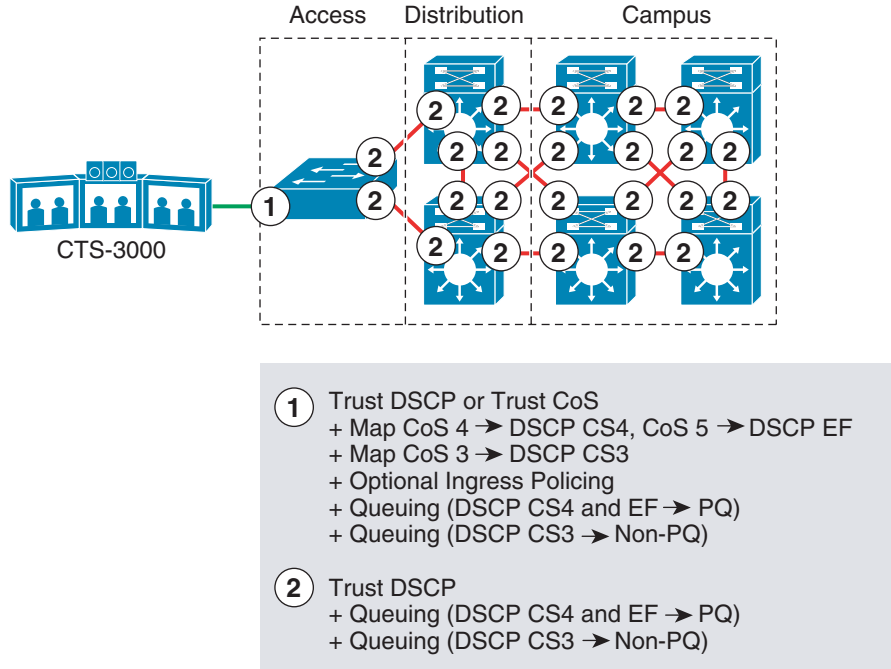
The reason it is recommended to trust DSCP and not CoS is two-fold: first, because marking granularity is lost every time a node is set to trust CoS. For example, if TelePresence endpoints are marking traffic to CS4 and Unified Video Advantage (or other Videoconferencing/Video Telephony endpoints) are marking their traffic to AF41 and the distribution-layer is set to trust CoS from the access-layer, then these flows both appear the same (as CoS 4) to the distribution-layer switch and are indistinguishable from each other from that node forward. Secondly, because trusting CoS implies using 802.1Q trunking between switches. Today, most enterprise campus networks are designed to be Layer 3 and thus 802.1Q is not used on inter-switch links.

Queuing is likewise recommended to be enabled on every node along the path. Note that this document generally focuses only on the QoS requirements for TelePresence. The actual QoS policies may be more complex than those shown here due to the myriad of other data, voice, and video applications on the network. It is recommended that customers use the information provided in this document in concert with

- *Enterprise QoS Solution Reference Network Design Guide*, Version 3.3, November 2005
- Szigeti, Tim and Hattingh, Christina. *End-to-End QoS Network Design: Quality of Service in LANs, WANs, and VPNs*. Indianapolis: Cisco Press, 2004. ISBN-10: 1-58705-176-1; ISBN-13: 978-1-58705-176-0.

A summary of the minimum QoS design requirements within an enterprise campus supporting TelePresence are illustrated in [Figure 5-3](#).

Figure 5-3 Enterprise Campus QoS Design Recommendations for TelePresence



TelePresence Campus QoS Designs

Now that campus-specific considerations have been addressed, we can look at how these policies can be configured on specific platforms. Before we identify the platforms, let's briefly review some of the key network and QoS-related features required by switch platforms at the access, distribution, and core layers of the enterprise campus. These include GE (including 10/100/1000) or 10GE connectivity, adequate port buffering to accommodate TelePresence traffic rates and micro-second bursts, granular policing, and IPxQyT queuing. Optionally, it would be preferred to have dedicated per-port buffers, as well as support for conditional trust, Per-Port/Per-VLAN policing, and DSCP-to-Queue mapping (as opposed to CoS-to-Queue mapping).

Given these requirements, the following currently-shipping Catalyst platforms have been validated by Cisco Enterprise Systems Technical Marketing for TelePresence-enabled campus networks include:

- Catalyst 3560G and 3750G (and by extension the 3650-E and 3750-E)
- Catalyst 4500 and 4948
- Catalyst 6500 (Although only certain linecards are recommended. Some older linecards do not have the requisite buffer to handle TelePresence traffic rate and burst requirements.)

Platform-specific configurations for each of these series of switches are provided in subsequent sections.

Catalyst 3560G/3750G and 3650-E/3750E

The Cisco Catalyst 3560G is a fixed-configuration switch that supports up to 48 10/100/1000 ports with integrated Power over Ethernet (PoE), plus 4 Small Form-Factor Pluggable (SFP) ports for uplinks. The 3560 has a 32 Gbps backplane, which is moderately oversubscribed (52 Gbps theoretical maximum input vs. 32 Gbps backplane yields an oversubscription ratio of 1.625:1 or 13:8). Additionally, the 3560G supports IP routing (including IPv6), multicast routing, and an advanced QoS and security feature-set.

The Catalyst 3750G is nearly identical, with only a few additional key features, including the support for a stackable configuration (via Stackwise technology), allowing for the 32 Gbps backplane (comprised of dual counter-rotating 16 Gbps rings) to be extended over multiple 3750G switches (up to 9). Additionally, the 3750G provides support for 10 Gigabit Ethernet (10GE) connectivity. Obviously, however, the more switches in the stack, as well as the use of 10 GE connectors, increases the oversubscription ratio accordingly.

The 3560-E and 3750-E represent the next evolution of these switches. As before, the 3560-E is a fixed configuration switch, but now with a 128 Gbps backplane and 10 GE port support. Similarly, the 3750-E supports a 128 Gbps backplane with dual 10GE port support, as well as the support for a stackable configuration (via Stackwise Plus technology, allowing a 64 Gbps interconnect between stacked switches).

As the 3560G, 3750G, 3560-E, and 3750-E share virtually identical feature parity (the main differences being the backplane throughput and uplink port speeds), we consider them as a single switch and abbreviate the reference to simply C3560G/3750G.

From a QoS perspective, some of the relevant features of the C3560G/3750G/E include conditional trust, Per-Port/Per-VLAN policers (via Hierarchical QoS policies), DSCP-to-Queue mapping, 2Q3T or 1P1Q3T ingress queuing, and 4Q3T or 1P3Q3T egress queuing. Additionally, these platforms provide (minimally) 750 KB of receive buffers and (up to) 2 MB of transmit buffers for each set of 4 ports. These buffers can be allocated, reserved, or dynamically borrowed from a common pool, on a port-port, per-queue basis, depending on the administrative configurations chosen.

Let's begin leveraging these features into the validated best-practice designs for this switch family for supporting TelePresence at the campus access-layer.

As QoS is disabled by default on these switches, the first step that we must take is to globally enable QoS. We can do this by issuing the global command:

```
mls qos
```

With QoS enabled, we can configure the access-edge trust boundaries. As discussed previously, we have three options: trust DSCP, trust CoS, or conditional trust. It is recommended that ports used for data and VoIP Telephony be configured to conditionally trust CoS, while ports used for TelePresence be configured to either trust DSCP, trust CoS or conditionally trust CoS. Trusting DSCP on these ports is the simplest operationally. The interface command to configure DSCP trust is fairly straightforward:

```
mls qos trust dscp
```



Note

While Cisco IOS allows the configuration of trust CoS and conditional trust on uplink ports, uplink ports should be set to trust DSCP (only). This is required on the C3560G/3750G/E for two reasons: first, to preserve marking granularity between switches (as previously discussed in [Campus Inter-Switch Link QoS Considerations](#)), as well as to activate the DSCP-to-Queue mapping (versus the CoS-to-Queue mapping) on the uplink switch ports.

If you choose to trust CoS or conditionally trust CoS, ensure that the fifth parameter in the global CoS-to-DSCP map—which corresponds to the DSCP mapping for CoS 4—is set to 32 (CS4). Additionally, to support IP Telephony properly, the global CoS-to-DSCP mapping table should be modified such that CoS 5 (the sixth parameter in the CoS-to-DSCP map) is mapped to 46 (EF)—which is not the default (the default setting is 40/CS5). These settings are achieved via the following global and interface commands:

```
mls qos map cos-dscp 0 8 16 24 32 46 48 56
interface Gigx/y
  mls qos trust cos
```

If you choose to implement conditional trust on the TelePresence ports, it can be enabled with the following interface command:

```
mls qos trust device cisco-phone
```


Note

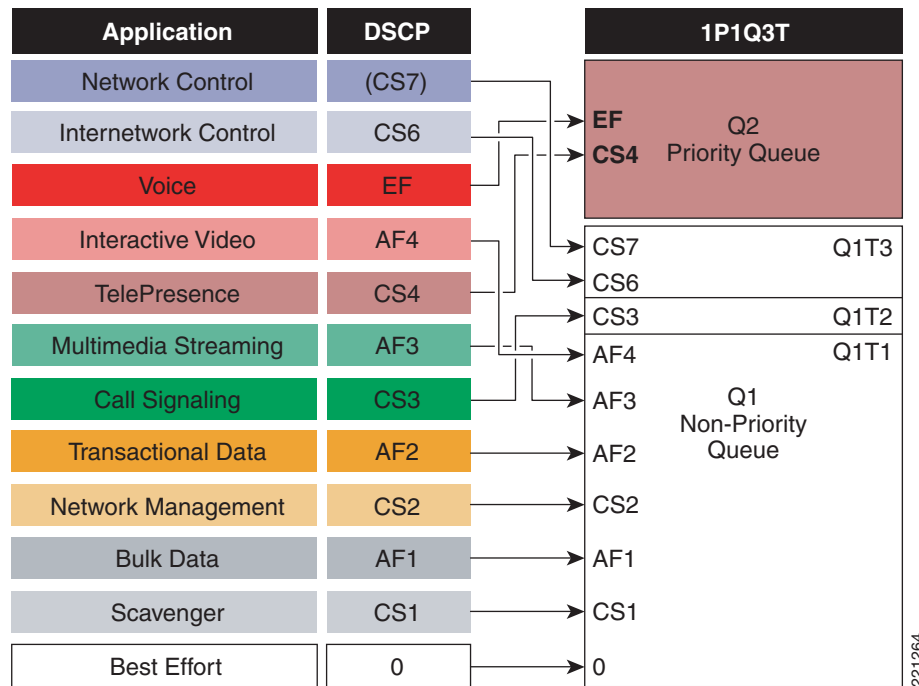
If conditional trust policies are to be used, then make sure that the TelePresence codec software is running version 1.1.0 (256D) or higher, as software version 1.0.1 (616D) incorrectly marks TelePresence audio traffic to CoS 5 (not CoS 4).

These configuration commands can be verified with the following commands:

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

Next, as the C3560G/3750G/E 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. As we've already established the requirement for strict-priority servicing of TelePresence (and VoIP) traffic, it is recommended to enable the 1P1Q3T ingress queuing structure with DSCP EF (VoIP) and CS4 (TelePresence) being mapped to the ingress PQ (Q2). The configurable thresholds in the non-priority queue can be used to protect control traffic. For example, Network Control traffic (such as Spanning Tree Protocol) associated with DSCP CS7 and Internetwork Control traffic (such as Interior Gateway Protocols, including EIGRP and OSPF) marked DSCP CS6 can be explicitly protected by assigned these to Q1T3. Additionally, a degree of protection can be offered to Call-Signaling traffic (which is essentially control traffic for the IP Telephony infrastructure), which is marked CS3. All other traffic types can be provisioned in Q1T1. The recommended ingress 1P1Q3T queuing configuration for the C3560G/3750G/E platforms is illustrated in [Figure 5-4](#).

Figure 5-4 Catalyst 3560G/3750G/E(1P1Q3T) Ingress Queuing Recommendations for TelePresence Deployments



Based on [Figure 5-4](#), the recommended configuration for ingress queuing on the C3560G/3750G/E for TelePresence deployments is as follows:

! This first section modifies the CoS-to-DSCP for VoIP

```
mls qos map cos-dscp 0 8 16 24 32 46 48 56
! Modifies CoS-to-DSCP mapping to map CoS 5 to DSCP EF
```

! This section configures the Ingress Queues and Thresholds for 1P1Q3T

```
mls qos srr-queue input buffers 70 30
! Configures the Ingress Queue buffers such that Q2 (PQ) gets 30% of buffers
mls qos srr-queue input priority-queue 2 bandwidth 30
! Configures the Ingress PQ (Q2) to be guaranteed 30% BW on stack ring
mls qos srr-queue input bandwidth 70 30
! Configures SRR weights between Ingress Q1 and Q2 for remaining bandwidth
mls qos srr-queue input threshold 1 80 90
! Configures Ingress Queue 1 Threshold 1 to 80% and Threshold 2 to 90%
! Ingress Queue 1 Threshold 3 remains at 100% (default)
! Ingress Queue 2 Thresholds 1, 2 and 3 remain at 100% (default)
```

! This section configures the Ingress CoS-to-Queue Mappings for TelePresence ports using trust-CoS

```
mls qos srr-queue input cos-map queue 1 threshold 1 0 1 2
! Maps CoS 0, 1, and 2 to Ingress Queue 1 (Q1T1)
mls qos srr-queue input cos-map queue 1 threshold 2 3
! Maps CoS 3 to Ingress Queue 1 Threshold 2 (Q1T2)
mls qos srr-queue input cos-map queue 1 threshold 3 6 7
! Maps CoS 6 and 7 to Ingress Queue 1 Threshold 3 (Q1T3)
mls qos srr-queue input cos-map queue 2 threshold 1 4 5
! Maps CoS 4 (TelePresence) and CoS 5 (VoIP) to Ingress-PQ Threshold 1 (Q2T1)
```

```
! This section configures the Ingress DSCP-to-Queue Mappings for TelePresence ports using
trust-DSCP
```

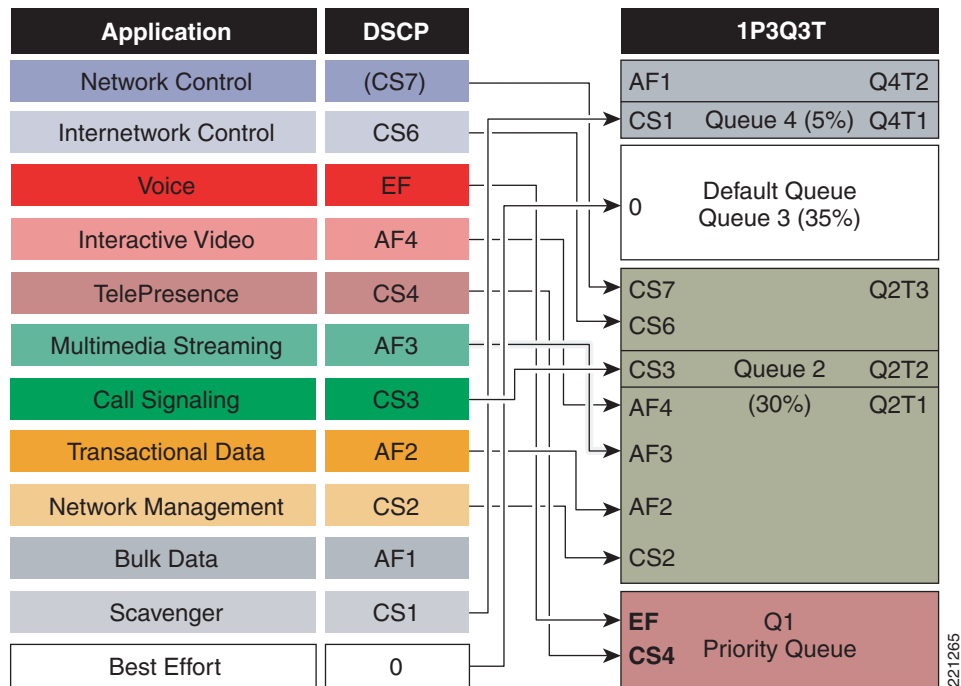
```
mls qos srr-queue input dscp-map queue 1 threshold 1 0 8 10 12 14
! Maps DSCP 0, CS1 and AF1 to Ingress Queue 1 Threshold 1 (Q1T1)
mls qos srr-queue input dscp-map queue 1 threshold 1 16 18 20 22
! Maps DSCP CS2 and AF2 to Ingress Queue 1 Threshold 1 (Q1T1)
mls qos srr-queue input dscp-map queue 1 threshold 1 26 28 30 34 36 38
! Maps DSCP AF3 and AF4 to Ingress Queue 1 Threshold 1 (Q1T1)
mls qos srr-queue input dscp-map queue 1 threshold 2 24
! Maps DSCP CS3 to Ingress Queue 1 Threshold 2 (Q1T2)
mls qos srr-queue input dscp-map queue 1 threshold 3 48 56
! Maps DSCP CS6 and CS7 to Ingress Queue 1 Threshold 3 (Q1T3)
mls qos srr-queue input dscp-map queue 2 threshold 1 32 46
! Maps DSCP CS4 (TelePresence)& EF (VoIP) to Ingress-PQ Threshold 1 (Q2T1)
```

**Note**

Non-Standard DSCP values can also be mapped to their respective queues (using the CoS-to-Queue Map as a reference); however, for the sake of simplicity, non-standard DSCP-to-Queue Mappings have not been shown in these configurations.

Following ingress queuing configuration, we can now proceed to configuring the egress queues. The C3560G/3750G/E supports either 4Q3T or 1P3Q3T egress queuing configurations. As the need for an EF PHB has already been established, both for VoIP and for TelePresence, it is recommended to enable the 1P3Q3T egress queuing configuration, with Q1 as the PQ. Then both VoIP (DSCP EF) and TelePresence (DSCP CS4) should be mapped to Q1 (the PQ). Default traffic can be assigned to Q3 and Q4 can be designated as a less than Best Effort queue, servicing Bulk (AF1) and Scavenger (DSCP CS1) traffic, being assigned to Q4T2 and Q4T1, respectively. Network Control (DSCP CS7) and Internetwork Control (DSCP CS6) can be mapped to the highest threshold of the preferential non-priority queue (Q2T3), while Call-Signaling (DSCP CS3) can be mapped to the second highest threshold in that queue (Q2T2). All other applications can be mapped to Q2T1. The recommended 1P3Q3T egress queuing configuration for the C3560G/3750G/E platforms is illustrated in [Figure 5-5](#).

Figure 5-5 Catalyst C3560G/3750G/E (1P3Q3T) Egress Queuing Recommendations for TelePresence Deployments



Based on [Figure 5-5](#), the recommended configuration for egress queuing on the C3560G/3750G/E for TelePresence deployments is as follows:

! This section configures the Output CoS-to-Queue Maps for TelePresence ports using trust-CoS

```

mls qos srr-queue output cos-map queue 1 threshold 3 4 5
! Maps CoS 4 (TelePresence) and CoS 5 (VoIP) to Egress Queue 1 Threshold 3 (PQ)
mls qos srr-queue output cos-map queue 2 threshold 1 2
! Maps CoS 2 to Egress Queue 2 Threshold 1 (Q2T1)
mls qos srr-queue output cos-map queue 2 threshold 2 3
! Maps CoS 3 (Call-Signaling) to Egress Queue 2 Threshold 2 (Q3T2)
mls qos srr-queue output cos-map queue 2 threshold 3 6 7
! Maps CoS 6 and CoS 7 (Net Control) to Egress Queue 2 Threshold 3 (Q2T3)
mls qos srr-queue output cos-map queue 3 threshold 3 0
! Maps CoS 0 (Best Effort) to Egress Queue 3 Threshold 3 (Q3T3)
mls qos srr-queue output cos-map queue 4 threshold 3 1
! Maps CoS 1 (Bulk/Scavenger) to Egress Queue 4 Threshold 3 (Q4T3)

```

! This section configures the Output DSCP-to-Queue Maps for TelePresence ports using trust-DSCP

```

mls qos srr-queue output dscp-map queue 1 threshold 3 32 46
! Maps DSCP CS4 (TelePresence) and EF (VoIP) to Egress Queue 1 (PQ)
mls qos srr-queue output dscp-map queue 2 threshold 1 16 18 20 22
! Maps DSCP CS2 and AF2 to Egress Queue 2 Threshold 1 (Q2T1)
mls qos srr-queue output dscp-map queue 2 threshold 1 26 28 30 34 36 38
! Maps DSCP AF3 and AF4 to Egress Queue 2 Threshold 1 (Q2T1)
mls qos srr-queue output dscp-map queue 2 threshold 2 24
! Maps DSCP CS3 to Egress Queue 2 Threshold 2 (Q2T2)
mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
! Maps DSCP CS6 and CS7 to Egress Queue 2 Threshold 3 (Q2T3)
mls qos srr-queue output dscp-map queue 3 threshold 3 0

```

```

! Maps DSCP DF to Egress Queue 3 Threshold 3 (Q3T3 - Default Queue)
mls qos srr-queue output dscp-map queue 4 threshold 1 8
! Maps DSCP CS1 to Egress Queue 4 Threshold 1 (Q4T1)
mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
! Maps DSCP AF1 to Egress Queue 4 Threshold 2 (Q4T2)

! This next section configures the WRED min and max thresholds for Q1

mls qos queue-set output 1 threshold 2 80 90 100 100
! Sets Egress Queue 2 Threshold 1 (Q2T1) to 80% and Threshold2 (Q2T2) to 90%
mls qos queue-set output 1 threshold 4 60 100 100 100
! Sets Egress Queue 4 Threshold 1 (Q4T1) to 60% and Threshold 2 (Q4T2) to 100%

! This section configures trust-DSCP and queuing on TelePresence access port and uplink
ports

interface GigabitEthernet1/0/1
description TelePresence or Uplink port
mls qos trust dscp
! Assigns the TelePresence port and/or uplink port to trust DSCP
queue-set 1
! Assigns interface to Queue-Set 1 (default)
srr-queue bandwidth share 1 30 35 5
! Q2 gets 30% of remaining BW (after PQ); Q3 gets 35% & Q4 gets 5%
priority-queue out
! Expedite queue is enabled for TelePresence and VoIP
!

! This section configures conditional-trust and queuing on TelePresence access ports

interface GigabitEthernet1/0/2
description IP Telephony and/or Data port
mls qos trust device cisco-phone
! Configures conditional trust based on the CDP advertisements of the TelePresence
system and attached 7975G IP phone
queue-set 1
! Assigns interface to Queue-Set 1 (default)
srr-queue bandwidth share 1 30 35 5
! Q2 gets 30% of remaining BW (after PQ); Q3 gets 35% & Q4 gets 5%
priority-queue out
! Expedite queue is enabled for TelePresence and VoIP
!

```

**Note**

As before, non-Standard DSCP values can also be mapped to their respective queues (using the CoS-to-Queue Map as a reference); however, for the sake of simplicity, non-standard DSCP-to-Queue Mappings have not been shown in these configurations.

These configuration commands can be verified with the following commands:

- **show mls qos queue-set**
- **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**
- **show mls qos interface buffers**
- **show mls qos interface queueing**
- **show controllers ethernet-controller port-asic statistics**

Catalyst 4500 and 4948

The Cisco Catalyst 4500 series switches are midrange modular platforms with chassis options to 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). 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 linecards that meet the requirements (at the time of writing) outlined in [TelePresence Campus QoS Designs](#) for the Catalyst 4500 include the 4448 and the 4548 series linecards (specifically, the WS-X4448-GB-RJ45 and the WS-X4524-GB-RJ45V or WS-X4548-GB-RJ45V).

On the other hand, the degree of oversubscription and buffering capabilities on the C4500 series linecards varies by linecard. Some linecards are entirely non-blocking, while others, such as the 4448 and the 4548, provision a single 1 Gbps uplink to the switch fabric for every 4 or 8 (10/100/1000) ports, which equates to an 4:1 (for the 4524) or an 8:1 (4448 and 4548) theoretical oversubscription ratio. As such, the 4448 and 4548 series linecards, while suitable at the campus access-edge, would not be recommended to be used as uplinks nor within the distribution and core layers of a TelePresence-enabled campus.

The Catalyst 4948 series provides an advanced feature set of intelligent network services, but is engineered and optimized to support high-performance wire-speed switching for data center server traffic. Thus the Catalyst 4948 has a completely non-blocking architecture and as such would be suitable at any layer (access, distribution, or core) within a TelePresence-enabled campus network. Specifically, the Catalyst 4948 provides 96 Gbps of switching fabric for its fixed configuration 48 x 10/100/1000 ports plus 4 SFP ports (which may be GE or 10 GE). Additionally, the Catalyst 4948 provides approximately 16 MB of buffering which is shared among all 48 ports.

As the Catalyst 4500 and 4948 share virtually identical feature parity (the main differences being the backplane throughput and buffer architectures), we consider them as a single switch and abbreviate the reference to simply C4500/4948.

From a QoS perspective, some of the relevant features of the C4500/4948 include conditional trust, an elegant Per-Port/Per-VLAN policer implementation, DSCP-to-Queue mapping, 4Q1T or 1P3Q1T queuing support, and an advanced congestion algorithm (Dynamic Buffer Limiting or DBL).

Let's begin leveraging these features into the validated best-practice designs for this switch family for supporting TelePresence at the campus access-layer.

The first thing to note is a minor syntactical difference when configuring QoS features on the C4500/4948; specifically, QoS commands on this platform do not include the mls prefix used on the C3560G/3750G/E and the C6500 series platforms. For example, to globally enable QoS on the C4500/4948 (which is disabled by default), the command is not mls qos, but simply:

```
qos
```

With QoS enabled, we can configure the access-edge trust boundaries. As discussed previously, we have three options: trust DSCP, trust CoS, or conditional trust. It is recommended that ports used for data and VoIP Telephony be configured to conditionally trust CoS, while ports used for TelePresence be configured to either trust DSCP, trust CoS or conditionally trust CoS. Trusting DSCP on these ports is the simplest operationally.

```
qos trust dscp
```

If you choose to trust CoS or conditionally trust CoS, then CoS 5 must be explicitly mapped to DSCP EF prior to the port being configured to trust CoS. All other CoS-to-DSCP mappings can be left at their respective default values. These functions can be achieved via the following global and interface commands:

```
qos map cos 5 to 46
!
interface Gigx/y
  qos trust cos
```

If you choose to implement conditional trust on the TelePresence ports, it can be enabled with the following interface command:

```
qos trust device cisco-phone
```


Note

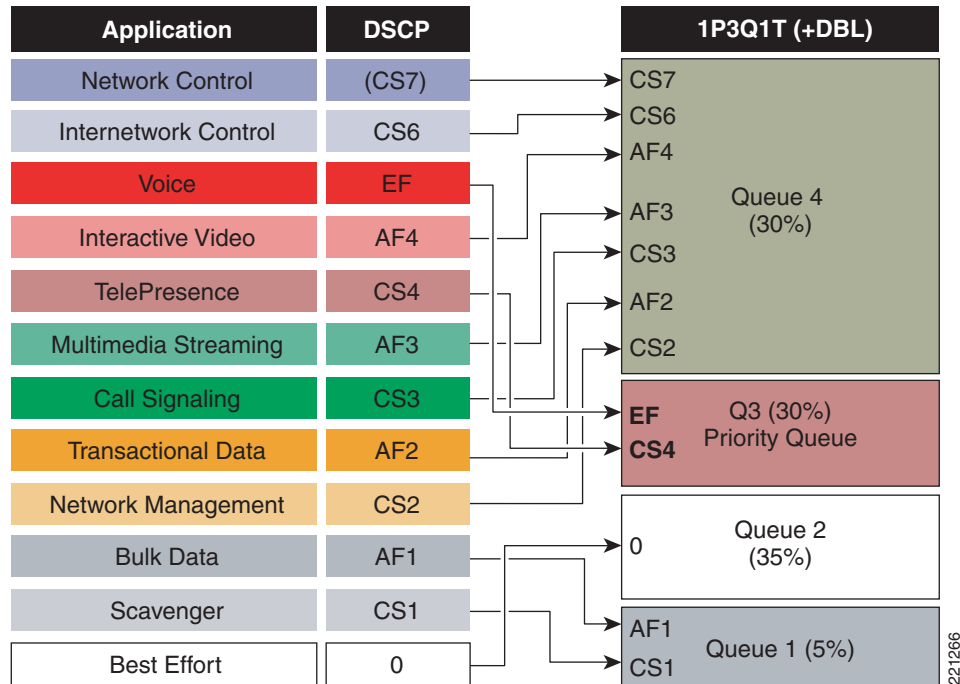
If conditional trust policies are to be used, then make sure that the TelePresence codec software is running version 1.1.0 (256D) or higher, as software version 1.0.1 (616D) incorrectly marks TelePresence audio traffic to CoS 5 (not CoS 4).

As with configuration commands, the C4500/4948 omits the mls prefix in the corresponding verification commands. These configuration commands can be verified with the following commands:

- **show qos**
- **show qos maps**
- **show qos interface**

As the C4500/4948 does not support ingress queuing (although it bears mentioning that the internal servicing architectures have been tested and found to be adequate in protecting TelePresence traffic even in the event of oversubscription), we can move on to configuring egress queuing. The C4500/4948 can be configured to operate in a 4Q1T mode or a 1P3Q1T mode, the latter of which is recommended for VoIP Telephony and TelePresence deployments. On the C4500, however, the strict priority queue, when enabled, is Q3. As the C4500/4948 supports DSCP-to-Queue mappings, we can distinguish between applications such as generic Videoconferencing/Video Telephony (AF4) and TelePresence (CS4), even though these share the same CoS and IP Precedence values (Cos/IPP 4). Given these abilities, it is recommended to enable 1P3Q1T queuing on the C4500/4948, with VoIP (EF) and TelePresence (CS4) assigned to the strict-priority queue (Q3). Q2 may be dedicated to service default traffic and Q1 can be used to service less than Best Effort Scavenger (CS1) and Bulk (AF1) traffic. All other applications can be mapped to Q4, the preferential queue. The recommended (1P3Q1T + DBL) egress queuing configuration for the C4500/4948 platform is illustrated in [Figure 5-6](#).

Figure 5-6 Catalyst C4500/4948 (1P3Q1T + DBL) Egress Queuing Recommendations for TelePresence Deployments



Note

As before, non-Standard DSCP values can also be mapped to their respective queues; however, for the sake of simplicity, non-standard DSCP-to-Queue Mappings have not been shown in these configurations.

As previously mentioned, the C4500/4948 supports an advanced congestion avoidance algorithm—Dynamic Buffer Limiting (DBL)—rather than Weighted Tail Drop (WTD) or Weighted-Random Early-Detect (WRED). Therefore no DSCP-to-Threshold mappings are required on the C4500/4948. However, to leverage DBL, it must be globally enabled (as it is disabled by default). This is achieved with the following global command:

```
qos db1
```

Optionally, DBL can be configured to operate to support RFC 3168 IP Explicit Congestion Notification (IP ECN or simply ECN), which utilizes the remaining 2 bits of the IPv4/IPv6 Type of Service (ToS) Byte (the DSCP value uses the first 6 bits of the ToS Byte). The following global command enables ECN for DBL:

```
qos db1 exceed-action ecn
```



Note

For more information on IP ECN, refer to RFC 3168 (at www.ietf.org/rfc/rfc3168) and Szigeti, Tim and Hattingh, Christina. *End-to-End QoS Network Design: Quality of Service in LANs, WANs, and VPNs*. Indianapolis: Cisco Press, 2004. ISBN-10: 1-58705-176-1; ISBN-13: 978-1-58705-176-0.

Additionally, to leverage DBL (with/without ECN) on a per-interface basis, a service policy applying DBL to all flows must be constructed and applied to each interface. This can be done by using the following basic policy-map:

```
policy-map DBL
```

```

class class-default
  db1
!
interface Gig x/y
  service policy output DBL

```

However, at this point, an important consideration pertaining to DBL much be taken into account, namely DBL (when enabled and configured as per the above recommendations) is active on all flows, including flows destined to the PQ (Q3)—which in our case includes VoIP and TelePresence traffic. As DBL introduces dynamic drops, especially on bursty, large-packet flows, this is detrimental to TelePresence call-quality. Therefore, to explicitly disable DBL on PQ traffic, the following amendments can be made to the previous policy:

```

class-map match-any PQ
  match ip dscp ef
  match ip dscp cs4
policy-map DBL
  class PQ
  class class-default
    db1
!
interface Gig x/y
  service policy output DBL

```

In this modified policy, the class-map PQ identifies traffic destined to the Priority Queue, specifically EF (VoIP) or CS4 (TelePresence) traffic. In the policy-map, the PQ-class receives no action (DBL or otherwise) and serves only to exclude these flows from the following class-default policy of applying DBL to all (other) flows. It is highly recommended to use this modified policy on C4500/4948 platforms supporting TelePresence in conjunction with DBL; otherwise DBL drops negatively impact TelePresence call-quality.

Piecing this together, the C4500/4948 egress queuing recommendation, shown in [Figure 5-6](#), is as follows:

```

!This section enables DBL globally and excludes DBL on PQ flows

qos db1
  ! Globally enables DBL
qos db1 exceed-action ecn
  ! Optional: Enables DBL to mark RFC 3168 ECN bits in the IP ToS Byte
class-map match-any PQ
  match ip dscp ef
  match ip dscp cs4
  ! Classifies traffic mapped to PQ for exclusion of DBL-policy
policy-map DBL
  class PQ
  ! No action (DBL or otherwise) is applied on traffic mapped to PQ
  class class-default
    db1
  ! Enables DBL on all (other) traffic flows

! This section configures the DSCP-to-Transmit Queue Mappings

qos map dscp 0 to tx-queue 2
  ! Maps DSCP 0 (Best Effort) to Q2
qos map dscp 8 10 12 14 to tx-queue 1
  ! Maps DSCP CS1 (Scavenger) and AF11/AF12/AF13 (Bulk) to Q1
qos map dscp 16 18 20 22 to tx-queue 4
  ! Maps DSCP CS2 (Net-Mgmt) and AF21/AF22/AF23 (Transactional) to Q4
qos map dscp 24 26 28 30 to tx-queue 4
  ! Maps DSCP CS3 (Call-Sig) and AF31/AF32/AF33 (MultiMedia) to Q4

```

```

qos map dscp 34 36 38 to tx-queue 4
! Maps DSCP AF41/AF42/AF43 (Interactive-Video) to Q4
qos map dscp 32 46 to tx-queue 3
! Maps DSCP CS4 (TelePresence) and EF (VoIP) to Q3 (PQ)
qos map dscp 48 56 to tx-queue 4
! Maps DSCP CS6 (Internetwork) and CS7 (Network Control) to Q4

! This section configures queues, activates the PQ and applies DBL

interface range GigabitEthernet1/1 - 48
  tx-queue 1
  bandwidth percent 5
  ! Q1 gets 5% BW
  tx-queue 2
  bandwidth percent 35
  ! Q2 gets 35% BW
  tx-queue 3
  priority high
  ! Q3 is PQ
  bandwidth percent 30
  ! Q3 (PQ) gets 30% BW
  shape percent 30
  ! Shapes/limits PQ to 30% BW
  tx-queue 4
  bandwidth percent 30
  ! Q4 gets 40%
  service-policy output DBL
  ! Applies DBL to all flows except VoIP & TelePresence
!
```

**Note**

As before, non-Standard DSCP values can also be mapped to their respective queues; however, for the sake of simplicity, non-standard DSCP-to-Queue Mappings have not been shown in these configurations.

These configuration commands can be verified with the following commands:

- **show qos dbl**
- **show qos maps dscp tx-queue**
- **show qos interface**

Catalyst 6500

The Cisco Catalyst 6500 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 edge. The Catalyst 6500 platform is available in 3, 4, 6, 9, or 13 slot combinations; these models include the 6503, 6504, 6506, 6509 (regular or Network Equipment Building System [NEBS] compliant), and 6513. Additionally, these chassis options are also available in Enhanced models, designated by a -E suffix (such as 6503-E, 6504-E, etc.) for additional feature functionality and performance (except the 6513 at the time of writing).

Overall, the Catalyst 6500 provides the highest performance switching plane, supporting a 720 Gbps switching fabric and the option to run either centralized or distributed forwarding to achieve optimal performance. Additionally, the Catalyst 6500 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.

The queuing and buffering capabilities of Catalyst 6500 supervisors and linecards vary according to type. A summary of the ingress and egress queuing structures, as well as ingress, egress, and total buffering capabilities for Catalyst 6500 supervisors and linecards that have been tested and validated for TelePresence campus network designs are presented in [Figure 5-7](#) through [Figure 5-9](#).

[Figure 5-7](#) presents queuing and buffering details by Catalyst 6500 supervisor types.

Figure 5-7 Catalyst 6500 Queuing and Buffering Details by Supervisor Types

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	WRR	1p2q2t	WRR	512 KB	73 KB	439 KB
WS-SUP720-3B							
WS-SUP720-3BXL							
WS-SUP32-10GE	2q8t	WRR	1p3q8t	SRR	1.3 MB	166 KB	1.2 MB
WS-SUP32-GE							

224698

[Figure 5-8](#) presents queuing and buffering details for Catalyst 6500 GigabitEthernet linecards (including 10/100/1000 linecards).

Figure 5-8 Catalyst 6500 Queuing and Buffering Details by GigabitEthernet (and 10/100/1000) Linecards

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-X6148A-GE-TX	1q2t	WRR	1p3q8t	WRR	5.5 MB	120 KB	5.4 MB
WS-X6148A-GE-45AF							
WS-X6516A-GBIC	1p1q4t	WRR	1p2q2t	WRR	1 MB	135 KB	946 /kB
WS-X6548-GE-TX	1q2t	WRR	1p2q2t	WRR	1.4 MB	185 KB	1.2 MB
WS-X6548V-GE-TX							
WS-X6548-GE-45AF							
WS-X6724-SFP with CFC	1q8t	WRR	1p3q8t	DWRR	1.3 MB	166 KB	1.2 MB
WS-X6724-SFP with DFC3	2q8t	WRR					
WS-X6748-GE-TX with CFC	1q8t	WRR					
WS-X6748-GE-TX with DFC3	2q8t	WRR					
WS-X6748-SFP with CFC	1q8t	WRR					
WS-X6748-SFP with DFC3	2q8t	WRR					

224699

It is important to note that the 6148A-GE and the 6548-GE are both engineered with 8:1 oversubscription ratios and, as such, while suitable at the access layer, these linecards would not be recommended to deploy as uplinks or within the distribution and core layers of the TelePresence-enabled campus network. Furthermore, it would be recommended to configure ingress queuing on these linecards, according to the recommendations presented in [Ingress Queuing Design—1Q2T](#).

In contrast, the 6748-GE is virtually non-blocking, supporting a dual 20 Gbps connection to the switch fabric for its 48 (10/100/1000) ports, which equates to a minimal 6:5 oversubscription ratio.

In turn, [Figure 5-9](#) presents queuing and buffering details for Catalyst 6500 TenGigabitEthernet (10GE) linecards.

Figure 5-9 Catalyst 6500 Queuing and Buffering Details by TenGigabitEthernet Linecards

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-X6704-10GE with CFC	1q8t	WRR	1p7q8t	DWRR	16 MB	2 MB	14 MB
WS-X6704-10GE with DFC3	8q8t	WRR					
WS-X6708-10G-3C	8q4t	DWRR	1p7q4t	DWRR SRR	198 MB	108 MB	90 MB
WS-X6708-10G-3CXL							

224700

It is important to note that, at the time of writing, all Catalyst 6500 linecards support only CoS-to-Queue configurations for both ingress and egress queuing (with the exception of the 6708 [WS-X6708-10GE] linecard family, which can be configured with either CoS-to-Queue mappings or with DSCP-to-Queue mappings). Considerations relating to CoS-to-Queue mapping limitations are covered shortly.

Finally, [Figure 5-10](#) summarizes egress queuing structures by supervisor and linecard families, as a quick reference to selecting the relevant recommended queuing configurations.

Figure 5-10 Catalyst 6500 Supervisors and Linecards By Egress Queuing Structures

1P2Q2T CoS-to-Queue	1P3Q8T CoS-to-Queue	1P7Q8T CoS-to-Queue	1P7Q4T* *DSCP-to-Queue
<ul style="list-style-type: none"> • WS-SUP720 • WS-SUP720-3B • WS-SUP720-3BXL • WS-6516A-GBIC • WS-X6548-GE-TX • WS-X6548V-GE-TX • WS-X6548-GE-45AF 	<ul style="list-style-type: none"> • WS-SUP32-GE • WS-SUP32-10GE • WS-X6148A-GE-45F • WS-X6148A-GE-TX • WS-X6724-SFP • WS-X6748-GE-TX • WS-X6748-SFP 	<ul style="list-style-type: none"> • WS-X6704-10GE 	<ul style="list-style-type: none"> • WS-X6708-10G-3C • WS-X6708-10G-3CXL

224701

From a QoS perspective, some of the relevant features of the C6500 include port-trust, linecard-dependant queuing options, and WRED support. Let's examine how these features can be leveraged into validated best-practice designs for the Catalyst 6500 (which we will abbreviate to C6500) at the access-edge.

As with the previously discussed switch platforms, QoS is disabled by default and must be explicitly enabled globally on the C6500 for any configured policies to take effect. The command to globally enable QoS on the C6500 is:

```
mls qos
```

With QoS enabled, we can configure the access-edge trust boundaries. At the time of writing, on the C6500, we have only two port-trust options: trust DSCP and trust CoS.

**Note**

While trusting IP Precedence is a configurable option, this functionality is superseded by trusting DSCP. Additionally, at the time of writing, conditional trust is not available on the C6500.

When considering which trust option to configure, there is an important relationship between trust and ingress queuing on the C6500 to consider, namely, if a port is set to trust CoS, then ingress queuing is automatically enabled. This becomes an especially relevant consideration on linecards with high oversubscription ratios, such as the 6148A and 6548 (both with 8:1 oversubscription ratios). Therefore, it is recommended to set the ports connecting to TelePresence systems to trust CoS. However, keep in mind that if CoS is to be trusted, then ensure that the fifth parameter in the global CoS-to-DSCP map—which corresponds to the DSCP mapping for CoS 4—is set to 32 (CS4). Additionally, to support IP Telephony properly, the global CoS-to-DSCP mapping table should be modified such that CoS 5 (the sixth parameter in the CoS-to-DSCP map) is mapped to 46 (EF), which is not the default (the default setting is 40/CS5). These settings are achieved via the following global and interface commands:

```
mls qos map cos-dscp 0 8 16 24 32 46 48 56
interface Gigx/y
  mls qos trust cos
```

However, on all inter-switch link ports (uplinks/downlinks, etc.) it is recommended to set port-trust to trust DSCP to preserve marking granularity and Diffserv Per-Hop Behaviors. For this same reason, it is highly recommended to set port-trust to trust DSCP (or better yet, to use service policies with policers) on ports connected to endpoints that may be generating generic Videoconferencing/Video Telephony traffic (marked AF41); otherwise the DSCP value (AF41) for these flows will be lost and will be remapped to CS4 (since the CoS-to-DSCP map maps CoS 4 to DSCP CS4), eliminating your ability to distinguish between them on subsequent switch/router hops along the network path. The interface command to configure DSCP trust on an interface is as follows:

```
mls qos trust dscp
```

These configuration commands can be verified with the following commands:

- **show mls qos**
- **show mls qos maps | begin Cos-dscp map**
- **show queuing interface gigabitethernet x/y | include trust**

Before we describe linecard-specific queue recommendations, it bears mentioning that the queuing structures on the C6500—both ingress and egress—are CoS-based (with the sole exception, at the time of writing, of the WS-X6708-10GE, which supports DSCP-based queuing). This presents a challenge to network administrators deploying both generic Videoconferencing/Video Telephony (marked AF41 per RFC 4594) and TelePresence (marked CS4 per RFC 4594), as these applications both share the same CoS value of 4. As such, TelePresence and generic Videoconferencing/Video Telephony traffic are indistinguishable from one another with a CoS-based queuing scheme, with both applications always being mapped to the same queue. Since TelePresence requires an Expedited Forwarding Per-Hop Behavior (as explained in [Chapter 4, “Quality of Service Design for TelePresence”](#) and as allowed for by RFC 4594), both TelePresence and Videoconferencing/Video Telephony can be assigned to the strict-priority hardware queue on C6500 linecards (along with VoIP). Therefore, while this technical limitation exists, network administrators are encouraged to configure the hardware strict-priority queues on their C6500 platforms to adequately provision for their VoIP, TelePresence, and Videoconferencing/Video Telephony traffic.

While this may sound a bit complicated or excessive, in practice it is not that difficult to do, especially when considering that VoIP is such a lightweight application. For example, consider a Catalyst 6513 with redundant supervisors and 11 x 48-port linecards. If each of these ports supported VoIP, a total of only 68 Mbps of PQ would be required on the uplink (6.8% of a GE link). Additionally, if a generic

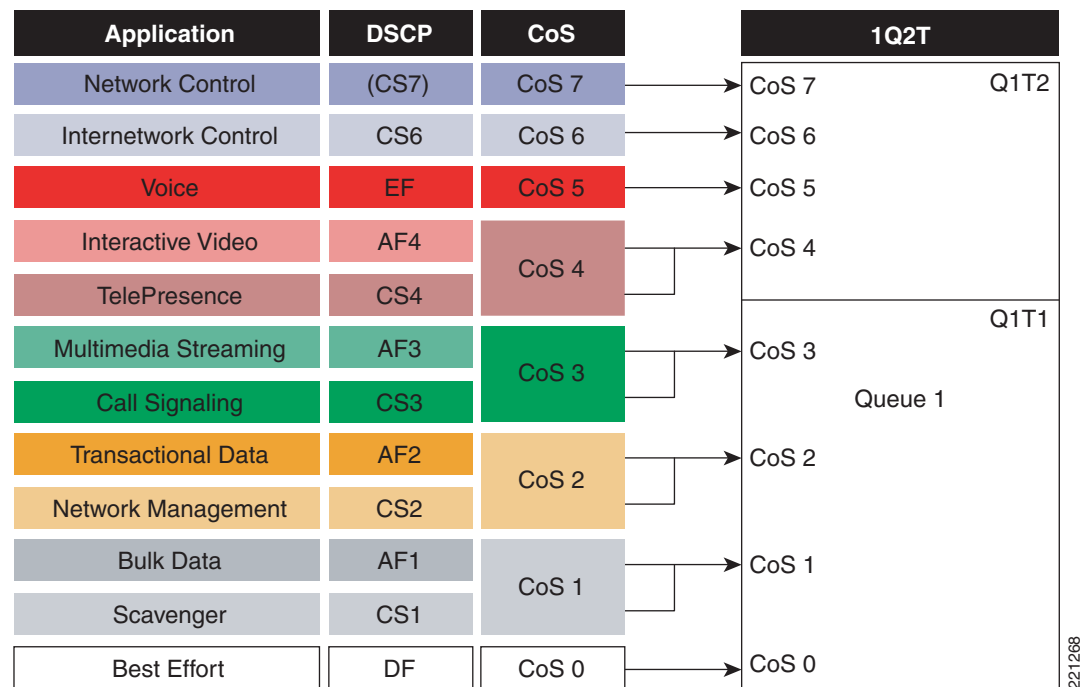
Videoconferencing/Video Telephony application was provisioned on each port that would permit 384 Kbps of AF41 video traffic per port, the combined total would be 270 Mbps (27% of a GE uplink). This leaves enough PQ traffic to support 2 separate CTS-3000 systems connected to the same chassis and still be at a theoretical maximum of only 30% of a single GE uplink.

With this in mind, let's now consider the best practice ingress and egress queuing configurations for the 6148A, 6548 and 6748 linecards.

Ingress Queuing Design—1Q2T

As shown in [Figure 5-10](#), both the 6148A and 6548 linecards support a CoS-based ingress queuing structure of 1Q2T which can be leveraged to offset their oversubscription ratios. The 1Q2T ingress queuing structure uses Tail-Drop thresholds, which by default are set at 80% of the queue (Q1T1) and at 100% of the queue (Q1T2 which, incidentally, is non-configurable). By default, CoS values 0 through 4 are mapped to Q1T1 and CoS values 5 through 7 are mapped to Q1T2. The only improvement we can make on this default configuration to optimize TelePresence traffic on these oversubscribed linecards is to map CoS 4 (TelePresence) to the second threshold (Q1T2), along with CoS 5 (VoIP) and CoS 6 and 7 (Network Control traffic). The recommending ingress 1Q2T queuing configuration for C6500 6148A and 6548 linecards is illustrated in [Figure 5-11](#).

Figure 5-11 Catalyst 6500 (1Q2T) Ingress Queuing Recommendations for TelePresence Deployments



The configuration for the C6500 1Q2T ingress queuing structure (for the 6148A and 6548 linecards) illustrated in [Figure 5-11](#) is as follows:

```
interface GigabitEthernet x/y
  rcv-queue cos-map 1 1 0 1 2 3
    ! Maps CoS values 0-3 to Q1T1
  rcv-queue cos-map 1 2 4 5 6 7
    ! Maps CoS values 4-7 to Q1T2
```

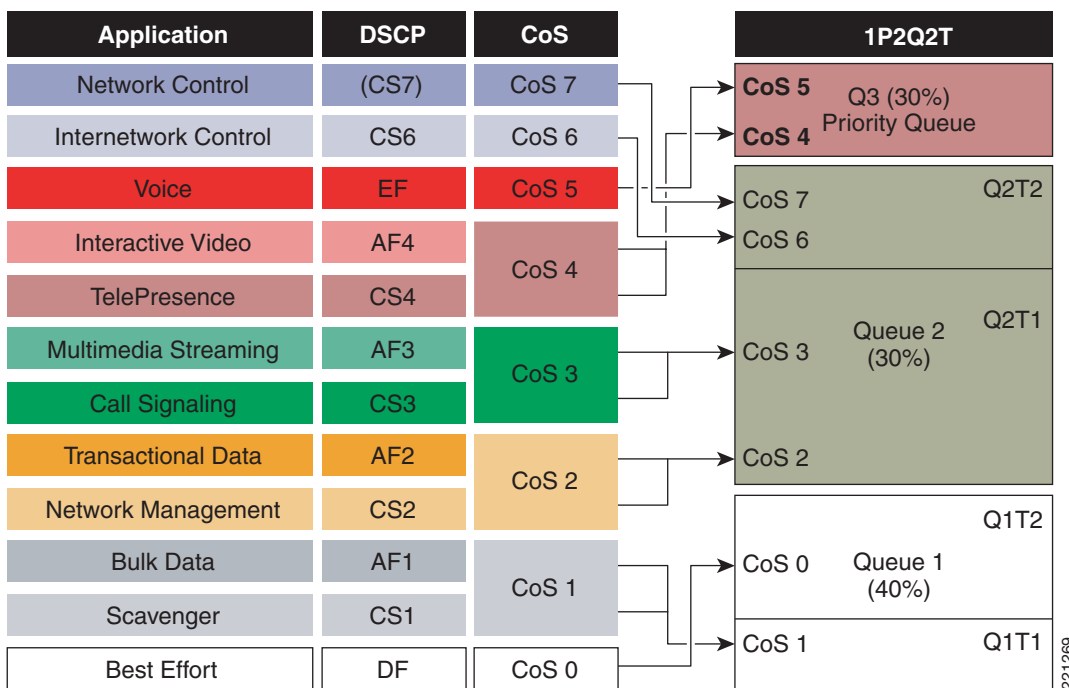
These configuration commands can be verified with the following command:

- **show queueing interface GigabitEthernet x/y**

Egress Queuing Design—1P2Q2T

As shown in [Figure 5-10](#), the Sup720s and 6548 linecards support a CoS-based egress queuing structure of 1P2Q2T, which uses WRED as a congestion avoidance mechanism. Under such a queuing structure, TelePresence traffic, along with VoIP, is recommended to be mapped to the strict-priority queue. Furthermore, because the egress queuing is CoS-based, Videoconferencing/Video Telephony (AF41) will also be assigned to the strict-priority queue. For non-realtime classes, the per-queue WRED thresholds can be configured to allow for granular QoS within a given queue. Specifically, Q1T1's minimum WRED threshold is set to 40% and its maximum threshold to 80%; then by mapping CoS 1 (Scavenger/Bulk) to Q1T1, we are restricting such traffic within Q1, with the remaining buffers (Q1T2) being exclusively reserved for CoS 0 (Best Effort traffic). Similarly, we can set Q2T1's minimum WRED threshold to 70% and maximum threshold to 80%; then by mapping CoS 2 (Transactional/Network Management) and CoS 3 (Call-Signaling/Multi-Media Streaming) to Q2T1, we are restricting these flows to a maximum of 80% of Q2, with the remaining buffers (Q2T2) being exclusively reserved for CoS 6 and 7 (Network Control traffic). The recommending egress 1P2Q2T queuing configuration for C6500 6548 linecards is illustrated in [Figure 5-12](#).

Figure 5-12 Catalyst 6500 (1P2Q2T) Egress Queuing Recommendations for TelePresence Deployments



The configuration for the C6500 1P2Q2T egress queuing structure (for the 6548 linecards) illustrated in [Figure 5-12](#) is as follows:

```
!
interface GigabitEthernet x/y

! This section sets the queue limits and bandwidth allocations
wrr-queue queue-limit 40 30
```

```

! Sets the buffer allocations to 40% for Q1 and 30% for Q2
! Also implicitly sets PQ (Q3) to 30%)
wrr-queue bandwidth 40 30
! Sets the WRR weights for 40:30 (Q1:Q2) bandwidth servicing

! This section sets the Min and Max WRED thresholds for Q1
wrr-queue random-detect min-threshold 1 40 80
! Sets Min WRED Thresholds for Q1T1 and Q1T2 to 40 and 80, respectively
wrr-queue random-detect max-threshold 1 80 100
! Sets Max WRED Thresholds for Q1T1 and Q1T2 to 80 and 100, respectively

! This section sets the Min and Max WRED thresholds for Q2
wrr-queue random-detect min-threshold 2 70 80
! Sets Min WRED Thresholds for Q2T1 and Q2T2 to 70 and 80, respectively
wrr-queue random-detect max-threshold 2 80 100
! Sets Max WRED Thresholds for Q2T1 and Q2T2 to 80 and 100, respectively

! This section maps the CoS values to the Queues/Thresholds
wrr-queue cos-map 1 1 1
! Maps CoS 1 (Scavenger/Bulk) to Q1 WRED Threshold 1
wrr-queue cos-map 1 2 0
! Maps CoS 0 (Best Effort) to Q1 WRED Threshold 2
wrr-queue cos-map 2 1 2 3
! Maps CoS 2 (Trans-Data & Mgmt) and CoS 3 (Call-Sig + Multimedia) to Q2T1
wrr-queue cos-map 2 2 6 7
! Maps CoS 6 (Routing) and CoS 7 (STP) to Q2 WRED Threshold 2
priority-queue cos-map 1 4 5
! Maps CoS 4 (TelePresence & Interactive-Video) and CoS 5 (VoIP) to the PQ
!

```

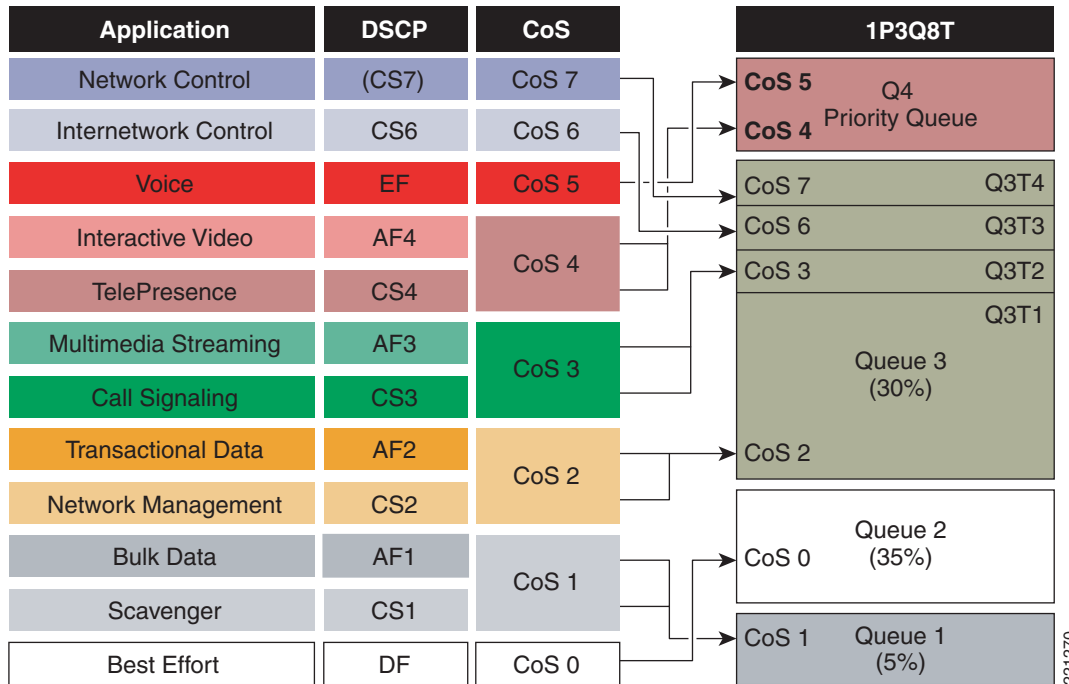
These configuration commands can be verified with the following command:

- **show queueing interface GigabitEthernet x/y**

Egress Queuing Design—1P3Q8T

As shown in [Figure 5-10](#), both the 6148A and 6748 linecards support a CoS-based egress queuing structure of 1P3Q8T, which uses WRED as a congestion avoidance mechanism. Under such a queuing structure, TelePresence traffic, along with VoIP, is recommended to be mapped to the strict-priority queue. Furthermore, because the egress queuing is CoS-based, Videoconferencing/Video Telephony (AF41) is also assigned to the strict-priority queue. CoS 1 (Scavenger/Bulk) can be constrained to a less than Best Effort queue: Q1. Q2 can then be dedicated for the default class (CoS 0). To minimize TCP global synchronization, WRED can be enabled on the non-realtime queues for congestion avoidance (technically, the congestion avoidance behavior is RED, as only one CoS weight is assigned to each queue). However, in Q3 the WRED thresholds can be set to give incremental preference to Network control traffic (CoS 7 and 6), followed by Call-Signaling traffic (CoS 3), and finally by Network Management traffic (CoS 2). The recommended egress 1P3Q8T queuing configuration for C6500 6148A and 6748 linecards is illustrated in [Figure 5-13](#).

Figure 5-13 Catalyst 6500 (1P3Q8T) Egress Queuing Recommendations for TelePresence Deployments



The configuration for the C6500 1P3Q8T egress queuing structure (for the 6548 linecards) illustrated in Figure 5-13 is as follows:

```

interface GigabitEthernet x/y

! This section sets the queue limits and bandwidth allocations
wrr-queue queue-limit 5 35 30
! Allocates 5% for Q1, 35% for Q2 and 30% for Q3
priority-queue queue-limit 30
! Allocates 30% for the Strict-Priority Queue (Q4)
wrr-queue bandwidth 5 35 30
! Sets the WRR weights for 5:35:30 (Q1:Q2:Q3) bandwidth servicing

! This section enables WRED on Q1, Q2 and Q3
wrr-queue random-detect 1
! Enables WRED on Q1
wrr-queue random-detect 2
! Enables WRED on Q2
wrr-queue random-detect 3
! Enables WRED on Q3

! This section sets Q1T1 WRED Thresholds to 80% (min) and 100% (max)
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
! Sets Min WRED Threshold for Q1T1 to 80% and all others to 100%
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
! Sets Max WRED Threshold for Q1T1 to 100% and all others to 100%

! This section sets Q2T1 WRED Thresholds to 80% (min) and 100% (max)
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
! Sets Min WRED Threshold for Q2T1 to 80% and all others to 100%
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
! Sets Max WRED Threshold for Q2T1 to 100% and all others to 100%

```

```

! This section sets Q3T1 to 60:70, Q3T2 to 70:80, Q3T3 to 80:90 and Q3T4 to 90:100
wrr-queue random-detect min-threshold 3 60 70 80 90 100 100 100 100
! Sets Min WRED Threshold for Q3T1 to 60%, Q3T2 to 70%, Q3T3 to 80%
! Q3T4 to 90%, and all others to 100%
wrr-queue random-detect max-threshold 3 70 80 90 100 100 100 100 100
! Sets Max WRED Threshold for Q3T1 to 70%, Q3T2 to 80%, Q3T3 to 90%
! and all others to 100%

! This section maps CoS values to egress Queues/Thresholds
wrr-queue cos-map 1 1 1
! Maps CoS 1 (Scavenger/Bulk) to Q1 WRED Threshold 1
wrr-queue cos-map 2 1 0
! Maps CoS 0 (Best Effort) to Q2 WRED Threshold 1
wrr-queue cos-map 3 1 2
! Maps CoS 2 (Net-Mgmt and Transactional Data) to Q3 WRED T1
wrr-queue cos-map 3 2 3
! Maps CoS 3 (Call-Signaling and Mission-Critical Data) to Q3 WRED T2
wrr-queue cos-map 3 3 6
! Maps CoS 6 (Routing) to Q3 WRED T3
wrr-queue cos-map 3 4 7
! Maps CoS 7 (Spanning Tree) to Q3 WRED T4
priority-queue cos-map 1 4 5
! Maps CoS 4 (TelePresence & Int-Video) and CoS 5 (VoIP) to the PQ
!

```

These configuration commands can be verified with the following command:

- **show queuing interface GigabitEthernet x/y**

Egress Queuing Design—1P7Q8T

As shown in [Figure 5-10](#), the 6704-10GE linecards support a CoS-based egress queuing structure of 1P7Q8T, which uses WRED as a congestion avoidance mechanism. Such a queuing structure allows for a dedicated queue for each CoS value.

Normally, TelePresence traffic (represented by CoS 4, along with Videoconferencing) would be mapped to the strict priority queue in campus queuing models; however, as a dedicated hardware queue exists on this platform, an exception could be made to this rule. As long as both TelePresence and Videoconferencing traffic are bounded by Call-Admission Control (CAC) mechanisms, these can adequately be provisioned in a dedicated, non-priority queue and still meet the tight service levels required by TelePresence. Remember, in the campus, latency and jitter are so minimal that they are considered negligible (typically in the microseconds on a per-switch basis). However loss is the most important service level consideration within the campus, since packet loss can easily occur within a campus due to instantaneous buffer congestion. The key is to ensure that switching buffers are never overrun during such instantaneous bursts. This can be achieved either by assigning TelePresence traffic to a (sufficiently deep) priority hardware queue or by assigning TelePresence to a non-priority queue, which has a limit to how much traffic can be admitted to the queue (via a CAC mechanism).

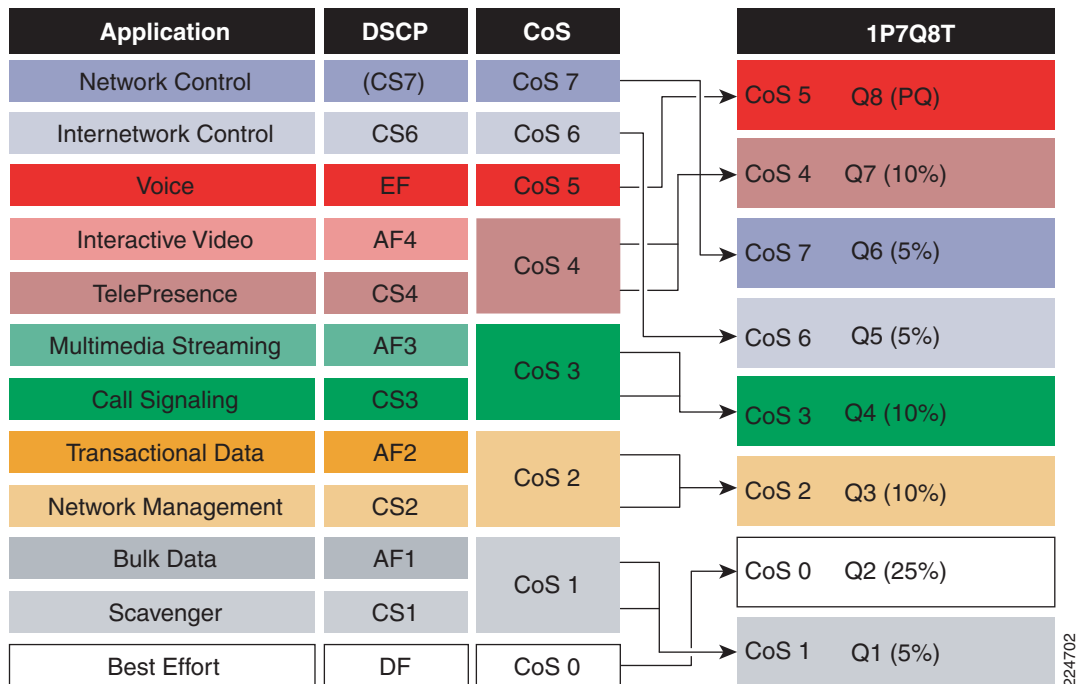
In this example, we present the option of assigning TelePresence (along with Videoconferencing) traffic to a dedicated non-priority queue (Q7), whereas VoIP continues to be mapped to the strict-priority queue (Q8). All other CoS-based traffic classes can be mapped to their dedicated queues.

To minimize TCP global synchronization, WRED can be enabled on the non-realtime queues for congestion avoidance (technically, the congestion avoidance behavior is RED, as only one CoS weight is assigned to each queue). This can be done by setting the minimum threshold for WRED to 80% and the maximum threshold to 100% (the tail of the queue) on a per-queue basis. However, there is no value—and only harm—from enabling WRED on the dedicated, non-priority queue for

TelePresence/Videoconferencing (Q7), as you never want to aggressively drop TelePresence traffic for any reason. As such, WRED can be disabled by setting the minimum and maximum WRED threshold to 100% (the tail of the queue).

The recommended egress 1P7Q8T queuing configuration for C6500 6704-10GE linecards is illustrated in Figure 5-14.

Figure 5-14 Catalyst 6500 (1P7Q8T) Egress Queuing Recommendations for TelePresence Deployments



The configuration for the C6500 1P7Q8T egress queuing structure (for the 6704-10GE linecards) illustrated in Figure 5-14 is as follows:

```

!
interface range TenGigabitEthernet4/1 - 4
  wrr-queue queue-limit 5 25 10 10 5 5 10
    ! Allocates 5% to Q1, 25% to Q2, 10% to Q3, 10% to Q4,
    ! Allocates 5% to Q5, 5% to Q6 and 10% to Q7
  wrr-queue bandwidth 5 25 10 10 5 5 10
    ! Sets the WRR weights for 5:25:10:10:5:5:10 (Q1 through Q7)
  priority-queue queue-limit 30
    ! Limits PQ to 30%

  wrr-queue random-detect 1
    ! Enables WRED on Q1
  wrr-queue random-detect 2
    ! Enables WRED on Q2
  wrr-queue random-detect 3
    ! Enables WRED on Q3
  wrr-queue random-detect 4
    ! Enables WRED on Q4
  wrr-queue random-detect 5
    ! Enables WRED on Q5
  wrr-queue random-detect 6
    ! Enables WRED on Q6
  wrr-queue random-detect 7

```

```

! Enables WRED on Q7

wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
! Sets Min WRED Threshold for Q1T1 to 80% and all others to 100%
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
! Sets Max WRED Threshold for Q1T1 to 100% and all others to 100%

wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
! Sets Min WRED Threshold for Q2T1 to 80% and all others to 100%
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
! Sets Max WRED Threshold for Q2T1 to 100% and all others to 100%

wrr-queue random-detect min-threshold 3 80 100 100 100 100 100 100 100
! Sets Min WRED Threshold for Q3T1 to 80% and all others to 100%
wrr-queue random-detect max-threshold 3 100 100 100 100 100 100 100 100
! Sets Max WRED Threshold for Q3T1 to 100% and all others to 100%

wrr-queue random-detect min-threshold 4 80 100 100 100 100 100 100 100
! Sets Min WRED Threshold for Q4T1 to 80% and all others to 100%
wrr-queue random-detect max-threshold 4 100 100 100 100 100 100 100 100
! Sets Max WRED Threshold for Q4T1 to 100% and all others to 100%

wrr-queue random-detect min-threshold 5 80 100 100 100 100 100 100 100
! Sets Min WRED Threshold for Q5T1 to 80% and all others to 100%
wrr-queue random-detect max-threshold 5 100 100 100 100 100 100 100 100
! Sets Max WRED Threshold for Q5T1 to 100% and all others to 100%

wrr-queue random-detect min-threshold 6 80 100 100 100 100 100 100 100
! Sets Min WRED Threshold for Q6T1 to 80% and all others to 100%
wrr-queue random-detect max-threshold 6 100 100 100 100 100 100 100 100
! Sets Max WRED Threshold for Q6T1 to 100% and all others to 100%

wrr-queue random-detect min-threshold 7 100 100 100 100 100 100 100 100
! Sets all Min WRED Thresholds for Q7 to 100 (Disables WRED on Q7)
wrr-queue random-detect max-threshold 7 100 100 100 100 100 100 100 100
! Sets all Max WRED Thresholds for Q7 to 100 (Disables WRED)
! WRED is disabled on TelePresence Queue (Q7); Tail-drop only

wrr-queue cos-map 1 1 1
! Maps Scavenger/Bulk to Q1T1
wrr-queue cos-map 2 1 0
! Maps Best Effort to Q2T1
wrr-queue cos-map 3 1 2
! Maps Trans-Data & Net-Mgmt to Q3T1
wrr-queue cos-map 4 1 3
! Maps MM-Streaming & Call-Signaling to Q4T1
wrr-queue cos-map 5 1 6
! Maps IP Routing to Q5T1
wrr-queue cos-map 6 1 7
! Maps Spanning Tree to Q6T1
wrr-queue cos-map 7 1 4
! Maps TelePresence to Q7T1 (WRED is Disabled)
priority-queue cos-map 1 5
! Maps Voice to the PQ (Q8)
mls qos trust dscp
! Sets interface to trust DSCP
!

```

**Note**

Due to WRED defaults and/or settings, the order of entering WRED min/max threshold commands may need to be reversed for some queues. The commands have been shown in the order presented to emphasize and simplify the logic of the configuration, rather than to appease the IOS parser.

These configuration commands can be verified with the following command:

- **show queueing interface GigabitEthernet x/y**

Egress Queuing Design—1P7Q4T (DSCP-to-Queue)

As shown in [Figure 5-10](#), the 6708-10GE linecards support a CoS- or DSCP-based egress queuing structure of 1P7Q4T, which uses WRED as a congestion avoidance mechanism.

The increased granularity in queue-mapping that is presented by the support of DSCP-to-Queue mapping allows—for the first time—TelePresence traffic to be separated from Videoconferencing traffic. As such, TelePresence can be assigned to its own dedicated, non-priority queue (Q7), separate from the dedicated, non-priority queue for Videoconferencing (Q5). Again, if a non-priority queue is to be used to service TelePresence traffic, then it is essential that a Call Admission Control mechanism be used to limit TelePresence traffic, inline with the bandwidth assigned to this queue.

VoIP (DSCP EF) will continue to be mapped to the strict-priority queue (Q8).

Queue 6 can be used to service control traffic, including (in respective order):

- Network Control traffic (DSCP CS7)
- Internetwork Control traffic (DSCP 6)
- Call Signaling traffic (DSCP CS3)
- Network Management traffic (CS2)

This respective order can be enforced by leveraging the 4 drop thresholds that can be assigned per non-priority queue. Specifically, CS2 can be mapped to Q6T1, CS3 can be mapped to Q6T2, CS6 can be mapped to Q6T3, and CS7 can be mapped to Q6T4 (the tail of the queue).

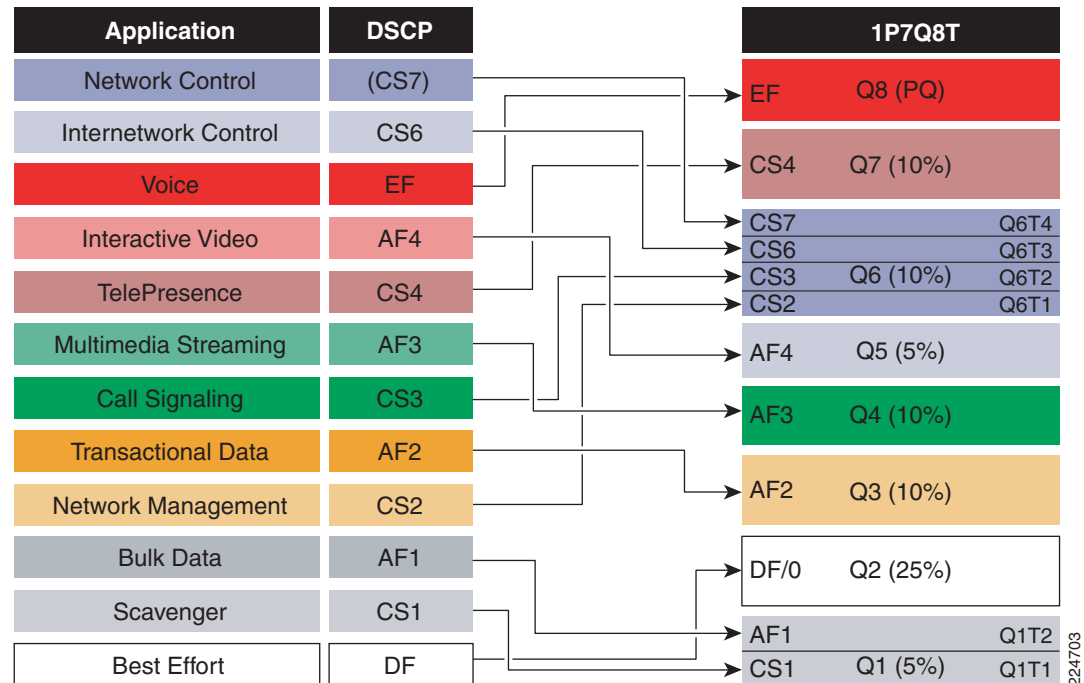
Multimedia Streaming (AF3) traffic can be assigned to a dedicated, non-priority queue (Q4), as can be Transactional Data (AF2) traffic (to Q3). This leaves Q2 to serve as a default/Best Effort queue and Q1 to serve as a “less than Best Effort” queue to service Bulk Data (AF1) and Scavenger (CS1) traffic, respectively. Again, the respective priority of Bulk Data over Scavenger traffic in Q1 can be enforced by leveraging the configurable WRED drop thresholds for the queue, such that Scavenger is mapped to Q1T1 and Bulk Data to Q1T2.

To minimize TCP global synchronization, WRED can be enabled on the non-realtime queues for congestion avoidance. Additionally, as this platform supports DSCP-to-Queue/Threshold mappings, the Assured Forwarding Per-Hop Behavior (AF PHB) can be fully-implemented on all AF traffic classes by enabling DSCP-based WRED (such that AFx3 drops before AFx2, which-respectively-is dropped before AFx1).

However, as noted before, there is no value—and only harm—from enabling WRED on the dedicated, non-priority queue for TelePresence (Q7), as you never want to aggressively drop TelePresence traffic for any reason. As such, WRED can be disabled by setting the minimum and maximum WRED threshold to 100% (the tail of the queue).

The recommended egress 1P7Q4T (DSCP-to-Queue) queuing configuration for C6500 6708-10GE linecards is illustrated in [Figure 5-15](#).

Figure 5-15 Catalyst 6500 (1P7Q4T—DSCP-to-Queue) Egress Queuing Recommendations for TelePresence Deployments



224703

The configuration for the C6500 1P7Q4T (DSCP-to-Queue) egress queuing structure (for the 6708-10GE linecards) illustrated in Figure 5-15 is as follows:

```

!
interface range TenGigabitEthernet4/1 - 8
  wrr-queue queue-limit 5 25 10 10 10 10 10
    ! Allocates 5% to Q1, 25% to Q2, 10% to Q3, 10% to Q4,
    ! Allocates 5% to Q5, 5% to Q6 and 10% to Q7
  wrr-queue bandwidth 5 25 10 10 10 10 10
    ! Sets the WRR weights for 5:25:10:10:10:10:10 (Q1 through Q7)
  priority-queue queue-limit 20
    ! Limits PQ to 20%

  wrr-queue random-detect 1
    ! Enables WRED on Q1
  wrr-queue random-detect 2
    ! Enables WRED on Q2
  wrr-queue random-detect 3
    ! Enables WRED on Q3
  wrr-queue random-detect 4
    ! Enables WRED on Q4
  wrr-queue random-detect 5
    ! Enables WRED on Q5
  wrr-queue random-detect 6
    ! Enables WRED on Q6
  wrr-queue random-detect 7
    ! Enables WRED on Q7

  wrr-queue random-detect min-threshold 1 60 70 80 90 100 100 100
    ! Sets Min WRED Thresholds for Q1T1 to 60%, Q1T2 to 70%, Q1T3 to 80%, Q1T4 to 90%
  wrr-queue random-detect max-threshold 1 70 80 90 100 100 100 100
    ! Sets Max WRED Thresholds for Q1T1 to 70%, Q1T2 to 80%, Q1T3 to 90%, Q1T4 to 100%

  wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100

```

```

! Sets Min WRED Threshold for Q2T1 to 80% and all others to 100%
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100
! Sets all Max WRED Thresholds for Q2 to 100%

wrr-queue random-detect min-threshold 3 70 80 90 100 100 100 100
! Sets Min WRED Thresholds for Q3T1 to 70%, Q3T2 to 80%, Q3T3 to 90%, Q3T4 to 100%
wrr-queue random-detect max-threshold 3 80 90 100 100 100 100 100
! Sets Max WRED Thresholds for Q3T1 to 80%, Q3T2 to 90%, Q3T3 & Q3T4 to 100%

wrr-queue random-detect min-threshold 4 70 80 90 100 100 100 100
! Sets Min WRED Thresholds for Q4T1 to 70%, Q4T2 to 80%, Q4T3 to 90%, Q4T4 to 100%
wrr-queue random-detect max-threshold 4 80 90 100 100 100 100 100
! Sets Max WRED Thresholds for Q4T1 to 80%, Q4T2 to 90%, Q4T3 & Q4T4 to 100%

wrr-queue random-detect min-threshold 5 70 80 90 100 100 100 100
! Sets Min WRED Thresholds for Q5T1 to 70%, Q5T2 to 80%, Q5T3 to 90%, Q5T4 to 100%
wrr-queue random-detect max-threshold 5 80 90 100 100 100 100 100
! Sets Max WRED Thresholds for Q5T1 to 80%, Q5T2 to 90%, Q5T3 & Q5T4 to 100%

wrr-queue random-detect min-threshold 6 60 70 80 90 100 100 100
! Sets Min WRED Thresholds for Q6T1 to 60%, Q6T2 to 70%, Q6T3 to 80%, Q6T4 to 90%
wrr-queue random-detect max-threshold 6 70 80 90 100 100 100 100
! Sets Max WRED Thresholds for Q6T1 to 70%, Q6T2 to 80%, Q6T3 to 90%, Q6T4 to 100%

wrr-queue random-detect min-threshold 7 100 100 100 100 100 100 100
! Sets all Min WRED Thresholds for Q7 to 100 (Disables WRED on Q7)
wrr-queue random-detect max-threshold 7 100 100 100 100 100 100 100
! Sets all Max WRED Thresholds for Q7 to 100 (Disables WRED)
! WRED is disabled on TelePresence Queue (Q7); Tail-drop only

mls qos trust dscp
! Sets interface to trust DSCP

mls qos queue-mode mode-dscp
! Enables DSCP-to-Queue mapping mode

wrr-queue dscp-map 1 1 8
! Maps Scavenger (CS1) to Q1T1
wrr-queue dscp-map 1 2 14
! Maps Bulk Data (AF13) to Q1T2 - AF13 PHB
wrr-queue dscp-map 1 3 12
! Maps Bulk Data (AF12) to Q1T3 - AF12 PHB
wrr-queue dscp-map 1 4 10
! Maps Bulk Data (AF11) to Q1T4 - AF11 PHB

wrr-queue dscp-map 2 1 0
! Maps Best Effort to Q2T1

wrr-queue dscp-map 3 1 22
! Maps Transactional Data (AF23) to Q3T1 - AF23 PHB
wrr-queue dscp-map 3 2 20
! Maps Transactional Data (AF22) to Q3T2 - AF22 PHB
wrr-queue dscp-map 3 3 18
! Maps Transactional Data (AF21) to Q3T3 - AF21 PHB

wrr-queue dscp-map 4 1 30
! Maps Multimedia Streaming (AF33) to Q4T1 - AF33 PHB
wrr-queue dscp-map 4 2 28
! Maps Multimedia Streaming (AF32) to Q4T2 - AF32 PHB
wrr-queue dscp-map 4 3 26
! Maps Multimedia Streaming (AF31) to Q4T3 - AF31 PHB

wrr-queue dscp-map 5 1 38
! Maps Multimedia Conferencing (AF43) to Q5T1 - AF43 PHB

```

```
wrr-queue dscp-map 5 2 36
! Maps Multimedia Conferencing (AF42) to Q5T2 - AF42 PHB
wrr-queue dscp-map 5 3 34
! Maps Multimedia Conferencing (AF41) to Q5T3 - AF41 PHB

wrr-queue dscp-map 6 1 16
! Maps Net-Mgmt (CS2) to Q6T1
wrr-queue dscp-map 6 2 24
! Maps Call-Signaling (CS3) to Q6T2
wrr-queue dscp-map 6 3 48
! Maps IP Routing (CS6) to Q6T3
wrr-queue dscp-map 6 4 56
! Maps Spanning-Tree (CS7) to Q6T4

wrr-queue dscp-map 7 1 32
! Maps TelePresence (CS4) to Q7T1

priority-queue dscp-map 1 46
! Maps Voice (EF) to the PQ (Q8)
!
```

**Note**

Due to WRED defaults and/or settings, the order of entering WRED min/max threshold commands may need to be reversed for some queues. The commands have been shown in the order presented to emphasize and simplify the logic of the configuration, rather than to appease the IOS parser.

These configuration commands can be verified with the following command:

- **show queueing interface GigabitEthernet x/y**

