



## CHAPTER 6

# Branch QoS Design for TelePresence

---

## TelePresence Branch QoS Design Overview

The primary business advantages of TelePresence systems include:

- Reduced travel time and expense
- Improved collaboration and productivity
- Improved quality of work/life (due to reduced travel)
- The green advantage of a reduced carbon footprint

However, these business advantages are not fully realized if TelePresence systems are connected solely via an Intra-Campus Deployment Model (as illustrated in [Figure 3-1](#)); rather, gaining these advantages requires TelePresence systems to be deployed over wide area networks, whether these are private WANs or Virtual Private Networks.

WANs or VPNs may be used to interconnect large campuses to each other or may be used to connect one or more large campuses with smaller branch offices (as illustrated in [Figure 3-3](#)). To simplify these permutations, we refer to all TelePresence connections over a wide area as Branch Places-in-the-Network (PINs).

Branch PINs serve as boundary points between local area and wide area networks and, as such, these are often the most bottlenecked PINs and therefore have the most critical QoS requirements within the network infrastructure. To help select the best policies to be used at these critical PINs, it is beneficial to review some important considerations, which we discuss next.

## LLQ versus CBWFQ Considerations

Probably the most controversial decision relating to TelePresence deployments is whether to provision TelePresence traffic over the WAN/VPN in a strict-priority Low-Latency Queue (LLQ) or in a dedicated bandwidth-guaranteed Class-Based Weighted-Fair Queue (CBWFQ).

In campus networks, placing TelePresence in the strict-priority hardware queues yielded superior results during testing, especially in terms of protection against packet loss during momentary periods of congestion, which occur regularly in campus networks even under normal operating conditions. Additionally, placing TelePresence traffic in these strict-priority queues does not involve any incremental or ongoing monetary expense (beyond initial configuration), as this potential for strict-priority servicing already exists within the campus network infrastructure and the exercise simply becomes a matter of re-configuring existing queuing structures to enable strict-priority queuing for TelePresence.

Therefore, the decision to service TelePresence with strict-priority queues within the campus is relatively straightforward.

However, the corresponding decision becomes more complicated over the WAN/VPN due to three main considerations:

- The cost of subscribing to realtime SP services
- The “33% LLQ Rule”
- The potential effect of TelePresence on VoIP

Let us look at each of these considerations in turn. The first and foremost consideration is the ongoing cost of subscribing to realtime services from a service provider. Service providers generally charge enterprise customers premium rates for the amount of traffic they want serviced within a realtime class. At times, these additional premiums may make it cost prohibitive to provision TelePresence traffic within a realtime SP class. At the very least, such expensive premiums could diminish the overall business cost savings that TelePresence can provide an enterprise (versus employee travel expenses).

The second consideration is the potential impact of the “33% LLQ Rule” (referenced in [Chapter 4, “Quality of Service Design for TelePresence”](#) in the section [Queuing TelePresence](#)). At times, administrators cannot provision adequate amounts of bandwidth for TelePresence and remain within this conservative design recommendation. This is generally the case when dealing with (45 Mbps) T3/DS3 links. According to the “33% LLQ Rule,” no more than 15 Mbps of traffic of such a link should be assigned for strict-priority servicing. However, if a network administrator already has VoIP provisioned (quite properly) in an LLQ on such a link, and is looking to also provision TelePresence with strict-priority servicing, then they have a decision to make. For example, if they wish to deploy a CTS-3000 at 1080p-Best (requiring 15 Mbps just for TelePresence), then they either have to upgrade the link’s bandwidth capacity (which is often cost-prohibitive, as generally the next tier of bandwidth is OC3) or they violate this design rule to accommodate all of their realtime traffic.

At this point, it bears repeating that the “33% LLQ Rule” is a conservative design recommendation, with the intent of reducing the variance in application response times of non-realtime applications during periods that the realtime classes are being utilized at maximum capacity. This is an exceptionally relevant concern when dealing with a high-bandwidth realtime application such as TelePresence.

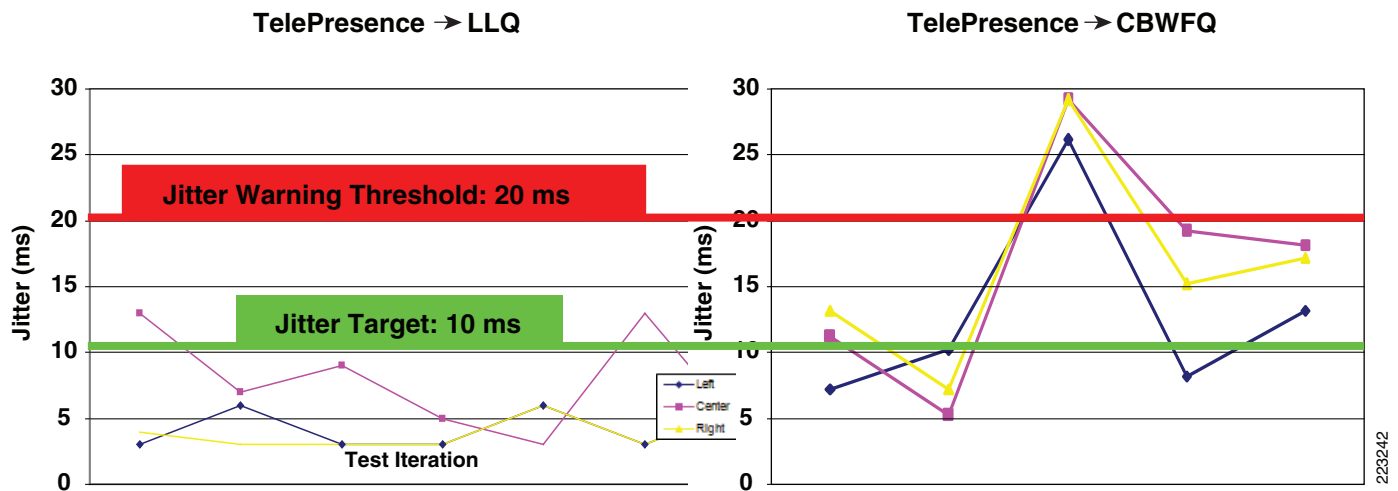
For example, let us reconsider a (45 Mbps) T3/DS3 link configured to support two separate CTS-3000 calls, both configured to transmit at full 1080p-Best resolution. Each such call requires 15 Mbps of realtime traffic. Prior to TelePresence calls being placed, non-realtime applications would have access to 100% of the bandwidth (to simplify the example, we are assuming there are no other realtime applications, such as VoIP, on this link). However, once these TelePresence calls are established, realtime TelePresence calls would suddenly dominate more than 66% of the link and all non-realtime applications would just as suddenly be contending for less than 33% of the link. TCP windowing for many of these non-realtime applications would begin slow-starting, resulting in many data applications hanging, timing out, or becoming stuck in a non-responsive state. Such network behavior, changing from one minute to the next, generally translates into users calling the IT help desk complaining about the network (which happens to be functioning properly, albeit in a poorly-configured manner).

That being said, it bears repeating that the “33% LLQ Rule” rule is not to be viewed as a mandate, but is simply a best practice design recommendation. There may be cases where specific business objectives cannot be met while holding to this recommendation. In such cases, enterprises must provision according to their detailed requirements and constraints. However, it is important to recognize the tradeoffs involved with over-provisioning realtime traffic classes in conjunction with the negative performance impact this has on non-realtime-application response times.

Quite naturally then, to make such provisioning decisions, a network administrator might wonder about the tradeoffs involved in TelePresence application performance when TelePresence is placed in a LLQ versus a CBWFQ. In such a comparison, the most sensitive service level attribute is jitter, as both

policies can be configured to completely prevent packet loss. As one might suspect, provisioning TelePresence in a LLQ results in lower peak-to-peak jitter values, as compared to provisioning TelePresence in a CBWFQ, which is shown in Figure 6-1.

Figure 6-1 Jitter Comparisons Between LLQ and CBWFQ WAN Edge Queuing Policies



Cisco Enterprise System Engineering testing showed that a 12-class RFC 4594-based QoS policy on a fully-congested T3 link—with TelePresence being serviced in a LLQ—yielded between 3-13 ms of peak-to-peak jitter to TelePresence; whereas an identical test—but with TelePresence being serviced in a CBWFQ—yielded between 5-29 ms of peak-to-peak jitter to TelePresence. As detailed in [Chapter 4, “Quality of Service Design for TelePresence,”](#) TelePresence has a one-way peak-to-peak jitter target of 10 ms and a warning threshold of 20 ms of peak-to-peak jitter, which if exceeded over an extended period can generate warning messages on the screen. During some of the CBWFQ tests, this warning message was observed on the screen, indicating that network congestion was affecting the TelePresence call quality.



**Note**

These tests were performed using TelePresence codec software version 1.1.0 (256D). Newer versions of CTS software have superior traffic smoothing capabilities as well as deeper de-jitter buffering, both of which amount to less overall sensitivity of TelePresence to jitter. Therefore the advantage of LLQ over CBWFQ queuing policies is less with newer versions of CTS software.

Therefore, while a moderate performance advantage to TelePresence can be observed when it is provisioned in a LLQ versus a CBWFQ, the advantage is not so great as to preclude recommending provisioning TelePresence in a CBWFQ when it is not viable to be provisioned with a LLQ. In other words, from a purely technical standpoint, **the best performance levels for TelePresence can be achieved when it is provisioned in a LLQ.** However, when other factors (such as additional ongoing costs or over-provisioning constraints for realtime bandwidth, etc.) need to be taken into account and render provisioning TelePresence in an LLQ unviable, then **the next best levels of service can be achieved by provisioning TelePresence in a dedicated CBWFQ.**

The third main consideration of whether to use LLQ or CBWFQ is the potential effect of TelePresence traffic on VoIP traffic if both are to be serviced in a strict-priority queue. To better understand how these realtime applications can be provisioned with strict-priority servicing and protected from interfering with each other, we must take a closer look at Cisco’s IOS LLQ/CBWFQ mechanisms. To do so, let us consider a simple LLQ/CBWFQ policy.

**Example 6-1 Simple LLQ/CBWFQ Policy**

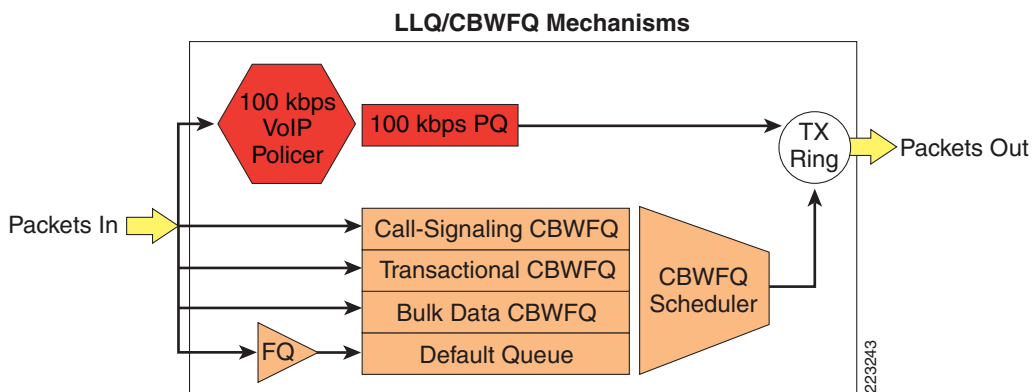
```

policy-map WAN-EDGE
  class VOIP
    priority 100
  class CALL-SIGNALING
    bandwidth percent 5
  class TRANSACTIONAL
    bandwidth percent 20
  class BULK
    bandwidth percent 10
  class class-default
    fair-queue

```

The underlying mechanisms for this LLQ/CBWFQ policy are graphically represented in [Figure 6-2](#).

**Figure 6-2 Cisco IOS LLQ/CBWFQ Mechanisms—Part 1**

**Note**

For the sake of simplicity, some Layer 2 subsystems (including Link Fragmentation and Interleaving) have been omitted from [Figure 6-2](#), as these mechanisms simply are not relevant at the link speeds required by TelePresence.

In [Figure 6-2](#), we see a router interface that has been configured with a 5-class LLQ/CBWFQ policy, with VoIP assigned to a 100 kbps LLQ and additional three explicit CBWFQs defined for Call-Signaling, Transactional Data, and Bulk Data respectively, as well as a default queue that has a Fair-Queuing pre-sorter assigned to it. There are two additional underlying mechanisms that may not be obvious from the configuration, but are shown in [Figure 6-2](#):

- An implicit policer attached to the LLQ
- A final output buffer called the Tx-Ring

Let us first take a look at the implicit policer attached to the LLQ. The threat posed by any strict priority-scheduling algorithm is that it could completely starve lower priority traffic. To prevent this, the LLQ mechanism has a built-in policer. This policer (like the queuing algorithm itself) engages only when the interface is experiencing congestion. Therefore, it is important to provision the priority classes properly. In this example, if more than 100 kbps of VoIP traffic was offered to the interface and the interface was congested, the excess VoIP traffic would be discarded by the implicit policer. However, traffic admitted by the policer gains access to the strict priority queue and is handed off to the Tx-Ring ahead of all other CBWFQ traffic.

The Tx-Ring is a final output buffer that serves the purpose of always having packets ready to be placed onto the wire so that link utilization can be driven to 100%. It is actually a full Tx-Ring that signals the Cisco IOS software to indicate that an interface is experiencing congestion and as such the LLQ/CBWFQ algorithms need to be engaged.

Now, let us consider the case of servicing not just VoIP with strict-priority queuing, but also TelePresence.

**Note**

For the sake of example and illustration simplicity, let us assume TelePresence only requires 400 kbps of traffic for these next two examples only.

Two options exist to the network administrator. The first is to admit both VoIP and TelePresence to the same LLQ. Thus our example policy becomes:

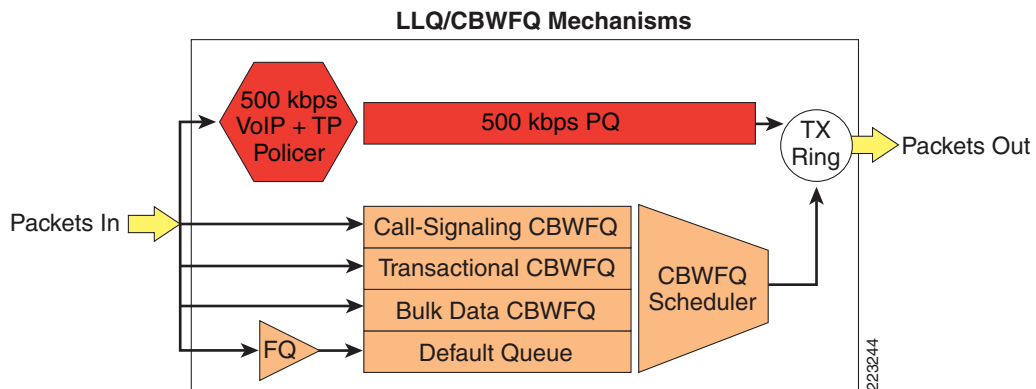
**Example 6-2 VoIP and TelePresence in a Single LLQ Policy**

```
class-map match-any REALTIME
  match dscp ef          ! Matches VoIP
  match dscp cs4        ! Matches TelePresence
  ...

policy-map WAN-EDGE
  class REALTIME
    priority 500          ! 100 kbps for VoIP + 400 kbps for TelePresence
  class CALL-SIGNALING
    bandwidth percent 5
  class TRANSACTIONAL
    bandwidth percent 20
  class BULK
    bandwidth percent 10
  class class-default
    fair-queue
```

The corresponding IOS mechanisms for [Example 6-2](#) are illustrated in [Figure 6-3](#).

**Figure 6-3 Cisco IOS LLQ/CBWFQ Mechanisms—Part 2**



In [Figure 6-3](#), we can see that not only has the LLQ been expanded in size (to 500 kbps), but also the implicit policer (for the combined VoIP and TelePresence class) has been increased to 500 kbps. Such a policy continues to protect VoIP from data as well as TelePresence from data. However, this policy does

potentially allow TelePresence to interfere with VoIP. This is because traffic offered to the LLQ class is serviced on a first-come, first-serve basis. Therefore, should TelePresence traffic suddenly burst, then it is possible—even likely—that VoIP traffic would be dropped.

At this point, we can realize another benefit of the implicit policer for the LLQ: not only does this mechanism protect non-realtime queues from bandwidth-starvation, but also it allows for Time-Division Multiplexing (TDM) of the LLQ. TDM of the LLQ allows for the configuration and servicing of “multiple” LLQs, while abstracting the fact that there is only a single LLQ “under-the-hood,” so to speak. Pertinent to our example, by configuring two LLQs, not only are VoIP and TelePresence protected from data applications, but VoIP and TelePresence are also protected from interfering with each other.

Let us take a look at our final policy example to cover this point. In [Example 6-3](#), a dual-LLQ design is used, one each for VoIP and TelePresence.

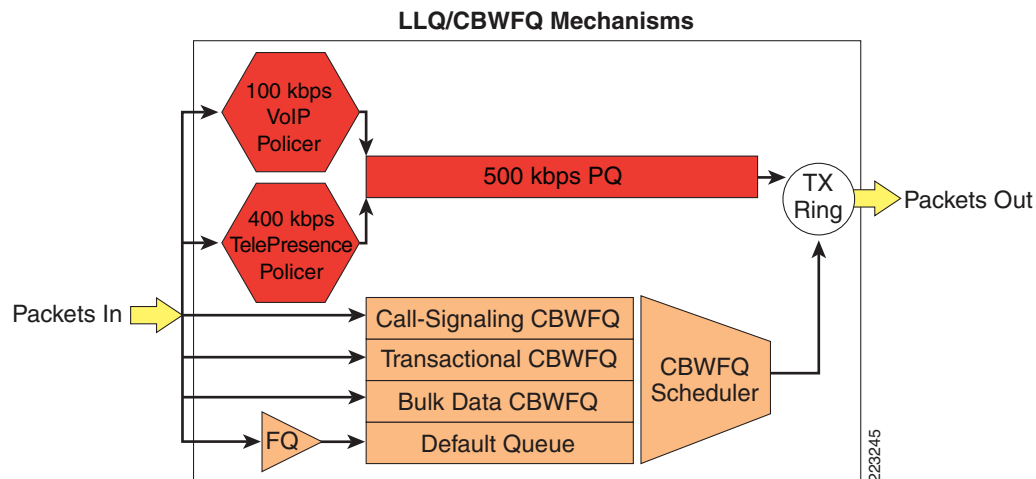
### Example 6-3 VoIP and TelePresence in a Dual-LLQ Policy

```
class-map match-all VOIP
  match dscp ef                ! Matches VoIP
class-map match-all TELEPRESENCE
  match dscp cs4              ! Matches TelePresence
...

policy-map WAN-EDGE
  class VOIP
    priority 100              ! 100 kbps LLQ for VoIP
  class TELEPRESENCE
    priority 400              ! 400 kbps LLQ for TelePresence
  class CALL-SIGNALING
    bandwidth percent 5
  class TRANSACTIONAL
    bandwidth percent 20
  class BULK
    bandwidth percent 10
  class class-default
    fair-queue
```

The corresponding IOS mechanisms for [Example 6-3](#) are illustrated in [Figure 6-4](#).

Figure 6-4 Cisco IOS LLQ/CBWFQ Mechanisms—Part 3



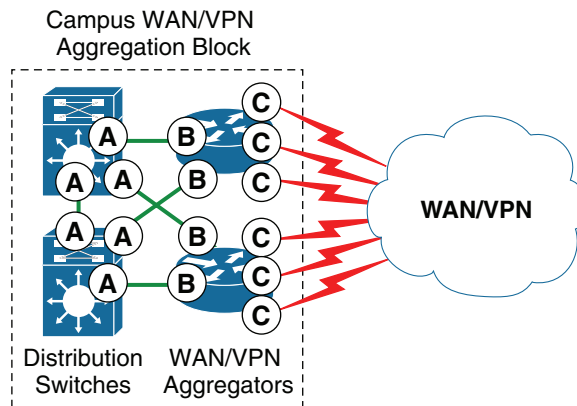
In [Figure 6-4](#), we see that two separate implicit policers have been provisioned, one each for the VoIP class (to 100 kbps) and another for the TelePresence class (to 400 kbps), yet there remains only a single strict-priority queue, which is provisioned to the sum of all LLQ classes, in this case to 500 kbps (100 kbps + 400 kbps). Traffic offered to either LLQ class is serviced on a first-come, first-serve basis until the implicit policer for each specific class has been invoked. For example, if TelePresence attempts to burst beyond a 400 kbps rate (remember, this rate has been reduced in order to simplify this example, both textually and also graphically), then it is dropped.

Therefore, to sum up this final consideration regarding whether or not to use LLQ for TelePresence: **if strict priority servicing for TelePresence is desired and viable, and if another realtime class (such as VoIP) has already been configured with a LLQ, then a dual-LLQ design would be recommended** in order to protect VoIP and TelePresence from interfering with each other.

## Campus WAN/VPN Block Considerations

Typically the first step in connecting a branch to a campus is to build out a WAN/VPN aggregation block at the main campus site. An example enterprise campus WAN/VPN aggregation block is illustrated in [Figure 6-5](#).

**Figure 6-5** Enterprise Campus WAN/VPN Aggregation Block QoS Design Recommendations for TelePresence



- A Interswitch Link Policies:**  
Trust DSCP  
+ Queuing (CoS 4 and 5 → PQ)  
+ Queuing (CoS 3 → Non-PQ)
- B Router LAN Edge Policies:**  
Trust DSCP (default)  
+ LLQ for VoIP (EF)  
+ LLQ for TelePresence (CS4)  
+ CBWFQ for Call-Signaling (CS3)
- C WAN/VPN Edge QoS Policies:**  
Trust for VoIP (EF)  
+ LLQ or CBWFQ for TelePresence (CS4)  
+ CBWFQ for Call-Signaling (CS3)

223246

Interswitch links are shown in [Figure 6-5](#) as points labeled A. While technically-speaking some of these links are interconnecting switches to routers, however their role and configuration are the same as the interswitch links described and defined in [Chapter 5, “Campus QoS Design for TelePresence.”](#) To be completely technically accurate, we could refer to these links as LAN-to-LAN non-edge links, but this term becomes a bit wordy and unwieldy, and as such we continue to use the simpler term interswitch links, but with a broadened meaning to include these switch-to-router links as well.

**Note**

As noted above, the term used here as interswitch links in this context refers to LAN-to-LAN non-edge links, not (necessarily) trunked links encapsulated with Cisco InterSwitch Link (ISL) trunking protocol.

As previously detailed (in [Chapter 5, “Campus QoS Design for TelePresence”](#)), all interswitch links should be configured to trust DSCP and perform hardware queuing, such that CoS 4 (TelePresence) and CoS 5 (VoIP) are assigned to the strict priority hardware queue and CoS 3 (Call-Signaling) is assigned to a non-priority queue within the platform/linecard’s 1PxDyT queuing structure.

Next, it would be recommended to enable LLQ/CBWFQ (or hardware queuing policies, if supported) on the router’s LAN edges, labeled as points B in [Figure 6-5](#). This is recommended when the levels of WAN/VPN aggregation may make it theoretically possible to oversubscribe these WAN-to-LAN links. For example, if the WAN/VPN aggregation router was homing seven individual OC3 circuits (totaling 7 \* 155 Mbps or 1.085 Gbps), but connecting to the distribution switches via GigabitEthernet links, then the potential for oversubscription on these WAN-to-LAN links would exist. Therefore, these links should be protected with a queuing policy, as a queuing policy would be the only way to provide service level **guarantees** on these links, regardless of how rarely these queuing policies would engage. As discussed in the previous section, if LLQ is to be used for VoIP and TelePresence, then these should be configured with a dual-LLQ policy, similar to the simplified example provided in [Example 6-3](#) (but with different bandwidth values for both VoIP and TelePresence, based on how many calls of each type were being supported over these links).

Finally, WAN/VPN edge QoS policies would be required on all points labeled C in [Figure 6-5](#). The specifics of these WAN/VPN edge policy permutations are discussed in detail in this chapter.

## TelePresence Branch LAN Edge

The LAN edge of the branch PIN performs essentially the same QoS services as does the campus access edge, namely the enforcement of a trust boundary, CoS-to-DSCP mapping (if required), optional TelePresence policing (to prevent network abuse of a trusted switch port), and queuing. However, there are some design considerations unique to the branch LAN edge discussed below.

### TelePresence Branch LAN Edge QoS Design Considerations

Depending on the platform(s) used at the branch, QoS functions may be performed in hardware, in software, or in a combination of both. This is because Cisco IOS-based routers perform QoS in software, while Cisco Catalyst switches perform QoS in hardware. Additionally, some devices, such as the Cisco Integrated Services Routers, combine functionality from both product families within a single platform (for example, a Cisco ISR equipped with a Cisco EtherSwitch network module).

A general rule of thumb relating to QoS design is to **always enable QoS policies in hardware, rather than in software**, whenever a choice exists. This is because QoS policies performed in software require (marginal) incremental CPU loads to enforce (the actual incremental load varies according to platform, line rates, policy complexity, traffic patterns, and other variables). However, QoS policies performed in

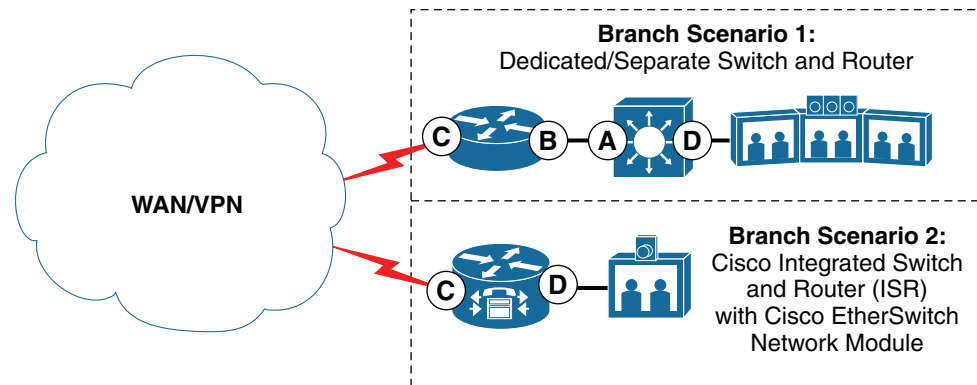
hardware are performed at line rates (GE or 10GE) without **any** incremental load to the CPU. Therefore, it is generally more efficient and effective to design QoS policies to be performed within hardware whenever possible.

Applying this rule of thumb to the branch LAN edge would typically result in one of two scenarios:

- The first scenario consists of a Cisco router on the WAN edge connecting to a Cisco Catalyst switch, which then connects to the branch endpoints, including the TelePresence system.
- The second scenario consists of a Cisco Integrated Services Router equipped with an EtherSwitch module performing all QoS functions within a single box.

These scenarios and their respective placements of QoS policies are illustrated in [Figure 6-6](#).

**Figure 6-6 Enterprise Branch QoS Design Recommendations for TelePresence**



#### Branch QoS Policies:

- |   |  |
|---|--|
| <p><b>(A) Interswitch Link Policies:</b><br/>Trust DSCP<br/>+ Queuing (CoS 4 and 5 → PQ)<br/>+ Queuing (CoS 3 → Non-PQ)</p>   | <p><b>(C) WAN/VPN Edge QoS Policies:</b><br/>Trust for VoIP (EF)<br/>+ LLQ or CBWFQ for TelePresence (CS4)<br/>+ CBWFQ for Call-Signaling (CS3)</p>  |
| <p><b>(B) Router LAN Edge Policies:</b><br/>Trust DSCP (default)<br/>+ (Optional) LLQ for VoIP (EF)<br/>+ (Optional) LLQ for TelePresence (CS4)<br/>+ (Optional) CBWFQ for Call-Signaling (CS3)</p> | <p><b>(D) Branch Access Edge Policies:</b><br/>Trust for DSCP or Trust CoS<br/>+ Map CoS 4 → DSCP CS4<br/>+ Map CoS 5 → DSCP EF and CoS 3 → DSCP CS3<br/>+ (Optional) Ingress Policing<br/>+ Queuing (CoS 4 and 5 → PQ)<br/>+ Queuing (CoS 3 → Non-PQ)</p> |

223247

We can see that for the most part the QoS policies deployed in the branch are reflective of the policies deployed at the campus WAN block. For example, the WAN/VPN edge QoS policies are applied to the branch router's WAN interface (to complement the WAN edge policies on the WAN/VPN aggregator's WAN edge). These are shown as points labeled C in [Figure 6-6](#). The considerations and details of these WAN/VPN edge policies are discussed in detail throughout this chapter.

In the case of a branch PIN using dedicated/separate switch(es) and router(s), the administrator has two additional policy points to configure: the router's LAN edge (shown the point labeled B in [Figure 6-6](#)) and the switch-to-router link (shown as the point labeled A in [Figure 6-6](#)). A queuing policy on the router's LAN edge (point B) is optional, as in many cases it may be theoretically impossible to congest this interface in the WAN-to-LAN direction: this interface would likely be a GigabitEthernet interface, and as such, well above the access rate of the WAN/VPN link. If it is to be configured with a queuing

policy, then a dual-LLQ policy would be recommended, such that VoIP and TelePresence are assigned to separate LLQs and Call-Signaling is protected with a CBWFQ. Such a policy would be similar to the simplified example provided in [Example 6-3](#) (but with different bandwidth values for both VoIP and TelePresence, based on how many calls of each type were being supported over these links).

Next, the administrator would need to configure the switch-to-router links, labeled as point A. These ports are essentially serving the same roles as campus interswitch links (as previously discussed in [Chapter 5, “Campus QoS Design for TelePresence”](#)). These ports should be configured to trust DSCP and perform hardware queuing, such that CoS 4 (TelePresence) and CoS 5 (VoIP) are assigned to the strict priority hardware queue and that CoS 3 (Call-Signaling) is assigned to a non-priority queue within the platform/linecard’s 1PxQyT queuing structure.

The final policy points are the branch access edges, which are shown as points labeled D in [Figure 6-6](#). These ports should be configured with either static or conditional trust of either DSCP or CoS (as described in detail in [Chapter 5, “Campus QoS Design for TelePresence”](#)” in [Access Switch Port QoS Considerations](#)).

If CoS is trusted, then the necessary CoS-to-DSCP mappings must be in place, such that CoS 4 (TelePresence) is mapped to (the default value of) DSCP CS4 (32), CoS 5 (VoIP) is mapped to (the non-default value of) DSCP EF (46), and CoS 3 (Call-Signaling) is mapped to (the default value of) DSCP CS3 (24).

An optional recommendation for the branch 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. This recommendation helps prevent an unknowing and/or disgruntled individual that gains physical access to the TelePresence switch port from sending rogue traffic over the network that can hijack voice or video queues and easily ruin voice 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 - DSCP 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**

Access Edge policers are described in detail on a platform-by-platform basis in Chapter 2, “Campus QoS Design” of the QoS SRND at [www.cisco.com/go/srnd](http://www.cisco.com/go/srnd).

Additionally, the recommended burst parameter for TelePresence policers is discussed in detail in [TelePresence Branch WAN Edge LLQ Policy](#).

Finally, it is recommended to enable queuing on these branch access edge links in the case of congestion. While the likelihood of such an event is rare, these may occur during DoS/worm attacks; therefore, provisioning queuing policies on these links is mandatory to provide service level guarantees in **any** event.

In the case of a branch PIN using an ISR with an EtherSwitch module, the only real twist is that the administrator would configure the WAN/VPN edge policies in the router console mode, and then switch to the EtherSwitch console mode (using the **service-module gigabitEthernet module/number session** IOS command) to configure hardware QoS policies on the EtherSwitch module (which is essentially a Cisco Catalyst 3750 switch), using the Catalyst 3750 configuration recommendations provided in [Chapter 5, “Campus QoS Design for TelePresence.”](#)

## TelePresence Branch LAN Edge QoS Designs

The configuration details for branch LAN edge policies are identical to the platform- and linecard-specific policies used in the campus access edge, which have already been covered in detail in [Chapter 5, “Campus QoS Design for TelePresence”](#) and as such it would be redundant to again detail these designs in this chapter.

## TelePresence Branch WAN Edge

The WAN edge of the branch is likely the most congested node within the network and as such requires the most attention from a QoS perspective. Let us now recap some of the considerations of the WAN edge and then delve into design detail.

## TelePresence Branch WAN Edge Design Considerations

In a private WAN design, the principal decision that the administrator needs to make has already been covered in detail, namely whether to service TelePresence with a LLQ or a CBWFQ. A point to keep in mind with respect to private WAN scenarios is that additional costs are typically not incurred when provisioning additional traffic with strict priority servicing and, as such, these scenarios are generally more conducive to provisioning TelePresence in a LLQ than other VPN scenarios.

Once the LLQ versus CBWFQ decision is made, then the WAN edge policies are fairly straightforward and are a function of this decision coupled with the number of traffic classes that have been defined by the enterprise’s strategic business QoS objectives (as discussed in [Chapter 4, “Quality of Service Design for TelePresence”](#)).

## TelePresence Branch WAN Edge QoS Design

To recap, LLQ provides superior levels of service for TelePresence as compared to CBWFQ, yet may entail additional costs or other constraints. Therefore, each administrator must make an informed decision as to which WAN edge queuing strategy (LLQ or CBWFQ) to employ. Both configuration options are presented in detail in the following sections.

### TelePresence Branch WAN Edge LLQ Policy

If TelePresence is to be assigned to an LLQ, then in addition to adequately provisioning priority bandwidth to the LLQ, one additional design parameter needs to be calculated: the burst parameter of the implicit policer of the LLQ.

The implicit policer for the LLQ is a token-bucket algorithm policer (like any other IOS or Catalyst policer) and as such needs a burst parameter to be defined in order to police to a sub-line rate. To better understand why this is so, let us briefly recap how token-bucket policers work.

To regulate transmissions at sub-line rates, the concept of an interval must be applied. An interval is a sub-second period of time during which an application may send traffic. For example, if a policer was to limit an application to 15 Mbps of a 45 Mbps circuit, then the policer would allow the application to transmit for a total of 333 ms per second (15 / 45 Mbps) and it would drop any packets offered during the remaining 667 ms per second. Now, if the application sent all its traffic in a single burst, this could tie up the circuit for up to one-third of a second, which may cause excessive jitter and/or drops to other applications. Therefore, rather than allowing a single interval per second for an application to send

traffic, it is generally more efficient to configure policers that allow for transmission over multiple sub-second intervals. The amount of traffic that an application can transmit during a sub-second interval is called the committed burst or Bc. The time interval itself is referred to as the time constant or Tc. The relationship between the burst, the interval, and the overall policing rate is:

$$\text{Bc} = \text{Policing Rate} * \text{Tc}$$

Now let us apply this theory to TelePresence traffic patterns so that we can define an optimal value for the burst parameter of the LLQ's implicit policer.

TelePresence codecs, whether operating at 720p or 1080p resolution, display 30 frames per second. Put differently, TelePresence codecs send information representing one frame every 33 ms (1 second/30 frames-per-second). We can use this information as a starting point, as it directly correlates to our interval (Tc).

Now, if TelePresence had a fixed packet size and a constant packetization rate (like VoIP), then we could simply divide the per second bandwidth requirements (shown in Table 4-1 in Chapter 4, “Quality of Service Design for TelePresence”) by 30 (fps) to arrive at our burst parameter. For example, under this assumption, a CTS-3000 transmitting at 1080p-Best would need a burst parameter of 62,500 Bytes (15 Mbps / 30 fps / 8 bits).

**Note**

However, while a fixed packet size and a constant packetization rate might make burst calculations a bit simpler, these would result in exponentially higher bandwidth requirements for TelePresence. As discussed in Chapter 4, “Quality of Service Design for TelePresence,” if TelePresence were uncompressed, it would result in 1.5 Gbps of bandwidth per display, rendering TelePresence virtually undeployable—especially over wide area networks.

However, TelePresence does not have a fixed packet size, nor a constant packetization rate, as it utilizes advanced video compression techniques to achieve compression rates of over 99%, thus massively reducing the bandwidth requirements for TelePresence and rendering it more deployable, even over WANs. Notwithstanding, these high compression algorithms within TelePresence systems do have a direct impact in burst calculations for policers, such as the implicit policer within LLQ.

Therefore, to configure the policing burst such that it does not drop TelePresence traffic, we have to analyze what would be the maximum transmission (in Bytes) within a 33 ms interval—in other words, the worst-case scenario per frame of TelePresence video. In H.264 video, which TelePresence systems utilize, this worst-case scenario would be the full screen of (spatially-compressed) video, which is periodically sent, known as the Instantaneous Decoding Refresh (IDR) frame. The IDR frame is the key frame that subsequent video frames reference, sending only differential information between subsequent frames and the IDR frame, rather than the full-picture again.

**Note**

For more information about H.264 video encoding, refer to RFC 3964 “RTP Payload Format for H.264 Video” at <http://www.ietf.org/rfc/rfc3984>.

The maximum IDR frame sizes observed during extensive testing of TelePresence systems (using CTS software version 1.1.0 [256D]) was 64 KB. Therefore, the LLQ burst parameter should be configured to permit up to 64 KB of burst per frame per screen. In the case of a triple-display CTS-3000 systems, we should allow for 192 KB of burst (3 \* 64 KB) in the rare event of a “triple-IDR storm,” where all three codecs send IDR frames simultaneously.

**Note**

If Cisco design recommendations for TelePresence room lighting and other environmental variables are not followed, then IDR frame sizes may vary in size beyond 64 KB, which may in turn affect the network QoS policies.

However, it bears mentioning that the version of CTS software used in this phase of testing (1.1.0 [256D]) did not support an auxiliary video stream. If newer versions of CTS software are being used or if the use of an auxiliary video stream (for sharing PowerPoint presentations, etc.) is planned, then a larger value of TelePresence burst would be required. Subsequent testing has shown that a value of 256 KB is sufficient to support TelePresence with an auxiliary video stream (192 KB for worst-case primary video + 64 KB for worst-case auxiliary video). Therefore, the examples that follow utilize this higher burst value to adequately provision for the use of TelePresence with an auxiliary video stream; if, on the other hand, such use is not planned, then a value of 192 KB is sufficient for TelePresence burst provisioning.

Now let us put this all together into a configuration. To quickly recap, the full syntax of the LLQ command in Cisco IOS is:

```
priority { bandwidth-kbps | percent percentage } [burst]
```

As can be seen, the burst parameter is an optional parameter that can be explicitly defined as part of the **priority** command. If the burst is not explicitly defined, then it defaults to a value computed as 200 ms of traffic at the configured LLQ bandwidth rate. However, it is important to note that the burst value is expressed in Bytes (not bits).

For example, if **priority 1000** was configured for a class, then the default burst parameter would be set to 25000 Bytes (1000 kbps \* 200 ms / 8 bits). This value would not appear in the configuration, but could be verified with a **show policy map interface** verification command.

Let us look at the worst-case burst for a CTS-1000 system. Applying the IDR as the worst-case burst scenario for TelePresence primary video (at 64 KB) coupled with an allowance of auxiliary video bursting of the same amount (64 KB), the configuration to provision a branch WAN edge queuing policy that provisions a TelePresence CTS-1000 system running at 1080p-Best (with the optional support of an auxiliary video stream) to a LLQ with an optimal burst parameter (of 128 KB) is shown in [Example 6-4](#).

**Example 6-4 Dual-LLQ Branch WAN Edge Policy for VoIP and TelePresence (CTS-1000 at 1080p-Best with Auxiliary Video)**

```
policy-map WAN-EDGE
  class VOIP
    priority percent 10           ! LLQ for VoIP (example amount of BW)
  class TELEPRESENCE
    priority 5500 128000        ! LLQ for CTS-1000 (1080p-Best + aux video)
  class DATA
    ...
```

Likewise, the configuration to provision a branch WAN Edge queuing policy that provisions a TelePresence CTS-3000 system running at 1080p-Best (with the optional support of an auxiliary video stream) to a LLQ with an optimal burst parameter (of 256 KB) is shown in [Example 6-5](#).

**Example 6-5 Dual-LLQ Branch WAN Edge Policy for VoIP and TelePresence (CTS-3000 at 1080p-Best with Auxiliary Video)**

```
policy-map WAN-EDGE
  class VOIP
    priority percent 10           ! LLQ for VoIP (example amount of BW)
  class TELEPRESENCE
    priority 15000 256000       ! LLQ for CTS-3000 (1080p-Best + aux video)
  class DATA
    ...
```

These configurations can be verified with the following command:

- **show policy-map interface**

## TelePresence Branch WAN Edge CBWFQ Policy

If, on the other hand, TelePresence is to be assigned to a CBWFQ, then in addition to adequately provisioning guaranteed bandwidth to the CBWFQ, one additional design parameter needs to be considered, the length of the CBWFQ.

By default, Class-Based Weighted Fair Queues are 64 packets deep. Extensive testing has shown that this default queue-depth has at times resulted in tail-drops when provisioned to protect TelePresence flows. Therefore, on most interfaces it is recommended to increase the default queue-depth for the TelePresence queue to 128 packets, using the **queue-limit 128** command in conjunction with the CBWFQ **bandwidth** command.

An example policy provisioning a TelePresence CTS-3000 system running at 1080p-Best (with the optional support of an auxiliary video stream) to a CBWFQ, with an extended queue-depth to 128 packets, is shown in [Example 6-6](#).

### Example 6-6 CBWFQ Branch WAN Edge Policy for TelePresence (CTS-3000 at 1080p-Best with Auxiliary Video)

```
policy-map WAN-EDGE
  class VOIP
    priority percent 10          ! LLQ for VoIP (example amount of BW)
  class TELEPRESENCE
    bandwidth 15000             ! CBWFQ for CTS-3000 (1080p-Best + aux video)
    queue-limit 128            ! Extended queue-limit for TelePresence CBWFQ
  class DATA
    ...
```

This configuration can be verified with the following command:

- **show policy-map interface**

## TelePresence Branch T3/DS3 WAN Edge Design

When configuring a WAN edge policy for TelePresence, there are a couple of additional considerations that need to be taken into account when using T3 interfaces, namely, adjusting the hold-queue size (if needed) to accommodate all LLQ/CBWFQs and tuning the Tx-Ring to minimize TelePresence jitter on converged links.

The total number of buffers that the IOS software allocates for queuing per interface (regardless of whether the interface is configured with FIFO, LLQ, or CBWFQ) is called the output queue or hold-queue. The size of the output queue and can be adjusted with the **hold-queue** interface command to a value between 0 and 4096 packets; the default output queue size of a T3 serial interface is 1000 packets.

Normally the default hold-queue size is sufficient for a T3 interface. Consider our worst-case example, where we have a 12-class RFC 4594-based QoS policy and let us choose the option with TelePresence assigned to a CBWFQ. Besides TelePresence, there are 10 CBWFQs, each a default queue-depth of 64 packets, for an output queue depth of, so far, 640 packets. Additionally, there is the 128 packet extended queue-depth for the TelePresence CBWFQ, bringing our running total output queue depth to 768 packets which, even factoring a moderate allowance for LLQ queue depth, is well below our default value of 1000 packets for this T3 interface.

However, should the administrator—for whatever reason—require expanding the queue depths of each CBWFQ to 128, then the output queue depth requirement would be at least (11 \* 128) 1408 packets, not even factoring a moderate allowance for the LLQ. In such a scenario, LLQ traffic could be impacted if

the output queue size was not expanded accordingly. For instance, in such a case, the network administrator could increase the size of the output queue to 1500 packets by using the **hold-queue 1500** interface command.

Earlier in this chapter we introduced the Tx-Ring. To quickly recap, the Tx-Ring represents the size of the final output buffer (a FIFO queue) that maximizes physical link bandwidth utilization by matching the outbound packet rate on the router with the physical interface rate. The Tx-Ring also serves to indicate interface congestion to the IOS software. Prior to interface congestion, packets are sent on a FIFO basis to the interface via the Tx-Ring. However, when the Tx-Ring fills to its queue-depth/limit, then it signals to the IOS software to engage any LLQ/CBWFQ policies that have been attached to the interface. Subsequent packets are then queued within IOS according to these LLQ/CBWFQ policies, dequeued into the Tx-Ring, and then sent out the interface in a FIFO manner. These operations are illustrated in [Figure 6-2](#), [Figure 6-3](#), and [Figure 6-4](#).

The Tx-Ring can be configured on certain platforms, such as the Cisco PA-T3+ port adapter interface, with the **tx-ring-limit** interface command. The value of the **tx-ring-limit** number can be from 1 to 32,767 packets. The default for serial interfaces on the PA-T3+ is 64 packets.

During testing it was observed that the default tx-ring-limit limit of 64 packets was shown to cause somewhat higher jitter values to TelePresence traffic during fully-congested scenarios. The reason for this is the bursty nature of TelePresence traffic. Even though TelePresence traffic is prioritized when LLQ/CBWFQ policies are active, if there are no TelePresence packets to send, the FIFO Tx-Ring is filled with other traffic. When a new TelePresence packet arrives, even if it gets priority treatment from the Layer 3 LLQ/CBWFQ queuing system, the packet are dequeued into the FIFO Tx-Ring when space is available. However, with the default settings, there can be as many as 63 (non-TelePresence) packets in the Tx-Ring in front of that TelePresence packet. In such a worst-case scenario it would take as long as 17 ms to transmit these non-TelePresence packets out of the T3 interface. This 17 ms of instantaneous delay (i.e., jitter) exceeds the jitter target for TelePresence.

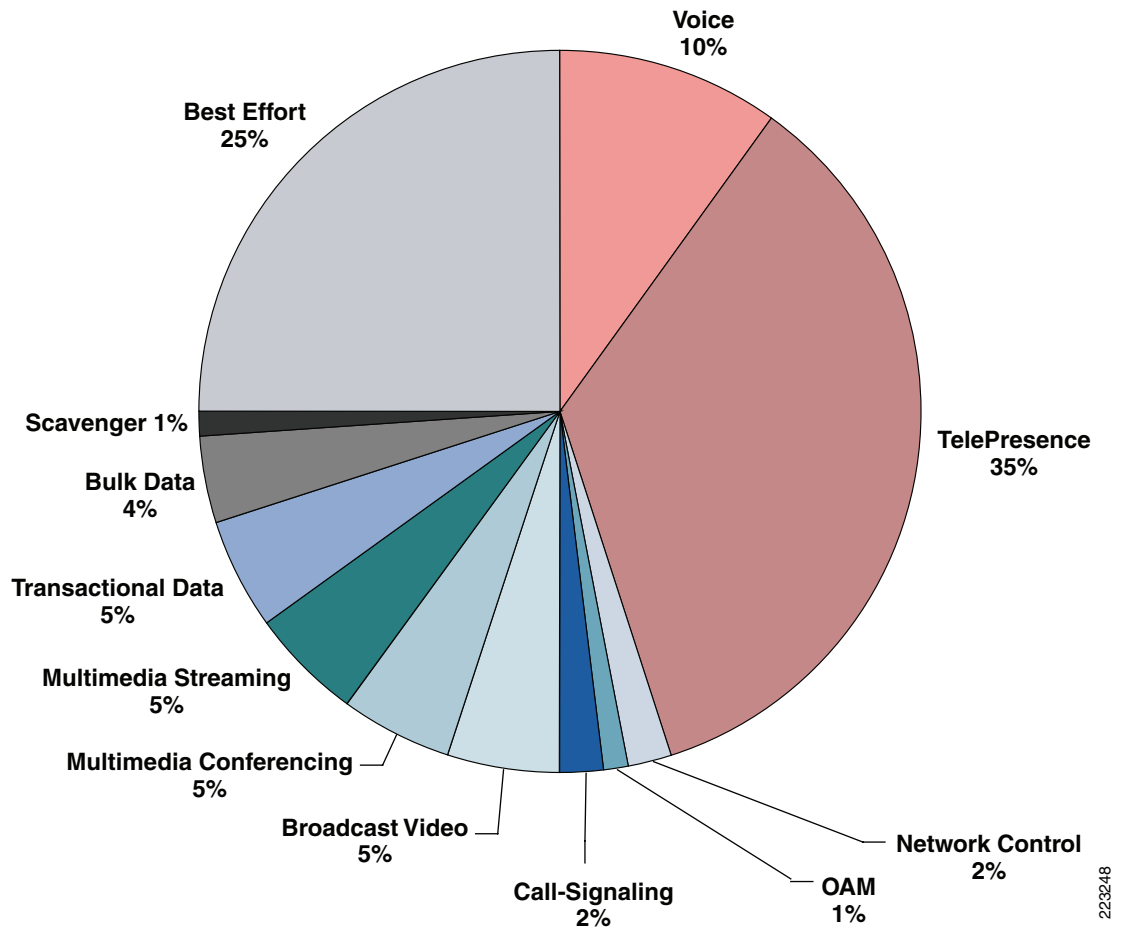
Additionally, a lower Tx-Ring results in the IOS software engaging congestion management policies sooner and more often, resulting in lower overall jitter values for priority traffic, such as TelePresence. On the other hand, setting the value of the Tx-Ring too low may result in significantly higher CPU utilization rates, as the processor is being continually interrupted to engage queuing policies, even when congestion rates are just momentary bursts and not sustained rates. Thus when tuning the Tx-Ring, a trade-off setting is required such that jitter is minimized, but not at the expense of excessive CPU utilization rates. Therefore, extensive testing has shown that setting the Tx-Ring to a value of 10 packets is optimal on converged T3 links supporting TelePresence and other applications (such as voice, data, and other video applications). This can be achieved by using the **tx-ring-limit 10** interface command.

**Note**

As explained above, tuning the Tx-Ring to value of 10 is only required on converged links that support additional applications beyond TelePresence. On T3 circuits that are dedicated to TelePresence, lowering the Tx-Ring to a non-default value is not required; in fact, such tuning can actually deteriorate the quality of TelePresence calls on such dedicated T3 circuits. It is important to keep in mind that in the case of dedicated circuits, it is not other applications that could potentially fill the Tx-Ring, but rather other TelePresence flows. Furthermore, when properly provisioned, dedicated links should not generate sustained congestion scenarios.

Now let us put this all together into a full example. In this 12-class RFC 4594-based case-study example, it was decided to service TelePresence traffic over the branch T3 WAN edge in a CBWFQ, as provisioning both VoIP and TelePresence in a dual-LLQ design would, in this case, require 45% of priority queuing, which would cause excessive variations in application response times to the other 10 application classes. The WAN edge bandwidth allocations for this case-study example are shown in [Figure 6-7](#).

**Figure 6-7** Case Study Example Bandwidth Allocations of a RFC 4594-Based LLQ/CBWFQ Policy Over a Branch T3 WAN Edge



The corresponding configuration for this case study example is shown in [Example 6-7](#).

**Example 6-7** Case Study Example Configuration of a RFC 4594-Based LLQ/CBWFQ Policy (with TelePresence in a CBWFQ) Over a Branch T3 WAN Edge

```

!
class-map match-all VOICE
  match dscp ef
class-map match-all TELEPRESENCE
  match dscp cs4
class-map match-all NETWORK-CONTROL
  match dscp cs6
class-map match-all OAM
  match dscp cs2
class-map match-all CALL-SIGNALING
  match dscp cs3
class-map match-all BROADCAST-VIDEO
  match dscp cs5
class-map match-all MULTIMEDIA-CONFERENCING
  match dscp af41 af42 af43
class-map match-all MULTIMEDIA-STREAMING
  match dscp af31 af32 af33
class-map match-all TRANSACTIONAL-DATA
  match dscp af21 af22 af23

```

```

class-map match-all BULK-DATA
  match dscp af11 af12 af13          ! Bulk-Data markings
class-map match-all SCAVENGER
  match dscp cs1                    ! Scavenger marking
!

!
policy-map WAN-EDGE-T3
  class VOICE
    priority percent 10              ! LLQ for VoIP
  class TELEPRESENCE
    bandwidth percent 35           ! CBWFQ for TP (CTS-3000)
    queue-limit 128                ! Expanded Queue-Limit for TP
  class NETWORK-CONTROL
    bandwidth percent 2              ! CBWFQ for Routing
  class OAM
    bandwidth percent 1              ! CBWFQ for Ops/Admin/Mgmt
  class CALL-SIGNALING
    bandwidth percent 2              ! CBWFQ for Call-Signaling
  class BROADCAST-VIDEO
    bandwidth percent 5              ! CBWFQ for Broadcast Video
  class MULTIMEDIA-CONFERENCING
    bandwidth percent 5              ! CBWFQ for IP/VC
    random-detect dscp-based         ! DSCP-WRED for IP/VC
  class MULTIMEDIA-STREAMING
    bandwidth percent 5              ! CBWFQ for Streaming-Video
    random-detect dscp-based         ! DSCP-WRED for Stream-Video
  class TRANSACTIONAL-DATA
    bandwidth percent 5              ! CBWFQ for Trans-Data
    random-detect dscp-based         ! DSCP-WRED for Trans-Data
  class BULK-DATA
    bandwidth percent 4              ! CBWFQ for Bulk Data
    random-detect dscp-based         ! DSCP-WRED for Bulk Data
  class SCAVENGER
    bandwidth percent 1              ! Minimum CBWFQ for Scavenger
  class class-default
    bandwidth percent 25             ! CBWFQ for Best Effort
    random-detect                    ! WRED for Best Effort
!
...
!
interface Serial6/0
  description BRANCH-TO-CAMPUS-T3
  ip address 192.168.2.9 255.255.255.252
  tx-ring-limit 10                ! Tuned T3 Tx-Ring
  dsu bandwidth 44210
  framing c-bit
  cablelength 10
  serial restart-delay 0
  max-reserved-bandwidth 100      ! LLQ/CBWFQ BW Override
  service-policy output WAN-EDGE-T3 ! Attaches policy to T3 int
!

```

**Note**

Since this policy supports a less-than Best Effort Scavenger-class, it requires an explicit CBWFQ to be configured on class-default (the Best Effort class); otherwise, the queuing algorithm robs bandwidth from class-default to service Scavenger traffic. Additionally, when class-default is configured with an explicit **bandwidth** command, then the **max-reserved-bandwidth** interface command must also be configured on the outgoing interface. Additional details on this behavior can be found in the QoS SRND at [www.cisco.com/go/srnd](http://www.cisco.com/go/srnd) on pages 3-8 and 3-9.

Optionally, if the network administrator chooses to use LLQ instead of CBWFQ, then the only change to the above policy would be to the TELEPRESENCE class within the WAN-EDGE-T3 policy-map, as shown in [Example 6-8](#).

**Example 6-8 Policy Amendment to Example 6-7 to Provision TelePresence in a Dual-LLQ Policy Over a T3 Branch WAN Edge**

```
!
policy-map WAN-EDGE-T3
  class VOICE
    priority percent 10           ! LLQ for VoIP
  class TELEPRESENCE
    priority percent 35 256000   ! LLQ for CTS-3000 (1080p-Best + aux video)
  class NETWORK-CONTROL
...

```

These configurations can be verified with the following commands:

- **show interface**
- **show policy-map interface**
- **show controllers Serial *module/interface* | include tx\_limited**

## TelePresence Branch OC3-POS WAN Edge Design

When configuring a WAN edge policy for TelePresence, there are a couple of additional considerations that need to be taken into account when using OC3-POS interfaces, such as the Cisco 7600 SPA-2XOC3-POS. Both of these considerations relate to Low-Latency Queuing:

- There is a 35% hard limit to the amount of traffic that can be configured with priority queuing.
- There is different configuration syntax for enabling LLQ on these interfaces.

Let us discuss these considerations in more detail.

The first consideration is fairly straightforward: on these OC3-POS interfaces, there is a hard limit of 35% for the sum of all traffic that can be configured with LLQ. This means that either a single LLQ class can be configured with a maximum of 54.25 Mbps or the sum of all LLQ classes can be configured for a combined maximum of 54.25 Mbps. This hard-limit, incidentally, is quite consistent with Cisco's "33% LLQ Rule."

The second consideration has to do with a change in syntax for configuration of LLQ on these OC3-POS interfaces. The configuration syntax for strict priority queuing on an OC3 POS interface is such the **priority** command does not include a bandwidth parameter, either as an absolute value (defined in kbps) or as a percentage of the link's bandwidth. That being said, there is correspondingly no implicit policer within the LLQ **priority** command, but rather an explicit policer must be configured on the policy-map class and then the **priority** command can be applied to the class. This difference is not limited to syntax only, but also affects behavior, as an implicit policer only engages when the LLQ is active (i.e., during periods of congestion), but an explicit policer, such as required for a LLQ class on an OC3-POS interface, is always on.

**Note**

The always-on nature of explicit policers is advantageous from an Admission Control perspective. For example, without a comprehensive, network-aware Call Admission Control system in place, there would be no way to always enforce limits on TelePresence traffic without explicit policers (remember, implicit policers, like those included within LLQ, are only active during congestion scenarios). Admission Control considerations and designs are discussed in more detail in [Chapter 8, “Capacity Planning and Call Admission Control.”](#)

The configured explicit LLQ policer may be either a single-rate or a dual-rate policer. When applied to TelePresence traffic, only a single-rate policer is relevant (as we are not interested in marking down excess TelePresence traffic, which we could do with two levels of granularity via a dual-rate policer). Additionally, testing has shown that using a dual-rate policer offers no performance advantages whatsoever; therefore it is recommended to configure a single-rate policer on the TelePresence LLQ class.

As with an implicit policer, a committed burst parameter is required when defining an explicit policer. As discussed in [TelePresence Branch WAN Edge LLQ Policy](#), the recommended value for TelePresence committed burst for a CTS-3000 system running 1080p-Best (with optional auxiliary video support) is 256 KB.

Reflecting the foregoing points, the configuration syntax for creating a single-rate explicit policer and applying it to a TelePresence LLQ class (for a CTS-3000 system running at 1080p-Best with auxiliary video) LLQ is shown in [Example 6-9](#).

**Example 6-9 TelePresence LLQ Policy Over a OC3-POS Branch WAN Edge**

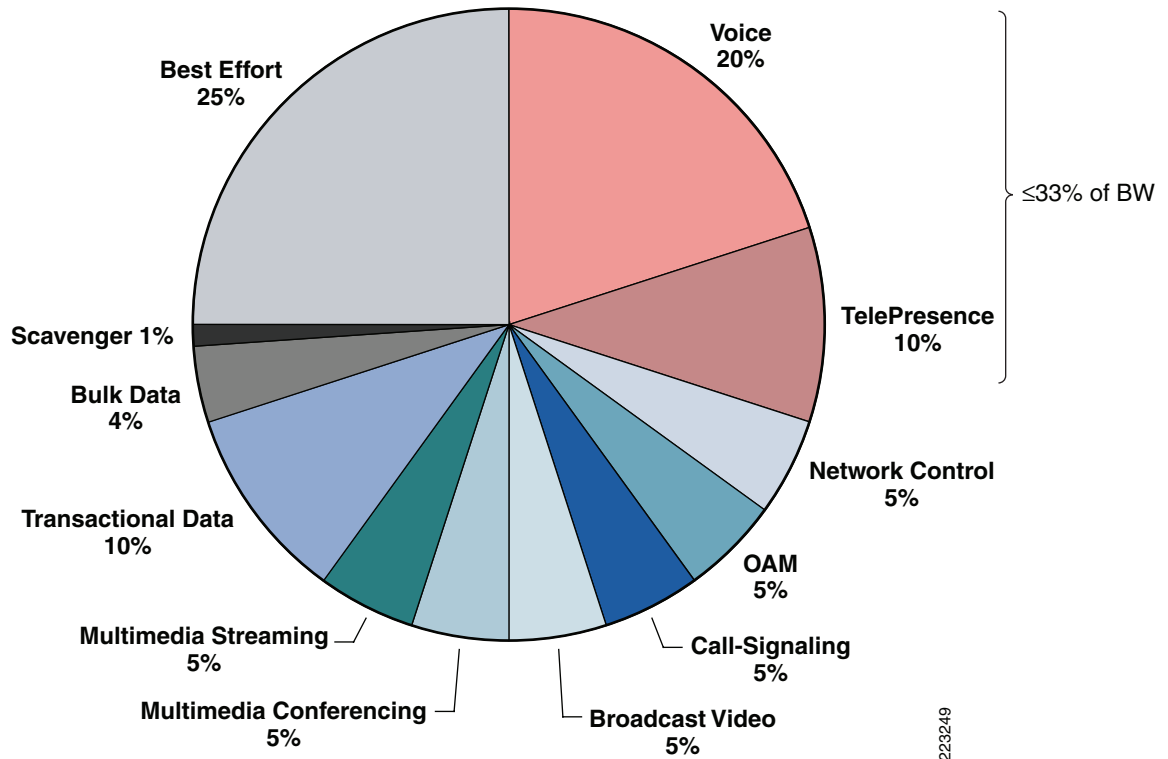
```
!
policy-map WAN-EDGE
class TELEPRESENCE
  police cir 15000000           ! TP is policed to 15 Mbps
    bc 256000                 ! Bc is 256 KB
    conform-action transmit    ! Conforming action --> transmit
    exceed-action drop         ! Single-Rate Policing action
  priority                    ! LLQ command for OC3-POS
!
```

This configuration can be verified with the following command:

- **show policy-map interface**

Now let us again put this all together into a full example. In this 12-class RFC 4594-based case-study example, it has been decided to service TelePresence traffic over the branch OC3-POS WAN edge in a dual-LLQ design, along with voice. The WAN edge bandwidth allocations for this case-study example are shown in [Figure 6-8](#).

**Figure 6-8 Case Study Example Bandwidth Allocations of a RFC 4594-Based LLQ/CBWFQ Policy Over a Branch OC3-POS WAN Edge**



The corresponding configuration for this second case study example is shown in [Example 6-10](#) (the class-maps are not repeated, as these do not change).

**Example 6-10 Case Study Example Configuration of a RFC 4594-Based LLQ/CBWFQ Policy (with TelePresence in a Dual-LLQ) Over a Branch OC3-POS WAN Edge**

```

!
policy-map WAN-EDGE-OC3-POS
  class VOICE
    police cir 31000000           ! Voice is policed to 31 Mbps (20%)
      bc 15500                   ! Bc is 15.5 KB
      conform-action transmit    ! Conforming action --> transmit
      exceed-action drop        ! Single-Rate Policing action
      priority                   ! LLQ command for OC3-POS
  class TELEPRESENCE
    police cir 15000000         ! TP is policed to 15 Mbps
      bc 256000                 ! Bc is 256 KB
      conform-action transmit    ! Conforming action --> transmit
      exceed-action drop        ! Single-Rate Policing action
      priority                   ! LLQ command for OC3-POS
  class NETWORK-CONTROL
    bandwidth percent 5        ! CBWFQ for Routing
  class OAM
    bandwidth percent 5        ! CBWFQ for Network Management
  class CALL-SIGNALING
    bandwidth percent 5        ! CBWFQ for Call-Signaling
  class BROADCAST-VIDEO
    bandwidth percent 5        ! CBWFQ for Broadcast Video
  class MULTIMEDIA-CONFERENCING
    bandwidth percent 5        ! CBWFQ Video-Conferencing

```

```

    random-detect dscp-based                ! DSCP-WRED for Video-Conferencing
class MULTIMEDIA-STREAMING
    bandwidth percent 5                     ! CBWFQ for Streaming-Video
    random-detect dscp-based                ! DSCP-WRED for Streaming-Video
class TRANSACTIONAL-DATA
    bandwidth percent 10                   ! CBWFQ for Transactional Data
    random-detect dscp-based                ! DSCP-WRED for Transactional Data
class BULK-DATA
    bandwidth percent 4                     ! CBWFQ for Bulk Data
    random-detect dscp-based                ! DSCP-WRED for Bulk Data
class SCAVENGER
    bandwidth percent 1                     ! Minimum CBWFQ for Scavenger
class class-default
    bandwidth percent 25                     ! CBWFQ for Best Effort
    random-detect ! WRED for Best Effort
!
...
interface POS3/0/1
description BRANCH-TO-CAMPUS-OC3-POS
ip address 192.168.5.1 255.255.255.252
clock source internal
service-policy output WAN-EDGE-OC3-POS ! Attaches policy to OC3-POS
!

```

**Note**

No Tx-Ring tuning is required on the OC3-POS link; neither is a **max-reserved-bandwidth 100** interface command required.

Optionally, if the network administrator chooses to use CBWFQ instead of LLQ for TelePresence, then the only change to the above policy would be to the TELEPRESENCE class within the WAN-EDGE-OC3-POS policy-map, as shown in [Example 6-11](#).

**Example 6-11 Policy Amendment to Example 6-10 to Provision TelePresence in a CBWFQ Over an OC3-POS Branch WAN Edge**

```

!
policy-map WAN-EDGE-OC3-POS
class VOICE
    police cir 31000000                     ! Voice is policed to 31 Mbps (20%)
    bc 15500                               ! Bc is 15.5 KB
    conform-action transmit                 ! Conforming action --> transmit
    exceed-action drop                     ! Single-Rate Policing action
    priority                               ! LLQ command for OC3-POS
class TELEPRESENCE
    bandwidth percent 10                   ! CBWFQ for TelePresence
class NETWORK-CONTROL
...

```

**Note**

Testing has shown that extending the TelePresence CBWFQ queue-limit beyond the default value of 64 packets is not required on OC3-POS interfaces because of the extremely fast serialization rate—relative to TelePresence transmission rates—of these interfaces.

These configurations can be verified with the following command:

- **show policy-map interface**

## TelePresence Branch IPSec VPN Edge

In the initial releases of Cisco TelePresence software, native encryption within the codecs was not supported. Nonetheless, for certain enterprises, encryption is a business requirement for all IP communications. This business requirement can be achieved by performing IPSec encryption and decryption within the network infrastructure, such as at the branch WAN/VPN edges over a private WAN or an MPLS VPN infrastructure. However, it is good to review some design considerations relating to IPSec and QoS interaction prior to examining the configuration details required for these deployments.



### Note

Cisco does not recommend deploying TelePresence over the Internet—with or without IPSec encryption—as critical service level parameters, such as latency, jitter, and loss, cannot be guaranteed over the Internet. These IPSec designs for TelePresence are intended for use over private WAN and/or MPLS VPN scenarios.

## TelePresence Branch IPSec VPN Edge Considerations

One of the first considerations is that IPSec adds network overhead to the packets. How much overhead depends on the encryption and tunneling options defined within the security associations. For example, a typical IPSec configuration uses IPSec Tunnel Mode with **esp-3des** and **esp-md5-hmac**, which results in an overhead of 56 bytes, accounted for as shown in [Table 6-1](#).

**Table 6-1**      *IPSec Network Overhead Breakdown*

Component	Overhead in Bytes
IPSec header (bytes)	20
ESP Header (bytes SPI)	4
ESP Header (bytes Sequence)	4
IOS ESP-DES/3DES (bytes IV)	8
ESP-DES/3DES 64-bit (bytes pad)	6
ESP Trailer (byte PAD length)	1
ESP Trailer (byte Next Header)	1
ESP MD5 96 digest (bytes)	12
Total IPSec Overhead	56

This being the case, since the average packet size for TelePresence is around 1200 Bytes, encryption overhead is typically <5%. [Table 6-2](#) shows a detailed breakdown of the respective encrypted bandwidth requirements for all TelePresence motion-handling and resolution options.

**Table 6-2** *TelePresence Bandwidth Requirements with IPSec Encryption*

Motion Handling	Best	Better	Good	Best	Better	Good
Resolution	1080p	1080p	1080p	720p	720p	720p
CTS 1000						
Max with IPSec overhead (Kbps)	5,792	5,194	4,596	4,596	3,400	2,204
CTS 3000						
Max with IPSec overhead (Kbps)	15,360	13,566	11,772	11,772	8,184	4,596

Another important consideration is the interaction of IPSec and QoS, particularly with respect to Anti-Replay. In order to understand this interaction implication, it is beneficial to briefly recap the purpose and function of IPSec Anti-Replay.

IPSec offers inherent message-integrity mechanisms to provide a means to identify whether an individual packet is being replayed by an interceptor or hacker. This concept is called connectionless integrity. IPSec also provides for partial sequence integrity, preventing the arrival of duplicate packets.

**Note**

Anti-Replay concepts are outlined in RFC 2401, “Security Architecture for the Internet Protocol” at [www.ietf.org/rfc/rfc2401](http://www.ietf.org/rfc/rfc2401).

When ESP authentication is configured in an IPSec transform set, for each security association, the receiving IPSec peer verifies that packets are received only once. Because two IPSec peers can send millions of packets, a 64-packet sliding window is implemented to bind the amount of memory required to tally the receipt of a peer’s packets. Packets can arrive out of order, but they must be received within the scope of the window to be accepted. If they arrive too late (outside the window), they are dropped.

The operation of the Anti-Replay window protocol is as follows:

1. The sender assigns a unique sequence number (per security association) to encrypted packets.
2. The receiver maintains a 64-packet sliding window, the right edge of which includes the highest sequence number received.
3. The receiver evaluates the received packet’s sequence number:
  - If a received packet’s sequence number falls within the window and was not received previously, the packet is accepted and marked as received.
  - If the received packet’s sequence number falls within the window and previously was received, the packet is dropped and the replay error counter is incremented.
  - If the received packet’s sequence number is greater than the highest sequence in the window, the packet is accepted and marked as received, and the sliding window is moved “to the right.”
  - If the received packet’s sequence number is less than the lowest sequence in the window, the packet is dropped and the replay error counter is incremented.

While Anti-Replay is useful in validating message integrity, in a converged IPSec VPN implementation with QoS enabled, lower-priority packets are often delayed so that higher-priority packets receive preferential treatment, which has the unfortunate side effect of sufficiently reordering packets so they are out of sequence from an IPSec Anti-Replay perspective. Therefore, there is a concern that through the normal QoS prioritization process, the receiver might drop packets as Anti-Replay errors, when, in fact, they are legitimately sent or received packets.

Traffic assigned to CBWFQ classes is much more sensitive to Anti-Replay than traffic assigned to a LLQ. This is because LLQ traffic is always sent in order, with strict priority; but CBWFQ traffic may be delayed by other CBWFQ flows and be sent in gaps exceeding the receiver's 64-packet sliding Anti-Replay window. Furthermore, by default, each CBWFQ class receives a queue with a length of 64 packets. Meanwhile, the receiving IPSec peer has a single 64-packet Anti-Replay window (per IPSec Security Association) with which to process packets from **all** LLQ and CBWFQ bandwidth classes. Therefore, a mismatch is created between the queue depths on the sender's output interface (multiple queues of 64 packets each) as compared to the width of the receiver's Anti-Replay window (a single sliding window of 64 packets per SA). As more bandwidth classes are defined in the sender's policy map, this mismatch increases. This is an inefficient use of expensive WAN/VPN bandwidth, as many packets are transmitted only to be dropped before decryption.

Cisco IOS allows the Anti-Replay window to be expanded (up to a maximum value of 1024 packets) or, alternatively, to be disabled entirely.

During testing it was observed that when TelePresence was provisioned within a dual-LLQ design, with a default-sized Anti-Replay window (of 64 packets), Anti-Replay errors did not affect either TelePresence or voice flows; however, there were significant Anti-Replay errors occurring on CBWFQ classes, inline with the behavior described above. These errors were reduced as the Anti-Replay window was enlarged, to the maximum of 1024 packets, yet were only eliminated altogether when Anti-Replay was disabled.

Additionally, when TelePresence was provisioned with a CBWFQ, with a default-sized Anti-Replay window, significant replay errors occurred on TelePresence flows, resulting in unusable call-quality. Replay errors were still noticed even when the Anti-Replay sliding window was set to the maximum of 1024 packets and were only eliminated when the Anti-Replay feature was disabled.

**Therefore, when encrypting TelePresence over the private WAN and/or MPLS VPN, it is recommended to assign TelePresence to a LLQ and/or to disable Anti-Replay.**

## TelePresence Branch IPSec VPN Edge QoS Design

As discussed in the previous section, it is not recommended to deploy TelePresence over IPSec VPNs over the Internet, due to the lack of service level guarantees of the Internet in general. But rather, if required due to business reasons, IPSec encryption via the network infrastructure provides an additional security overlay to private WANs or MPLS VPNs for TelePresence calls.

If TelePresence is to be deployed over IPSec VPNs over private WANs or MPLS VPNs, then three additional points should be kept in mind:

- Provision the additional bandwidth required by encryption, according to [Table 6-2](#).
- Provision TelePresence traffic into a LLQ (or a dual-LLQ, along with voice).
- Either maximize or disable Anti-Replay.

The first bullet is straightforward and the second bullet has already been covered (see [Example 6-5](#) and [Example 6-9](#)). The third bullet, however, requires some new commands that we have not yet detailed.

To minimize Anti-Replay errors or to eliminate them completely, Cisco IOS introduced a pair of commands in 12.3(14)T that could either enlarge the Anti-Replay window (to a maximum of 1024 packets) or disable it entirely, the **set security-association replay window-size** and the **set security-association replay disable** commands, respectively.

Let us consider two examples to illustrate these options. In [Example 6-12](#), a dual-LLQ QoS policy (with modified priority bandwidth for the TelePresence class) is applied in conjunction with a native IPSec tunnel (including a maximized Anti-Replay window) to a branch T3 WAN/VPN edge interface.

**Example 6-12 Dual-LLQ Design with Native IPsec Tunnel and Maximized Anti-Replay Window**

```

!
policy-map WAN-EDGE-IPSEC
  class VOIP
    priority percent 10          ! LLQ for VoIP (example amount of BW)
  class TELEPRESENCE
    priority 15360 256000      ! LLQ for CTS-3000 with IPsec (1080p-Best + aux video)
  class DATA
    ...
!
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key CTS address 192.168.2.10
crypto isakmp keepalive 10
!
!
crypto ipsec transform-set CTS-IPSEC esp-3des esp-md5-hmac
!
crypto map CMAP local-address Serial6/0
crypto map CMAP 10 ipsec-isakmp
  set peer 192.168.2.10
  set security-association replay window-size 1024! Maximizes A/R
  set transform-set CTS-IPSEC
  match address BRANCH-TO-CAMPUS
  qos pre-classify
!
!
interface Serial6/0
  description BRANCH-TO-CAMPUS-T3
  ip address 192.168.2.9 255.255.255.252
  tx-ring-limit 10          ! Tunes T3 Tx-Ring
  dsu bandwidth 44210
  framing c-bit
  cablelength 10
  serial restart-delay 0
  crypto map CMAP
  max-reserved-bandwidth 100      ! LLQ/CBWFQ BW Override
  service-policy output WAN-EDGE-IPSEC  ! Attaches Dual-LLQ Policy
!
!
ip access-list extended BRANCH-TO-CAMPUS
  permit ip 10.16.0.0 0.0.255.255 10.17.0.0 0.0.255.255
!

```

In [Example 6-13](#), a dual-LLQ QoS policy (with modified priority bandwidth for the TelePresence class) is applied in conjunction with a GRE IPsec tunnel (with a disabled Anti-Replay window) to a branch T3 WAN/VPN edge interface.

**Example 6-13 Dual-LLQ Design with GRE IPsec Tunnel and Disabled Anti-Replay Window**

```

!
policy-map WAN-EDGE-IPSEC
  class VOIP
    priority percent 10          ! LLQ for VoIP (example amount of BW)
  class TELEPRESENCE
    priority 15360 256000      ! LLQ for CTS-3000 with IPsec
  class DATA
    ...
!

```

```

crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key telep address 192.168.2.10
crypto isakmp keepalive 10
!
!
crypto ipsec transform-set CTS-IPSEC esp-3des esp-md5-hmac
!
crypto map CMAP local-address Serial6/0
crypto map CMAP 10 ipsec-isakmp
set peer 192.168.2.10
set security-association replay disable ! Disables Anti-Replay
set transform-set CTS-IPSEC
match address BRANCH-TO-CAMPUS
qos pre-classify
!
!
interface Tunnel0
  ip address 10.18.1.1 255.255.255.252
  tunnel source 192.168.2.9
  tunnel destination 192.168.2.10
!
interface Serial6/0
  description BRANCH-TO-CAMPUS-T3
  ip address 192.168.2.9 255.255.255.252
  tx-ring-limit 10 ! Tunes T3 Tx-Ring
  dsu bandwidth 44210
  framing c-bit
  cablelength 10
  serial restart-delay 0
  crypto map CMAP
  max-reserved-bandwidth 100 ! LLQ/CBWFQ BW Override
  service-policy output WAN-EDGE-IPSEC ! Attaches Dual-LLQ Policy
!
!
ip access-list extended BRANCH-TO-CAMPUS
  permit gre host 192.168.2.9 host 192.168.2.10
!

```

These configurations can be verified with the following command:

- **show policy-map interface**
- **show crypto engine accelerator statistic** *module*

## TelePresence Branch MPLS VPN

MPLS VPN architectures are comprised of customer edge (CE) routers, provider-edge (PE) routers, and provider (P) routers. MPLS VPNs provide fully-meshed Layer 3 virtual WAN services to all interconnected CE routers. Let us discuss some of the critical QoS design considerations pertaining to MPLS VPNs and then translate these considerations into configuration examples.



### Note

MPLS VPN architectures are defined in RFC 2547 “BGP/MPLS VPNs” at [www.ietf.org/rfc/rfc2547](http://www.ietf.org/rfc/rfc2547).

## TelePresence Branch MPLS VPN Edge Considerations

The advent of MPLS VPN service offerings that inherently offer full-mesh connectivity has shifted the QoS administration paradigm. Under traditional hub-and-spoke Layer 2 WAN designs, the enterprise network administrator controlled all the QoS policies by configuring these on the WAN aggregator routers' and branch routers' WAN edges, as previously discussed. However, under a full-mesh topology, it is the service provider's QoS policies on the PE edges routers that ultimately determine how traffic enters a branch and these SP policies may be different from the enterprise's policies on the (unmanaged) CE edges.

Therefore, to ensure end-to-end service levels, enterprise administrators must choose service providers that offer compatible policies to meet their business objectives; furthermore, enterprises must fully understand the SP's QoS policies and map their policies to match in a complementary manner.

First, let us briefly discuss service provider selection based on SLA requirements. As brought out in [Chapter 4, "Quality of Service Design for TelePresence,"](#) the bandwidth and service level requirements of TelePresence (including latency, jitter, and loss requirements) are very high—some are even higher than the SLAs of VoIP. Therefore, to achieve these tight end-to-end SLAs, it is mandatory that the SP be able to guarantee a subset of these SLAs from PE-edge-to-PE-edge.

In the past, to facilitate VoIP deployments over MPLS VPNs, Cisco initiated a Cisco Powered Network (CPN) "IP Multiservice Service Provider" designation that required SPs to offer (independently-verified) PE-to-PE SLA guarantees that would enable enterprise customers to fulfill the end-to-end SLA requirements of VoIP. This initiative was well received by both enterprise customers and SPs, as enterprise customers did not have to do as much research and testing to validate potential SP networks to support their VoIP deployments and SPs, in turn, received a competitive advantage in marketing their networks that could meet VoIP SLAs.

Subsequently, this initiative has similarly been applied to TelePresence. Cisco is in the process of validating various service providers that can meet a subset of the stringent SLAs required by TelePresence, so that enterprise customers can provide end-to-end SLA guarantees for TelePresence.

Second, let us turn our attention to enterprise-to-service provider mapping, which usually involves three main points to consider:

1. The number of enterprise traffic classes versus the number of service provider traffic classes; and if collapsing is required, how to perform this efficiently.
2. Marking or remarking requirements on CE egress to gain admission to the desired SP traffic class; and (optional) remarking requirements on CE ingress to restore enterprise traffic markings for provisioning, accounting, or management purposes.
3. Non-traditional WAN access media, such as sub-line-rate Ethernet access, and the QoS implications these pose.

Let us discuss each of these in turn, beginning with the number of traffic classes.

The number of traffic classes within an enterprise network is a function of its business objectives (as discussed in detail in Chapter 1 of the QoS SRND at [www.cisco.com/go/srnd](http://www.cisco.com/go/srnd)). As an informational guide, RFC 4594 "Configuration Guidelines for DiffServ Service Classes" ([www.ietf.org/rfc/rfc4594](http://www.ietf.org/rfc/rfc4594)), outlines up to 12 classes of traffic that may be present within an enterprise. This is not to say that it is mandatory for enterprises to have 12 traffic classes today; but rather that the potential exists in enterprise networks for up to 12 traffic classes and—given the trends in emerging new applications and evolving business objectives—even if enterprises are not deploying 12 class models today, they may need to in the near future. For configuration and testing purposes, we can use this 12-class enterprise model as a worst-case scenario, providing maximum disparity between the number of enterprise traffic classes versus the number of service provider traffic classes. The Cisco-modified 12-class RFC 4594 enterprise model is shown in [Figure 6-9](#).

**Figure 6-9 Cisco-Adapted 12-Class RFC 4594-based Enterprise Classification and Marking Model**

Application	L3 Classification		IETF
	PHB	DSCP	RFC
Network Control	CS6	48	RFC 2474
VoIP Telephony	EF	46	RFC 3246
Broadcast Video	CS5	40	RFC 2474
Multimedia Conferencing	AF41	34	RFC 2597
Real-Time Interactive/TelePresence	CS4	32	RFC 2474
Multimedia Streaming	AF31	26	RFC 2597
Call Signaling	CS3	24	RFC 2474
Low-Latency/Transactional Data	AF21	18	RFC 2597
Operations/Administration/Management	CS2	16	RFC 2474
High-Throughput/Bulk Data	AF11	10	RFC 2597
Best Effort	DF	0	RFC 2474
Low-Priority/Scavenger Data	CS1	8	RFC 3662

223250

**Note**

Some of the application class names show both the RFC 4594 names as well as the better known, but less wordy, Cisco QoS Baseline application class names. For example, “Low Latency Data,” “High Throughput Data,” and “Low Priority Data” are generally more easily referred to as “Transactional Data,” “Bulk Data,” and “Scavenger,” respectively. Nonetheless, the names can be viewed as synonymous.

Now let us look at service provider class models. At the time of writing, in North America, most service providers offer 3- or 4-class QoS models, although some are planning 6-class models. In EMEA or Asia Pacific, some providers offer even more classes. Rather than presenting models for each number of SP classes, we consider just two, a 4-class and a 6-class model, and the principles applied to these enterprise-to-SP mapping examples can be extended to other traffic class models. These 4-class and 6-class examples are graphically illustrated in [Figure 6-10](#).

Figure 6-10 Example 4-Class and 6-Class MPLS VPN SP QoS Models

4-Class SP Model		6-Class SP Model	
EF CS5	SP-Real-Time (RTP/UDP) 30%	EF CS5	SP-Real-Time (RTP/UDP) 20%
CS6 AF3 CS3	SP-Critical 1 (TCP) 20%	CS4	SP-Critical 1 (TelePresence) 10%
AF2 CS2	SP-Critical 2 (UDP) 20%	CS6 AF3 CS3	SP-Critical 2 (TCP) 20%
DF	SP-Best Effort 30%	AF1 CS1	SP-Scavenger 5%
		DF	SP-Best Effort 25%

223098

Comparing Figure 6-7 to Figure 6-8 highlights the fact that, more often than not, there are generally fewer SP traffic classes than enterprise classes, and thus there are times when more than one enterprise traffic class is assigned to the same SP class. When such collapsing has to be done, it is recommended to **avoid mixing TCP-based applications with UDP-based applications within a single service provider class**. This is due to the behavior of these respective protocols during periods of congestion.

Specifically, due to TCP transmission guarantees and its windowing behavior, TCP transmitters throttle back flows when drops are detected. In contrast, most UDP transmitters are completely oblivious to drops and, therefore, never lower transmission rates because of dropping. When TCP flows are combined with UDP flows within a single service provider class and the class experiences congestion, TCP flows continually lower their transmission rates, potentially giving up their bandwidth to UDP flows that are oblivious to drops. This effect is called TCP starvation/UDP dominance. Even if WRED is enabled on the service provider class, the same behavior would be observed because WRED (for the most part) manages congestion only on TCP-based flows. Granted, it is not always possible to separate TCP-based flows from UDP-based flows, but it is beneficial to be aware of this behavior when making such application-mixing decisions within a single service provider class.

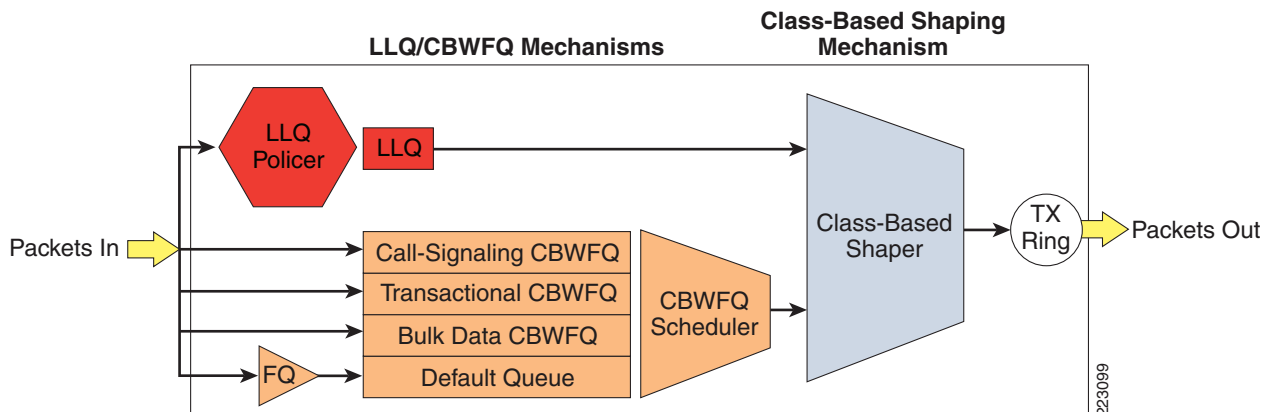
Now let us look at traffic marking and remarking requirements. As can be seen in Figure 6-8, DSCP values serve as the admission criteria per SP class. These DSCP values likely vary from one provider to another, therefore it is important for the enterprise subscriber be fully informed of the DSCP admission criteria for each SP class. At times applications may need to be remarked in order to gain admission to the desired SP class. When such is the case, **marking should be done as the final operations on the (unmanaged) CE egress edge**. Otherwise, if remarking is done at an earlier node, say the campus access edge, then changes to the SP QoS policies or migration to another SP would be much more difficult to manage, as would using multiple SPs for redundancy (each with its own marking scheme).

Also, there may be times when the enterprise has a business requirement to maintain DSCP markings in the branch, perhaps for traffic accounting purposes or for other reasons. In such cases, the enterprise subscriber may choose to make the MPLS VPN appear DSCP-transparent by **restoring enterprise DSCP markings on the CE ingress edge**.

Additionally, each SP class is likely policed on the PE ingress edge. Excess traffic may either be remarked or dropped. Again, it is important for the enterprise subscriber to know exactly how excess traffic is treated on a per-class basis. Understanding SP policing policies is an especially important consideration for the TelePresence class. As we have already discussed in [TelePresence Branch WAN Edge LLQ Policy](#), TelePresence requires 256 KB of committed burst from a policer. **Therefore, it is essential to confirm with the service provider that whatever class TelePresence traffic is assigned to is being policed with at least 256 KB of burst.**

And finally, let us discuss the QoS implications of non-traditional WAN access-media, such as Ethernet. As previously discussed, queuing policies only engage when the physical interface is congested (as is indicated to IOS software by a full Tx-Ring). This means that queuing policies never engage on media that has a contracted sub-line rate of access, whether this media is Frame Relay, ATM, or Ethernet. In such a scenario, **queuing can only be achieved at a sub-line rate by introducing a two-part policy**, sometimes referred to a **Hierarchical QoS (HQoS) policy** or nested QoS policy, **wherein 1) traffic is shaped to the sub-line rate, and 2) traffic is queued according to the LLQ/CBWFQ policies within the sub-line rate.** With such an HQoS policy, it is not the Tx-Ring that signals IOS software to engage LLQ/CBWFQ policies, but rather it is the Class-Based Shaper that triggers software queuing when the shaped rate has been reached. Such an HQoS policy is graphically illustrated in [Figure 6-11](#).

**Figure 6-11 Hierarchical QoS Policy—Shaping to a Sub-Line Rate with Queuing within the Shaped Rate**

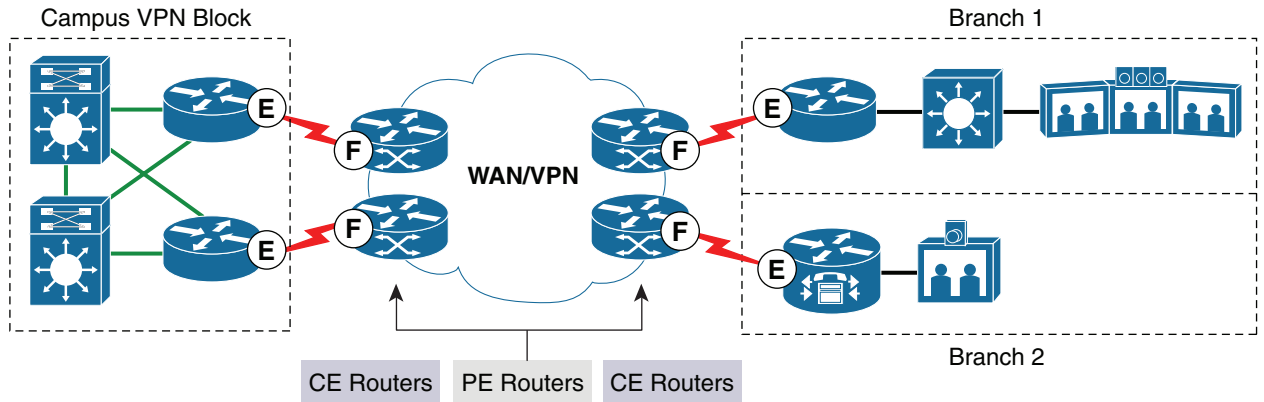


Let us consider a practical example in which an SP offers an enterprise subscriber a GigabitEthernet handoff, but with a (sub-line rate) contract for only 50 Mbps. Normally, queuing policies only engage on this GE interface when the offered traffic rate exceeds 1000 Mbps. However, the enterprise administrator wants to ensure that traffic within the 50 Mbps contracted rate is properly prioritized prior to PE handoff. Therefore, they configure an HQoS policy, such that the interface shapes all traffic to the contracted 50 Mbps rate and attaches a nested queuing policy to the shaping policy, such that traffic is properly prioritized within this 50 Mbps sub-line rate.

The only other consideration an administrator should keep in mind with HQoS policies is their potential performance impact. When performed in IOS software on routers, then these policies generate a marginal CPU load; the actual amount of the load depends on platforms, speeds, policy complexity, traffic rates, and other factors. A rule of thumb, however, is to always keep CPU levels below 75% during normal operating conditions, as this allows some cycles to always be available to process network events. Some platform guidance for HQoS policies are presented, along with detailed configurations, in [TelePresence Branch MPLS VPN QoS Designs](#).

Finally, let us take a look at how all these QoS policies fit together for a TelePresence-enabled branch subscribing to a MPLS VPN, as illustrated in [Figure 6-12](#).

**Figure 6-12 Enterprise and Service Provider MPLS VPN QoS Design Recommendations for TelePresence**



**Enterprise Subscriber (Unmanaged CE Routers):**

<p><b>(E) Outbound Policies:</b>                  HQoS Shaper (if required)                  ≤33% of BW {                  + LLQ for VoIP (EF)                  + LLQ or CBWFQ for TelePresence (CS4)                  + Remark TelePresence (if necessary)                  + CBWFQ for Call-Signaling (CS3)                  + Remark Call-Signaling (if necessary)</p>	<p><b>Inbound Policies:</b>                  Trust for DSCP                  + Restore TelePresence to CS4 (if necessary)                  + Restore TelePresence to CS4 (if necessary)</p>
---	---

**Service Provider:**

<p><b>(F) Outbound Policies:</b>                  + LLQ for Real-Time                  + CBWFQ for Critical Data</p>	<p><b>Inbound Policies:</b>                  Trust for DSCP                  Police on a per-Class Basis</p>
--	--

223100

As shown in Figure 6-12, the enterprise subscriber provisions LLQ/CBWFQ policies for VoIP and TelePresence (in conjunction with HQoS sub-line rate shapers, if required) and performs any application-class remarking on the CE egress edges. Optionally, if required, the enterprise may restore their markings on the CE ingress edges for any traffic that required remarking over the MPLS VPN.

In turn, the service provider polices traffic on a per-class basis on their PE ingress edges and provisions LLQ/CBWFQ policies according to their class-models on the PE egress edges. They may also perform QoS and/or MPLS Traffic Engineering within their core; however, such policies are beyond the scope of our enterprise-centric designs.



**Note**

Due of the explicit ingress policing on PE edges of MPLS VPNs, it cannot be overemphasized that the enterprise subscriber needs a comprehensive Call Admission Control system in place to limit the amount of TelePresence traffic over the MPLS VPN; otherwise the call-quality of **all** TelePresence calls over the MPLS VPN may degrade to the point of unusability.

## TelePresence Branch MPLS VPN QoS Designs

Having reviewed the many design considerations for MPLS VPNs, let us now put them into practice by constructing configuration policies to meet the requirements of several specific scenario examples, including a 4-class SP model, a 6-class SP model, and a sub-line rate access example.

### TelePresence 4-Class MPLS VPN SP Model QoS Design

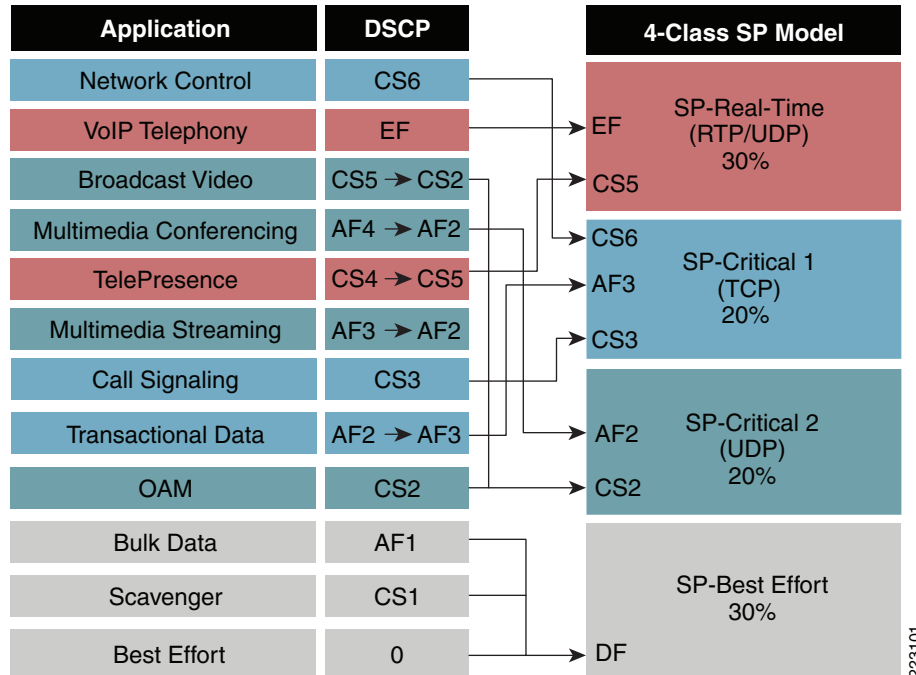
Let us begin by constructing a CE edge policy to support a 4-class SP model example. In this example, since there are so few classes to choose from, TelePresence may need to be combined with another application. **It is highly recommended not to combine TelePresence with any unbounded application** (i.e., an application without any admission control) **within a single SP class**, since this could lead to class congestion, resulting in TelePresence drops (with or without WRED enabled on the SP class), which would ruin TelePresence call quality. Therefore, in such a design two choices exist:

- Assign TelePresence into the SP-Realtime class along with voice.
- Assign TelePresence to a dedicated non-priority SP class.

We consider the option of assigning TelePresence into the SP-Realtime class for this example and then consider the option of assigning it to a non-priority class in the following (6-class SP model) example.

Given the 4-Class SP model illustrated in [Figure 6-10](#), we have a Realtime class, a default Best Effort class, and two additional non-priority traffic classes. In this case, the enterprise administrator may elect to separate TCP-based applications from UDP-based applications by using these two non-priority SP traffic classes. Specifically, if voice and TelePresence are the only applications to be assigned to the SP Realtime class, then Broadcast Video, Multimedia Conferencing, Multimedia Streaming, and Operations/Administration/Management (OAM) traffic (which is largely UDP-based) can all be assigned to the UDP SP-class (SP-Critical 2). This leaves the other non-priority SP class (SP-Critical 1) available for control plane applications, such as Network Control and Call-Signaling, along with TCP-based Transactional Data applications. [Figure 6-13](#) shows the per-class remarking requirements from the CE edge to gain access to the classes within the 4-class SP model, with TelePresence assigned to the SP-Realtime class, along with voice.

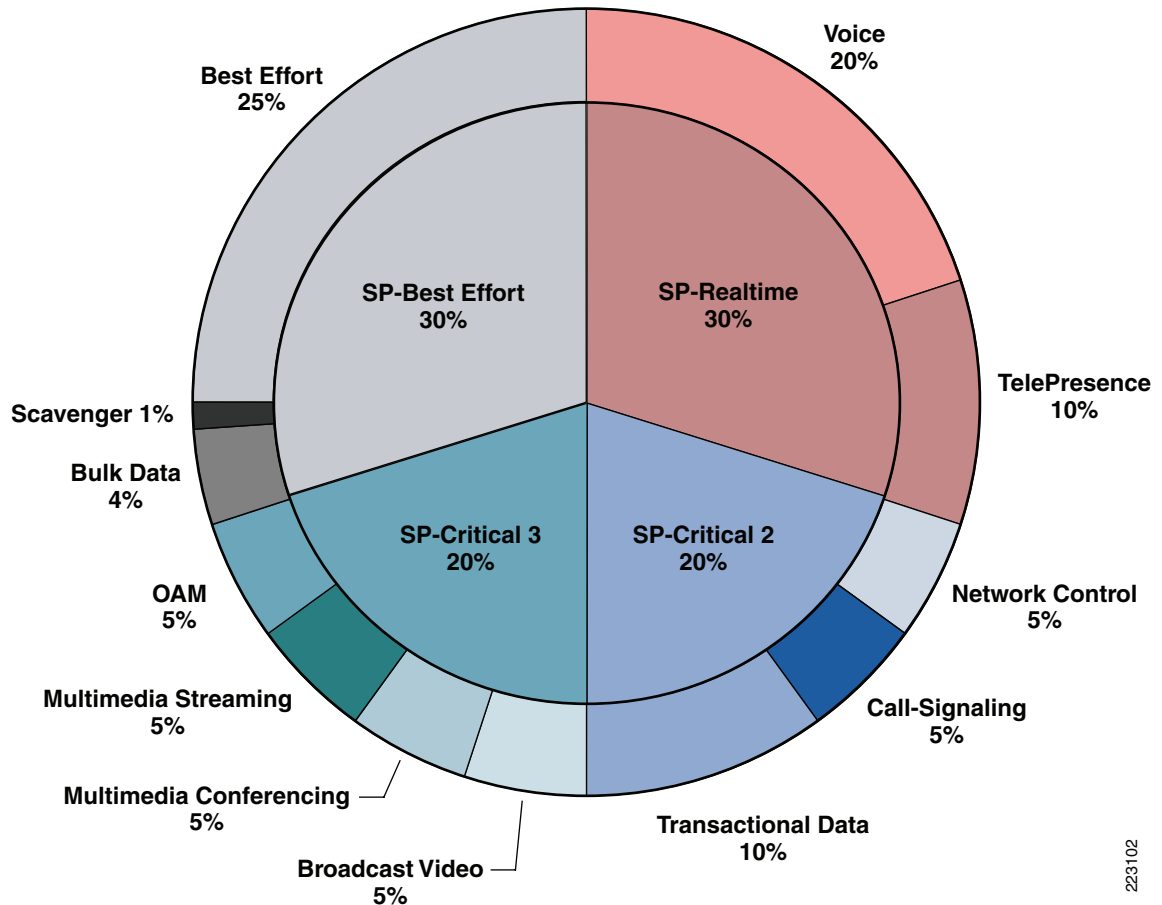
**Figure 6-13 Enterprise-to-SP Mapping—4-Class SP Model Example with TelePresence Assigned to the Realtime Class Along with Voice**



As shown in [Figure 6-13](#), in this example TelePresence traffic must be remarked on the CE egress edge to CS5 to gain access to the SP's Realtime class. Also, Broadcast Video must be remarked to CS2 to assign it to the UDP SP class (SP-Critical 2). Similarly, Multimedia Conferencing and Multimedia Streaming must be remarked to AF2 to assign these also to the UDP SP class. Correspondingly, Transactional Data traffic must be remarked to AF3 to gain access into the TCP SP class (SP-Critical 1). All other traffic does not require remarking to gain admission to the desired classes; this includes Bulk and Scavenger, as these default to the SP-Best Effort class without any explicit remarking.

Additionally, the relative per-class bandwidth allocations need to be aligned, such that the enterprise CE edge queuing policies are consistent with the SP's PE edge queuing policies to ensure compatible Per-Hop Behaviors (PHBs). Compatible bandwidth allocations are illustrated in [Figure 6-14](#), where the inner pie-chart represents the SP's per-class bandwidth allocations and the outer pie-chart represents the enterprise's per-class bandwidth allocations over an OC3 link.

**Figure 6-14** Enterprise-to-SP Bandwidth Allocation C4-Class SP Model Example with TelePresence Assigned to the Realtime Class Along with Voice



223102

The CE egress edge configuration for this policy is shown in [Example 6-14](#).

**Example 6-14** Enterprise-to-SP Mapping—4-Class SP Model Example with TelePresence Assigned to the Realtime Class Along with Voice

```

policy-map CE-EDGE-4CLASS-OC3-POS
  class VOICE
    police cir 31000000           ! Voice is policed to 31 Mbps (20%)
    bc 15500                     ! Bc is 15.5 KB
    conform-action transmit      ! Conforming action --> transmit
    exceed-action drop          ! Single-Rate Policing action
    priority                     ! LLQ command for OC3-POS
  class TELEPRESENCE
    police cir 15000000          ! TP is policed to 15 Mbps
    bc 256000                   ! Bc is 256 KB
    conform-action transmit      ! Conforming action --> transmit
    exceed-action drop          ! Single-Rate Policing action
    priority                     ! LLQ command for OC3-POS
    set dscp cs5                 ! Remark TelePresence to CS5
  class NETWORK-CONTROL
    bandwidth percent 5         ! CBWFQ for Routing
  class CALL-SIGNALING
    bandwidth percent 5         ! CBWFQ for Call-Signaling
  class TRANSACTIONAL-DATA
    bandwidth percent 10        ! CBWFQ for Transactional Data

```

```

    random-detect dscp-based                ! DSCP-WRED for Transactional Data
    set dscp af31                            ! Remark Transactional Data to AF31
class BROADCAST-VIDEO
    bandwidth percent 5                     ! CBWFQ for Broadcast Video
    set dscp cs2                             ! Remark Broadcast Video to CS2
class MULTIMEDIA-CONFERENCING
    bandwidth percent 5                     ! CBWFQ Video-Conferencing
    random-detect dscp-based                ! DSCP-WRED for Video-Conferencing
    set dscp af21                           ! Remark Video-Conferencing to AF21
class MULTIMEDIA-STREAMING
    bandwidth percent 5                     ! CBWFQ for Streaming-Video
    random-detect dscp-based                ! DSCP-WRED for Streaming-Video
    set dscp af21                           ! Remark Streaming-Video to AF21
class OAM
    bandwidth percent 5                     ! CBWFQ for Network Management
class BULK-DATA
    bandwidth percent 4                     ! CBWFQ for Bulk Data
    random-detect dscp-based                ! DSCP-WRED for Bulk Data
class SCAVENGER
    bandwidth percent 1                     ! Minimum CBWFQ for Scavenger
class class-default
    bandwidth percent 25                    ! CBWFQ for Best Effort
    random-detect                           ! WRED for Best Effort
!
...
...
interface POS3/0/1
    description BRANCH-CE-EDGE-OC3-POS
    ip address 192.168.5.1 255.255.255.252
    clock source internal
    service-policy output CE-EDGE-4CLASS-OC3-POS ! Attaches policy
!
```

Optionally, the original markings for TelePresence, Transactional Data, Broadcast Video, Multimedia Conferencing, and Multimedia Signaling can be restored on the CE ingress edges to make the MPLS VPN appear completely DSCP-transparent to the enterprise, despite the remarking requirements of the service provider. The TelePresence and Transactional Data remarking policies are 1:1 DSCP mappings (one DSCP is changed to another DSCP) and as such are easy to undo with a reversing 1:1 mapping operation. However, the remarking operations performed on Broadcast Video, Multimedia Conferencing, and Multimedia Signaling are 2:1 mappings (two DSCP values are changed to a single DSCP value); specifically Broadcast Video and OAM now share DSCP CS2 and Multimedia Conferencing and Multimedia Signaling share DSCP AF21. These 2:1 DSCP mappings require additional classification policies to identify the discrete applications now sharing a single codepoint. These additional classification policies can include NBAR or access-lists.

In [Example 6-15](#), TelePresence and Transactional Data are restored to their original enterprise-marked DSCP values via a simple 1:1 reverse-mapping. Broadcast Video is separated from OAM by referencing an ACL that identifies the source IP address of the Broadcast Video servers. Multimedia Streaming, in this example, consists of streaming RealAudio and VDO Live stateful protocols, both of which can be identified via NBAR and thus sifted apart from Multimedia Conferencing. An optional DSCP restoration policy for the CE ingress edge is shown in [Example 6-15](#).

#### **Example 6-15 Optional Enterprise DSCP Marking Restoration Policies for CE Ingress Edges**

```

class-map match-all SP-TELEPRESENCE
    match dscp cs5                            ! Remark value for TelePresence
class-map match-all SP-TRANSACTIONAL-DATA
    match dscp af31 af32 af33                ! Remark value(s) for Trans-Data
class-map match-all SP-BROADCAST-VIDEO
```

```

    match dscp cs2                                ! Shared DSCP for Bdcst Video + OAM
    match access-group name BROADCAST-VIDEO-SERVERS! References ACL
class-map match-all SP-MULTIMEDIA-STREAMING-REALAUDIO
    match dscp af21 af22 af32                    ! Shared DSCP for MM-Stream + Conf
    match protocol realaudio                    ! NBAR PDLM for RealAudio
class-map match-all SP-MULTIMEDIA-STREAMING-VDOLIVE
    match dscp af21 af22 af32                    ! Shared DSCP for MM-Stream + Conf
    match protocol vdolive                      ! NBAR PDLM for VDOLive
class-map match-all SP-MULTIMEDIA-CONFERENCING
    match dscp af21 af22 af32                    ! All other AF2 is MM-Conf only
!
policy-map CE-EDGE-IN
  class SP-TELEPRESENCE
    set dscp cs4                                ! Restores original marking for TP
  class SP-TRANSACTIONAL-DATA
    set dscp af21                                ! Restores original marking for TD
  class SP-BROADCAST-VIDEO
    set dscp cs5                                ! Restores original marking for BV
  class SP-MULTIMEDIA-STREAMING-REALAUDIO
    set dscp af31                                ! Restores original marking for MMS
  class SP-MULTIMEDIA-STREAMING-VDOLIVE
    set dscp af31                                ! Restores original marking for MMS
  class SP-MULTIMEDIA-CONFERENCING
    set dscp af41                                ! Restores original marking for MMC
!

interface POS3/0/1
  description BRANCH-CE-EDGE-OC3-POS
  ip address 192.168.5.1 255.255.255.252
  service-policy output CE-EDGE-OC3-POS        ! Attaches egress policy
  service-policy input CE-EDGE-IN              ! Attaches ingress policy
!
...
!
ip access-list extended BROADCAST-VIDEO-SERVERS! Reference ACL
permit ip any 10.200.200.0 0.0.0.255          ! Broadcast Video Server Subnet
!

```

These configurations can be verified with the following command:

- **show policy-map interface**

We can note a few important policy elements in [Example 6-15](#):

- It is important to give the remarked traffic classes unique names from the original enterprise-marked traffic classes, otherwise these interfere or overwrite each other. For example class “TELEPRESENCE” matches on the enterprise marking value for TelePresence (CS4), but class “SP-TELEPRESENCE” matches on the remarked value for TelePresence (CS5).
- Because of the default **match-all** operand on class-maps, we have to have two separate class maps to sift, [traffic marked AF2 and using the realaudio protocol] or [traffic marked Af2 and using the vdolive protocol]. If a single class-map was used, then the **match-all** operand would preclude any match (because the protocol in use cannot be both realaudio and vdolive at the same time; these protocols are unique and represent mutually exclusive criterion). Furthermore, a **match-any** operand on a combined class-map for AF2 **or** realaudio **or** vdolive would not work either, because this would fail the logical requirement that the traffic must be [marked AF2 **and** realaudio] or [marked AF2 **and** vdolive], and thus would result in false-positive matches on Multimedia Conferencing traffic marked AF2.
- It is important to keep in mind that classification logic, like ACL logic, is based on the first-true-match rule. Therefore, the order of classification and sifting must be given careful consideration in order to ensure that traffic marking gets restored properly.

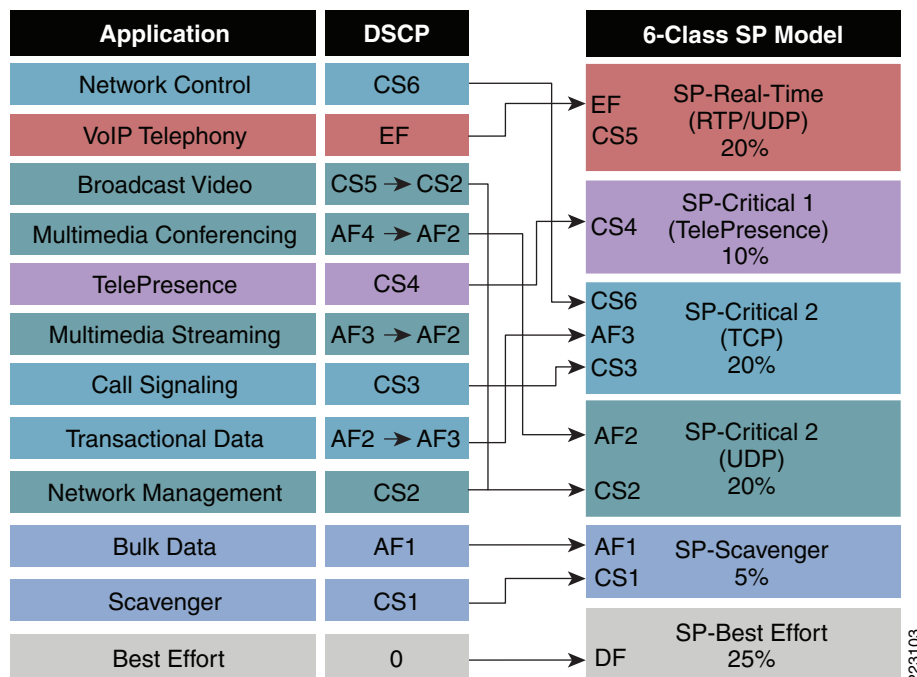
## TelePresence 6-Class MPLS VPN SP Model QoS Design

Now let us turn our attention to the 6-Class SP model, also illustrated in [Figure 6-10](#). In this model, we have a Realtime class, a default Best Effort class, a “less-than Best Effort” Scavenger class, and three additional non-priority traffic classes. Furthermore, to illustrate more design options, we assign TelePresence to a non-priority SP-class in this example; but of course TelePresence can also be assigned, in combination with voice, to the SP-Realtime class, as has already been detailed in the previous section.

In this case, the enterprise administrator can dedicate one of the non-priority classes (such as SP-Critical 1) for TelePresence. Again, it bears reiteration that it would not be recommended to assign TelePresence in conjunction with any unbounded application into a single SP class, as the other application could potentially cause the combined class to congest, resulting in TelePresence drops and loss of call-quality.

This leaves two additional non-priority classes, which again allows the administrator to separate TCP-based applications from UDP-based applications. Specifically, Broadcast Video, Multimedia Conferencing, Multimedia Streaming, and Operations/Administration/Management (OAM) traffic can all be assigned to the UDP SP-class (SP-Critical 3). This leaves the other non-priority SP class (SP-Critical 2) available for control plane applications, such as Network Control and Call-Signaling, along with TCP-based Transactional Data applications. [Figure 6-15](#) shows the per-class remarking requirements from the CE edge to gain access to the classes within the 6-class SP model, with TelePresence assigned to a non-priority SP class.

**Figure 6-15 Enterprise-to-SP Mapping—6-Class SP Model Example with TelePresence Assigned to a Dedicated, Non-Priority SP-Class**

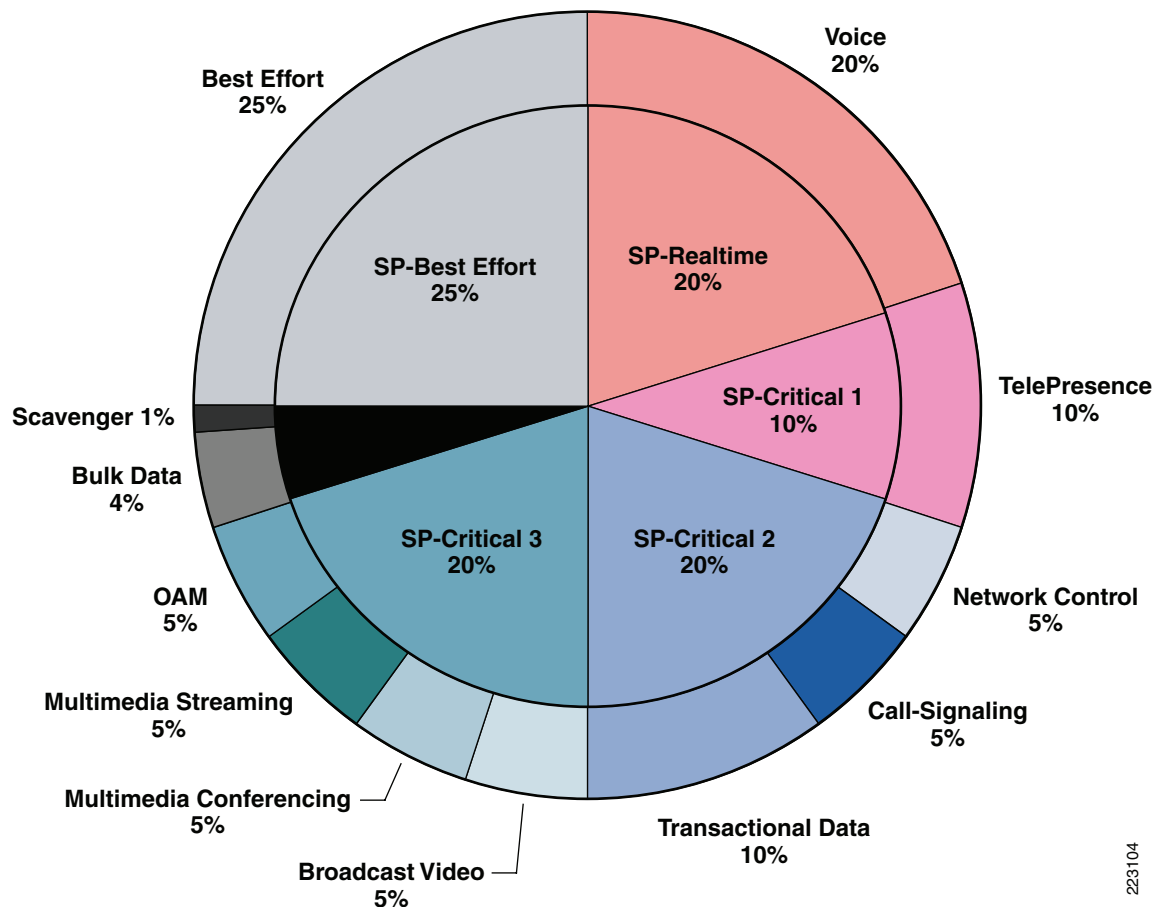


As shown in [Figure 6-15](#), in this second example TelePresence traffic does not need to be remarked to gain access to the dedicated, non-priority SP class to which it is assigned (SP-Critical 1). However as before, Broadcast Video must be remarked to CS2 to assign it to the UDP SP class (SP-Critical 3); Multimedia Conferencing and Multimedia Streaming must be remarked to AF2 to assign these also to the UDP SP class. Correspondingly, Transactional Data traffic must be remarked to AF3 to gain access into the TCP SP class (SP-Critical 2). All other traffic does not require remarking to gain admission to

the desired classes. However, it may be noted that Bulk and Scavenger no longer default to the SP-Best Effort class, but rather now default to the SP-Scavenger class, which is the desired policy to bind these potentially bandwidth-hogging applications.

Additionally, the relative per-class bandwidth allocations again need to be aligned, such that the enterprise CE edge queuing policies are consistent with the SP's PE edge queuing policies. Compatible bandwidth allocations are illustrated in Figure 6-16, where the inner pie-chart represents the SP's per-class bandwidth allocations and the outer pie-chart represents the enterprise's per-class bandwidth allocations over an OC3 link.

**Figure 6-16 Enterprise-to-SP Bandwidth Allocation—6-Class SP Model Example with TelePresence Assigned to a Dedicated, Non-Priority SP-Class**



The CE egress edge configuration for this policy is shown in Example 6-16.

**Example 6-16 Enterprise-to-SP Mapping—6-Class SP Model Example with TelePresence Assigned to a Dedicated, Non-Priority SP-Class**

```

policy-map CE-EDGE-6CLASS-OC3-POS
  class VOICE
    police cir 31000000          ! Voice is policed to 31 Mbps (20%)
    bc 15500                    ! Bc is 15.5 KB
    conform-action transmit      ! Conforming action --> transmit
    exceed-action drop          ! Single-Rate Policing action
    priority                    ! LLQ command for OC3-POS
  class TELEPRESENCE

```

```

    bandwidth percent 10                ! CBWFQ for TelePresence
class NETWORK-CONTROL
    bandwidth percent 5                 ! CBWFQ for Routing
class CALL-SIGNALING
    bandwidth percent 5                 ! CBWFQ for Call-Signaling
class TRANSACTIONAL-DATA
    bandwidth percent 10                ! CBWFQ for Transactional Data
    random-detect dscp-based           ! DSCP-WRED for Transactional Data
    set dscp af31                       ! Remark Transactional Data to AF31
class BROADCAST-VIDEO
    bandwidth percent 5                 ! CBWFQ for Broadcast Video
    set dscp cs2                        ! Remark Broadcast Video to CS2
class MULTIMEDIA-CONFERENCING
    bandwidth percent 5                 ! CBWFQ Video-Conferencing
    random-detect dscp-based           ! DSCP-WRED for Video-Conferencing
    set dscp af21                       ! Remark Video-Conferencing to AF21
class MULTIMEDIA-STREAMING
    bandwidth percent 5                 ! CBWFQ for Streaming-Video
    random-detect dscp-based           ! DSCP-WRED for Streaming-Video
    set dscp af21                       ! Remark Streaming-Video to AF21
class OAM
    bandwidth percent 5                 ! CBWFQ for Network Management
class BULK-DATA
    bandwidth percent 4                 ! CBWFQ for Bulk Data
    random-detect dscp-based           ! DSCP-WRED for Bulk Data
class SCAVENGER
    bandwidth percent 1                 ! Minimum CBWFQ for Scavenger
class class-default
    bandwidth percent 25                ! CBWFQ for Best Effort
    random-detect                       ! WRED for Best Effort
!

```

This configuration can be verified with the following command:

- **show policy-map interface**

Optionally, if the original markings for Transactional Data, Broadcast Video, Multimedia Conferencing, and Multimedia Signaling need to be restored, these can be done in a similar manner as demonstrated in [Example 6-15](#), with the exception of not requiring TelePresence traffic to be restored (as it does not get remarked in this 6-class model example).

## TelePresence Sub-Line Rate Ethernet Access QoS Designs

As previously discussed, to enforce CE edge queuing policies at sub-line rates, an HQoS policy must be used such that a shaper smooths out traffic to the sub-line rate and forces queuing to occur if this rate is exceeded.

As with policers, Cisco IOS shapers operate on a token-bucket principle, achieving sub-line rates by allowing traffic through in specified bursts (Bc) per sub-second intervals (Tc). As shaping introduces delay to packets above the burst value, it is important to properly size the bursts and intervals to minimize potential shaping jitter. For example, as previously discussed, 1080p sends 30 frames of video per second or, phrased differently, a frame's worth of information every 33 ms. However, if the shaping interval is set too low, say to 5 ms, then a frame's worth of information may be delayed over 3-5 shaping intervals (depending on the amount of frame information). Extensive lab testing has shown that configuring a shaping interval of 20 ms has resulted in the most consistent and minimal jitter values to support a CTS-3000 call.

The interval parameter cannot be set directly, but is set indirectly by explicitly configuring the burst parameter. The relationship between the interval, burst, and shaped rate is given as:

$$Tc = Bc / \text{Shaped Rate}$$

Or:

$$Bc = \text{Shaped Rate} * Tc$$

For example, on a FE or GE interface configured to support a sub-line rate of 50 Mbps, a burst value of 1 megabit (50 Mbps \* 20 ms) would result in an optimal shaping interval for TelePresence.



**Note**

For Cisco IOS Shapers, like the Class-Based Shaper, the burst is expressed in bits (not in Bytes, as is the case with policers).

Translating this into an HQoS policy yields the following configuration, as shown in [Example 6-17](#).

**Example 6-17 HQoS Policy to Queue and Shape TelePresence Traffic to a 50 Mbps Sub-Line Rate Over a GigabitEthernet Interface**

```

policy-map CE-EDGE                                ! CE Edge queuing policy
  class VOICE
  ...
  class TELEPRESENCE
  ...                                             ! Either a dual-LLQ or CBWFQ policy for TP
  ...
  !
policy-map HQoS-50MBPS                            ! CE Edge HQoS Shaping policy
  class class-default
    shape average 50000000 1000000              ! Rate=50Mbps; Bc=1Mb, Tc=20ms
    service-policy CE-EDGE                      ! Forces queuing at sub-line rate
  !
  ...
  !
interface GigabitEthernet0/1
  description CE-EDGE-GE
  ip address 192.168.1.50 255.255.255.252
  no ip redirects
  no ip proxy-arp
  duplex auto
  speed auto
  media-type rj45
  negotiation auto
  service-policy output HQoS-50MBPS            ! Attaches HQoS policy to GE int
  !

```

This configuration can be verified with the following command:

- **show policy-map interface**

The final point of consideration for this chapter is the performance impact of HQoS policies on various platforms. As previously discussed, when QoS is performed in hardware, such as on Catalyst switches, then there is no performance impact of QoS policies on the CPU. However, when QoS policies are performed in software, then there is a performance impact that depends on several factors, such as the platform, speed, traffic mix, QoS policy, etc.

At speeds up to 150 Mbps, Cisco IOS routers, like the 3800 series ISR or the Cisco 7200-VXR, may be used to enforce HQoS policies on Ethernet-based access-media. Before we look at the performance impact on these platforms, it bears mentioning that it is not primarily the speed that affects the CPU performance, but rather the Packets-per-second (PPS) rate.


That being said, [Table 6-3](#) shows the performance of a Cisco 3845 router enforcing HQoS policies at rates ranging from 50 Mbps through 150 Mbps.

**Table 6-3 Cisco 3845 Platform Performance of HQoS Policies at 50 Mbps Through 150 Mbps Traffic Rates (26KPPS Through 84 KPPS)**

Bidirectional Target Data Load	Actual Data Load	CPU	PPS	Video Quality
50 Mbps	50 Mbps Out / 48 Mbps In	31%	26,568	Near Perfect
100 Mbps	87 Mbps Out / 88 Mbps In	37%	53,633	Near Perfect
150 Mbps	145 Mbps Out / 147 Mbps In	57%	85,214	Near Perfect
150 Mbps	145 Mbps Out / 140 Mbps In	56%	83,879	Near Perfect

The corresponding performance graph for [Table 6-3](#) is shown in [Figure 6-17](#).

**Figure 6-17 Cisco 3845 Platform Performance of HQoS Policies at 50 Mbps Through 150 Mbps Traffic Rates (26KPPS Through 84 KPPS)**



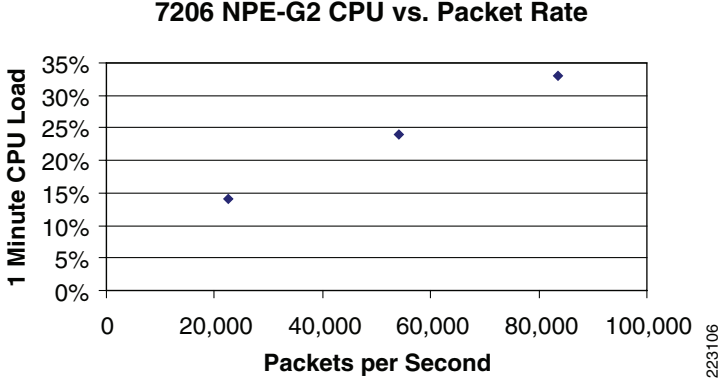
Additionally, [Table 6-4](#) shows the performance of a Cisco 7200VXR router with a Network Processing Engine (NPE) G2 enforcing HQoS policies at rates ranging from 50 Mbps through 150 Mbps.

**Table 6-4 Cisco 7200VXR with NPE-G2 Platform Performance of HQoS Policies at 50 Mbps Through 150 Mbps Traffic Rates (22KPPS Through 84 KPPS)**

Bidirectional Target Data Load	Actual Data Load	CPU	PPS	Video Quality
50 Mbps	50 Mbps Out / 45 Mbps In	14%	22,634	Near Perfect
100 Mbps	96 Mbps Out / 96 Mbps In	24%	54,011	Near Perfect
150 Mbps	142 Mbps Out / 147 Mbps In	33%	83,631	Near Perfect

The corresponding performance graph for [Table 6-4](#) is shown in [Figure 6-18](#).

Figure 6-18 Cisco 7200VXR with NPE-G2 Platform Performance of HQoS Policies at 50 Mbps Through 150 Mbps Traffic Rates (22KPPS Through 84 KPPS)



Both of these platforms are able to support HQoS policies for TelePresence at these speeds (50 Mbps through 150 Mbps). The goal, however, is to keep CPU levels below 75% during normal conditions, so that the router always has some cycles available to process network events.

For higher speeds, HQoS policies should be performed in hardware (such as on the Catalyst 3750-Metro switch with Enhanced Services modules) or on a hybrid hardware/software platform like the Cisco 7600 SIP/SPA combination.