



CHAPTER 12

Cisco TelePresence Multipoint Solution Circuit and Platform Recommendations

Recommendations for Multipoint over WAN Circuits

This section discusses the recommendations for support of multipoint and multiple point-to-point TelePresence over private WAN circuits. Recommendations are made with respect to supporting TelePresence in a dedicated circuit configuration and a converged circuit configuration. Dedicated circuit configurations are those in which TelePresence is the only traffic which exists on the private WAN circuit. In other words, the customer has deployed an overlay network designed exclusively for the support of TelePresence. Converged circuit configurations are those in which TelePresence is integrated with data, voice, and other video applications over a private WAN infrastructure.



Note

Recommendations for Multipoint over WAN circuits which involve a handoff to a service provider network, such as MPLS or Metro-Ethernet, and which involve shaping to limit the data rate to the service provider, are discussed in [Multipoint TelePresence over MPLS Circuits with Ethernet Handoff](#).

Cisco Router and Switch Platforms Tested

The following router and switch platforms were tested for the recommendations which follow:

WAN Catalyst 6500 switches:

- Catalyst 6506-E
 - IOS Version 12.2(18)SXF12
 - WS-SUP32-10GE-3B
 - WS-F6K-PFC3B
 - WS-F6K-MSFC2A
- 7600 SIP-400
 - SPA-2XOC3-POS
 - SPA-1XOC12-POS
 - SPA-1XOC48POS/RPR
- 7600-SIP-200
 - SPA-4XT3/E3

WAN Cisco 7200 routers:

- Cisco 7206VXR (NPE-G2) processor (rev A) with 917504K/65536K bytes of memory.
- IOS Version 12.4(4)XD7
- PA-POS-OC3MM
- PA-T3+

TelePresence over Dedicated WAN Circuits

The following sections discuss the recommendations for support of multipoint and multiple point-to-point TelePresence over dedicated private WAN circuits.

Dedicated T3, E3, and OC-3 POS Circuits

Table 12-1 summarizes the recommendations for supporting multipoint TelePresence over T3, E3, and OC-3 WAN circuits.

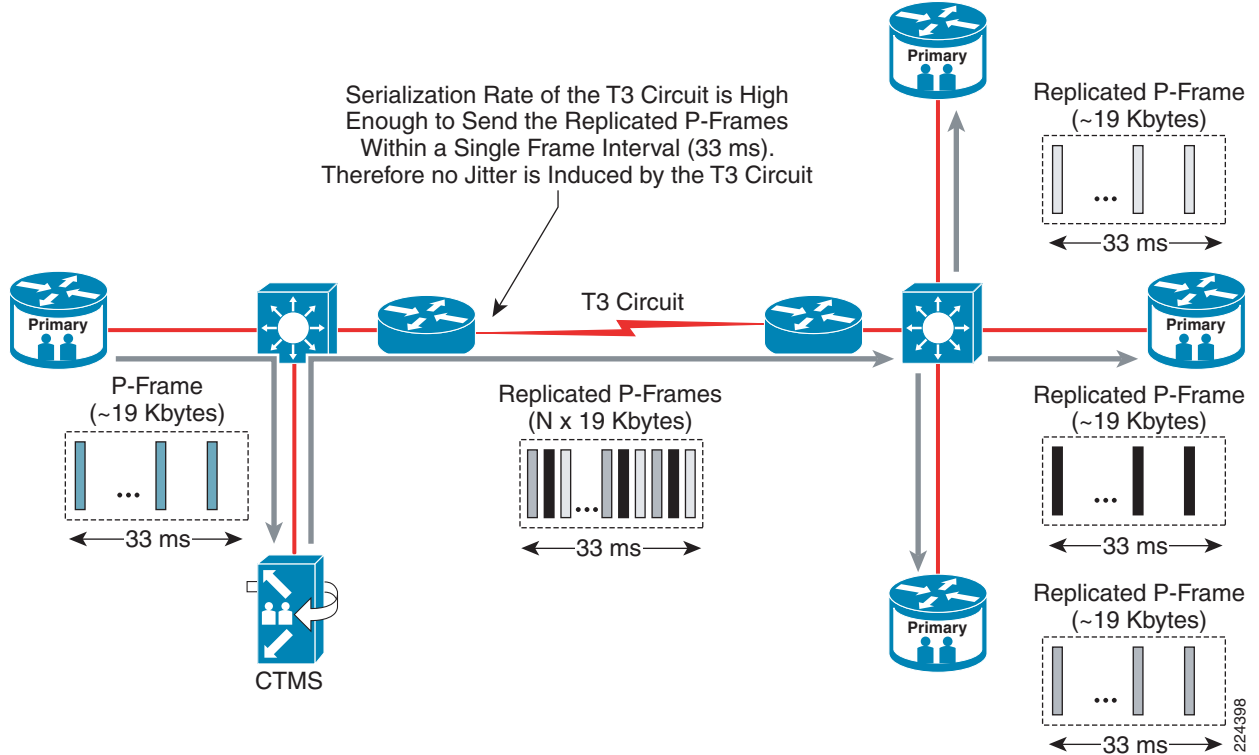
Table 12-1 Multipoint over Dedicated T3, E3, and OC-3 Circuit Results

Circuit	Bandwidth	Maximum Recommended CTS Table Segments	Recommended Starting Point for Tuning Output Hold Queue	TxRing Settings
T3	44.2 Mbps	7	840 (See discussion)	Default (See discussion)
E3	34.01 Mbps	5	600 (See discussion)	Default (See discussion)
OC-3 POS	155 Mbps	22 Estimated (See discussion)	2,640 (See discussion)	Default (See discussion)

Multipoint Recommendations

When configured in a single multipoint call, the maximum number of CTS table segments for lower speed WAN circuits, such as T3, E3, and OC-3 POS, is bounded not by the bandwidth of the circuits, but by the jitter induced by serialization delay across the circuits. The recommended maximum number of CTS table segments in a single multipoint call over such circuits is based upon the number of I-Frames and/or auxiliary video (PowerPoint slide) bursts simultaneously replicated by the CTMS and sent across the circuit. Figure 12-1 and Figure 12-2 help explain this using a T3 circuit as an example.

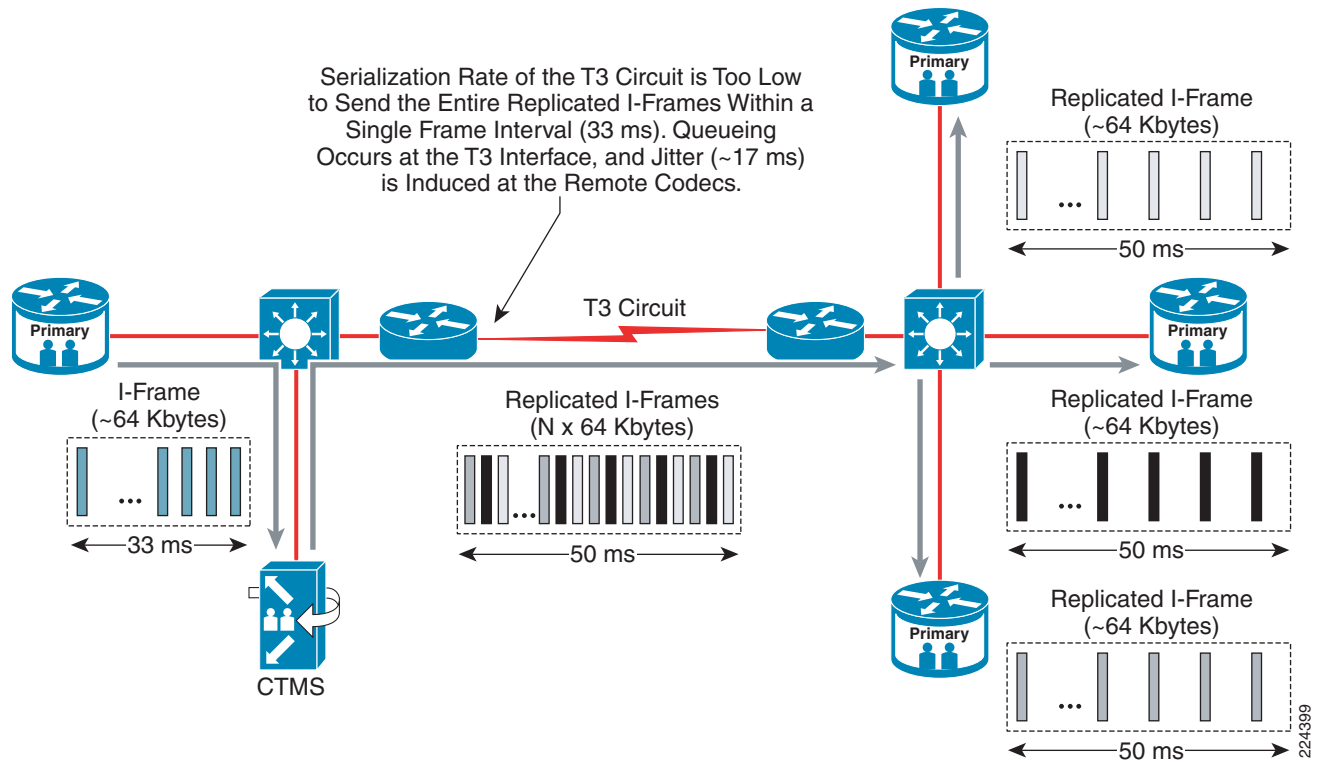
Figure 12-1 Example P-Frame Replication Across a T3 Circuit



As shown in [Figure 12-1](#), the transmitting codec sends a video frame from its camera every 33 ms frame interval. In this case, the video frame is a highly-compressed P-frame. The P-frame is replicated by the CTMS to the other three remote codecs across the T3 circuit. The replicated information is three times the size of the original P-frame. However, it must be received by each remote codec within the same 33 ms frame interval. In the example in [Figure 12-1](#), the serialization rate of the T3 circuit is high enough such that the entire replicated P-frames can still be sent across the circuit within 33 ms. Therefore, each remote codec receives its P-frame within the frame interval and no jitter is induced. This assumes that the latency due to physical distance of the circuit is constant.

[Figure 12-2](#) shows the same configuration with a I-frame replicated across the T3 circuit.

Figure 12-2 Example I-Frame Replication Across a T3 Circuit



As shown in Figure 12-2, the transmitting codec now sends a I-frame from its camera during the 33 ms frame interval. I-frames are about 3 to 4 times larger than P-frames, so there is much more information to be sent. The I-frame is again replicated by the CTMS to the other three remote codecs across the T3 circuit. The replicated information is three times the size of the original I-frame. It must still be received by each remote codec within the same 33 ms frame interval. In the example in Figure 12-2, the serialization rate of the T3 circuit is not fast enough to send the entire replicated I-frame within 33 ms. Therefore, queuing occurs at the T3 interface of the sending router. In the example above, the entire replicated I-frames are sent out in 50 ms. Since the CTMS replicates on a packet-by-packet basis, each remote CTS codec receives the full I-frame over a 50 ms interval. Therefore, approximately 50 ms - 33 ms = 17 ms of jitter is induced in each of the remote codecs when the I-frame is sent.

The tolerance of the CTS endpoints to jitter is based upon continuous jitter over a 50 second time period. Since jitter induced by an I-frame or auxiliary PowerPoint burst is a transient event, there is no concern that the call will terminate due to jitter exceeding the target of 40 ms for bad video quality. Also, the transient jitter can exceed the target of 20 ms for good video quality without causing any visible artifacts on the displays. However, the de-jitter buffers of the codecs are somewhat dynamic, with a target depth somewhere between 60-80 ms. If the transient jitter induced by I-frames and auxiliary bursts (due to PowerPoint slide transitions) exceeds the current depth of the de-jitter buffer, the codecs will mark packets as being "late." Late packets are not used in the decoding process of the codecs and cause visible artifacts on the displays. This degrades the video quality and the overall TelePresence experience.

Based upon ESE testing with CTS-1000s in a single multipoint call with low-speed (5 frames per second) auxiliary video input, the maximum number of endpoints which were observed to be supported on T3 and E3 circuits before late packets were reported on the CTS codecs is listed in Table 12-1. Note that these recommendations are based on testing across a single circuit.

**Note**

The number of CTS segments which can be supported with high-speed auxiliary video input may be lower due to the higher speed of the auxiliary video input and potentially larger bursts. ESE has not tested high-speed auxiliary video input currently.

The maximum number of segments supported over an OC-3 circuit is estimated based on the results from the T3 and E3 testing, as well as a tested configuration of 15 CTS-1000s and 3 CTS-3000s. Note that the number of replicated I-frames and auxiliary bursts in the tested configuration is somewhat lower than the number that would be generated by 22 CTS-1000s.

Multiple Point-to-Point Recommendations

For circuits which support multiple point-to-point TelePresence calls, the I-frames generated during each call are not time synchronized. In other words, there is no CTMS replication of I-frames to each site simultaneously. As a result of this, less jitter is seen by the codecs and the number of CTS table segments which can be supported across a circuit may be slightly higher. The maximum table segments which may be supported in this configuration should be calculated based on the bandwidth of the circuit and the bandwidth of each CTS endpoint—for example, using 15 Mbps per CTS-3000 and 5.5 Mbps per CTS-1000. Although these values represent maximum rates over a one-second period, it is not recommended to oversubscribe the circuit by provisioning CTS units such that the sum of their one second maximum rates exceeds the bandwidth of the circuit. Additionally, a wise rule-of-thumb for any network is to not exceed 75% utilization of any WAN link on a consistent basis, before considering increasing the bandwidth of the link.

Since the customer may at any time deploy a CTMS and begin supporting multipoint calls across the network, another wise rule-of-thumb is to design the network in such a manner that it is ready for multipoint. Therefore, it is recommended that regardless of whether the customer currently supports multipoint, the network should be designed in such a manner that it can be added without having to adjust network parameters on all circuits throughout the network when multipoint is supported.

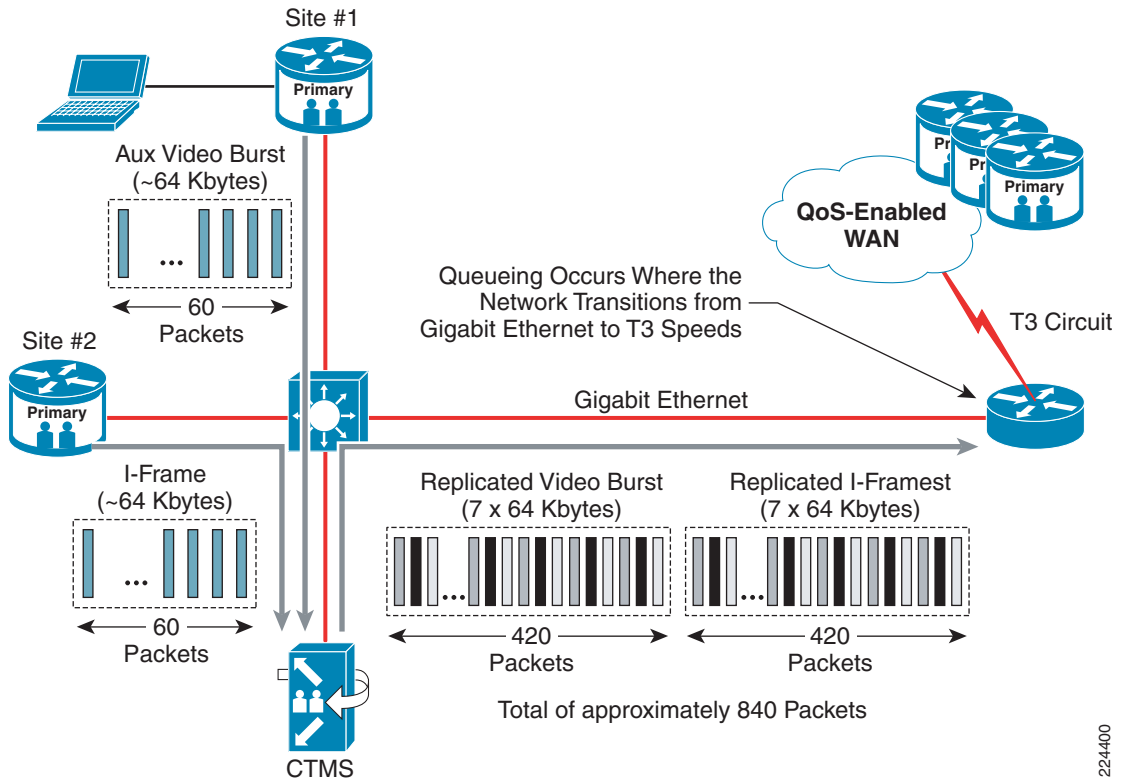
Output Hold Queue Recommendations

In multipoint configurations over dedicated circuits, the output hold queue of routers connected to the circuits may have to be increased to support the bursts due to I-frame replication and auxiliary bursts due to PowerPoint slide transitions.

Buffer size tuning is dependant upon size of the I-frame and auxiliary video (PowerPoint slide) bursts and their timing. Therefore, the output queue sizes listed in [Table 12-1](#) represent starting points for tuning the hold-queue out parameters on router WAN interfaces based upon testing within the ESE lab.

The following example demonstrates one simple method of estimating the output hold queue size based upon information in [Estimating Burst Sizes within Multipoint TelePresence Calls](#) in [Chapter 11](#), “[Cisco Multipoint Technology and Design Details](#).” Assume a multipoint TelePresence call consisting of 9 CTS-1000s. A maximum of seven of the CTS-1000s, totaling seven CTS table segments, are located across a T3 circuit from the CTMS, as shown in [Figure 12-3](#).

Figure 12-3 Output Hold Queue Estimation



224400

In the example in Figure 12-3, a slide show is being presented by Site #1. As the presenter transitions the slide, a question is being asked from Site #2, causing a speaker transition to occur simultaneously. As a result of the two simultaneous events, both an I-frame and an auxiliary video burst are sent simultaneously. Note that audio has been ignored to simplify the example. The I-frame and auxiliary video burst are replicated by the CTMS and sent to the seven CTS 1000s across the T3. (For simplicity of the example, replication of the I-frame and auxiliary video burst to the endpoints on the same side of the T3 circuit is not shown.)

Queuing primarily occurs at the transition from the higher speed Gigabit Ethernet interface to the lower speed T3 interface. In other words, queuing occurs at the output hold queue of the router's T3 interface. The following equations estimate the total burst size in packets:

60 Packets per I-frame * Replication to 7 Sites = 420 Packets

60 Packets per Aux Video Burst * Replication to 7 Sites = 420 Packets

Total Packets = 420 + 420 = 840 Packets

Based upon the example, a good starting point for tuning the output hold queue of the router's T3 interface would be approximately 840 packets. Actual test results for output hold queue sizes required for support of multipoint TelePresence over dedicated circuits have shown the equation above to be slightly conservative. The example above over-simplifies the arrival and transmission of the burst. It assumes that all 840 packets arrive at the same time. In reality, the I-frame is sent over a 33 ms frame interval and the auxiliary video burst is sent over a 200 ms frame interval with low-speed auxiliary video. Also, not shown is a single audio packet which arrives every 20 ms. Likewise, the example assumes that the entire burst arrives before any of it is serialized and sent onto the T3 circuit. In reality, as the information arrives, it is serialized. The higher the speed of the circuit, the more is serialized as it arrives,

and therefore the lower the maximum queue depth. It must also be noted that the size of I-frames and auxiliary video bursts is variable. The values shown above represent estimates for maximum sizes seen during testing.

Despite the over-simplification of the example, the method presented can still be used as a starting point for tuning output hold queues for supporting multipoint TelePresence. More generically, the starting point for tuning the output hold queue can be determined by the equation:

$N * 120 \text{ packets} + M * 240 \text{ packets} = \text{Starting point for tuning the output hold queue size for a dedicated circuit}$

The application of detailed queuing theory to obtain a more accurate assessment of the maximum queue depth is beyond the scope of this document. However, it should be noted that there is no apparent downside, in terms of TelePresence video quality, of simply increasing the output hold-queue on the router to its maximum value. TelePresence simply utilizes as much of the output queue space as needed. Bounding the maximum number of CTS table segments supported across the circuit and then allowing as much queuing space as needed is preferable to discarding TelePresence traffic.

TxRing Tuning Recommendations

For multipoint TelePresence over a dedicated circuit, it is recommended to leave the TxRing configuration of the WAN interface at its default value. Testing over a T3 circuit has indicated that decreasing the value to 10 packets—as was the recommendation for point-to-point TelePresence over a shared T3 circuit—resulted in output drops on the T3 interface. The length of the TxRing determines when queuing is enabled and disabled due to congestion on IOS routers. In a dedicated circuit configuration, TelePresence is the only traffic. Therefore, queuing is somewhat irrelevant and the TxRing can be left at its default value. Note that within many POS interfaces on IOS routers, the TxRing cannot be tuned anyway.

Multipoint Over Dedicated OC-12 and OC-48 POS Circuits

Testing the maximum number of CTS table segments which a dedicated OC-12 or OC-48 POS circuit is capable of handling is beyond the capability of the ESE lab. However both OC-12 and OC-48 POS have been tested with 15 CTS-1000s and 3 CTS-3000s in a single multipoint call with low-speed (5 frames per second) auxiliary video input. Additionally, another 8 simulated CTS-3000s were added to each circuit, for a total of 48 CTS table segments. It has been confirmed that dedicated OC-12 and OC-48 POS circuits are capable of handling the equivalent to the amount of TelePresence traffic that a single CTMS can currently generate across these circuits.

No tuning of output hold queues was necessary for the tested configuration on OC-12 or OC-48 circuits. Further, as mentioned previously, the TxRing cannot be tuned for the tested OC-12 and OC-48 interfaces.

Recommendations for Multipoint Over Converged WAN Circuits

The following sections discuss the recommendations for support of multipoint and multiple point-to-point TelePresence over converged WAN circuits. Converged circuit configurations are those in which TelePresence is integrated with data, voice, and other video applications over a private WAN infrastructure.

When configuring multipoint over a WAN edge design, the network administrator should follow the basic design guidance outlined in [Chapter 6, “Branch QoS Design for TelePresence.”](#) The following sections extend those discussions to multipoint TelePresence designs.

Multipoint TelePresence WAN Edge LLQ Policy

When configuring multipoint TelePresence over a converged private WAN infrastructure, it is often advantageous to place TelePresence in an LLQ. If multipoint TelePresence is to be assigned to an LLQ, then sufficient bandwidth must be allocated to the LLQ to support multiple CTS endpoints. The LLQ bandwidth is typically limited by either implicit policing or a configured explicit policer. In either case, the policed rate (CIR) must be sufficient for the number of CTS endpoints supported across the converged WAN circuit. This is regardless of whether the CTS endpoints are in a single multipoint call, multiple multipoint calls, or multiple point-to-point calls.

Bandwidth Provisioning

The amount of LLQ bandwidth provisioned via the policed rate can be calculated based upon the number of CTS endpoints to be supported across the WAN circuit and the resolution/quality configuration for the endpoints, as shown in [Table 4-1](#) in [Chapter 4](#), “Quality of Service Design for TelePresence.”

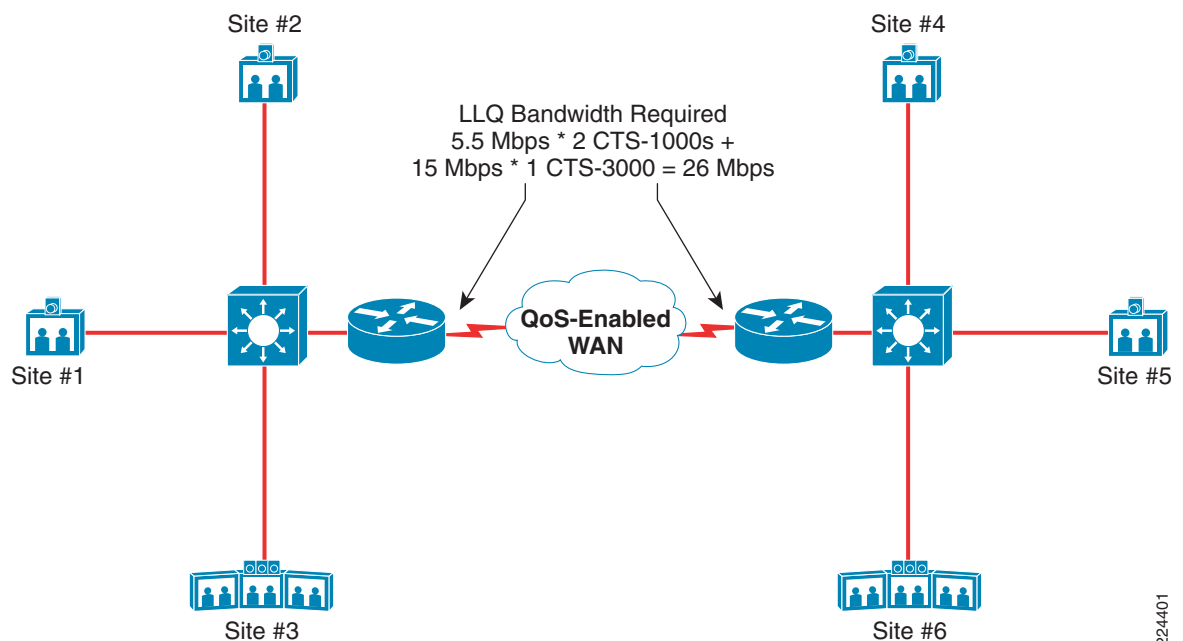
For example, with 1080p best quality the amount of provisioned LLQ bandwidth can be determined using the equation:

$$N * 5.5 \text{ Mbps} + M * 15 \text{ Mbps} = \text{Policed Rate (CIR)}$$

Where N is the number of CTS-1000s and M is the number of CTS-3000s to be supported across the circuit.

[Figure 12-4](#) shows this with an example TelePresence deployment.

Figure 12-4 Bandwidth Requirements for Multipoint and Multiple Point-to-Point TelePresence



In [Figure 12-4](#) there are a total of six CTS endpoints in the network—four CTS-1000s and two CTS-3000s. However each side of the WAN has only three CTS endpoints. The amount of bandwidth required to be provisioned in the LLQ policed rate across the WAN, in order to support having all TelePresence units in calls at the same time, is:

$$2 \text{ CTS-1000s} * 5.5 \text{ Mbps per} + 1 \text{ CTS-3000} * 15 \text{ Mbps} = 26 \text{ Mbps}$$

This holds regardless of whether the CTS units are in a single multipoint call or in multiple multipoint calls.



Note


These recommendations hold for low-speed (5 frames/second) auxiliary video input only. Additional bandwidth must be provisioned to support high-speed (30 frames/second) auxiliary video input.

It should be noted that with TelePresence in an LLQ configuration, it is recommended that the total amount of traffic allocated for the LLQ (Voice and TelePresence) be below approximately 1/3 of the circuit bandwidth in order to prevent possible decreases in responsiveness of applications which are not placed into the LLQ.

Burst Provisioning

The burst parameter of the LLQ policer must also be scaled to handle bursts from multiple CTS endpoints. This burst size is not directly dependent upon the resolution/quality configuration of the CTS endpoints. However, the required burst size is dependent upon whether CTS endpoints are in individual point-to-point calls, multiple multipoint calls, or one single multipoint call. [Figure 12-5](#) and [Figure 12-6](#) help to explain this.

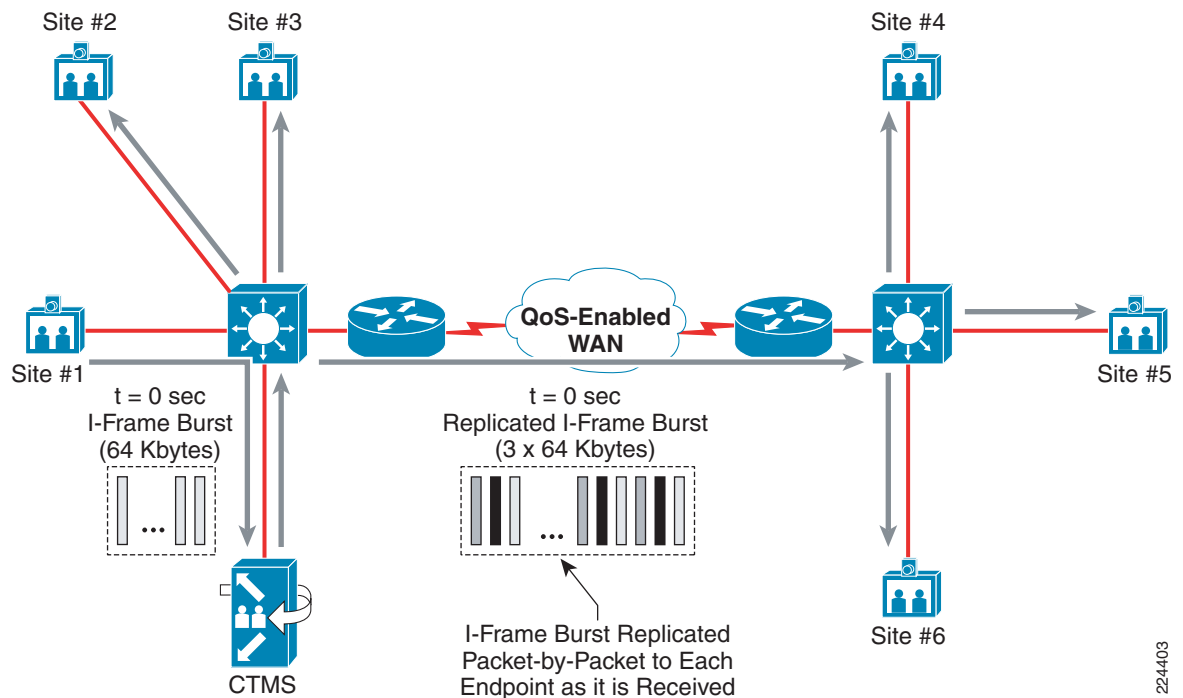
Figure 12-5 Multiple Point-to-Point TelePresence Calls



[Figure 12-5](#) shows an example of six CTS-1000 sites in three point-to-point calls. All three calls traverse the converged WAN. In the example, at time $t = 0$ seconds, Site#1 generates an I-frame burst, perhaps as part of normal video synchronization. Likewise at $t = 10$ seconds, Site #2 generates an I-frame burst, and at $t = 20$ seconds, Site #3 generates an I-frame burst. In the example, the I-frames are not time synchronized. In other words, bursts do not necessarily occur at the same time with separate point-to-point calls. Therefore the burst size of the LLQ policer of the WAN circuit only needs to handle a single I-frame burst in this example.

[Figure 12-6](#) shows the same six sites, in a single multipoint conference call this time.

Figure 12-6 Single Multipoint TelePresence Call



In this example, at time $t = 0$ seconds, Site #1 generates an I-frame burst, perhaps because it became the active site. The I-frame burst is immediately replicated by the CTMS to every other site. Since Sites #4, #5, and #6 are located across the WAN, the LLQ policer burst parameter (B_c) must be configured to handle three times the size of the initial I-frame burst generated by Site #1.

Since a single multipoint call represents a worst case scenario in terms of synchronization of bursts, it is recommended to provision the LLQ burst parameter (B_c) to support this scenario. Allocation of approximately 64 Kbytes of burst space per video camera, as well as another 64 Kbytes of burst space per auxiliary video input on each CTS endpoint, has been shown to be conservatively sufficient. Using this recommendation, each CTS-1000 would require approximately 128 Kbytes of burst space and each CTS-3000 would require approximately 256 Kbytes of burst space.

Using these recommendations, the amount of provisioned LLQ burst for a converged circuit which supports multipoint TelePresence can be calculated based upon the following equation:

$$N * 128 \text{ Kbytes} + M * 256 \text{ Kbytes} = \text{Burst Size } (B_c)$$

Where N is the number of CTS-1000s and M is the number of CTS-3000s to be supported across the circuit.

Going back to the example in Figure 12-3, the burst size required to be provisioned in the LLQ policer in order to support having all TelePresence units in a single multipoint call is:

$$2 \text{ CTS-1000s} * 128 \text{ Kbytes} + 1 \text{ CTS-3000} * 256 \text{ Kbytes} = 512 \text{ Kbytes}$$

This holds regardless of whether the CTS units are in a single multipoint call or in multiple multipoint calls.

Multipoint LLQ Configuration Examples

Example 12-1 and Example 12-2 show both an implicit and explicit policer configuration for support of multipoint TelePresence in an LLQ configuration along with voice.

Example 12-1 Implicit Policer

```

policy-map OC-3 WAN-EDGE
  class VOICE
    priority percent 10! LLQ for Voice (example amount of BW)
  class TELEPRESENCE
    priority 26000 512000! LLQ for TP (priority bandwidth)
  class DATA
  ...

```

Example 12-2 Explicit Policer

```

policy-map OC-3 WAN-EDGE
  class VOICE
    police cir 15500000! Explicit Policer for Voice (default burst)
      conform-action transmit
      exceed-action drop
      violate-action drop
    priority 15500! LLQ for Voice (explicit priority bandwidth)
  class TELEPRESENCE
    police cir 26000000 bc 512000 be 512000! Explicit Policer for TP
      conform-action transmit
      exceed-action drop ! Note: Excess burst (Be) not utilized.
      violate-action drop
    priority 26000! LLQ for TP (explicit priority bandwidth)
  ...

```

It should be noted that the combination of configuring a policed rate of 5.5 Mbps and a burst size 128 Kbytes per CTS-1000 ensures a policer time constant of approximately 186 ms. Likewise, configuring a policed rate of 15 Mbps and a burst size of 256 Kbytes per CTS-3000 ensures a policer time constant of approximately 137 ms. Therefore the recommendations above ensure that the time constant for the policer (given by the equation $T_c = CIR/Bc$) is always between 137 ms and 186 ms. Many service provider networks already utilize time constants of around 200 ms, implying a larger burst value configured for a given policed rate, than recommended above.

Empirical Test Results

Table 12-2 summarizes the test results for supporting multipoint TelePresence over shared T3, E3, and OC-3 WAN circuits, performed within the ESE lab.

Table 12-2 Multipoint over Shared T3, E3, and OC-3 POS Circuit Test Results

Circuit	CTS Unit Configuration	Policed Rate (CIR)	Actual Burst (Bc) Tested	Recommended Burst Size	Percentage of Recommended
E3	2 CTS-1000s	11 Mbps	160 Kbytes	256 Kbytes	63%
T3	3 CTS-1000s	16.5 Mbps	164 Kbytes	384 Kbytes	43%
OC-3	7 CTS-1000s and 1 CTS-3000	53.5 Mbps	680 Kbytes	1,152 Kbytes	59%
OC-3	9 CTS-1000s	49.5 Mbps	720 Kbytes	1,152 Kbytes	63%
OC-12	15 CTS-1000s and 3 CTS-3000s	127.5 Mbps	1,584 Kbytes	2,688 Kbytes	59%

The test results in [Table 12-2](#) show that allocating 128 Kbytes per CTS-1000 and 256 Kbytes per CTS-3000 are fairly conservative recommendations, which should be sufficient for most customer networks. Actual testing has shown that bursts are typically somewhere between 40% to 70% of this value. Several reasons for this wide range of results are:

- The size of I-frames (IDRs) is highly variable, with sizes observed up to 64 Kbytes.
- The size of auxiliary video bursts is also highly variable.
- The probability of part or all of an auxiliary video burst (PowerPoint slide transition) occurring during the exact same time as an I-frame (IDR) is low.

Multipoint TelePresence Branch WAN Edge CBWFQ Policy

If multipoint TelePresence is to be assigned to a CBWFQ, then sufficient bandwidth and queue size for bursts must be allocated to the CBWFQ to support multiple CTS endpoints.

Bandwidth Provisioning

Within a CBWFQ, bandwidth is typically limited via a bandwidth statement within the policy map, as shown in [Example 12-3](#).

Example 12-3 TelePresence CBWFQ Example

```
policy-map T3 WAN-EDGE
  class VOICE
    police cir 1184000! Explicit Policer for Voice (default burst)
      conform-action transmit
      exceed-action drop
      violate-action drop
      priority 1184! LLQ for Voice (explicit priority bandwidth)
  class TELEPRESENCE
    bandwidth percent 37! CBWFQ for TP (example amount of BW)
    queue-limit 120
  ...
```



Note

In [Example 12-3](#), TelePresence traffic can utilize more than 37% of the bandwidth if other CBWFQ classes and/or the voice LLQ class are not utilizing their complete share of the bandwidth of the circuit. This was also the case for TelePresence in the LLQ with an implicit policer.

When allocating CBWFQ bandwidth for multipoint TelePresence, it is recommended to again allocate approximately 5.5 Mbps per CTS-1000 and 15 Mbps per CTS-3000 (assuming low-speed auxiliary video). In the example above which supports three CTS-1000s, bandwidth is calculated as:

$$3 \text{ CTS-1000s} * 5.5 \text{ Mbps per CTS-1000} = 16.5 \text{ Mbps}$$

$$\text{Bandwidth} = 16.5 \text{ Mbps} / 44.21 \text{ Mbps T3 circuit bandwidth} = 37\%$$

Burst Provisioning

When configuring multipoint TelePresence in a CBWFQ, bursting is accommodated by a combination of the queue size configured within the policy map via the queue-limit command, as well as the output hold queue size of the WAN interface configured via the hold-queue out command.

Policy map Queue Size Recommendations

When multipoint TelePresence is placed within a CBWFQ on a converged circuit, the queue size for the TelePresence class within the policy map may need to be increased above the default limit of 64 packets to support the bursts.

There is no exact science as to the size of the queue limit to configure within the policy map for the TelePresence class based upon the number of CTS endpoints supported over the circuit. Tuning of the queue size is dependant upon the sizes of the I-frame and auxiliary video bursts and their timing, as well as the amount of congestion on the WAN circuit which causes queuing to occur in the first place on a converged circuit. Further, TelePresence bursts themselves may be the cause of the temporary congestion which results in queuing.

However, the method shown in [Figure 12-3](#) in [Output Hold Queue Recommendations](#) can be applied to multipoint TelePresence in a CBWFQ in order to provide the network administrator a starting point which to begin tuning the queue size for the TelePresence class.

As shown in [Figure 12-3](#), each I-frame from the camera of a CTS-endpoint can be as large as 64 Kbytes. Given a frame size of approximately 1,100 bytes, this translates to approximately 60 packets per I-frame per camera. Likewise the auxiliary video burst can be approximately 60 packets in size. Utilizing these rough estimates, and ignoring the audio for simplicity, the following equation can be used to provide a starting point for tuning the queue size within the policy map:

$N * 120 \text{ packets} + M * 240 \text{ packets} = \text{Starting point for tuning the queue size for the TelePresence class.}$

Based upon this equation, a starting point for tuning the queue size within the policy-map for the TelePresence class for a circuit that supports three CTS-1000s would be:

$3 \text{ CTS-1000s} * 120 \text{ packets per CTS-1000} = 360 \text{ packets}$

As with the empirical test results for policer burst sizes shown in [Table 12-1](#), actual test results for queue-limit sizes required for support of multipoint TelePresence in a CBWFQ design have shown the equation above to be somewhat conservative. This is partly because the equation above over-simplifies the arrival and transmission of the bursts. It assumes that all packets arrive at the same time. Likewise, the equation assumes that the entire burst arrives before any of it is serialized and sent onto the circuit. In reality, as the information arrives, it is serialized. The higher the speed of the circuit, the more is serialized as it arrives, and therefore the lower the required queue depth. Finally, it must also be noted that the size of I-frames and auxiliary video bursts is variable. The values shown above represent estimates for maximum sizes seen during testing.

Despite the over-simplification of the example, the method presented can still be used as a starting point for tuning queue-limit within the policy map for supporting multipoint TelePresence. If desired the network administrator can tune the queue-limit size down to the point where TelePresence drops still do not occur.

It should be noted that increasing the queue limit size of TelePresence traffic within a CBWFQ configuration can lead to higher jitter seen by the TelePresence codecs, since a higher queue limit implies more packets can be queued, rather than discarded. At some point, jitter can lead to “late” packets which are not used by the codecs in the decoding process and may cause visible artifacts in the TelePresence video. However, decreasing the queue limit size would result in “drops” seen by the codecs, which would also result in visible artifacts in the TelePresence video. In such cases, it may be appropriate to either reconfigure TelePresence in a LLQ configuration or increase the bandwidth of the circuit such that TelePresence no longer experiences long queuing delays due to congestion when in a CBWFQ configuration.

Output Hold Queue Size Recommendations

When multipoint TelePresence is placed within a CBWFQ on a converged circuit, the output hold queue size for the WAN interface may also need to be increased above the default limit to support the bursts.

As with the queue limit size within the policy map, there is no exact science as to the size of the output hold queue to configure on the WAN interface. Tuning of the output hold queue size is dependant upon the sizes of the I-frame and auxiliary video bursts and their timing as well as the amount of congestion on the WAN circuit which causes queuing to occur. TelePresence bursts themselves may again be the cause of the temporary congestion which results in queuing.

However, in a 12-class QoS model in which TelePresence is placed in a CBWFQ, it is recommended that the output hold queue size of the WAN interface be larger than the sum of the individual queue-limit sizes of the CBWFQ classes within the policy map. This allows each of the CBWFQ classes to reach the maximum size of their queue limits and provides room for LLQ traffic without exceeding the output hold queue size of the physical WAN interface. The queue limit sizes of each of the CBWFQ classes can be determined via the show policy-map command.

TxRing Limit Recommendations

For multipoint TelePresence over a converged circuit, it is recommended to leave the TxRing configuration of the WAN interface at its default value where possible.

Multipoint TelePresence over MPLS Circuits with Ethernet Handoff

This section discusses the recommendations for support of multipoint and multiple point-to-point TelePresence over a WAN infrastructure which consists of an Ethernet handoff to a service provider MPLS network. Both dedicated and converged network configurations are discussed.

When configuring multipoint over an WAN edge design, the network administrator should follow the basic design guidance outlined in [Chapter 6, “Branch QoS Design for TelePresence.”](#) Specifically, [TelePresence Branch MPLS VPN](#) discusses specific QoS parameters as they pertain to TelePresence deployments over an MPLS VPN. The following sections simply extend those discussions to multipoint TelePresence designs.

Cisco Router and Switch Platforms Tested

The following router and switch platforms were tested for the recommendations which follow:

Cisco 7200 routers:

- Cisco 7206VXR (NPE-G2) processor (rev A) with 917504K/65536K bytes of memory
- IOS Version 12.4(4)XD7
- PA-POS-OC3MM
- PA-T3+

Cisco 3800 ISR routers:

- Cisco 3845 (revision 1.0) with 478208K bytes of memory
- IOS Version 12.4(15)T

- C3845 Mother board 1GE(TX,SFP),1GE(TX), integrated VPN and 4W Port adapter, 2 ports

Cisco 2800 ISR routers:

- Cisco 2821 (revision 53.51) with 1032192K / 16384K bytes of memory
- IOS Version 12.4(15)T
- C2821 Motherboard with 2GE and integrated VPN Port adapter, 2 ports

Multipoint TelePresence over Dedicated MPLS Circuits

When deploying multipoint TelePresence over a WAN infrastructure consisting of a dedicated MPLS circuit with an Ethernet handoff to the service provider network, egress shaping may be applied to the CE edge device. Since the physical interface rate may be a 1 Gbps Ethernet connection, but the contracted data rate to the service provider network may only be 50 Mbps, 100 Mbps, or 150 Mbps; shaping smooths out traffic to meet the contracted sub-line rate.

However, it should be noted that networking gear is not designed to buffer large amounts of traffic. Shaping should only be used to temporarily smooth bursty traffic, where the overall traffic rate is limited by mechanisms such as call admission control or manual provisioning. When shaping is active for long periods of time, the buffering capacity of routers and switches will be overrun and traffic will drop, degrading the overall TelePresence experience. In such cases, additional bandwidth should be provisioned. Note that the higher the shaped data rate, the faster the buffering capacity is exhausted on a given router or switch platform when shaping becomes active.

In a dedicated or overlay deployment, queueing is not necessarily a requirement, since all the traffic on the circuit is TelePresence related. Therefore a single-level policy map applied to the egress interface of the CE device, such as shown in [Example 12-4](#), can be implemented.

Example 12-4 Egress CE Policy Map Which Enforces Shaping Only

```
policy-map DEDICATED-EGRESS-TOP-MAP
  class class-default
    shape average 50000000 1250000 6400000! All in bits/sec.
    shape max-buffers 4096
    ...
```

The network administrator needs to determine the appropriate values for the average rate (CIR), committed burst (Bc), and excess burst (Be) to best match the bursty characteristic of multipoint TelePresence. However, the network administrator is also constrained by the ingress policy map of the service provider PE device, which may police to the contracted rate. See [Example 12-5](#).

Example 12-5 Ingress PE Policy Map Which Enforces Policing Only

```
policy-map DEDICATED-TP-INGRESS
  class class-default
    police cir 50000000
      conform-action set-mpls-exp-transmit 5
      exceed-action drop
    ...
```

In the policer configuration in [Example 12-5](#), the average rate (CIR) matches the contracted rate of 50 Mbps. Note that the burst parameters (Bc and Be) are not specified explicitly within the configuration and therefore assume the default value. For most Cisco IOS devices the value corresponds to a time constant (Tc) of 250 ms. For example, in the configuration above, the default committed burst (Bc) is given by the equation:

$$Bc = Tc * CIR = 250 \text{ ms} * 50 \text{ Mbps} = 12,500,000 \text{ bits} = 1,562,500 \text{ bytes}$$

Note that for policers, the burst values are typically specified in bytes not bits. The configuration also specifies that any traffic that conforms to the committed burst (Bc) is transmitted, but any traffic that exceeds it is dropped. Therefore in the example above, the excess burst (Be) is effectively not utilized.

In order to keep the PE policer from dropping ingress TelePresence traffic, and therefore stay within contract, the network design engineer needs to choose the appropriate committed rate (CIR), committed burst (Bc), and excess burst (Be) value on the shaper of the egress CE device. Two possible methods of determining these values are discussed below.

Design Option #1

With design option #1, the network administrator sets the CIR of the shaper on the CE device to match the CIR of the policer of the PE device. This is typically done when contracting a sub-line rate from a service provider.

The network administrator must ensure that the total burst (Bc + Be) sent from the CE shaper is below the committed burst (Bc) of the PE policer. However, since the time constants may be different, this may be a challenge. Configuring a long time constant on the CE shaper, perhaps 250 ms, to match the time constant of the PE policer can result in late packets seen by the remote codecs if shaping is activated at all. This nullifies the need for the shaper. In such cases it may be better not to configure a shaper at all and let the policer discard the traffic. Therefore, it may be desirable to configure a smaller time constant on the CE shaper. However a smaller time constant results in a smaller committed burst size, given by the equation $Bc = Tc / CIR$.

TelePresence bursts are relatively rare occurrences from the standpoint of the overall video which is sent at 30 frames per second. ESE testing has shown that configuring a small committed burst which yields a time constant of around 25 ms and large excess burst (Be) to handle the rest of the overall burst can work effectively in a dedicated deployment.

An example of this method using 9 CTS-1000s over a 50 Mbps MPLS circuit is discussed below. The estimated burst from 9 CTS-1000s can be estimated based on the equation:

$$9 \text{ CTS-1000s} * 128 \text{ Kbytes per CTS-1000} = 1,152,000 \text{ bytes} = 9,216,000 \text{ bits}$$

Configuring a sustained rate (CIR) of 50 Mbps with a time constant (Tc) of 25 ms yields the following committed burst (Bc) size for the CE shaper:

$$Bc = Tc * CIR = 25 \text{ ms} * 50 \text{ Mbps} = 1,250,000 \text{ bits} = 156,250 \text{ bytes}$$

Therefore, in order to hold the entire burst from 9 CTS-1000s in one time constant, the following excess burst size would need to be configured:

$$Be = 9,216,000 \text{ bits} - 1,250,000 \text{ bits} = 7,966,000 \text{ bits} = 995,750 \text{ bytes}$$

This method relies on the fact that normal TelePresence traffic is well below the recommended bandwidth allocation of 5.5 Mbps per CTS-1000 or 15 Mbps per CTS-3000 and that the bursts are typically nowhere near a full 128 Kbytes per CTS-1000 or 256 Kbytes per CTS-3000. Therefore, the service provider PE policer committed burst size (Bc) is never exceeded and TelePresence traffic is never dropped.

[Table 12-3](#) summarizes the actual test results from ESE testing of multipoint TelePresence over a dedicated 50 Mps, 100 Mbps, and 150 Mbps Ethernet handoff to an MPLS network using this methodology.

Table 12-3 ESE Tested Results for Multipoint Over Dedicated MPLS Circuits

Shaped Rate	Max CTS Table Segments	Shaper Bc (Bytes)	Shaper Be (Bytes)	Total Shaper Burst (Bytes)	Percent of Theoretical Burst Size	Policer Be (Bytes)
50 Mbps	9 CTS-1000s	156,250	800,000	956,250	83%	1,562,500
100 Mbps	3 CTS-3000s and 10 CTS-1000s	312,500	800,000	1,112,500	54%	3,125,000
150 Mbps	3 CTS-3000s and 15 CTS-1000s (See discussion)	468,750	1,400,000	1,868,750	70%	4,687,500

For all testing, the policer on the PE device was set to its default time constant of $T_c = 250$ ms and the policed average rate matched the shaped average rate of the CE device. Note that the shaper time constant was 1/10th of the policer time constant. For each test, the excess burst (Be) values were tuned down just above the point at which drops began. As can be seen, the actual burst size required ranged from 54% to 83% of the theoretical required burst sizes.

For all of the testing, the physical line rate was 1 Gbps. Because of the high physical line rate, serialization delay was low. Therefore, even though the shaped rates were 50 Mbps, 100 Mbps, and 150 Mbps, the number of CTS table segments supported was not constrained by induced jitter causing late packets on the codecs, as was the case with dedicated E3, T3, and OC-3 circuits. Instead the number of CTS table segments is bounded by the shaped bandwidth and should follow the recommendations of allocating approximately 5.5 Mbps per CTS-1000 and 15 Mbps per CTS-3000 (assuming low-speed auxiliary video input only).

Output hold queue sized for the testing were increased to their maximum values of 4,096 for all MPLS testing over a dedicate circuit with Ethernet handoff. Since TelePresence is the only traffic in a dedicated configuration, no detrimental effects are expected from this. No adjustments were made to the TxRing for the Gigabit Ethernet interface.

Design Option #2

In situations where the policer time constant (T_c) of the service provider PE has been set to prevent bursts, it may be necessary to provision a greater amount of bandwidth to support multipoint TelePresence. However, determining just how much bandwidth is required is not easy to determine. ESE has not conducted any testing with this configuration currently. The following discussion highlights some of the main issues involved.

Using the same example of 9 CTS-1000s over a 50 Mbps MPLS Ethernet handoff, if the service provider also configures a time constant (T_c) of 25 ms, then the amount of burst that can be accommodated per time constant is given by the equation:

$$B_c = T_c * CIR = 25 \text{ ms} * 50 \text{ Mbps} = 12,500,000 \text{ bits} = 156,250 \text{ bytes}$$

In this case, if the shaper were also configured with a time constant (T_c) of 25 ms, then any excess burst (Be) configured on the shaper to try to accommodate the burstiness of multipoint TelePresence would immediately be dropped by the policer. Eliminating the excess burst ($B_e = 0$) and attempting to buffer traffic in the shaper for 8 time constants (1,152,000 bytes/156,250 bytes) or approximately 200 ms would result in late packets seen by the codecs and would likely overrun the buffering capacity of the router itself.

With the ability to only handle a 156.25 Kbyte burst, the provisioning of the network is sufficient to handle a single point-to-point CTS-1000 call. It may also be able to handle two of the CTS-1000s in a multipoint call. Two CTS-1000s in a multipoint call could result in a 256 Kbyte burst when a speaker transition occurs at the same time as a PowerPoint slide transition. During the first time constant ($t = 0$ to 25 ms) the shaper would send 156.25 Kbytes and during the second time constant ($t = 25$ to 50 ms) the shaper would send the remaining 100 Kbytes. There would be additional P-slice data and audio to send as well. However, a 50 ms delay may not be sufficient to cause late packets to be seen by the codecs.

Further, by relying on the statistical assumption that speaker transitions will not likely occur at exactly the same time as PowerPoint slide transitions, the network administrator may even be able to accommodate up to four CTS-1000s in a single multipoint call. If such events do occur, some packet loss will be seen. However the network is not provisioned to handle a multipoint call with all nine CTS-1000s in a single multipoint call. Such calls will in all likelihood fail.

Since the network administrator is usually not in control of how many devices across any given circuit can be in a single multipoint call, it is recommended to design for the worse case scenario. In such cases, the only alternative may be to increase the provisioned bandwidth such that the committed burst of both the shaper and policer can accommodate the multipoint TelePresence bursts. For example, to accommodate the theoretical burst size of 1,152,000 bytes from 9 CTS-1000s in a single multipoint call, the CIR would need to be increased as follows:

$$\text{CIR} = \text{Bc} / \text{Tc} = 1,152,000 \text{ bytes} / 25 \text{ ms} = 368.64 \text{ Mbps}$$

This translates to approximately 40 Mbps per CTS-1000. At this point the design engineer may need to accept the possibility of some packet losses occurring due to a PowerPoint slide transition occurring at the same time as a speaker transition and provision the network to accommodate either a slide transition or a speaker transition occurring, but not both simultaneously. This decreases the tolerance of the network to bursts by one-half.

Going back to the example of 9 CTS-1000s, the design engineer would now design the network to support the following bursts:

$$9 \text{ CTS-1000s} * 64 \text{ Kbytes per CTS-1000} = 576,000 \text{ bytes}$$

Recalculating the amount of bandwidth required to support this configuration yields the following:

$$\text{CIR} = \text{Bc} / \text{Tc} = 576,000 \text{ bytes} / 25 \text{ ms} = 184.32 \text{ Mbps}$$

This translates to approximately 20 Mbps per CTS-1000. Finally, if this amount of bandwidth is still too high, the design engineer may redesign the network further to try and buffer the burst over two time constants, if the router has sufficient capacity to handle the burst.

Multipoint TelePresence over Converged MPLS Circuits

When deploying multipoint TelePresence over a WAN infrastructure consisting of a converged MPLS circuit with an Ethernet handoff to the service provider network, a hierarchical QoS policy may need to be applied to the CE edge device. This policy map not only shapes the overall rate of traffic to meet the contracted sub-line rate, but also polices TelePresence traffic to a particular rate within the overall contracted rate.

As previously stated in [Chapter 6, “Branch QoS Design for TelePresence,”](#) policing is not a desirable thing for TelePresence traffic. Any packets discarded by a policer degrade the TelePresence experience. However, policing may be required to limit the amount of TelePresence and/or real-time traffic sent to the service provider network when mechanisms such as CAC are not available. Also, policing may be required on certain platforms when placing traffic into an LLQ configuration, as is recommended for TelePresence.

An example of an egress hierarchical policy map applied to the interface of a CE device is shown in [Example 12-6](#).

Example 12-6 Egress CE Policy Map Which Enforces Shaping and Policing

```

policy-map EGRESS-TOP-MAP
  class class-default
    shape average 50000000 1250000 1822000
    service-policy EGRESS-CE-MODEL
policy-map EGRESS-CE-MODEL
  class VOICE
    police 1000000
    priority percent 2
  class TELEPRESENCE
    set dscp cs5
    priority percent 33          ! Places TP in LLQ
    police cir 16500000 bc 384000! Policed for 3 CTS-1000s
  class NETWORK-CONTROL
    bandwidth percent 5
  class CALL-SIGNALING
    bandwidth percent 5
  class OAM
    bandwidth percent 5
  class MULTIMEDIA-CONFERENCING
    set dscp af21
    bandwidth percent 10
  class MULTIMEDIA-STREAMING
    set dscp af21
    bandwidth percent 5
  class BROADCAST-VIDEO
    set dscp cs2
    bandwidth percent 5
  class LOW-LATENCY-DATA
    set dscp af31
    bandwidth percent 5
  class HIGH-THROUGHPUT-DATA
    bandwidth percent 4
  class SCAVENGER
    bandwidth percent 1
  class class-default
    bandwidth percent 20
    random-detect
  ...

```

The policy map is similar to that which would be applied to support a single TelePresence device in a point-to-point configuration, however it has been scaled to support multipoint TelePresence.

Policer Details

The average rate (CIR) and committed burst (Bc) of the policer within the TelePresence traffic class have been scaled for support of multipoint TelePresence. Scaling of the average rate can be accomplished by utilizing 5.5 Mbps per CTS-1000 and 15 Mbps per CTS-3000. Keeping in mind that it is recommended to limit real-time traffic to approximately 1/3 of the bandwidth of any link, the configuration example above is designed to support a maximum of 3 CTS-1000s * 5.5 Mbps per CTS-1000 = 16.5 Mbps in the TelePresence traffic class. This is approximately 33% of the overall 50 Mbps shaped rate. In this example, if more than 3 CTS-1000s need to be supported in an LLQ configuration, a higher contracted rate to the service provider would be recommended.

Scaling the committed burst (Bc) of the policer within the TelePresence traffic class can also be accomplished by utilizing the burst estimates of 128 Kbytes per CTS-1000 and 256 Kbytes per CTS-3000. In the example above, which supports 3 CTS-1000s, this translates to 3 CTS-1000s * 128 Kbytes per CTS-1000 = 384 Kbytes.

Shaper Details

The configuration of the shaper follows the guidance presented in [Design Option #1](#) for support of multipoint TelePresence over a dedicated MPLS link. The average rate (CIR) is set to the overall contracted rate from the service provider. The time constant (Tc) is selected to be relatively low, 25 ms in the example above. This results in a relatively small committed burst (Bc) of 1,250,000 bits or 156,250 bytes. In order to accommodate bursts from the policer which has been configured for 384,000 bytes, an excess burst of $384,000 - 156,250 = 227,750$ bytes or 1,822,000 bits is configured.

Tested Results

In reality, both the shaper excess burst (Be) and the policer committed burst (Bc) may be tuned down somewhat. [Table 12-4](#) summarizes the actual test results from ESE testing of multipoint TelePresence over a shared 50 Mps, 100 Mbps, and 150 Mbps Ethernet handoff to an MPLS network using the methodology discussed above.

Table 12-4 ESE Tested Results for Multipoint over Shared MPLS Circuits

Shaper CIR	Shaper Bc	Shaper Be	Total Shaper Burst	Policer CIR	Policer Bc	Maximum CTS Table Segments	Output Hold Queue
50 Mbps	156.25 Kbytes	156.25 Kbytes	312.5 Kbytes	16.5 Mbps	192 Kbytes	3 CTS-1000s	40
100 Mbps	312.5 Kbytes	365.5 Kbytes	678 Kbytes	33 Mbps	384 Kbytes	6 CTS-1000s	80
150 Mbps	467.75 Kbytes	800 Kbytes (See discussion)	1,268.75 Kbytes	49.5 Mbps	805 Kbytes	9 CTS-1000s	4096 (See discussion)

As can be seen from [Table 12-4](#), testing has indicated that the policer committed burst (Bc) can normally be tuned below the 128 Kbyte per CTS-1000 recommendation. However, the shaper excess burst has to also account for the other CBWFQ classes of traffic, Voice LLQ class traffic, as well as the TelePresence LLQ class traffic. Therefore the overall burst capability of the shaper is above the policer burst capability in a converged circuit configuration.



Note

The Be value of 800 Kbytes for the 150 Mbps test was not tuned down to a minimum. It is believed that a value of 436 Kbytes would have worked as well.

Output Hold Queue and TxRing Tuning

In a converged configuration, the output hold queue will likely need to be tuned above the default of 40 packets for a Gigabit Ethernet interface. [Table 12-4](#) presents the output hold queue sizes required for the 50 and 100 Mbps tests. However, the results are highly dependant on the overall traffic on the interface. The output hold queue for the 150 Mbps test was simply increased to its maximum value of 4,096 to see if there were any detrimental effects on TelePresence quality by simply increasing the output hold queue to its maximum value. No adverse effects were observed. However, since TelePresence was placed in the LLQ, no adverse effects from queueing delays were expected either.

Platform and Linecard Test Results and Recommendations

This section summarizes the recommendations for support of multipoint TelePresence over various switch linecards and platforms within the LAN; as well as various switch and router platforms within the WAN. All recommendations apply to converged network designs.



Note

The platforms and linecards listed in the tables below have been tested by ESE. Just because a platform, linecard, or particular WAN media has not been tested, does not necessarily mean that it will not work in a multipoint TelePresence design.

Multipoint TelePresence over Switch Linecards and Platforms within the LAN

The reader is encouraged to review [Chapter 5, “Campus QoS Design for TelePresence”](#) for a detailed discussion of the queuing structure and per-port buffering of each switch platform and/or linecard, as well as detailed recommendations for configuring each switch platform and/or linecard to support TelePresence. The configuration recommendations presented in [Chapter 5, “Campus QoS Design for TelePresence”](#) apply to multipoint TelePresence as well.

[Table 12-5](#) and [Table 12-6](#) summarize the results from the platform and linecard testing for support of multipoint TelePresence, both for use at the LAN head-end (connected to the CTMS) and for the LAN remote-site (connected to CTS endpoints).



Note

ESE testing of all linecards and/or platforms at the head-end was performed up to 24 table segments or one-half the current capacity of a fully loaded CTMS.

Table 12-5 *Modular Switch Platforms*

Platform	Supervisor	Linecard	Connected to CTMS	Connected to CTS Endpoints
Catalyst 6500	Sup-32	WS-X6148A-GE-TX	Tested and Caveats (See discussion)	Tested and Recommended
	Sup-32	WS-X6548-GE-TX	Tested and Caveats (See discussion)	Tested and Recommended
	Sup-720	WS-X6748-GE-TX	Tested and Recommended	Tested and Recommended
Catalyst 4500	SupII+	WS-X4548-GB-RJ45V	Tested and Caveats (See discussion)	Tested and Recommended
	SupII+	WS-X4448-GB-RJ45	Tested and Caveats (See discussion)	Recommended

Table 12-6 Fixed Configuration Switch Platforms

Platform	Connected to CTMS	Connected to CTS Endpoints
Catalyst 4948-10GE	Tested and Recommended	Tested and Recommended
Catalyst 3750G-48PS-E	Tested and Caveats (See discussion)	Tested and Caveats (See discussion)

Both the WS-X6748-GE-TX linecard of the Catalyst 6500 platform and the Catalyst 4948-10GE platform are recommended to connect either a CTMS at the head-end of a multipoint TelePresence deployment or multiple CTS endpoints at the remote-end of a multipoint deployment, in a converged network design.

The WS-X6548-GE-TX and WS-X6148A-GE-TX linecards of the Catalyst 6500 platform are recommended to connect multiple CTS endpoints at the remote-end of a multipoint deployment in a converged network design. However, since these linecards are 8:1 oversubscribed, they are not recommended to connect a CTMS at the head-end of a multipoint deployment unless the network administrator dedicates the entire range of eight ports corresponding to the 1 Gbps uplink to the switch fabric shared by these eight ports. Port ranges are 1-8, 9-16, 17-24, 18-32, 33-40, or 41-48.

The WS-X4548-GB-RJ45V and WS-X4448-GB-RJ48 linecards of the Catalyst 4500 platform are recommended to connect multiple CTS endpoints at the remote-end of a multipoint deployment in a converged network design. However, since these linecards are again 8:1 oversubscribed, they are not recommended to connect a CTMS at the head-end of a multipoint deployment unless the network administrator dedicates the entire range of eight ports corresponding to the 1 Gbps uplink to the switch fabric shared by these eight ports. Port ranges are 1-8, 9-16, 17-24, 18-32, 33-40, or 41-48.

The Catalyst 3750G-48PS-E is recommended to connect either a CTMS at the head-end of a multipoint TelePresence deployment or multiple CTS endpoints at the remote-end of a multipoint deployment, in a converged network design. However, testing was only performed in a non-stacked configuration. Due to potential issues regarding oversubscription of the dual 16 Gbps ring backplane, it is not recommended to use the Catalyst 3750G-48PS-E in a stacked configuration either for connectivity to the CTMS or CTS endpoints.

Table 12-7 summarize the results from supervisor and linecard ports used as uplinks between switches when supporting multipoint TelePresence.

Table 12-7 Uplink Ports between Switches

Supervisor or Linecard Port	Results
Catalyst 6500 Sup-32 10GE Port	Tested and Recommended
Catalyst 6500 Sup-32 1 GE Port	Tested and Caveats (See discussion)
WS-X6704-10GE	Tested and Recommended
WS-X6708-10GE	Tested and Recommended
Catalyst 6500 Sup-720 GE Port	Tested and Caveats (See discussion)
WS-X6748-GE-TX	Tested and Caveats (See discussion)
Catalyst 4500 SupII+ 10 GE Port	Tested and Recommended
Catalyst 4948-10GE Platform 10 GE Port	Tested and Recommended
Catalyst 3750G-48PS-E Platform GE Port	Tested and Caveats (See discussion)

The 10 GE uplink ports on the Catalyst 6500 Sup-32 and Catalyst 4500 SupII+ are both recommended for uplinks between switches. Likewise ports on the Catalyst 6500 WS-X6704-10GE and WS-X6708-10GE linecards are recommended for use as uplink ports between switches. Finally the 10 GE uplink ports on the Catalyst 4948-10GE platform are also recommended as uplinks between switches.

The 1 GE uplink ports on the Catalyst 6500 Sup-32 and Sup-720, 1 GE ports on the WS-X6748-GE-TX, and 1 GE ports on the Catalyst 3750G-48PS-E platform can also be used as uplink ports between switches. However, caution should be applied when utilizing 1 gigabit Ethernet ports for connectivity to endpoint devices (including TelePresence CTS units) and only providing a 1 gigabit Ethernet uplink port between switches. The amount of oversubscription of the uplink port can be high, resulting in congestion. Multiple gigabit Ethernet ports can be implemented in a Gigabit EtherChannel configuration with some caveats.

[Example 12-7](#) shows part of the configuration of a Catalyst 3750G switch for a Gigabit EtherChannel configuration consisting of three ports.

Example 12-7 Catalyst 3750G Gigabit EtherChannel Configuration

```

!
port-channel load-balance src-dst-ip
!
!
interface Port-channel1
 no switchport
 ip address 10.16.3.1 255.255.255.0
 load-interval 30
!
!
interface GigabitEthernet1/0/33
 no switchport
 no ip address
 load-interval 30
 shutdown
 srr-queue bandwidth share 1 30 35 5
 priority-queue out
 mls qos trust dscp
 channel-group 1 mode auto
!
interface GigabitEthernet1/0/34
 no switchport
 no ip address
 load-interval 30
 shutdown
 srr-queue bandwidth share 1 30 35 5
 priority-queue out
 mls qos trust dscp
 channel-group 1 mode auto
!
interface GigabitEthernet1/0/35
 no switchport
 no ip address
 load-interval 30
 srr-queue bandwidth share 1 30 35 5
 priority-queue out
 mls qos trust dscp
 channel-group 1 mode auto
!

```

Depending upon the platform, load balancing on the port-channel group can be done in various ways—by source IP address, by destination IP address, by source MAC address, by destination MAC address, by source and destination IP address, or by source and destination MAC address.

There are two main issues with regard to the use of EtherChannel technology in this configuration. First, EtherChannel technology does not take into account the bandwidth of each flow. Instead, it relies on the statistical probability that, given a large number of flows of relatively equal bandwidths, the load is equally distributed across the links of the port-channel group. Since the CTMS represents a single IP address and a single MAC address, load balancing traffic based on source IP address or source MAC address on the port channel is not considered to be effective. All outbound traffic from the CTMS would take a single path. This could result in a lopsided distribution of traffic, in which one link could have more than 1 Gbps of traffic, since video flows are variable in nature, while other links have relatively little traffic. Depending upon the configuration of the network past the EtherChannel group, the destination MAC address in all likelihood corresponds to a router port. Therefore, load-balancing based on the destination MAC address or source and destination MAC address is also not considered to be effective. Load balancing based on the destination IP address, or source and destination IP address, are considered to be effective for this type of design, since the CTMS sends traffic to multiple CTS endpoint IP addresses. Finally, it should be noted that each source and destination IP address represents a video connection between the CTMS and one CTS endpoint. Therefore there is no chance of out-of-sequence packets being seen by the CTMS or CTS units in this configuration during normal operation.

The second issue with EtherChannel technology is that it does not take into account any QoS configuration on the individual Gigabit Ethernet links. Again, it relies on the statistical probability that, given a large number of flows of with different QoS markings, the load of those individual flows is equally distributed across the links of the port-channel group. Given a failover situation in which one of the links of an EtherChannel group fails, the sessions crossing that link would be re-allocated across the remaining links. Since EtherChannel technology has no awareness of QoS markings, it could easily re-allocate more real-time flows across any one of the links than the link is configured to accommodate. This would result in degraded real-time services such as TelePresence and voice, although there is sufficient real-time bandwidth within the EtherChannel group to easily accommodate the real-time traffic. This can happen in a non-failover situation as well. Therefore, caution should be used when deciding to utilize EtherChannel technology versus a single higher-speed uplink port.

Multipoint TelePresence over WAN Switch and Router Platforms

Table 12-8 summarizes the recommendations for WAN router and switch platforms along with associated media tested for support of multipoint TelePresence in a converged network design.

Table 12-8 Platforms and WAN Media

Platform	WAN Hardware	WAN Media	Results
Cisco 2821 ISR	Integrated Gigabit Ethernet Interface	MPLS Handoff with 50 Mbps Shaped Fast Ethernet	Tested and Caveats (See discussion)
Cisco 3845 ISR	Integrated Gigabit Ethernet Interface	MPLS Handoff with 50 Mbps Shaped Gigabit Ethernet	Tested and Recommended
Cisco 3845 ISR	Integrated Gigabit Ethernet Interface	MPLS Handoff with 100 Mbps Shaped Gigabit Ethernet	Tested and Recommended
Cisco 3845 ISR	Integrated Gigabit Ethernet Interface	MPLS Handoff with 150 Mbps Shaped Gigabit Ethernet	Tested and Recommended (See Discussion)

Table 12-8 **Platforms and WAN Media**

Platform	WAN Hardware	WAN Media	Results
Cisco 7206VXR with NPE-G2	Integrated Gigabit Ethernet Interface	MPLS Handoff with 50 Mbps Shaped Gigabit Ethernet	Tested and Recommended
Cisco 7206VXR with NPE-G2	Integrated Gigabit Ethernet Interface	MPLS Handoff with 100 Mbps Shaped Gigabit Ethernet	Tested and Recommended
Cisco 7206VXR with NPE-G2	Integrated Gigabit Ethernet Interface	MPLS Handoff with 150 Mbps Shaped Gigabit Ethernet	Tested and Recommended
Cisco 7206VXR with NPE-G2	PA-T3+	Clear Channel T3	Tested and Recommended
Cisco 7206VXR with NPE-G3	PA-POS-OC3MM	OC-3 POS	Tested and Recommended
Catalyst 6506-E with Sup-32	7600-SIP-200 with SPA-4XT3/E3	Clear Channel E3	Tested and Recommended
Catalyst 6506-E with Sup-32	7600-SIP-400 with SPA-2XOC3-POS	OC-3 POS	Tested and Recommended
Catalyst 6506-E with Sup-32	7600-SIP-400 with SPA-1XOC12-POS	OC-12 POS	Tested and Recommended
Catalyst 6506-E with Sup-32	7600-SIP-400 with SPA-1XOC48-POS	OC-48 POS	Tested and Recommended

For all tests involving handoff to an MPLS network via either a 50 Mbps, 100 Mbps, or 150 Mbps Ethernet link, an integrated Gigabit Ethernet interface on the Cisco 2821 ISR, Cisco 3845 ISR, or Cisco 7206VXR was utilized.

Significant input errors (overruns) were observed on the interface of the Cisco 2821 ISR at 1 Gbps physical line speeds during the testing when both background traffic and TelePresence traffic were applied. Based on these results, it is not recommended to utilize the Cisco 2821 ISR for a 50 Mbps shaped Ethernet handoff to a service provider MPLS network when the service provider physical line rate is 1 Gbps in a converged network design. Instead it is recommended to use the Cisco 2821 ISR for a 50 Mbps shaped Ethernet handoff to the service provider MPLS network only when the service provider physical line rate is 100 Mbps. Due to these overrun errors, the Cisco 2821 is also not recommended for Ethernet handoff to an MPLS network above a 50 Mbps shaped rate in a converged network design.

The Cisco 3845 ISR experienced a small number of input errors (overruns) during testing at 150 Mbps shaped Ethernet handoff to the MPLS network. However, the input errors were well below the 0.1% threshold for the drop rate of TelePresence. Therefore the Cisco 3845 ISR is recommended for shaped 50 Mbps, 100 Mbps, and 150 Mbps Ethernet MPLS handoff to a service provider network.

The Cisco 7206VXR experienced no input errors during any of the tests with a shaped Ethernet handoff to the MPLS network. Therefore the Cisco 7206VXR is recommended for shaped 50 Mbps, 100 Mbps, and 150 Mbps Ethernet MPLS handoff to a service provider network as well.

