Design Considerations for Cisco TelePresence over a PIN Architecture

At-A-Glance

Benefits of Cisco TelePresence

Cisco TelePresence enables organizations to realize the benefits of business transformation through more effective meetings. Increased productivity through reduced time in making decisions can accelerate time-to-market and implementation of corporate strategies. The collaborative nature of TelePresence helps facilitate cohesive internal teams and speed the integration of employees into the corporate culture. Customer intimacy can also be greatly enhanced though frequent face-to-face meetings. Enterprise organizations can scale and optimize the use of expertise throughout the organization, resulting in a more nimble business, while simultaneously reducing travel expenses. Reduced travel also helps maintain a positive corporate image by decreasing an organizations carbon footprint—as social and government pressures mount to address the growing concern of global warming.

Benefits of Enabling the Network for Cisco TelePresence

In order to maximize the benefits of business transformation through Cisco TelePresence, organizations must enable their network infrastructure to support such technology. This helps to minimize or eliminate disruptions of service and periods of degraded quality that can detract from the productivity and effectiveness of meetings. Enabling the network to support Cisco TelePresence also helps minimize or eliminate disruptions to other business-critical applications as demand for collaborative technologies increases.

The Medianet Framework

The medianet (Figure 1) provides a network-wide framework for enterprise network designs with video services in mind. It is designed for customers that are considering various types of video applications in use today or that might be used in the future—such as TelePresence, desktop video collaboration, digital media systems (DMS), and IP video surveillance (IPVS). The medianet starts with an end-to-end network infrastructure designed and built to achieve high availability. It also defines the sets of services that the network can provide to video applications, including the following:

- Access services—Provide access control, identity, mobility, and location services to video clients.
- Transport services—Provide packet delivery, ensuring service levels with Quality of Service (QoS) and delivery optimization.
- Bridging services—Include transcoding, conferencing, and recording services.
- Storage services—Include content capture, storage, retrieval, distribution, and management services.
- Session control services—Include signaling and control to setup and tear-down sessions, as well as gateway services.

Figure 1 The Medianet Framework

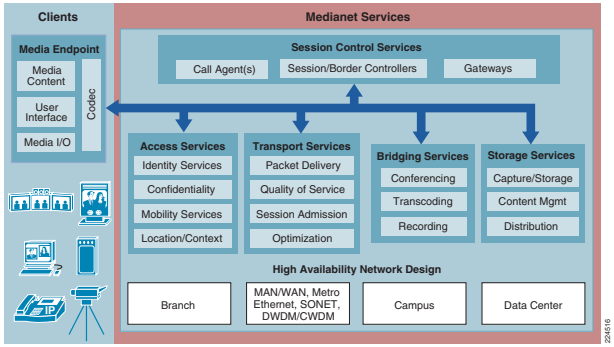

These services are applied to the network infrastructure components and technologies defined within the ESE Places-in-the-Network (PIN)—branch, WAN, campus, and data center.

Components of Cisco TelePresence

The components of the Cisco TelePresence solution consist of the following (Figure 3):

- Cisco TelePresence Systems (CTS) Endpoints—Include the CTS-3200, designed for large meetings of up to 18 participants on each side; the CTS-3000, designed medium-sized meetings of up to six participants on each side; the CTS-1000, designed for small meetings of up to two participants on each side; and the CTS-500, designed for personal conference rooms and public areas.
- Cisco TelePresence Multipoint Switch (CTMS)—Provides switching of the audio and video streams in a multipoint TelePresence meeting.
- Cisco Unified Communications Manager—Provides call signaling services for TelePresence end points.
- Cisco TelePresence Manager—Provides management and call scheduling services for the TelePresence endpoints.
- Meeting Scheduling Components—Include Microsoft Outlook or IBM Notes clients, LDAP/Microsoft Active Directory, and Microsoft Exchange or IBM Domino servers which provide scheduling services for TelePresence meetings.

SLA Determination

In order to properly enable Cisco TelePresence within a network, it is essential to first characterize the key service-level agreement (SLA) parameters to be met. These parameters include the following:

- Bandwidth—The bandwidth utilized per Cisco TelePresence endpoint varies based on factors that include the model deployed, the desired video resolution, interoperability with legacy video conferencing systems, and whether high or low speed auxiliary video input is deployed for the document camera or slide presentation. For example, when deploying 1080p best video resolution with high speed auxiliary video input and interoperability, the bandwidth requirement can be as high as 20.4 Mbps for a CTS-3200 and CTS-3000, or 10.8 Mbps for a CTS-1000 and CTS-500.
- Packet Loss—Due to the high amount of compression and motion-compensated prediction utilized by TelePresence video codecs, even a small amount of packet loss can result in visible degradation of the video quality. The packet loss target for Cisco TelePresence should be below 0.05 percent on the network.
- Jitter—All packets that comprise a frame of video must be delivered to the TelePresence end point before the replay buffer is depleted. Otherwise degradation of the video quality can occur. The peak-to-peak jitter target for Cisco TelePresence is under10 msec.
- Latency—The one-way network latency target for Cisco TelePresence is below 150 msec. This does not include latency induced by encoding and decoding at the CTS endpoints.
- Bursts—TelePresence video appears as a variable bit-rate stream with somewhat random bursts. Bursts are generally the result of a reference frame—also known as an I-Frame or IDR. These occur periodically throughout the video transmission, during speaker transitions in a multipoint meeting, after packet loss has occurred, or during slide transitions within a presentation. A general rule-of-thumb

is that the network infrastructure must account for bursts up to 64 KB per TelePresence video input (camera or auxiliary) when I-Frames occur. Video quality will degrade if packets associated with bursts are dropped or delayed excessively by any traffic policing or shaping within the network. Buffer capacity within LAN switches must also be sufficient to handle such bursts.

Branch PIN Design Considerations

Because CTS endpoints are managed by Cisco Unified Communications Manager in the same manner that Cisco IP Phones are managed, TelePresence end points can be trusted from a QoS perspective. TelePresence ports on the branch LAN switch, such as the Cisco Catalyst 3750E Series, need only to be configured to trust either the Differentiated Services Code Point (DSCP) or Class of Service (COS) value of inbound packets. These switches also provide both ingress and egress queueing to ensure that TelePresence traffic is allocated the appropriate bandwidth on the branch LAN.

High-availability designs within the branch include the use of redundant Cisco 3800 or Cisco 2800 Series Integrated Services Routers (ISR) with circuits to separate service provider networks. This provides a resilient path for audio and video media in the event of a circuit failure. It also provides a resilient path for call signaling to the Cisco Unified Communications Manager cluster. Optimization technologies, such as Cisco's Performance Routing (PfR) may be used to ensure that TelePresence media takes the optimal path through the WAN based on SLA requirements.

Cisco 3800 or Cisco 2800 Series ISRs also provide software defined QoS for multiple service classes. This allows the network administrator to separate TelePresence from other video traffic classes and even business critical traffic—effectively protecting each from affecting the other. The placement of TelePresence traffic in a dual-Low Latency Queuing (LLQ) configuration—along with voice traffic—can help reduce end-to-end jitter across the network.

WAN PIN Design Considerations

When Cisco TelePresence is deployed across a service provider (SP) managed WAN infrastructure—such as a Multiprotocol Label Switching (MPLS) network—the network administrator must ensure that the SLA parameters defined are met by the service provider. In addition, the enterprise might define more QoS service classes than available within the SP MPLS network. For example, the Cisco 12-class QoS model (based on IETF RFC 4594) defines a real-time interactive service class, which might be used for applications such as Cisco TelePresence —marking the audio and video media as CS4. However, the SP network might only implement four service classes. In such cases, the network administrator must take into account the mapping of the real-time interactive service class into the appropriate MPLS service class as traffic enters the MPLS network. Traffic might also need to be identified and re-mapped back to the appropriate service class as it exits the MPLS network.

When TelePresence is deployed across an enterprise-managed WAN infrastructure, the network administrator is responsible for ensuring SLA requirements are met. Jitter can be the result of long queue lengths, indicating that additional bandwidth might be required. In some cases parts of the network design might need to be modified to reduce the number of hops, which can potentially add latency and jitter. The choice of routing protocol implemented across the WAN and tuning of the timers should be done with rapid convergence in mind. At the WAN aggregation point, high-availability features—including built-in hardware and processor redundancy, In-Service Software Upgrade (ISSU) and nonstop forwarding (NSF) with stateful switchover (SSO)—can be deployed within Cisco Aggregation Services Router (ASR) platforms.

Figure 2 Cisco TelePresence PIN Overlay

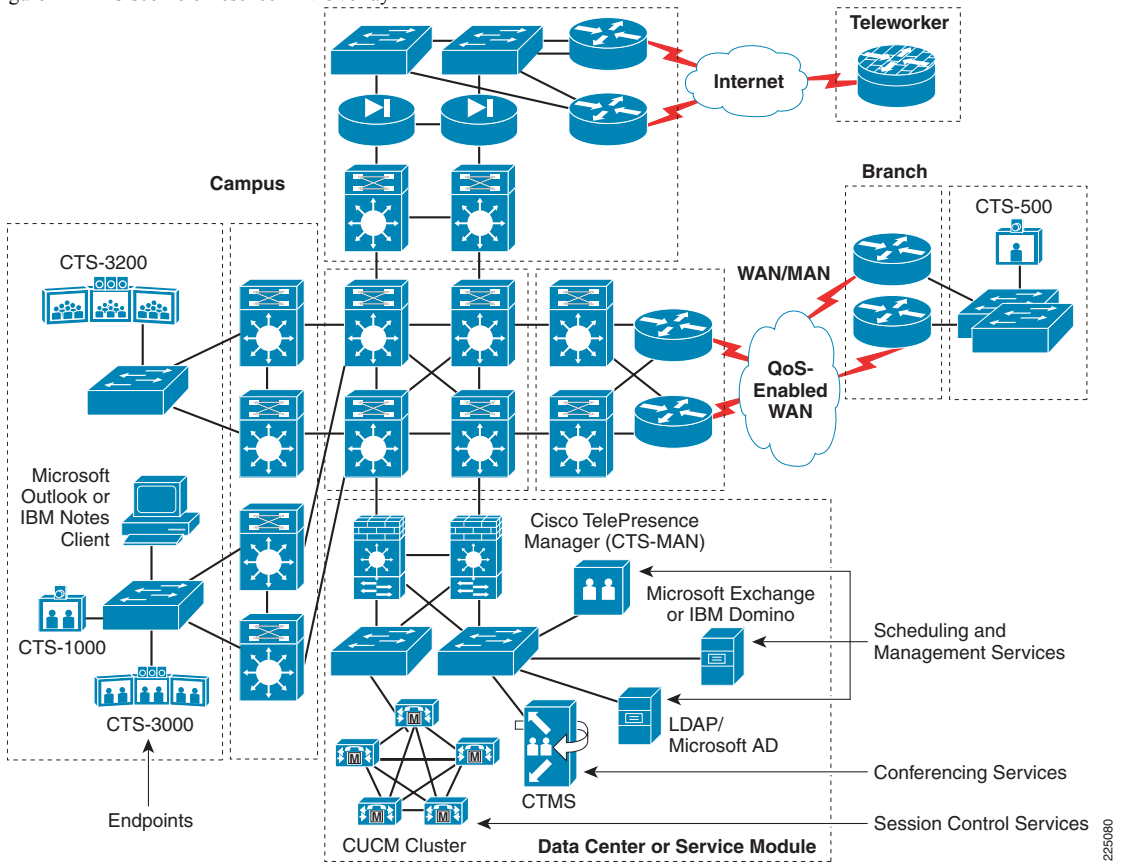

Campus PIN Design Considerations

A hierarchical campus design consisting of core, distribution, and access layers—utilizing redundant one and 10 Gigabit Ethernet links—provides bandwidth and high availability (HA) for Cisco TelePresence. Technologies such as Gateway Load Balancing Protocol (GLBP), Hot Standby Routing Protocol (HSRP), and Virtual Router Redundancy Protocol (VRRP) can provide rapid convergence times for video applications operating across the LAN. Additional HA features—including built-in hardware and processor redundancy, in-service software upgrade (ISSU), nonstop forwarding (NSF), and stateful switchover (SSO) can be deployed within Cisco Catalyst 6500 Series platforms.

Finally, as with the branch PIN, the access Cisco Catalyst 6500, Catalyst 4500, or Catalyst 3750E switch within the campus can simply trust the QoS marking from the TelePresence endpoint—as well as provide ingress and egress queueing—to ensure TelePresence traffic is allocated the appropriate

bandwidth on the campus LAN. Additional mechanisms, such as the Cisco Catalyst 6500 Programmable Intelligent Services Accelerator (PISA), allow deep-packet inspection, providing granular classification of TelePresence traffic as it enters the campus LAN.

Summary

Organizations are increasingly turning to collaborative technologies, such as Cisco TelePresence, to transform businesses and enhance business productivity. However the increasing demand for collaborative tools places increased strain on networking resources. The medianet provides a framework for enterprise organizations to design networks to support multiple video applications in use today or that might be deployed in the future. In order to properly enable Cisco TelePresence within a network, it is essential to first characterize key SLA parameter requirements—such as bandwidth, packet loss, latency, jitter, and bursts. Once these SLA parameters are understood, the network design engineer can identify critical services—such as QoS and high availability—that must be provided to support Cisco TelePresence. These services are then applied to the various PIN architectures of the branch, WAN, and campus in order to minimize disruptions to TelePresence sessions—as well as to prevent disruptions to business critical applications that could result from excessive video traffic.

Figure 3 TelePresence Solution

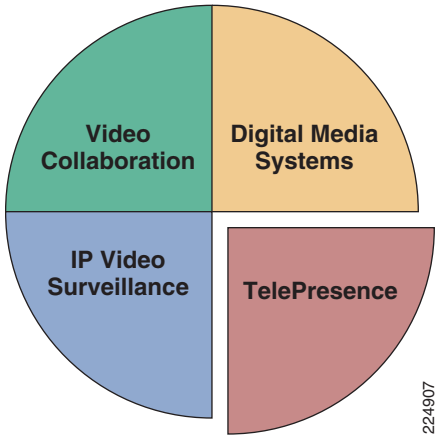

For more information, refer to the following URL:
http://www.cisco.com/go/designzone

For PIN architecture design considerations regarding Cisco Multipoint TelePresence, refer to Design Considerations for Cisco Multipoint TelePresence over a PIN Architecture At-a-Glance found at the above URL.