Design Considerations for Desktop Video Collaboration over a PIN Architecture—Solution Overview

March 30, 2009

Executive Summary

Video is exploding onto corporate networks, exponentially increasing bandwidth utilization and radically shifting traffic patterns. The development of newer video codecs, faster PC processors, and the emergence of web conferencing have converged to create a dramatic increase in the use of desktop video collaboration. Enterprise organizations are increasingly turning to collaborative technologies to increase productivity, scale expertise, and reduce travel costs. In order to take a proactive approach to the growth of video, the network administrator must understand the various services that a medianet can provide to help facilitate the convergence of desktop video collaboration onto the IP network infrastructure. This solution document provides a high level summary of the various services that a medianet provides for desktop video collaboration, based on Enterprise Solutions Engineering (ESE) Places-in-the-Network (PIN) architectures.

Desktop Video Collaboration Key Drivers

Although desktop video conferencing has been available for nearly a decade, recent advances in technology are driving a resurgence of interest in its use. These advancements include new video codec standards, advancements in personal computer hardware, and the integration of desktop video conferencing within collaborative software suites. Because of this integration, the term desktop video collaboration will be used interchangeably with desktop video conferencing within this document—in order to highlight the fact that video is rapidly becoming one component of a growing suite of collaborative applications that enhance business productivity.

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Video Codec Standards Development

Early video codecs were typically designed for room-based video conferencing systems that operate over relatively high bandwidths. When applied to desktop video collaboration, video frame rates were often decreased and smaller video formats utilized—in order to fit the available bandwidth. The end result was often postage stamp-sized video with poor resolution and unnatural motion. However, the recent development of the H.264 video codec standard has begun to allow higher quality video at lower bandwidths for a given video resolution. This combination of superior video quality at lower bandwidth utilization per call has made it more feasible to support desktop video collaboration over a wider range of IP networking infrastructures.

Integration of Video Support into PCs

The support of the newer H.264 video codec standard within applications running on generic personal computers—along with the steady increase in processor speeds and the integration of webcams into laptops—has helped to drive the growth of desktop video collaboration. The net result is that more and more video conferencing-capable endpoints are being deployed across enterprise networks. This is also driving an increased use of available bandwidth for desktop video collaboration across enterprise networks.

Integration with Collaborative Software Applications

The overall trend in desktop video conferencing endpoints has been a steady evolution in functionality, beginning with standalone hardware-based video phones, then to separate video-conferencing software applications running on personal computers, and finally to integrated collaborative software suites that bundle text messaging, voice mail access, and presence—along with audio and video conferencing capabilities. Another recent direction has been the integration of desktop video conferencing capabilities into web conferencing which again offers text messaging, presence, recording capabilities, desktop/application sharing capabilities, and so on—all targeted for multipoint meeting environments.

This steady evolution has resulted in a greater number of desktop video collaboration-enabled endpoints on the network. Today, any PC with a web browser and webcam can potentially participate in desktop video collaboration. Furthermore, whereas desktop video conferencing was at one time primarily a point-to-point application, it is rapidly becoming a multipoint application with the development of web conferencing

Business Benefits of Desktop Video Collaboration

The adoption of desktop video collaboration technology has enabled enterprises to capture the business benefits of increased productivity, through more effective meetings. This can translate to more rapid implementation of corporate strategy and faster time-to-market through increased decision making resulting from those meetings.

As globalization continues to increase, enterprise organizations are often challenged with developing and sustaining expertise. Desktop video collaboration offers an additional means of scaling the expertise of key personnel around the world—without them ever having to leave their desks. In addition, enterprises can benefit from reduced travel expense and increased productivity due to reduced downtime associated with business travel.

Finally, desktop video collaboration offers a means for enterprises to maintain a positive corporate image by reducing carbon emissions through decreased travel as social and governmental pressures mount to address the growing concern of global warming.

Benefits of Enabling the Network for Video Collaboration

In order to maximize the business benefits of increased productivity through the use collaborative tools (including desktop video conferencing), enterprises must enable their network infrastructures to support such technology. Enabling the network infrastructure to support desktop video collaboration helps to minimize or eliminate any disruption of service or periods of degraded quality—both of which detract from the productivity and effectiveness of meetings. Furthermore, enabling the network to support desktop video collaboration also helps to minimize or eliminate disruptions to other business critical applications as network demands for collaborative tools continue to increase.

Desktop Video Collaboration Components

Desktop video collaboration can be divided into the following components: *endpoints*, *session controllers* (which provide session control services), and *conferencing servers* (which provide bridging services). Additional services might be deployed for meeting scheduling or might be integrated within conferencing servers.

Endpoints

Desktop video collaboration endpoints include hardware-based devices, custom-built software applications, and generic software applications (web browsers with plugins). Custom-built software may be standalone video conferencing applications or integrated with additional collaborative tools such as presence, text messaging, and application / desktop sharing. Generic software applications such as web browsers typically access web conferences which provide collaborative tools such as presence, text messaging, application/desktop sharing, and recording—in addition to audio and video conferencing support. Table 1 highlights some of the Cisco desktop video collaboration endpoints and their uses.

Table 1 Cisco Desktop Video Collaboration Endpoints

Endpoint	Type	Key Benefits
Cisco IP Phone 7985G	Hardware endpoint	Can be used for standalone point-to-point or multipoint audio and video conferencing. Can also be integrated with Cisco Unified MeetingPlace 7.0 or MeetingPlace Express 2.0 to provide audio and video conferencing support within a web conference. Separate PC is required for desktop/application sharing, text messaging/chat, presence, white boarding, meeting recording and so on within the MeetingPlace web conference.
Cisco Unified Video Advantage (CUVA) with IP Phone or IP Communicator	Hardware and software endpoint	Can be used for standalone point-to-point or multipoint audio and video conferencing. Can also be integrated with Cisco Unified MeetingPlace 7.0 or MeetingPlace Express 2.0 to provide audio and video conferencing support within a web conference. Desktop/application sharing, text messaging/chat, presence, white boarding, meeting recording, and so on within the MeetingPlace web conference is provided through the PC that also runs CUVA.
Cisco Unified Personal Communicator (CUPC) with soft phone	Software endpoint	Can be used for standalone point-to-point or multipoint audio and video conferencing. Also provides presence, text messaging, directory services, and voice mail access; integrated within the application suite. Can also be integrated with Cisco Unified MeetingPlace 7.0 or MeetingPlace Express 2.0 to provide audio and video conferencing support within a web conference. Desktop/application sharing, text messaging/chat, presence, white boarding, message recording, and so on within the MeetingPlace web conference is provided through the PC that also runs CUPC.
Cisco Unified Videoconferencing Manager (CUV-M) Desktop 5.6	Software endpoint	Can be used for multipoint audio and video conferencing (via webcam and QuickTime plug-in). Requires CUV-M, Cisco Unified Video (CUV) Desktop Server, and Multipoint Control Unit (MCU) in order to provide audio and video conferencing as well as presence, text messaging/chat, and H.239 data sharing.
Cisco WebEx web conferencing	Web browser (generic software) with webcam	Can be used for audio and video conferencing (via webcam) through a WebEx web conference. Practical applications send video through the browser session, but implement audio callback to a separate IP phone—or dial-in from a separate IP phone—for business quality audio. Desktop/application sharing, text messaging/chat, presence, white boarding, meeting recording, and so on provided within the WebEx web conference. Integration with Cisco Unified MeetingPlace 7.0 allows combined internally and externally hosted web conferences with audio.

Session Controllers

From the perspective of desktop video conferencing, *session controllers* refer to devices that handle the call signaling which initiate and terminate audio and video sessions. These include traditional VoIP call-signaling platforms, such as the Cisco Unified Communications Manager (CUCM)—as well as H.323 gatekeeper functionality (either integrated within Cisco IOS routers, or other platforms, such as the CUV-M).

It should also be noted that, from the perspective of collaborative applications integrated with desktop video conferencing, session controllers may also include control of data sessions such as presence, text messaging/chat, and desktop/application sharing. These are not discussed within this document.

Conferencing Servers

Conferencing servers provide support for multipoint desktop video conferencing—either as standalone calls, or within a web conference. Often referred to as Multipoint Control Units (MCU), they include devices that provide audio and video switching, and potentially transcoding and transrating services for desktop video conferencing endpoints that support different frame rates and video formats. Conferencing servers may also include scheduling functions—either accessed directly through web interfaces on the conferencing server or via integration with E-mail and calendaring systems (such as those found in Microsoft Exchange/Outlook or IBM Lotus Domino/Notes).

Table 2 lists the Cisco products that provide video conferencing server functionality.

Table 2 Cisco Desktop Video Collaboration Conferencing Servers

Conferencing Server	Key Benefits
Cisco Unified Video Conferencing (CUVC) 35xx Series	Provides support for scheduled and ad hoc multipoint desktop and room-based video conferencing with transcoding and transrating. Basic scheduling functionality provided via web interface when deployed as a standalone device.
Cisco Unified MeetingPlace 7.0	Integrated rich-media solution consisting of application servers, integration servers, and media servers that provide scheduled multipoint audio and video conferencing and web-conferencing services in medium-to-large sized enterprise deployments. Scheduling functionality is integrated with E-mail/calendaring systems, such as Microsoft Exchange/Outlook and IBM Lotus Domino/Notes. Supports integration with Cisco WebEx for combined internally and externally hosted audio, video, and web conferences.
Cisco Unified MeetingPlace Express 2.0	Provides reservationless and scheduled support for multipoint audio and video conferencing, as well as web conferencing services for small-to-medium sized enterprise deployments. Does not provide transcoding or transrating. Scheduling functionality integrated with E-mail and calendaring systems, such as Microsoft Exchange/Outlook.

Table 2 Cisco Desktop Video Collaboration Conferencing Servers

Conferencing Server	Key Benefits
Cisco Unified MeetingPlace Express VT 2.0	Provides support for ad hoc multipoint and audio and video conferencing for small enterprise deployments. Does not provide transcoding or transrating.
Cisco Unified Videoconferencing Manager (CUV-M) 5.6	For customers that do not need the full rich-media capabilities of Cisco Unified MeetingPlace, CUV-M can be deployed along with the CUVC 35xx Series MCUs in order to provide support for scheduled and ad hoc multipoint desktop audio and video conferencing through the CUV-M Desktop web application and webcam.

Deployment Models

As a result of the ongoing evolution of desktop video collaboration from standalone hardware to part of an integrated software suite, various deployment models currently exist. Cisco offers two deployment models discussed in the following sections:

- Cisco Unified Video Conferencing, page 6
- Cisco WebEx Web Conferencing, page 10
- Cisco MeetingPlace and WebEx Integration, page 10

Cisco Unified Video Conferencing

Cisco Unified Video Conferencing deployments are targeted for customers planning to deploy and support voice, video and web conferencing services entirely within their organizations. CUVC deployments can be implemented with or without integration with Cisco Unified MeetingPlace. When deployed without Cisco Unified MeetingPlace, CUVC deployments provide point-to-point and multipoint video conferencing with basic scheduling support. An example of this type of deployment is shown in Figure 1.

Figure 1 Cisco Unified Video Conferencing without MeetingPlace Integration

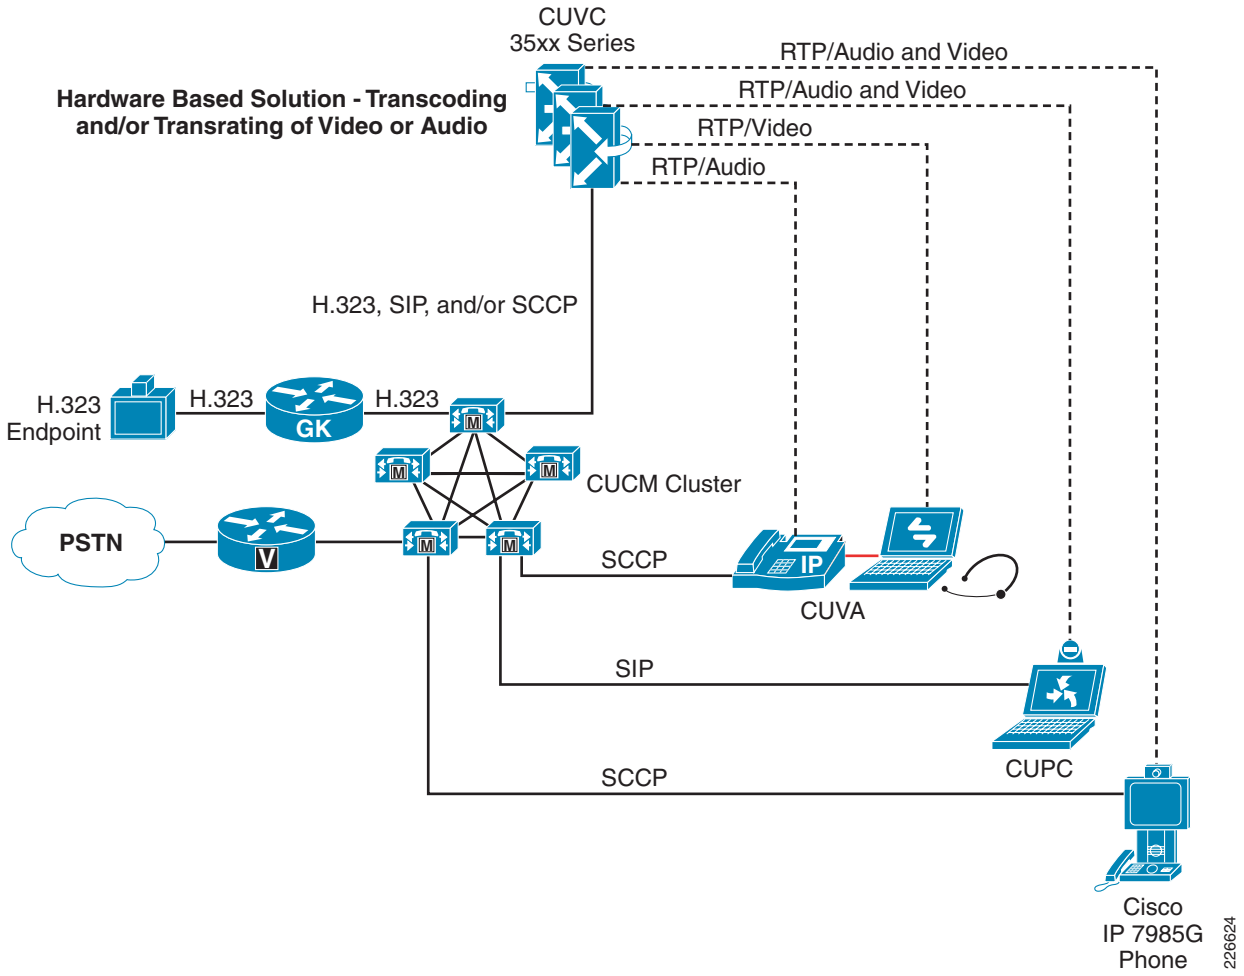

Cisco Unified Video Conferencing deployments rely on hardware endpoints, such as the Cisco IP Phone 7985G, or custom-built software endpoints, such as Cisco Unified Video Advantage (CUVA) with Cisco IP Communicator or IP Phone, and Cisco Unified Personal Communicator (CUPC). H.323 endpoints are also supported, typically via the gatekeeper functionality running on a Cisco IOS router platform. A Cisco Unified Communications Manager (CUCM) cluster provides resilient call signaling services. The Cisco Unified Video Conferencing (CUVC) 35xx Series functions as the MCU—providing multipoint services with transrating and transcoding support for the audio and video, if needed. This type of deployment provides no scheduling integration with E-mail/calendaring systems such as Microsoft Outlook/Exchange or IBM Lotus Domino/Notes and no web conferencing support.

An alternative for basic multipoint videoconferencing is the deployment of the Cisco Unified Videoconferencing Manager (CUV-M) 5.6 solution as shown in Figure 2.

Figure 2 Cisco Unified Videoconferencing Manager (CUV-M) 5.6 Solution

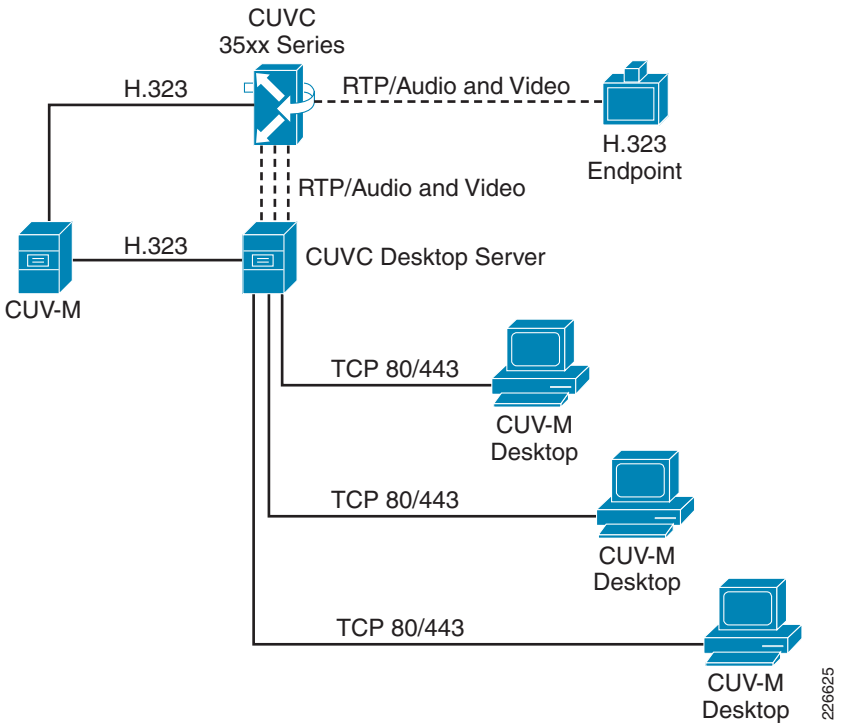

The CUV-M 5.6 solution does not utilize hardware endpoints, such as the Cisco IP Phone 7985G, or custom-built software endpoints, such as CUVA with IP Communicator or an IP Phone, or the Cisco Unified Personal Communicator (CUPC). These endpoints are designed to provide high quality audio and video. Instead, it relies on the CUV-M Desktop web application running on a PC for basic voice, video and data sharing capabilities. The CUV Desktop Server functions as an H.323 proxy between the CUV-M Desktop clients and the MCU with the CUV-M providing both scheduling/management functionality and H.323 gatekeeper functionality.

For customers that require full web conferencing and rich-media capabilities, the Cisco Unified Video Conferencing deployment can be integrated with Cisco Unified MeetingPlace 7.0—as shown in Figure 3.

Figure 3 Cisco Unified Video Conferencing with MeetingPlace 7.0 Integration

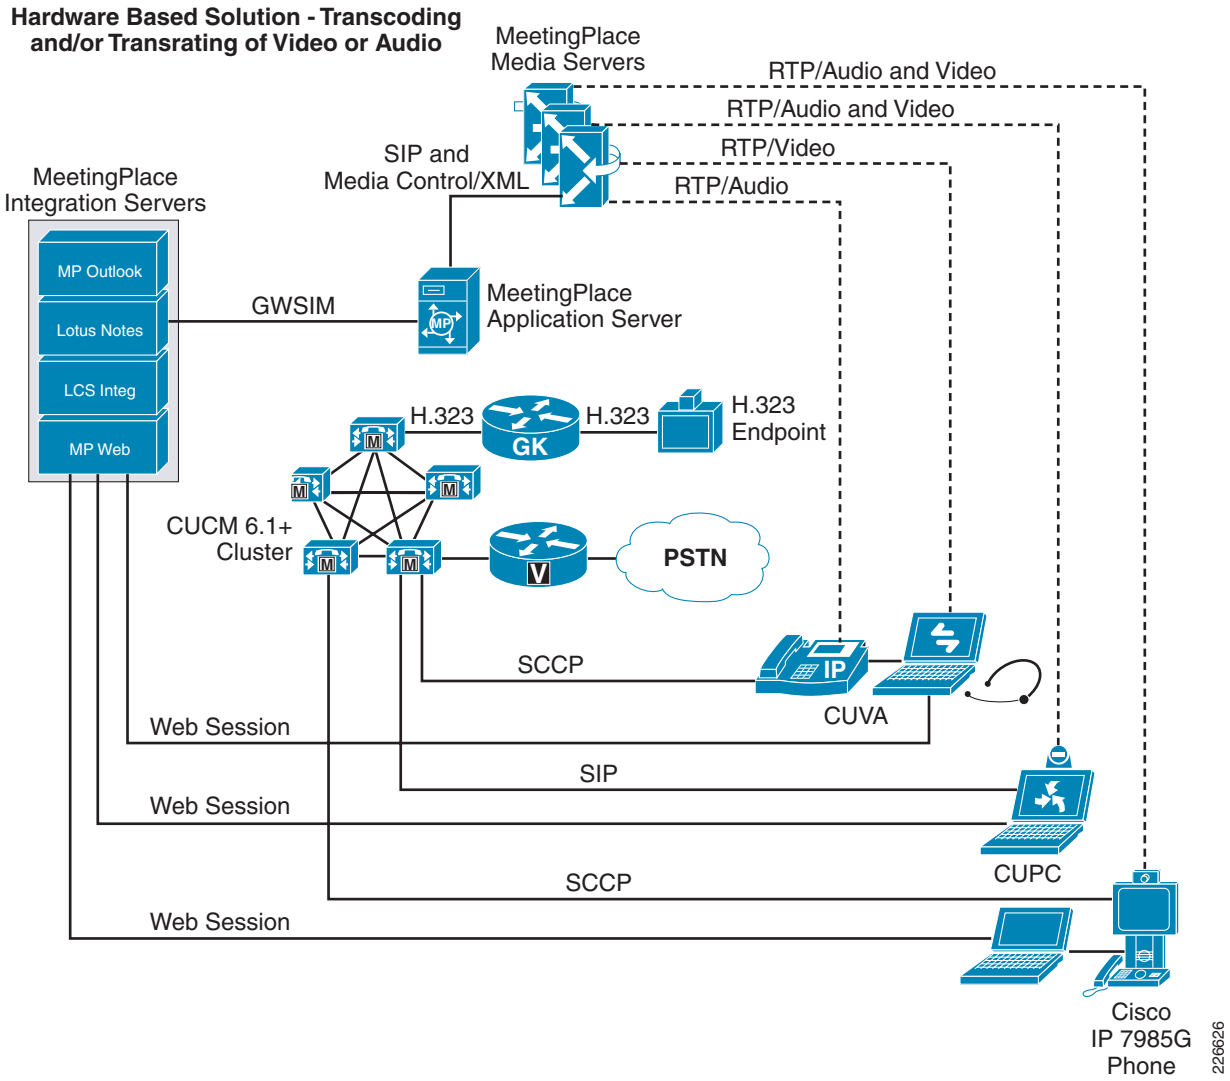

The Cisco Unified MeetingPlace 7.0 components consist of the following:

- *MeetingPlace Application Server*—Provides web-based administration and conferencing control for the MeetingPlace solution.
- *MeetingPlace Media Servers*—Provide audio and video MCU services (mixing, transcoding, and so on) for the MeetingPlace solution.
- *MeetingPlace Integration Services*—Provides services to integrate the MeetingPlace solution to other applications, such as Microsoft Outlook, IBM Lotus Domino, web conferencing, and so on. May be integrated together on one or more physical servers.

The Cisco Unified Video Conferencing with MeetingPlace 7.0 is targeted for large enterprise deployments. For medium-sized enterprises, Cisco Unified MeetingPlace Express 2.0 can be separately deployed. This type of deployment still provides scheduled audio, video, and web conferencing, but does not rely on the MeetingPlace Media Servers. Since audio and video switching is done in software within MeetingPlace Express, transcoding or transrating of the audio or video is not supported. However, for medium sized enterprise customers that implement a single type of desktop video conferencing endpoint, this might completely meet business requirements.

For very small deployments, Cisco offers Cisco Unified Video Conferencing with MeetingPlace Express VT 2.0. This type of deployment provides ad hoc audio and video conferencing. Again, because it does not rely on the MeetingPlace Media Servers, transcoding or transrating of the audio or video is not supported. However, for a very small enterprise that deploys a single type of desktop video conferencing endpoint, this might completely meet their business requirements as well.

Cisco WebEx Web Conferencing

Cisco WebEx web conferencing is targeted for customers who do not wish to deploy and support web conferencing services within their organizations. All conferencing services are located within the WebEx network. Customers simply schedule and access the WebEx network for web conferences on a fee-based model. In this model, desktop video conferencing support is offered through a PC webcam as part of an overall web conference. No specialized software or hardware client is used for desktop video conferencing. Instead, video is embedded within the web browser to the WebEx conference. In order to provide business quality audio, WebEx conferences provide an audio dial-back or dial-in function, which can be accessed by a separate IP phone or traditional telephone.

Cisco MeetingPlace and WebEx Integration

Cisco provides integration of MeetingPlace 7.0 with WebEx in order to support internally hosted and externally hosted web conferences with audio. In this model, the web conference component is hosted on WebEx servers within the WebEx MediaTone Network. Desktop video support is offered through a PC webcam with the video embedded within the web browser of the WebEx web conference. Integrated WebEx/MeetingPlace meetings can be scheduled either via MeetingPlace or via WebEx.

SLA Determination

From the perspective of enabling desktop video collaboration across the network infrastructure, the video media itself presents the greatest challenge. In order to properly enable desktop video collaboration within a network, it is essential to first characterize key service-level agreement (SLA) parameters to be met for the particular video component. The following list summarizes the key considerations in determining the appropriate SLA for a given implementation:

- *Bandwidth*—The amount of bandwidth utilized per desktop video conference call is typically at the discretion of the network administrator and based on the business requirements of the enterprise. Typical video rates can be as much as 1.5 Mbps for standard definition desktop video conferencing systems. High frame rates (up to 30 fps) and larger video formats (CIF and 4CIF) typically lead to better video quality, with the trade-off being higher bandwidth utilization. Keep in mind that application sharing components bundled with collaboration tools require additional bandwidth on top of the requirements for the audio and video media.
- *Packet loss*—Due to the high amount of compression and motion-compensated prediction utilized by video codecs, even a small amount of packet loss can result in visible degradation of the video quality. Packet loss tolerances for good video quality is highly subjective and can depend on a variety of factors, such as video resolution, frame rate, configured data rate, codec implementation, and even the specific PC upon which the video conferencing application is running. However, values between 0.1 and 1 percent often yield acceptable video quality. For example, testing shows that packet loss up to 1 percent might be acceptable for CUVA video configured for CIF (352 x 288 pixel) resolution and a 1.5 Mbps data rate. The resulting video might appear slightly more “jumpy” than

normal because the video pauses due to lost frames. However, because of the small resolution, this might not be significantly more noticeable than normal jumpiness caused by background processes running on the PC itself.

- *Jitter*—The tolerance of video codecs to jitter is often variable, depending upon the video codec deployed (H.263 or H.264) and the depth of the replay buffer of the codec. All packets that comprise a video frame must be delivered to the desktop video conferencing endpoint before the replay buffer is depleted. Otherwise degradation of the video quality can occur. The network should be designed to minimize jitter.
- *Latency*—The latency requirement for desktop video conferencing is in line with requirements for VoIP, which is based on the recommendations of the International Telecommunications Union-Telecommunication Standardization Sector (ITU-T) G.114 standard. When one-way latency begins to exceed approximately 200 msec, there is a noticeable degradation in the overall quality of the conversation. Desktop video conferencing includes the additional requirement of voice and audio synchronization. This often requires that audio and video media have the same service level across the network.
- *Bursts*—Video traffic on the network appears as a series of video frames spaced at regular intervals. Each video frame consists of multiple IP packets. The size of each frame is variable and partially determined by how the video is encoded. Therefore, video on the network often appears as a variable bit-rate stream with somewhat random bursts. The network infrastructure must be able to accommodate these bursts. Video quality will degrade if packets associated with video bursts are dropped or delayed excessively by any traffic policing or shaping within the network.

The Medianet Framework

A medianet (Figure 4) provides a network-wide framework for enterprise network designs with video services in mind. It is designed for enterprise customers who are considering various types of video applications in use today, or that might be used in the future. A 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. These services include the following:

- *Access services*—These provide access control and identity to video clients, as well as mobility and location services.
- *Transport services*—These provide packet delivery, ensuring the service levels with QoS and delivery optimization.
- *Bridging services*—These include transcoding, conferencing, and recording services.
- *Storage services*—These include content capture, storage, retrieval, distribution, and management services.
- *Session control services*—These include signaling and control to set up and tear down sessions, as well as gateway services.

Figure 4 The Medianet Framework

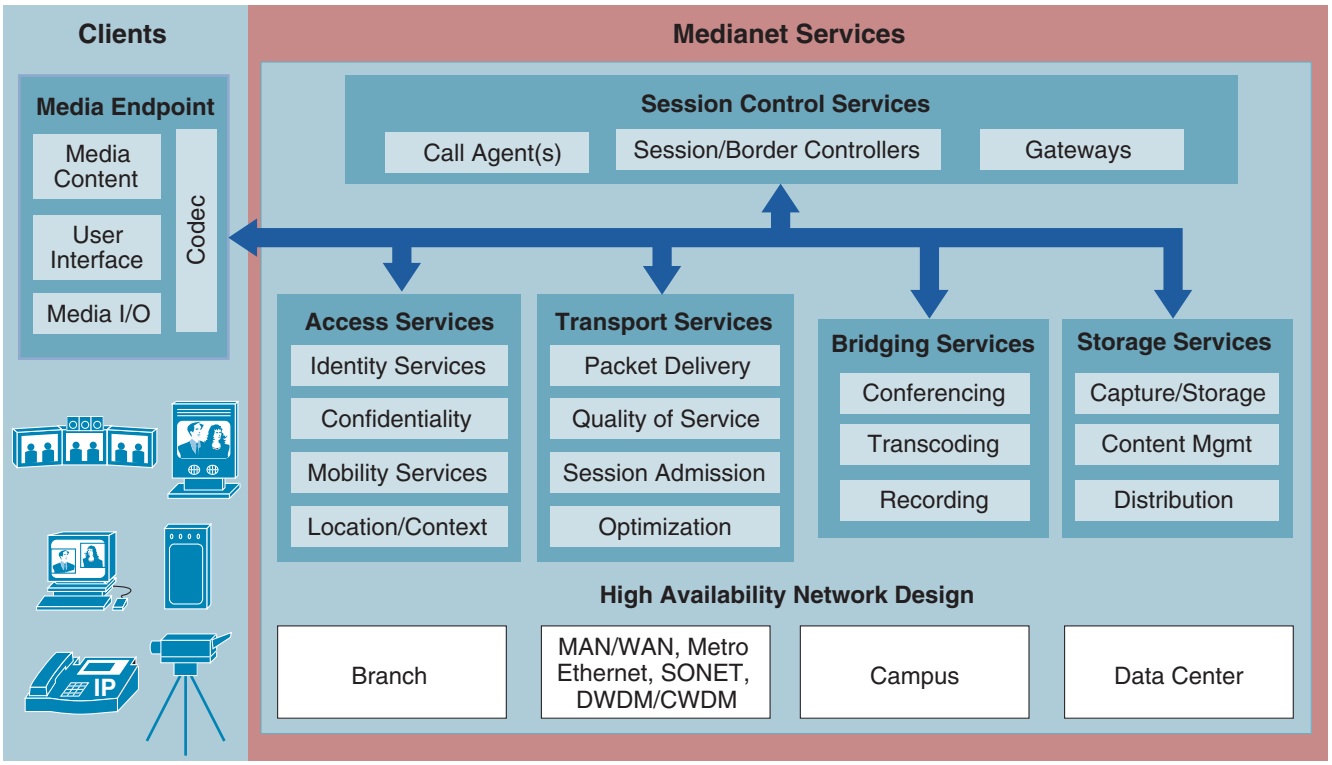

These services are then applied to the network infrastructure components and technologies defined within the Enterprise Solutions Engineering (ESE) Places-in-the-Network (PIN)—*Branch*, *WAN*, *Campus*, and *Data Center*.

PIN Architecture Design Considerations

Overlaying desktop video collaboration over a PIN architecture presents a unique challenge from other video technologies of a medianet, such as TelePresence, Digital Media Systems, and IP Video Surveillance. This is due to the wide diversity of products and technologies that can be provide the necessary services for desktop video collaboration. This diversity is a result of the ongoing evolution of desktop video collaboration from standalone hardware to part of an integrated software suite.

Figure 5 shows the overlay of the Cisco Unified Video Conferencing with MeetingPlace integration deployment over a PIN architecture, from a high level. This environment will be used as the basis for the discussion of the services required by each of the PIN areas to enable desktop video collaboration.

Figure 5 Desktop Video Collaboration PIN Overlay

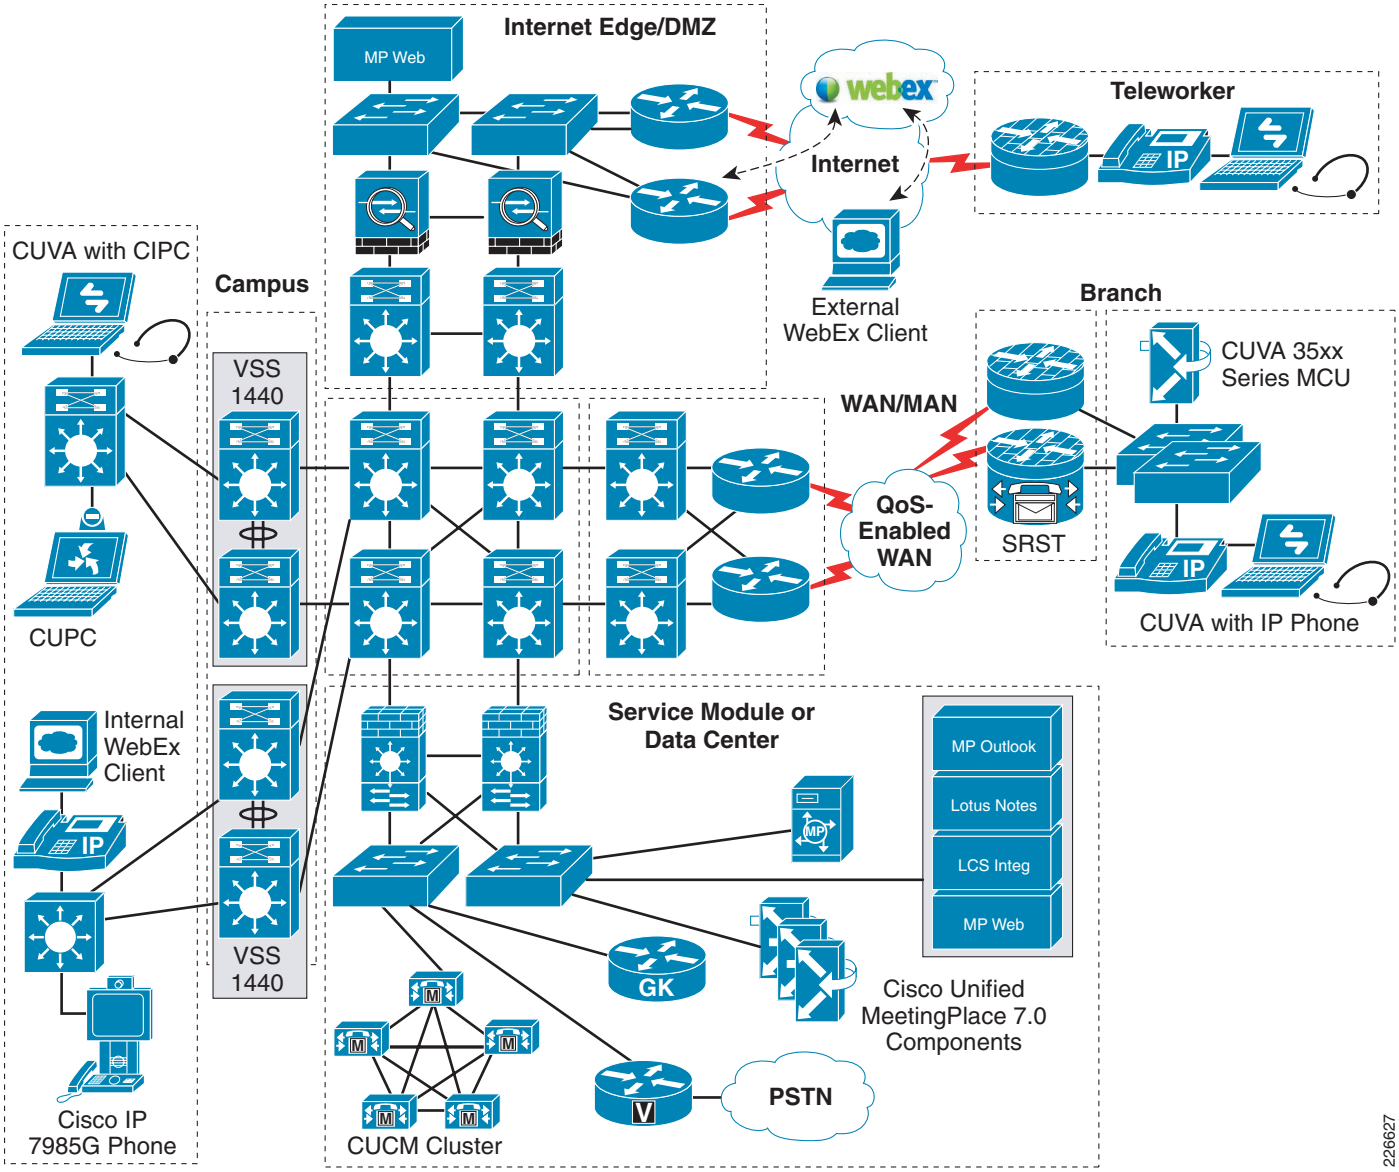

Branch PIN Design Considerations

From the perspective of desktop video collaboration, the Branch PIN design must provide the following key services in order to enable desktop video collaboration:

- WAN Bandwidth and Optimization, page 15
- Branch WAN Edge QoS, page 22
- Branch High Availability, page 22

These are discussed in the referenced sections and illustrated in Figure 6. Table 3 provides a legend that describes the numbered labels presented in Figure 6.

Figure 6 Branch PIN Services for Enabling Desktop Video Collaboration

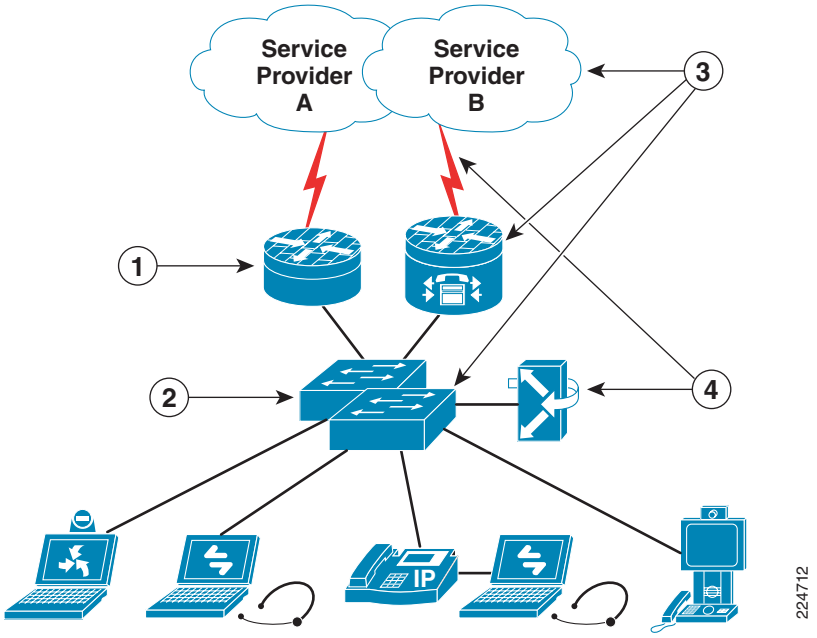

Table 3 Legend for Figure 6

Item	Description
1	WAN QoS • Ingress & Egress Queueing • Rate Limiting of Desktop Video Collaboration Traffic
2	LAN QoS • Classification & Marking of Desktop Video Collaboration Traffic • VLAN Assignment • QoS Trust Boundary Establishment • Ingress & Egress Queueing
3	High Availability • Redundant Circuits to Separate Service Providers • Redundant Cisco ISR Routers • Possible SRST for Additional Resilience of Call Control • Stacked Configuration of Catalyst 3750E Switches
4	WAN Bandwidth and Optimization • Bandwidth Provisioning Based on Desired Number of Simultaneous Calls Supported • Performance Routing for Optimal Utilization of Dual Circuits • Possible Deployment of Local MCU Resources for Multipoint Calls

WAN Bandwidth and Optimization

The amount of bandwidth utilized per desktop video conference call is typically at the discretion of the network administrator, based on the business requirements of the enterprise and the capabilities of the desktop video conferencing endpoints. For business quality desktop video conferencing, typical rates are above 384 Kbps per call.

Provisioning and Controlling Bandwidth Utilization

For Cisco Unified Video Conferencing endpoints (CUVA, CUPC, and the Cisco IP Phone 7985G), the amount of bandwidth utilized per desktop video conferencing call is relatively easily controlled via the region and location configurations with CUCM. Within the region configuration, the network administrator can select the maximum amount of bandwidth utilized for video per call as well as the audio codec supported between devices within the region itself and between regions. Within the location configuration, the network administrator can select the total amount of video allowed for all video calls to and from that location. Therefore, one method is to define each branch as a separate location and a separate region within CUCM. Alternatively, a single region can be defined for all branches if the amount of bandwidth allocated per desktop video conferencing call is the same across all branches. However, each branch can still be defined as a separate location in order to control the total amount of desktop video conferencing traffic to and from each branch.

In order to determine the amount of WAN bandwidth the network administrator must provision to a branch location to support desktop video conferencing, multiply the bandwidth by the number of simultaneous calls supported. For example, if 20 CUVA desktop video conferencing endpoints are deployed within a branch, but the network administrator decides that at most five simultaneous video conferencing calls need to be supported—and that 384 Kbps per call provides sufficient video quality—then the amount of bandwidth provisioned for desktop video collaboration support can be estimated as follows:

*5 simultaneous calls * 384 Kbps per call = 1,920 Kbps or 1.92 Mbps*

Note that this does not include network overhead which can conservatively be estimated to add approximately 20 percent more to the preceding bandwidth estimate. The bandwidth must then be provisioned for the QoS service class that contains desktop video conferencing traffic—in addition to the bandwidth requirements for the other service classes. The "Branch WAN Edge QoS" section on page 22 addresses this topics utilizing an Enterprise 12-class QoS model. Because detailed Erlang tables do not really exist for desktop video conference calls, the network administrator might need to estimate the bandwidth required and then monitor performance to determine the number of desktop video conferencing calls that fail or that are retried as audio calls due to insufficient branch WAN bandwidth.

Multipoint Call Considerations and MCU Placement

Care should be taken when considering multipoint desktop video conference calls. The placement of the video conference bridge or MCU will determine the number of video sessions that can traverse the branch WAN. An example is shown in Figure 7. Table 4 provides a legend that describes the numbered labels presented in Figure 7.

Figure 7 Branch Multipoint Desktop Video Conference Call Example

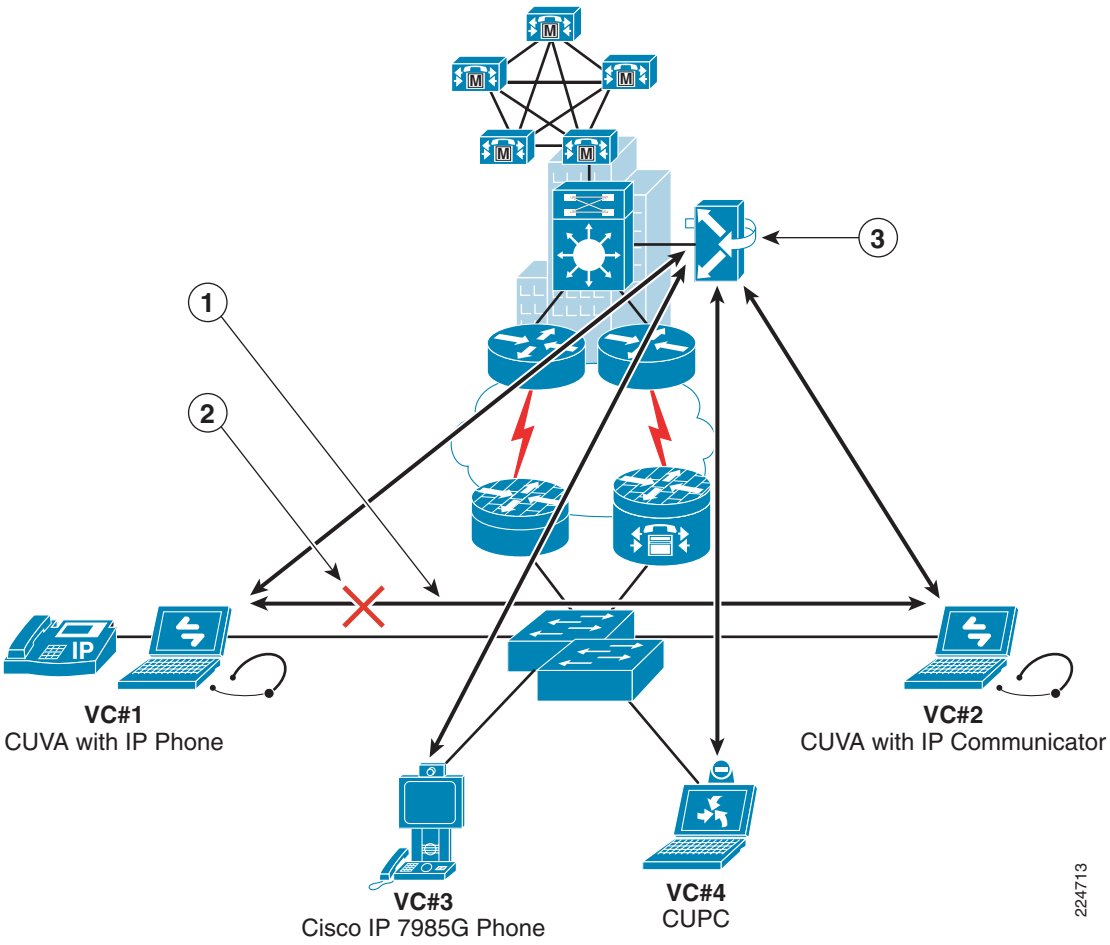

Table 4 Legend for Figure 7

Item	Description
1	VC#1 calls VC#2. Audio and video media local to the branch LAN.
2	VC#1 initiates an ad-hoc conference, adding VC#3 and VC#4.
3	Audio and video media now cross the branch WAN to centralized MCU.

In this example, VC#1 initiates a point-to-point desktop video conference call to VC#2. All audio and video media traverses the branch LAN only at this point. However, at some point during the call, VC#1 decides to initiate an ad hoc video conference to add in VC#3 and VC#4. Since no MCU exists at the branch, a MCU resource within the campus is automatically selected based upon the CUCM configuration. At that point, audio and video media streams from all four devices are re-directed to the centrally located MCU. Even though all of the endpoints are within the branch location, the WAN bandwidth is being utilized for the multipoint call.

The same situation can also occur for scheduled desktop video conference calls and for web conferences that include desktop video conferencing support. Therefore, it is critical for the network administrator to understand the call patterns within the branch to determine the amount of bandwidth that will be

required to support multipoint desktop video conferencing calls—and whether or not there might be benefits to placing MCU resources within a large branch. Careful placement of MCU resources can help optimize the use of branch WAN bandwidth during multipoint calls.

Bandwidth Optimization Techniques

The deployment of redundant circuits to the branch not only offers high availability of WAN connectivity to the branch, it also offers the possibility of optimizing the utilization of the dual circuits for the support of desktop video conferencing. A best practice is that each circuit be provisioned from a different service provider. Since each service provider network has different service-level parameters—which might in part depend on the amount and type of traffic between the branch and campus crossing the circuit at a given time—it might be possible to utilize bandwidth optimizing features such as Cisco Performance Routing (PfR) to select the best path for support of desktop video conferencing based on parameters such as the lowest amount of jitter, packet loss, or end-to-end delay. This might be beneficial for minimizing any degradation or disruption of service to desktop video collaboration calls.

Branch LAN QoS

Within the branch LAN, the QoS features that facilitate desktop video conferencing include establishment of the QoS trust boundary, classification and marking, VLAN assignment, and ingress and egress queueing. Platforms such as the Cisco Catalyst 3750E Series switches support these features.

QoS Trust Boundary Establishment, Classification & Marking, and VLAN Assignment of Desktop Video Collaboration Traffic

One of the unique challenges of desktop video collaboration is identifying all the audio, video, and application flows that originate from an endpoint and mark them appropriately. The following steps outline a methodology for accomplishing this.

1. Determining the QoS Model

The first step is to have a clear understanding of QoS model deployed within the enterprise and where each of the flows within a desktop video collaboration session fits within that model. Cisco currently recommends the RFC 4594-based QoS model shown in Figure 8 for deployment within Enterprise networks.

Figure 8 RFC 4594 - Based Enterprise QoS Model

Application Class	PHB	Admission Control	Queueing and Dropping	Application Examples
VoIP Telephony	EF	Required	Priority Queue (PQ)	Cisco IP Phones
Broadcast Video	CS5	Required	(Optional) PQ	Cisco IP Video Surveillance, Cisco Enterprise TV
Real-Time Interactive	CS4	Required	(Optional) PQ	Cisco TelePresence
Multimedia Conferencing	AF4	Required	BW Queue + DSCP WRED	Cisco Unified Personal Communicator
Multimedia Streaming	AF3	Recommended	BW Queue + DSCP WRED	Cisco Digital Media System (VoD)
Network Control	CS6		BW Queue	EIGRP, OSPF, BGP, HSRP, IKE
Call Signaling	CS3		BW Queue	SCCP, SIP, H.323
OAM	CS2		BW Queue	SNMP, SSH, Syslog
Transactional Data	AF2		BW Queue + DSCP WRED	Cisco WebEx, Cisco MeetingPlace, ERP Apps
Bulk Data	AF1		BW Queue + DSCP WRED	Email, FTP, Backup Apps, Content Distribution
Best Effort	default		Default Queue + RED	Default Class Traffic
Scavenger	CS1		Min BW Queue (deferential)	YouTube, iTunes, BitTorrent, Xbox Live

226628

Within this model, the audio and video components of desktop video collaboration sessions are recommended to be classified as multimedia (conferencing) and given an AF41 marking. However, if a call goes through as an audio only call, it should be classified as *VoIP telephony* and given an EF marking. The session control flows which initiate and terminate the audio and video should be classified as *call signaling* and given a CS3 marking. The final determination is to identify how to mark the various other data flows which make up an overall collaborative session or web conference because desktop video conferencing might only be a part of a larger collaborative session. For practical purposes, it might be necessary to classify these flows as *best effort* with a default marking.

2. Identifying Traffic Flows

The second step is to determine how the various audio, video, and signaling flows within a desktop video collaboration session can be identified. Flows coming from a desktop video collaboration endpoint can be identified based upon various criteria, which include the following:

- *802.1Q VLAN tag*—If the device is capable of sending packets that are encapsulated within an 802.1Q VLAN tag, the ingress switch port can classify and mark traffic based upon the VLAN on which it appears.
- *802.1p class of service (CoS) value*—The switch can be configured to trust the 802.1p CoS marking within the 802.1Q VLAN tag. This assumes the end device is trusted to send traffic with the correct CoS marking.
- *DSCP value*—The switch can be configured to simply use the DSCP values within the traffic sent by the desktop video collaboration endpoint. However, this leaves open the possibility that the device will send everything as high priority traffic. Access control list (ACL) entries might be needed to match certain protocol ranges as well as their DSCP values, while remarking all other traffic to default.

- *Source and/or destination address range*—Since a single switch interface may have a Voice and Data VLAN defined, and each VLAN has a different IP subnet, traffic can be classified based upon which IP subnet it arrives. It should be noted that classifying on individual source and/or destination IP addresses might be administratively unfeasible for desktop video collaboration because of the sheer number of endpoints involved.
- *Protocol range*—In many cases, the protocol or port range of audio and video traffic is known in advance. Using this pre-knowledge, the switch can identify audio and video traffic and mark it appropriately.

It should be noted that these methods are not mutually exclusive. Combinations of methods can be utilized and a combined criteria approach is recommended.

To further complicate things, desktop video collaboration endpoints are not required to behave identically. The following notes discuss how audio, video, and call signaling flows are sent from various Cisco endpoints:

- *CUVA with IP Phone*—Cisco Discovery Protocol (CDP) packets sent by the IP phone can be used to indicate to the switch that the device connected to the port is a Cisco IP Phone. The switch informs the IP phone of the VLAN tag it is to use for voice traffic. Therefore audio traffic appears on the Voice VLAN. Voice traffic also utilizes a UDP port range of 16384 to 32767. The PC that runs CUVA does not send VLAN tagged packets. Video traffic appears on the Data VLAN. Video traffic utilizes UDP port 5445. Therefore it can be distinguished from audio traffic via the difference in port range. CUCM controls the DSCP marking of audio and video traffic generated from the endpoint. By default CUCM instructs the audio and video traffic to be marked as AF41 for a video call, but instructs the audio traffic to be marked as EF for an audio-only call from the IP phone. Call signaling is generated only by the IP phone and therefore appears only on the Voice VLAN. However, there is additional signaling between the IP Phone and the PC that utilizes UDP port 4224. CUCM controls the DSCP marking of the call signaling. By default, CUCM instructs the call signaling traffic to be marked as CS3.
- *CUVA with IP Communicator*—Since there is no IP phone, neither the voice nor video traffic is VLAN encapsulated and therefore appears on the Data VLAN only. Voice traffic utilizes a UDP port range of 16384 to 32767. Video traffic utilizes UDP port 5445. Therefore, video can be distinguished from audio traffic via the difference in port range. CUCM controls the DSCP marking of audio and video traffic. By default CUCM instructs the audio and video traffic to be marked as AF41 for a video call, but instructs the audio traffic to be marked as EF for an audio-only call from IP Communicator. Call signaling is generated by the PC and therefore appears on the Data VLAN only. CUCM controls the DSCP marking of the call signaling. By default, CUCM instructs the call signaling traffic to be marked as CS3.
- *Cisco IP Phone 7985G*—CDP sent by the Cisco IP Phone 7985G can be used to indicate to the switch that the device connected to the port is a Cisco IP Phone. The switch informs the Cisco IP Phone 7985G of the VLAN tag it is to use for all traffic. Therefore, audio and video traffic both appear on the Voice VLAN. Audio and video traffic utilize a UDP port range of 16384 to 32767. Therefore, video cannot easily be distinguished from audio traffic via a difference in port range. CUCM controls the DSCP marking of audio and video traffic. By default CUCM instructs the audio and video traffic to be marked as AF41 for a video call, but instructs the audio traffic to be marked as EF for an audio-only call. Call signaling also appears on the Voice VLAN only. CUCM controls the DSCP marking of the call signaling. By default CUCM instructs the call signaling traffic to be marked as CS3.
- *CUPC with soft phone*—Since there is no IP phone, neither the voice nor video traffic is VLAN encapsulated and therefore appears on the Data VLAN only. Voice and video traffic utilize a UDP port range of 16384 to 32767. As a result, video cannot easily be distinguished from audio traffic via a difference in port range. CUCM controls the DSCP marking of audio and video traffic. By default, CUCM instructs the audio and video traffic to be marked as AF41 for a video call, but

should instruct the audio traffic to be marked as EF for an audio-only call from the soft phone. Call signaling is generated by the PC and appears on the Data VLAN. CUCM controls the DSCP marking of the call signaling. By default, CUCM instructs the call signaling traffic to be marked as CS3.

- *WebEx Web Conference (web browser)*—Video is embedded within the web conference session stream on the Data VLAN. This might be encrypted as well using SSL/TLS over HTTP (TCP port 443). It might not be possible to separately classify and mark the video traffic within a WebEx web conference. Typically, this means video traffic will be classified as Best Effort and marked default, along with all other traffic from the PC. There is no call signaling with regard to the video traffic to the endpoint. Audio is typically handled through dial-back or dial-in through a separate IP phone and is handled like a normal IP telephony call.

3. Defining the Policy Map

The third step is to define the policy map to apply to the ingress switch interface. For Cisco Unified Video Conferencing endpoints (CUVA, CUPC, and the Cisco IP Phone 7985G) audio, video, and call signaling traffic can be classified through ACLs applied to a generic policy map that is then applied to the ingress interfaces of the Cisco Catalyst switch ports. Within the policy map, traffic can be classified based on a combination of source IP subnet (based on its VLAN assignment), DHCP value, and protocol/port range. DSCP values of the voice, video, and signaling are trusted once they are matched by the ACLs. To further protect the network, each traffic type can be rate-limited via an inbound policer, if desired.

Note All other inbound traffic on the switch port may be classified as *best effort* and given default treatment. However, this methodology can be extended to identify specific collaborative application components (such as text messaging, data/application sharing, and so on), and to classify and mark them to something other than *best effort*, if desired.

The benefit of a single generic policy map versus multiple specific policy maps for each type of video endpoint is that it allows the enterprise to more easily migrate from one type of endpoint to another (for instance CUVA to CUPC) or to deploy multiple types of endpoints (for instance Cisco IP Phone 7985Gs and CUVA) without having to administratively determine which endpoint is connected to the particular switch port. Also, as companies move to the concept of *hoteling* (in which virtual cubicles are shared among employees in order to reduce expenses), it might not be possible to pre-determine which type of video endpoint will be connected to the switch port. A single generic policy map provides a convenient solution for this requirement.

Ingress Queueing

When the total amount of traffic from all ports on a switch or switch stack has the potential to exceed the backplane speed of the switch or switch stack, traffic can be momentarily dropped at the ingress queue before it enters the switch. Multiple ingress queues on switch ports can provide benefits in this scenario. The Cisco Catalyst 3750E Series switch, which is recommended for branch LAN deployments, provides ingress queueing for this purpose.

Ingress queueing on the Cisco Catalyst 3750E series consists of two queues, each with three thresholds (1P2Q3T)—although the third threshold is pre-defined for the queue-full state. If the network administrator has some concerns about the possibility of oversubscribing the stack-rings of a Cisco Catalyst 3750E Series switch stack (resulting in input queue drops), voice traffic can be placed into the ingress priority queue based on its DSCP marking of EF. Similarly, desktop video conferencing audio and video traffic may also be placed into the ingress priority queue based on its DSCP marking of AF41, if desired. An alternative method is to leave the traffic in the non-priority queue, placing it into the higher

drop threshold of that queue. Care should be applied when selecting traffic for the priority queue, since it is limited to 40 percent of the bandwidth of the switch port. Also, placing everything into the priority queue will nullify the effects of having a priority queue.

Egress Queueing

When the total amount of outbound traffic on switch port (or switch stack) has the potential to exceed the speed of the port, traffic can be momentarily dropped at the egress queue before it exits the switch. Multiple egress queues on switch ports can provide benefits in this scenario. The Cisco Catalyst 3750E Series switch also provides egress queueing for this purpose.

Egress queueing on the Cisco Catalyst 3750E series consists of four queues (one of which can be a priority queue); each has three thresholds (1P3Q3T). When applying egress queueing to branch PIN deployments that include desktop video endpoints, it is important to determine the best way to map the recommended RFC 4594-based Enterprise 12-class QoS model presented in Figure 8 onto the architecture of the branch LAN switch. Figure 9 shows an example mapping onto the 1P3Q3T architecture of the Cisco Catalyst 3570E Series switch.

Figure 9 Mapping the Enterprise QoS Model to the Branch LAN Switch

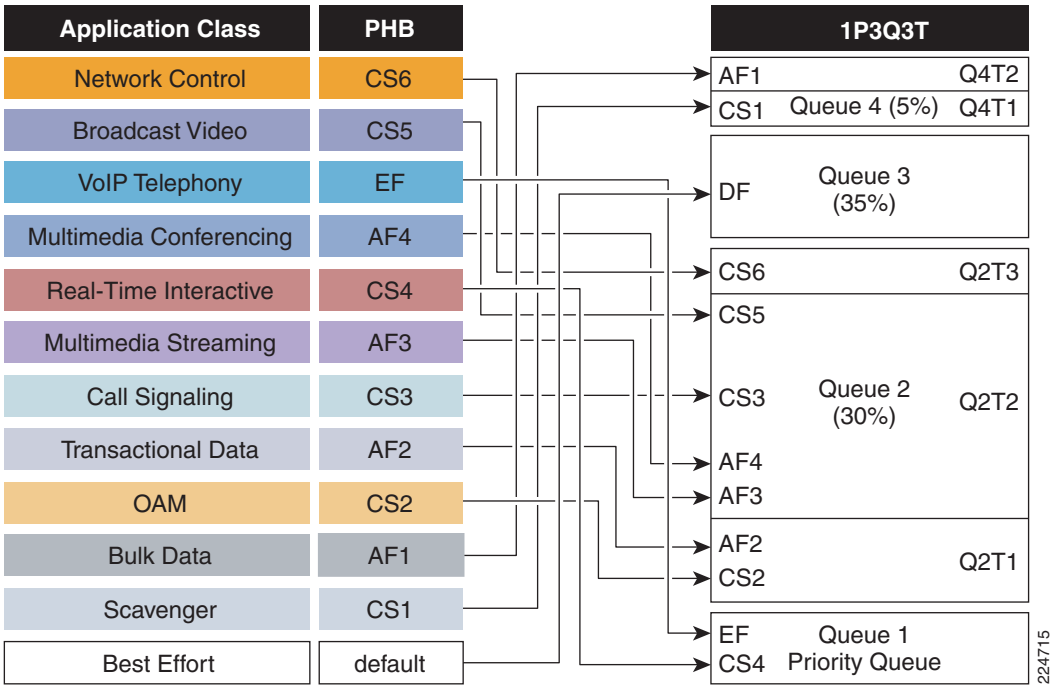

This mapping places desktop video conferencing traffic in the second drop threshold of Queue 2 (Q2 T2), along with various other traffic types. It should be noted by the network administrator that in any design in which the number of service classes exceeds the number of queues, the potential exists to mix multiple traffic types within a single queue. In such situations, the amount of traffic entering the particular queue might not be bounded by any call admission control mechanisms. Therefore, in order to prevent degradation of desktop video conferencing quality due to output queue drops on the Catalyst LAN switch, the queues should be regularly monitored.

Branch WAN Edge QoS

Within the branch WAN, the QoS features that facilitate desktop video conferencing include ingress and egress queueing, and rate limiting.

Ingress and Egress Queueing

Ingress queueing on Cisco 3800 and Cisco 2800 Series ISR platforms is usually not enabled, since the limiting factor of the branch WAN edge design is typically the WAN speed itself and not the router's ability to process packets. However monitoring of both the LAN and WAN interfaces should be performed regularly to quickly identify any interface input drops because these will degrade the quality of desktop video conferencing sessions.

Egress queueing should be configured because the amount of traffic generated by devices connected to the branch LAN can easily exceed the bandwidth of the branch WAN. One key benefit of Cisco 3800 and Cisco 2800 Series ISRs is that queueing is handled in software. A policy-map can be constructed with a queueing structure that matches the RFC 4594-based Enterprise 12-class QoS model shown in Figure 8.

Within the Enterprise 12-class QoS model, it is recommended to map desktop video conferencing to the Multimedia Conferencing traffic class on the WAN circuit. Current best practice recommendations are to configure the Multimedia Conferencing traffic class as a class-based weighted fair queue (CBWFQ) class. The percentage of bandwidth allocated for this class must be equal to or greater to the amount needed based on the number of desired simultaneous calls and the desired bandwidth per call. All desktop video conferencing calls might be degraded if insufficient bandwidth is allocated for the traffic class. Note also that, per RFC 4594, DSCP-based WRED may be enabled on the class. WRED might provide a benefit if the video conferencing application exhibits elasticity—meaning that it will decrease its transmission rate due to congestion. In such cases, dropping some packets before the queues are full will cause some video conferencing sessions to lower their transmission rate, potentially alleviating the congestion.

The policy map may also include a Broadcast Video class for IP video surveillance traffic, a Multimedia Streaming class for digital media systems, and a Real-Time Interactive class for TelePresence. This design guarantees a percentage of the WAN bandwidth for each traffic class. This effectively isolates each video traffic class. Since separate traffic classes exist for data traffic, it also effectively isolates business critical data applications from being disrupted by excessive video traffic. Regular monitoring of all of the traffic classes should be performed in order to quickly identify classes that are dropping traffic. In such situations, it might be necessary to re-allocate bandwidth across the various traffic classes, or possibly increase the overall circuit bandwidth.

Branch High Availability

High availability within the branch PIN design must be considered within the branch WAN edge and the branch LAN design. Since desktop video collaboration relies on a remote CUCM cluster, additional considerations for high availability of call signaling must also be assessed.

High Availability at the Branch WAN Edge

High availability at the branch WAN edge can be provided by dual WAN circuits, each terminating on redundant Cisco 3800 or Cisco 2800 Series ISRs. Additionally, the WAN circuits should be provisioned from different service providers and utilize different entrance facilities into the branch itself. This is consistent with current Branch PIN best practice design guidelines for large enterprise branches.

Ideally, the percentage of bandwidth utilization on any one circuit should be below 75 percent in order to allow for temporary spikes in utilization without dropping traffic or incurring excessive queueing delay. Unless Resource Reservation Protocol (RSVP)-based call-admission control (CAC) is

implemented for voice and desktop video conferencing calls, CUCM has no way of knowing that one of the redundant circuits has failed and that a reduced capacity exists to handle desktop video conference calls. Therefore, in this type of scenario, the utilization of each circuit during normal operation should optimally be below 37.5 percent (75/2) in order to allow for one circuit to smoothly handle the entire traffic load in case of a failure of the other WAN circuit. Even with an RSVP-based CAC mechanism, with insufficient bandwidth on a single circuit, fewer desktop video conferences could be accommodated. The only way to ensure the same number of simultaneous desktop video conferencing calls go through during a circuit outage is to provision each circuit with sufficient bandwidth to handle the entire traffic load. Further, the policy map defined on each circuit must have sufficient bandwidth allocated for the Multimedia Conferencing traffic class to handle the entire number of desired desktop video conference calls—in order to provide a true high-availability design for desktop video conferencing.

Dual WAN circuits terminating on redundant Cisco 3800 or Cisco 2800 Series ISRs also provides a high availability design for call signaling from branch IP telephony and desktop video conferencing endpoints to a remote CUCM cluster—provided that sufficient bandwidth exists on each circuit to handle all of the traffic within the call-signaling traffic class. Even if insufficient bandwidth exists for branch-to-campus video conferencing calls, intra-branch video conference calls and intra-branch IP telephony calls will still go through. However, because it is often not feasible from a cost perspective to provide dual entrance facilities to a branch location, the network administrator may consider the deployment of the Survivable Remote Site Telephony (SRST) feature of Cisco 3800 or Cisco 2800 Series ISRs as an added layer of call-signaling resiliency. In the event that the entrance cable to the branch is cut, severing both WAN circuits, the branch would be able to place intra-branch calls. If a PSTN gateway is deployed within the branch, voice calls can also be placed to campus locations via the PSTN. However, before deploying SRST for desktop video conferencing, the network administrator should verify that the desktop video conferencing endpoints deployed within the branch location are supported by SRST. For example, SRST currently supports audio-only on the Cisco IP Phone 7985G. However, providing the added layer of resilience for IP telephony devices and desktop video conferencing endpoints operating in audio-only mode might be sufficient justification for adding SRST to the existing branch Cisco ISR platform from a business perspective.

High Availability within the Branch LAN

Because all of the Cisco Unified Video Conferencing endpoints support a single Ethernet interface, and because most PCs are typically connected to a single Ethernet network interface, there really is no high availability in the case of a LAN switch failure for devices directly connected to the Catalyst LAN switch.

However, the Cisco Catalyst 3750E switch stack does provide high availability for the Layer-3 switching function. This is often referred to as Layer-3 Non-Stop Forwarding (L3 NSF). Layer-3 switching within larger branch LAN implementations might be something to consider. As discussed previously, various desktop video conferencing endpoints send audio and video traffic on either the voice or data VLAN, depending on the endpoint. With the deployment of multiple types of video conferencing endpoints within the LAN, audio and/or video traffic might need to be routed between VLANs through the branch ISR. Although some routing of traffic between VLANs is normal, as the number of endpoints increases, so does the amount of branch-to-branch LAN traffic passing through the router. Since the primary purpose of the WAN edge ISR is to route packets to and from the WAN, at some point the network administrator might desire to implement Layer-3 switching within the Cisco Catalyst 3750E switch stack. Within the Cisco Catalyst 3750E switch stack, one switch is elected the stack master, which controls the rest of the switch stack. Upon failure of the stack master, any one of the other switches can be elected as the new switch master and Layer-3 switching can continue.

For Cisco Catalyst 3750E switch stacks operating in Layer-2 mode and connected to redundant Cisco 3800 or Cisco 2800 Series ISR via Gigabit Ethernet connections, a best practice is to connect the uplink ports to different switches across the switch stack. Multiple HSRP groups can be run across the ISR

interfaces to provide high availability and VLAN-based load balancing across the interfaces. Although some disruption in an ongoing desktop video collaboration call will likely be noticed by the end users during a link failure, the HSRP parameters can be tuned to balance the length of the disruption against the increased network overhead.

WAN PIN Design Considerations

From the perspective of desktop video collaboration, the key services that the WAN PIN design must provide in order to enable desktop video collaboration are as follows:

- WAN bandwidth and the optimization of the available bandwidth for desktop video conferencing
- QoS within the WAN in order to meet the desired SLA parameters for desktop video conferencing
- High availability

The specific requirements vary depending upon whether the WAN is completely managed by the enterprise or managed by a service provider.

Service Provider-Managed WAN Considerations

Figure 10 shows an example of a deployment of desktop video collaboration over a service provider-managed WAN, such as an MPLS network. Table 5 provides a legend that describes the numbered labels presented in Figure 10.

Figure 10 Service Provider-Managed WAN PIN Services

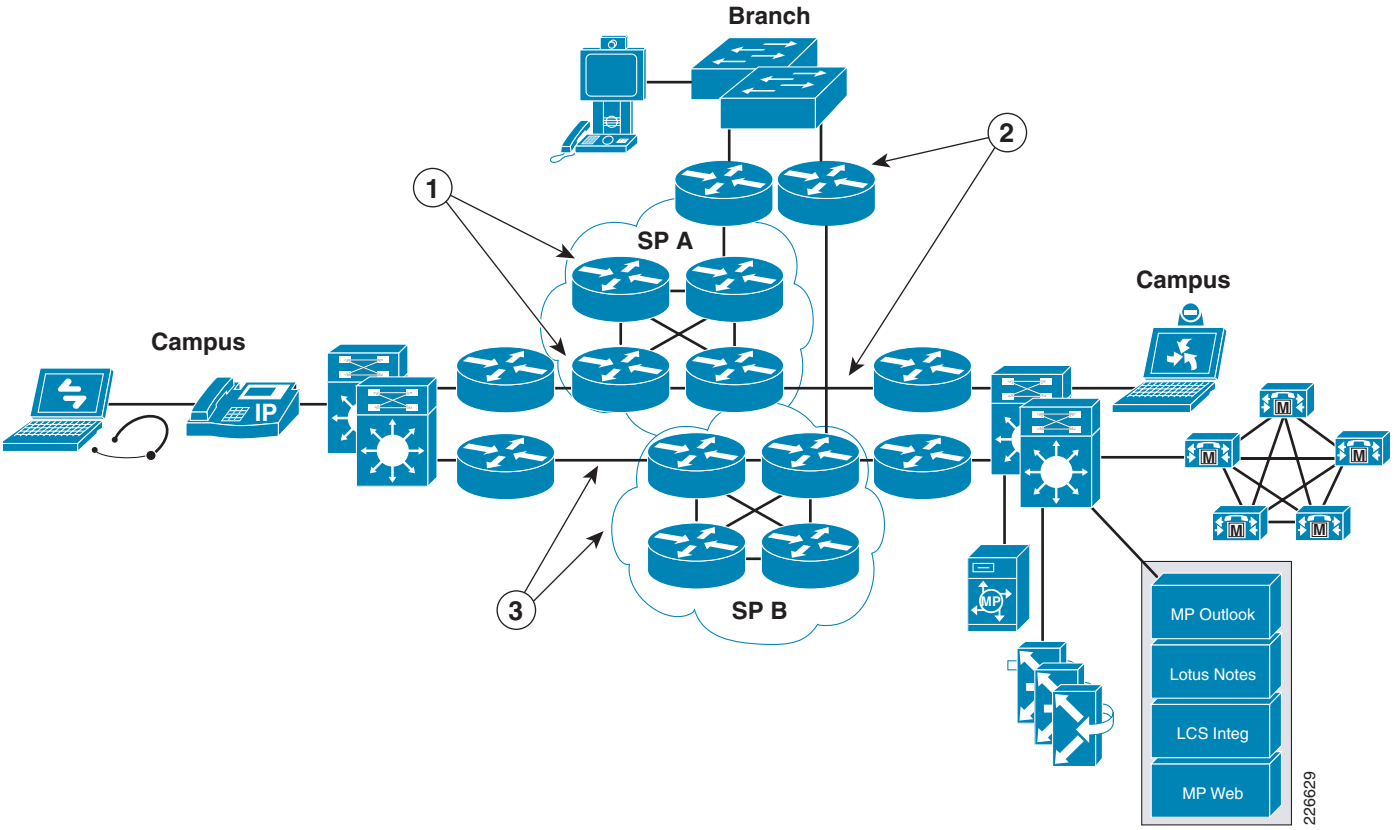

Table 5 Legend for Figure 10

Item	Description
1	WAN QoS • Egress Queueing and Re-Marking to Service Provider Classes • Rate Limiting of Desktop Video Collaboration Traffic • SLA to Meet Requirements of Desktop Video Collaboration
2	WAN High Availability • Redundant WAN Services From Separate Service Providers to Each Campus and Branch Location • Redundant WAN Edge Routers at Campus and Branch Locations
3	WAN Bandwidth and Optimization • Bandwidth Provisioning Based on Desired Number of Simultaneous Calls Supported • Performance Routing for Optimal Utilization of Multiple Paths

For a high-availability, managed-service WAN deployment, each campus and branch should have redundant connectivity to two service providers via dual WAN routers. This ensures continued operation of desktop video collaboration sessions across the enterprise in the event of a failure of one of the service provider networks. Each of the service providers must be able to meet the target SLA parameters for desktop video collaboration. This should be negotiated with the service provider prior to enabling desktop video collaboration across the network.

Depending upon the type of managed service being provisioned—such as an MPLS service—the service provider may provide multiple service classes within their network. However, the enterprise might still define more QoS service classes than available within the service provider network. For example, the Cisco RFC 4594-based QoS model defines 12 service classes. However, the service provider network might only implement six service classes. In such cases, the network administrator must take into account the mapping of all the enterprise service classes into the appropriate service provider service classes as traffic enters the MPLS network. Figure 11 provides an example of such a mapping.

Figure 11 Mapping of Enterprise Service Classes to Service Provider Classes

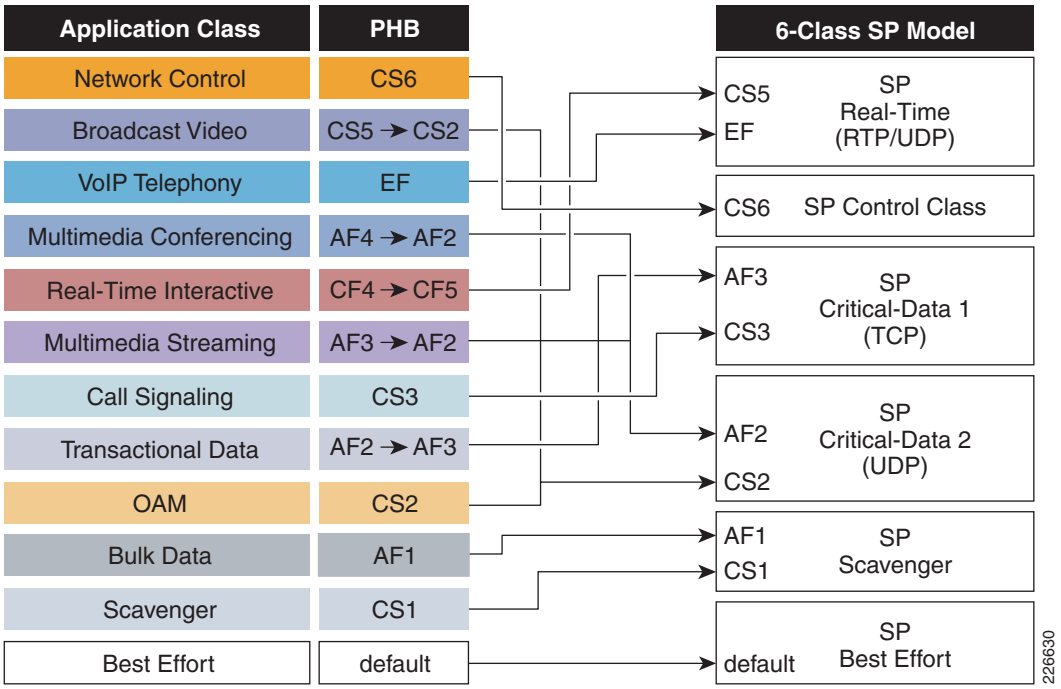

Mapping can be accomplished via a hierarchical policy map applied to the egress interface of the enterprise CE router, as the traffic exits the enterprise network. The top level of the policy map can be used to rate limit the overall amount of traffic sent to the service provider to the overall negotiated rate.

It is important for the network design engineer to understand that when mapping multiple enterprise service classes into a single service provider class, sufficient bandwidth must be provisioned per service provider service class. In Figure 11, the *broadcast video* service class (which handles IP video surveillance), the *multimedia conferencing* service class (which handles desktop video conferencing), the *multimedia streaming* service class (which handles digital media systems), and the *OAM* service class are all mapped into a single *critical-data 2* service provider class. Therefore this service class must be provisioned with sufficient bandwidth to handle the sum of all of the enterprise service classes into which it is mapped. Depending on how the policy map is configured, each of the enterprise service classes may be allowed exceed the bandwidth percentage allocated for the service class, if extra bandwidth is available. This is done in order to optimize the use of the available bandwidth, but might result in an oversubscription of any of the service provider classes.

Another important point to understand is the policy of the service provide for handling out-of-contract traffic per service class. If the service provider remarks out-of-contract traffic per service class down to *best effort*, then desktop video conferencing might be unaffected if the *critical-data 2* service class is temporarily oversubscribed. However, if the service provider polices out-of-contract traffic per service class, then desktop video collaboration—as well as digital media systems traffic and IP video surveillance traffic—will all be degraded if the *critical-data 2* service class is temporarily oversubscribed. This is one disadvantage of mapping many enterprise service classes into fewer service provider service classes.

Desktop video conferencing traffic may also need to be identified and re-mapped back to the *multimedia conferencing* service class as it exits the MPLS network. The difficulty lies with being able to differentiate desktop video conferencing traffic from other video traffic types such as IP video surveillance and digital media systems traffic. All three types of traffic may utilize RTP for audio and

video traffic, and may therefore have the same UDP port range of 16384 to 32767. CUVA desktop video conferencing endpoints make it relatively easy to differentiate the video component through the use of UDP port 5445; however, the audio component still uses the UDP range of 16384 through 32767.

One possible solution for this dilemma is to segment the UDP port range for each category of devices, so that each can be identified coming back out of the MPLS network. Another possible method is to differentiate based on source or destination IP address or IP subnet. Although this might not be feasible for desktop video conferencing endpoints, it might be feasible for IP video surveillance endpoints to be on different VLANs and therefore different IP subnets for traffic isolation purposes. In some cases, it simply might not be possible to properly distinguish some types of video traffic from others, as it exits MPLS network. In such cases, the network administrator might end up with multiple video traffic types marked the same at the LAN on the far side of the MPLS network.

Enterprise Managed WAN Considerations

Figure 12 shows an example of a deployment of desktop video collaboration over an enterprise-managed WAN. Table 6 provides a legend that describes the numbered labels presented in Figure 12.

Figure 12 Enterprise-Managed WAN PIN Services

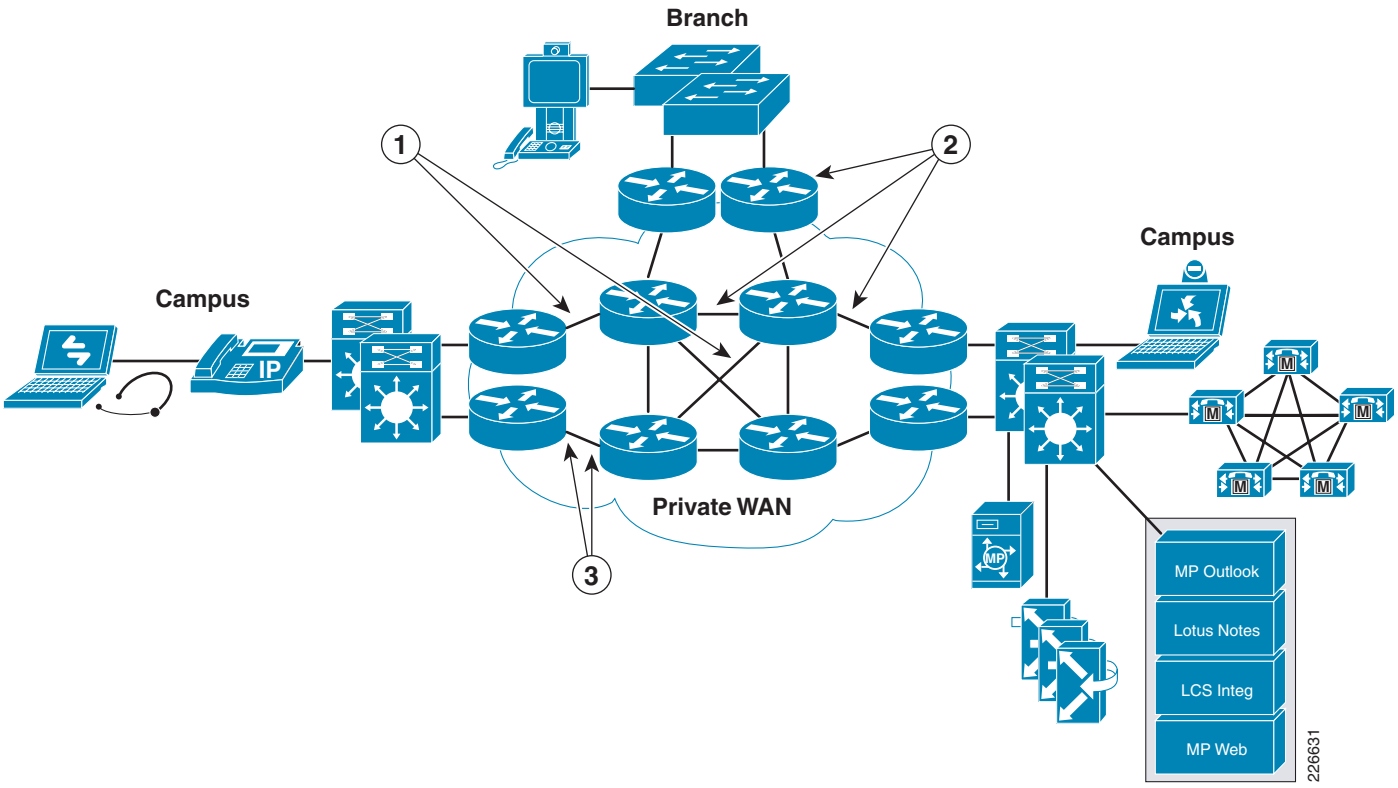

Table 6 Legend for Figure 12

Item	Description
1	WAN QoS • Egress Queueing • Rate Limiting of Desktop Video Collaboration Traffic • Minimize Hops To Minimize Latency and Jitter
2	WAN High Availability • Redundant Circuits From Separate Service Providers to Each Campus and Branch Location • Redundant WAN Edge Routers at Campus and Branch Locations • Redundant Paths Through the WAN Core
3	WAN Bandwidth and Optimization • Bandwidth Provisioning Based on Desired Number of Simultaneous Calls Supported • Tuning of Routing Protocols for Rapid Reconvergence • Performance Routing for Optimal Utilization of Multiple Paths

Within a private WAN, the enterprise itself is responsible for providing high availability throughout the network infrastructure. Dual circuits provisioned from different service providers and terminated on different routers at campus and branch locations are still necessary for high availability. At the WAN aggregation point, high-availability features including built-in hardware and processor redundancy, In-Service Software Upgrade (ISSU), and Cisco nonstop forwarding (NSF) with stateful switchover (SSO) can be deployed within Cisco Aggregation Services Router (ASR) platforms as well. In addition, the WAN core itself must be designed for both redundancy in terms of a failure of either a circuit or router platform.

The network administrator is also responsible for ensuring SLA requirements are met for desktop video conferencing across the private WAN infrastructure. 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 also require modification to reduce the number of hops within the network—which can potentially add latency and jitter. Finally, tuning of routing protocol timers might require evaluation in order to optimize the reconvergence time of the WAN to minimize disruptions to ongoing video conferencing calls during a circuit failure.

One of the advantages of the enterprise-managed WAN is the ability to run as many service classes as needed throughout the enterprise because the network administrator has complete control over the deployment of service classes. As a result, the RFC 4594-based enterprise QoS model shown in Figure 8 can be applied across the entire WAN. No issues arise with this type of deployment with regard to marking and remarking of various video traffic types, such as digital media systems, IP video surveillance, and desktop video conferencing.

Campus PIN Design Considerations

From the perspective of desktop video collaboration, the key services that the Campus PIN design must provide in order to enable desktop video collaboration are as follows:

- Bandwidth and the optimization of the available bandwidth for desktop video conferencing
- High availability
- QoS within both the LAN and WAN components of the campus
- Visibility and manageability for troubleshooting desktop video conferencing flows

These are shown in Figure 13. Table 7 provides a legend that describes the numbered labels presented in Figure 13.

Figure 13 Campus PIN Services for Enabling Desktop Video Collaboration

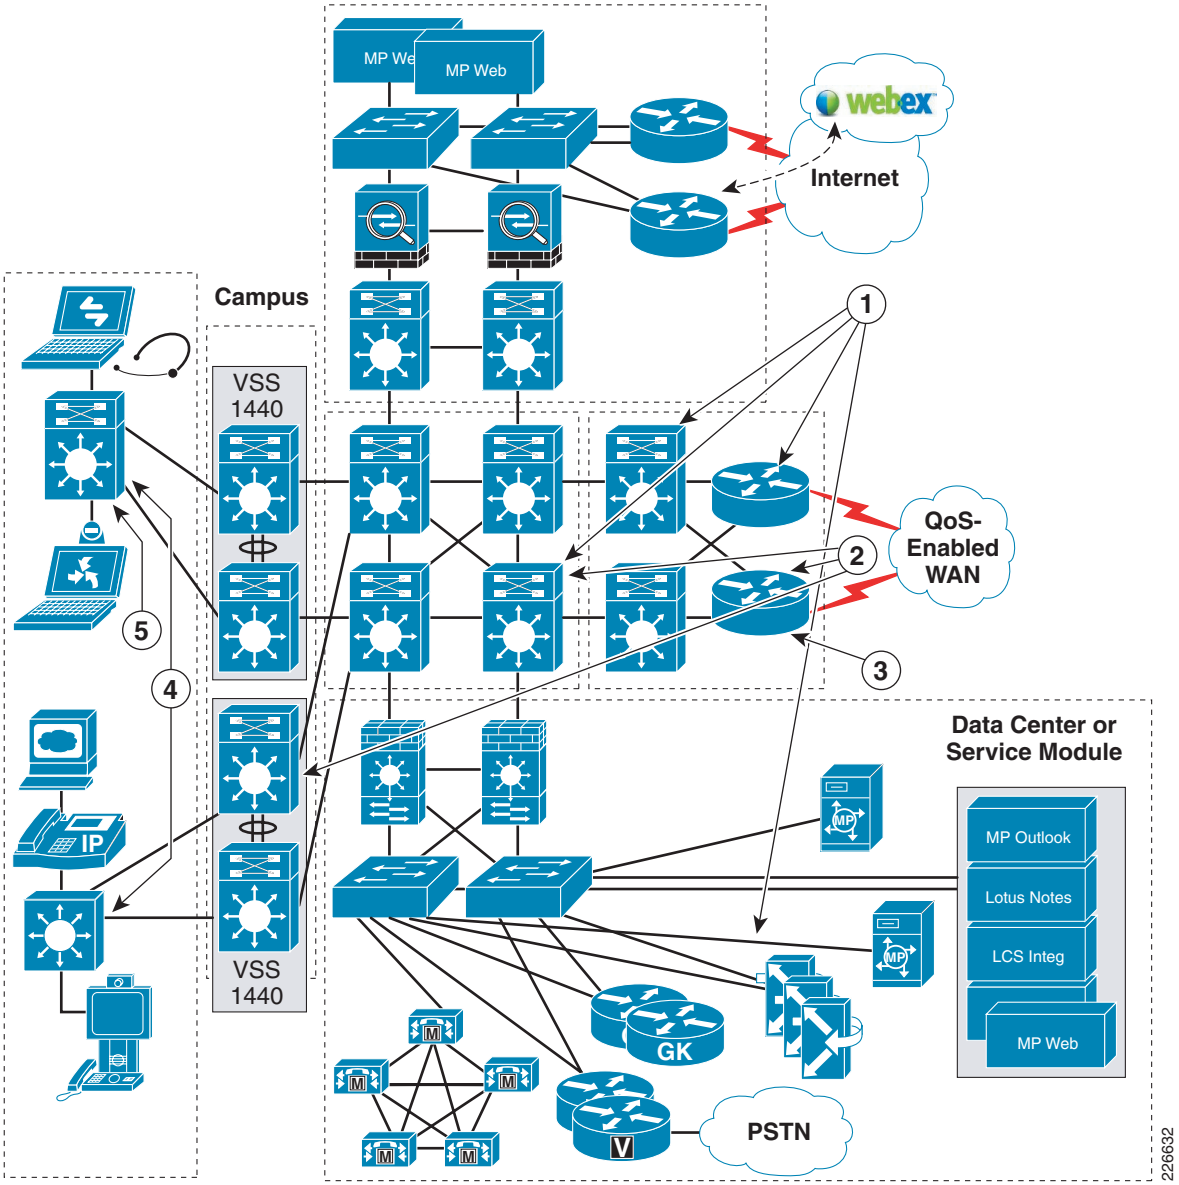

Table 7 Legend for Figure 13

Item	Description
1	High Availability • Redundant Circuits to Separate Service Providers at the WAN Edge • Redundant Connectivity at the Internet Edge • Redundant Ethernet Links in the Core, Distribution, and Access Layers • Built-in hardware and processor redundancy, ISSU, and NSF/SSO Throughout the Campus • Redundant MeetingPlace, Gatekeeper, and PSTN Gateway Components
2	Bandwidth and Optimization • Hierarchical Campus Design • 1 Gbps and 10 Gbps Ethernet Links within the Core, Distribution, and Access Layers • Performance Routing at the WAN Edge for Optimal Utilization of Dual Circuits • Virtual Switching System (VSS) for Scalability, Load Balancing, and Ease of Administration
3	WAN Edge QoS • Ingress and Egress Queueing • Rate Limiting of Desktop Video Collaboration Traffic
4	LAN QoS • Classification & Marking of Desktop Video Collaboration Traffic • VLAN Assignment • QoS Trust Boundary Establishment • Ingress and Egress Queueing
5	Visibility & Manageability • Cisco IP Service Level Agreement (IPSLA) for Network Diagnostics • NetFlow for Audio and Video Traffic Flow Statistics • Sup32 PISA / NAM for Deep Packet Inspection and Analysis

Bandwidth and Optimization

The same bandwidth and optimization services that apply to the branch and WAN PIN designs also apply to the WAN edge of the campus PIN design. Since the number of desktop video collaboration endpoints within the campus might be far greater than a branch, a hierarchical campus design consisting of core, distribution, and access layers is recommended for scalability. Redundant Gigabit Ethernet and 10 Gigabit Ethernet links combine to provide scalable bandwidth and high availability for desktop video collaboration within the campus itself. Additional features, such as the Cisco Catalyst 6500 VSS can be deployed within the core and distribution layers of the campus—as well as within the data center/services module. VSS provides scalability and high availability through cross-chassis EtherChannel load-balancing and ease of administration.

Campus High Availability

Besides redundant Cisco Catalyst switch platforms and Ethernet links, additional high-availability features, including built-in hardware and processor redundancy, ISSU, and NSF/SSO can be deployed within Cisco Catalyst 6500 and Cisco Catalyst 4500 Series platforms. These features help minimize disruptions in ongoing desktop video collaboration sessions that traverse the campus. Deployment of specific components can be within the campus services module or data center to promote high availability, bandwidth scalability, security, and manageability. The components include the following: CUCM cluster; redundant Cisco Unified MeetingPlace 7.0 Application, Integration, and Media Servers; redundant gatekeepers; and redundant PSTN gateways. Additional features, such as Uni-Directional Link Detection (UDLD), help to quickly isolate breaks in fiber optic cabling between core and distribution switches or between distribution and access switches within the campus. This allows spanning-tree protocol—or alternatively Cisco FlexLink technology—to rapidly reconverge the campus network, thereby minimizing the disruption in desktop video collaboration sessions. The Gateway Load Balancing Protocol (GLBP) can be run across Cisco Catalyst switch interfaces to provide high availability and load balancing across the campus LAN.

Campus QoS

As with the branch PIN, the access Cisco Catalyst 6500, Cisco Catalyst 4500, or Cisco Catalyst 3750E switch within the campus establishes the QoS trust boundary, classification, and marking of audio and video media as it enters the network—as well as ingress and egress queueing to ensure desktop video conferencing traffic is allocated the appropriate bandwidth with the campus LAN. Additional mechanisms, such as the Cisco Catalyst 6500 Sup32 Programmable Intelligent Services Accelerator (PISA) allows deep-packet inspection (DPI), thereby providing more granular classification of desktop video collaboration traffic as it enters the campus LAN.

Campus Visibility and Manageability

The addition of the Catalyst 6500 Network Analysis Module (NAM) can add detailed traffic analysis capability on top of the deep-packet inspection and classification capability of the Cisco Catalyst 6500 Sup32 PISA. Additionally, NetFlow statistics captured from Cisco Catalyst 6500 platforms can provide real-time monitoring for all application traffic flows—allowing visibility into voice and video flows within the campus network. Finally, the Cisco IPSLA capability of Cisco Catalyst 6500 platforms allow network administrators to proactively verify the network operation and accurately measure network performance.

Conclusions

Enterprise organizations are increasingly turning to collaborative technologies, such as desktop video conferencing, to increase productivity, scale expertise, and reduce travel costs. This has led to a dramatic increase in desktop video collaboration traffic across enterprise networks. This solution overview has provided an overview of the components required for desktop video collaboration. It has also provided guidance for choosing services that the various PINs (Branch, WAN, and Campus) can provide to facilitate desktop video collaboration deployments. In order to take a proactive approach to the growth of video, the network administrator must understand the various services that a medianet can provide to facilitate the convergence of desktop video collaboration onto the IP network infrastructure.

Terms and Acronyms

Acronyms	Definition
4CIF	4 x Common Intermediate Format - 704 x 576 pixel video resolution
10 GE	10 Gigabit Ethernet
ACL	Access Control List
AF	Assured Forwarding class of service
ASR	Advanced Services Router
CBWFQ	Class-Based Weighted Fair Queue
CDP	Cisco Discovery Protocol
CE	Customer Edge - MPLS Customer Edge Router
CIF	Common Intermediate Format - 352 x 288 pixel video resolution
Codec	Coder/Decoder
CoS	Class of Service
CS	Class Selector class of service
CUPC	Cisco Unified Personal Communicator
CUVA	Cisco Unified Video Advantage
CUVC	Cisco Unified Video Conferencing
CUV-M	Cisco Unified Video Manager
DPI	Deep Packet Inspection
DMS	Digital Media System
DSCP	Differentiated Services Code Point
EF	Expedited Forwarding class of service
GE	Gigabit Ethernet
GLBP	Gateway Load Balancing Protocol
H.323	Call Signaling Standard
H.263	Video compression standard, also know as MPEG-2
H.264	Video compression standard, also known as MPEG-4 AVC
HA	High Availability
HD	High Definition video resolution
HSRP	Hot-Standby Routing Protocol
IP	Internet Protocol
IPVS	IP Video Surveillance
ISR	Integrated Services Router
ISSU	In-Service Software Upgrade
LAN	Local Area Network
MAN	Metropolitan Area Network
MCU	Multipoint Control Unit

Acronyms	Definition
MPEG-4	Moving Pictures Expert Group 4 standard
MPLS	Multi Protocol Label Switching
NAM	Network Analysis Module
NSF	Non-Stop Forwarding
OAM	Operations, Administration, and Maintenance
PC	Personal Computer
PfR	Performance Routing
PHB	Per-Hop Behavior
PIN	Places-in-the-Network
PISA	Programmable Intelligent Services Accelerator
PQ	Priority Queue
PSTN	Public Switched Telephone Network
QoS	Quality of Service
RSVP	Resource Reservation Protocol
RTP	Real-time Transport Protocol
SD	Standard Definition video resolution
SIP	Session Initiation Protocol
SLA	Service Level Agreement
SOAP	Simple Object Access Protocol
SONET	Synchronous Optical Networking
SP	Service Provider
SRST	Survivable Remote Site Telephony
SSO	Stateful-Switchover
TCP	Transmission Control Protocol
UC	Unified Communications
UDP	User Datagram Protocol
UDLD	Uni-Directional Link Detection
VLAN	Virtual Local Area Network
VSS	Virtual Switching System
WAN	Wide Area Network
WRED	Weighted Random Early Discard
XML	Extensible Markup Language

Related Documents

System Reference Network Designs

Enterprise 3.0 Campus Architecture Overview and Framework
http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/campover.html

Transport Diversity: Performance Routing (PfR)
http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/Transport_diversity/Transport_Diversity_PfR.html

Branch Office Architecture Overview
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns171/c649/ccmigration_09186a00807593b7.pdf

Data Center Infrastructure Design Guide
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns107/c649/ccmigration_09186a008073377d.pdf

Telepresence Network System Design Guide
http://www.cisco.com/en/US/docs/solutions/TelePresence_Network_Systems_1.1_DG.pdf

Websites

Campus Solutions
http://www.cisco.com/en/US/netsol/ns340/ns394/ns431/networking_solutions_packages_list.html

WAN and Aggregation Services Solutions
http://www.cisco.com/en/US/netsol/ns483/networking_solutions_packages_list.html

Branch Office Solutions
http://www.cisco.com/en/US/netsol/ns477/networking_solutions_packages_list.html

Data Center 3.0 Solutions
http://www.cisco.com/en/US/netsol/ns708/networking_solutions_solution_segment_home.html

Video Solutions
http://www.cisco.com/en/US/netsol/ns340/ns394/ns158/networking_solutions_packages_list.html

Telepresence Solutions
http://www.cisco.com/en/US/netsol/ns669/networking_solutions_solution_segment_home.html

Unified Communications Solutions
http://www.cisco.com/en/US/netsol/ns340/ns394/ns165/ns152/networking_solutions_package.html

Wide Area Application Services Solutions
http://www.cisco.com/en/US/products/ps5680/Products_Sub_Category_Home.html

The Medianet Solutions
http://www.cisco.com/go/designzone