



CHAPTER 8

Cisco Video Infrastructure Components

Last revised on: October 30, 2009

This chapter describes the Cisco Unified Videoconferencing infrastructure components and related network design considerations. This chapter focuses on traditional (room-based) IP videoconferencing. Therefore, it is specifically intended for those environments that have yet to migrate to an end-to-end Cisco Unified Communications Manager Video Telephony model or Cisco Unified MeetingPlace Video Integration.

For information on integrating the Cisco Unified Videoconferencing infrastructure components with Cisco Unified Communications Manager, refer to the *IP Video Telephony* section in the *Cisco Unified Communications Solution Reference Network Design (SRND)*, available at

<http://www.cisco.com/go/ucsrnd>



Note

MCU and gateway features may change with new software releases, and those changes might not be represented in this document. Refer to the latest product documentation and release notes for details.

What's New in This Chapter

Table 8-1 lists the topics that are new in this chapter or that have changed significantly from previous releases of this document.

Table 8-1 **New or Changed Information Since the Previous Release of This Document**

New or Revised Topic	Described in:
Collaboration	Collaboration for Desktop Sharing, page 8-21
Connecting to external networks	Connecting the Video Network with External Untrusted Networks, page 8-27
Integration with other collaboration applications	Integration and Interoperability, page 8-25
Multipoint conferencing	Multipoint Conferencing, page 8-2
Recording conferences	Conference Recording, page 8-24
Scheduling video conferences with Cisco WebEx	Video Scheduling with Cisco WebEx, page 8-8
Streaming conference video	Conference Streaming, page 8-23

Design and Deployment Overview

Video infrastructure design is a very important element in an H.323, Session Initiation Protocol (SIP), or Skinny Client Control Protocol (SCCP) videoconferencing network. In the H.320 circuit-switched network, Multipoint Control Units (MCUs) and H.320 endpoints are connected directly to the PSTN network. In the past, an MCU could have multiple PRI connections into the switched network. The switched network supplied a dedicated transport with guaranteed bandwidth and predictable delay. Now that video is being moved onto IP networks that share bandwidth with data, placement of video infrastructure components becomes very important. Installing a central MCU and/or gateway in an IP environment does not always work. Bandwidth in an IP network is not dedicated to each video device on the network, therefore it is important to design the network accordingly.

The design of systems with Cisco Unified Videoconferencing components is best illustrated by reviewing design scenarios step-by-step. This method provides an understanding of how best to deploy Cisco Unified Videoconferencing for your particular environment. The scenarios in this chapter are designed to simplify the end-user experience and conserve bandwidth. Therefore, no matter what stage your videoconferencing network is in today, you can benefit from deploying one or more of the videoconferencing scenarios described in this chapter.

Multipoint Conferencing

Whenever three or more parties join a videoconference, a Multipoint Control Unit (MCU) is required. The MCU can mix all video streams together and transmit a composite image of all participants back to the originating sources. This composite view (called *continuous presence*) is necessary for all participants to see each other simultaneously. The continuous presence view can display from 2 to 32 windows (participants) in a variety of different layouts. Each layout offers the ability to make one of the windows voice-activated, which is useful if there are more participants in the conference than there are windows to display them all in the composite view. More than 32 participants can be in a single videoconference, and the MCU will allow the last active sites to be displayed. Moreover, as participants join or leave the conference, the Cisco Unified Videoconferencing MCU can automatically adjust the continuous presence view by dynamically changing the layout.

Continuous presence can have multiple participants using different video and audio codecs. In this situation, the MCU must reconstitute all the video streams into a single image while at the same time preserving the video and audio codec of each participant. In addition, the MCU can save network bandwidth by avoiding the need for the sources to transmit large video streams to all sources. Although MCUs can be implemented as either hardware or software, hardware-based MCUs guarantee superior video quality by performing advanced transcoding, transrating, and composition features.

The Cisco Unified Videoconferencing solution consists of the following main components:

- Cisco IOS Gatekeeper

This component provides address resolution and H.323 audio and video setup and tear down.

- Cisco Unified Videoconferencing MCU

This component performs call signaling and multipoint processing of all audio. In addition, the MCU hosts the web interface and controls one or more Enhanced Media Processors (EMPs).

- EMP

The EMP is a dedicated multipoint processor for video. The EMP cannot function without a controlling MCU.

- H.320 Gateway

This component bridges ISDN video endpoints into the IP H.323 infrastructure. There are several ways to use the ISDN gateways in a video solution.

- Desktop Streaming Server

The Desktop streaming server streams the conference to user desktop PCs or laptops. The Desktop server provides the collaboration features for video conferences by allowing users to share desktop or laptop screens and by providing enhanced roster and conference moderation functionality. It also enables conferences to be streamed using Real Time Streaming Protocol (RTSP), to be viewed through a media player or browser plug-in as a webcast.

- Desktop Recording Server

The Recording server can provide recording of the videoconference and desktop sharing. It can also be a repository for the recorded meetings to be viewed later.

- Management and Integration Server

The Management server can be the central entity for configuration and management of the various video devices and the videoconferencing elements. Thus, it can be used to configure, manage, and monitor devices such as MCUs and gatekeepers, among other. The server can also perform conference scheduling, be a gatekeeper, and provide cascading logic for video conferences between different locations to optimize WAN utilization if possible.

The Integration servers allow the enterprise calendaring systems to assign and reserve resources such as MCU ports for the conference, thus ensuring that sufficient resources are available during the conference.

The recommended multipoint hardware platforms are the Cisco Unified Videoconferencing 5000 Series, 3545 and 3515 MCUs. With the Cisco Unified Videoconferencing 3545 and 3515 MCUs, video processing is done by the EMP module, and the MCU performs audio processing and conference control.

The Cisco Unified Videoconferencing 3515 MCU is designed to be a self-contained unit for smaller deployments. Therefore, it has an integrated EMP along with an integrated MCU inside a fixed chassis that cannot be upgraded in the field. By contrast, the Cisco Unified Videoconferencing 5200 or 3545 System is modular and is designed for maximum deployment flexibility. Therefore, MCUs, EMPs, or H.320 gateway modules can be inserted into the Cisco Unified Videoconferencing 5200 or 3545 chassis in order to process audio and video, and this system can be upgraded in the field.

The Cisco Unified Videoconferencing 5000 Series and 3500 Series MCUs support H.323 (standard definition and high definition), SCCP, and SIP endpoint interoperability. The encoder-per-port hardware architecture allows any supported connection speed from 64 kbps up to 2 Mbps, with any supported codec and at any standard-definition resolution. Supported protocols include H.261, H.263, and H.264 video codecs, as well as G.711, G.722, G.722.1, G.723.1, G.728, and G.729A audio codecs.


High-definition video calls provide better video to users. Using traditional or desktop-based high-definition endpoints can enhance the conference experience. High-definition video also requires that the network must be able to process the greater volume of traffic for the calls and must provide Quality of Service (QoS), traffic classification, and call admission control where network resources are not sufficient. For additional information on multipoint conferencing, continuous presence, and voice activation, refer to the *IP Video Telephony* section in the *Cisco Unified Communications Solution Reference Network Design (SRND)*, available at

<http://www.cisco.com/go/ucsrnd>

Initiating a Call

To initiate a multipoint call using a Cisco Unified Videoconferencing MCU, the endpoint dials the appropriate service prefix followed by a conference ID. (The conference ID can be up to 9 digits long.) If a service on an MCU is 40880 for a 384-kbps call, the user might dial 4088011223, the call would be routed to the MCU using the 40880 service prefix, and the MCU would initiate an ad hoc conference with an ID of 4088011223. Users can also initiate a call and invite the other participants by dialing the conference ID, the invite string **, and the E.164 address of the another participant. [Figure 8-1](#) shows the dial sequence of an MCU call with Video 1 initiating an MCU call to 40880123 and inviting Video 2 and Video 3.

Figure 8-1 Example Dial Sequence for Initiating a Call (MCU Invite Call)




In addition, it is possible for an H.323 endpoint to dial the MCU by IP address. The MCU will answer the call with an audio IVR and an on-screen menu.

Cascading MCUs

Cascading MCUs allows larger conferences to be created by combining resources from multiple MCU blades or units. Cascading is also used in distributed MCU environments to save bandwidth on low-speed WAN links (see [Distributed MCUs](#), page 8-5). Cisco Unified Videoconferencing MCUs support cascading. An administrator can cascade MCUs by inviting conference calls on different MCUs to join in a single combined call. In [Figure 8-2](#), a conference was started on each of the four MCUs. To cascade the MCUs in this example, an administrator accessed the web interface on MCU 1 and invited conferences 4088112 on MCU 2, 4088632 on MCU 3, and 4088552 on MCU 4.

Figure 8-2 Cascaded MCU Conference



Note

There is no physical connection made between the cascaded devices. Cascading occurs over the LAN or WAN, allowing MCUs to be distributed across a network. An MCU can invite H.323 endpoints, H.320 endpoints through a gateway, or other MCUs through the web interface on the MCU.

Distributed MCUs

With the ability to cascade multipoint conferences, administrators can build an H.323 video network with distributed MCU services. A distributed MCU architecture saves WAN bandwidth when a conference includes multiple participants on two or more campuses connected by a WAN. By distributing MCUs across the network, it is possible to have multipoint conferences across WAN links without limiting the number of users at remote sites.

Centrally located MCU services require all conference participants to place a call across the WAN to the MCU, while distributed MCUs allow users to call to their local MCU and join the other MCUs into a cascaded conference across the WAN. [Figure 8-3](#) illustrates centralized MCU services, with three users from Campus B joining a conference hosted at Campus A. In this model, all three calls must traverse the WAN link.



Figure 8-3 *Centralized MCU Services, with All Calls Traversing the WAN*


Figure 8-4 illustrates a distributed MCU model with two video terminals at Campus A calling into a local MCU, and three video terminals at Campus B calling into a local MCU. In this model, there is just a single call cascading the two conferences across the WAN. In most distributed networks, a Cisco Unified Videoconferencing 5000 or 3545 System will be located at the headquarters site and a Cisco Unified Videoconferencing 3515 MCU will be located at each large remote site.

Figure 8-4 *Distributed MCU Services, with a Single Call Traversing the WAN*

Dynamic Cascading of MCUs

Dynamic cascading improves upon normal cascading by allowing conferences to be distributed automatically to local MCUs, with no administrator or participant intervention. By using single number reach, dynamic cascading hides the complexity from participants and automatically combines distributed MCUs from multiple locations. Dynamic cascading works through automatic scheduling and intelligent assignment of distributed MCU resources. Participants join a single combined call that reduces the network bandwidth consumption to only a single video call across the WAN.

As the conference grows in size, dynamic cascading will intelligently expand the conference as long as sufficient network resources are available. The Cisco Unified Videoconferencing 5000 Series or 3500 Series MCUs support dynamic cascading. In Figure 8-5, a conference was automatically started on each of the four MCUs. The same number (4088222) was dialed by all eight participants, but dynamic cascading automatically assigned the local terminals to the respective local MCUs. The option is available for the MCUs to dial out directly to the terminals, thereby simplifying the user experience.

Figure 8-5 *Dynamically Cascaded MCU Conference***Note**

Dynamic cascading requires a Cisco IOS Gatekeeper and a videoconferencing scheduler using either Cisco Unified MeetingPlace Video Integration or Cisco Unified Videoconferencing Manager.

**Note**

When deployments have a combination of Cisco Unified Videoconferencing 5000 Series or 3500 Series MCUs, cascading between the two series types of MCUs must be done manually per conference. However, dynamic cascading is supported between MCUs of the same series type.

Scheduled Conferencing

Scheduled videoconferencing provides the ability to reserve resources in advance, as well as to simplify the videoconferencing experience for end users. Administrators can grant users the ability to schedule conferences through email integration or a web interface. The scheduling interface hides all videoconferencing complexity from the end user.

There are three applications used to provide scheduling videoconferencing in a Cisco Unified Videoconferencing environment:

- Cisco Unified MeetingPlace with Video Integration
- Cisco Unified Videoconferencing Manager
- Cisco WebEx with Cisco Unified Videoconferencing integration

Video Scheduling with Cisco Unified MeetingPlace

Cisco Unified MeetingPlace is the recommended method for scheduling video conferences. The participants schedule video resources through email (Microsoft Outlook or Lotus Notes) or a simple-to-use web interface. Web conferencing can easily be added to provide integrated voice, video, and web conferencing solutions that fit any enterprise.

For additional information on Cisco Unified MeetingPlace Video Integration, including detailed video call flows, refer to the *Cisco Unified MeetingPlace* section of the *Cisco Unified Communications Solution Reference Network Design (SRND)*, available at

<http://www.cisco.com/go/ucsrnd>

Video Scheduling with Cisco WebEx

Enterprises that use Cisco WebEx for collaboration can use the WebEx scheduling features for their conferences. WebEx scheduling is done through the cloud services provided by WebEx. WebEx scheduling provides integration for various email systems through plug-ins. The plug-ins enable users to schedule conferences using the enterprise calendaring systems and to choose users and resources that then get updated in the WebEx cloud through the users' WebEx accounts.

WebEx integrations with the MCU and with Cisco Unified Videoconferencing Desktop allows the enterprise to provide the needed QoS for the scheduled calls and for integration with user endpoints. For additional details on the integration, see the section on [Integration and Interoperability](#), page 8-25.

Video Scheduling with Cisco Unified Videoconferencing Manager

For environments without Cisco Unified MeetingPlace, scheduling is accomplished with Cisco Unified Videoconferencing Manager, which is a standalone videoconferencing option only. It enables users to schedule voice and video conferences. Cisco Unified Videoconferencing Manager can be used with the Outlook integration option or with web-based scheduling. Moreover, synchronization with LDAP can minimize user setup and maintenance as well as providing authentication. In addition, LDAP can be used to integrate video terminals into the LDAP directory and associate Class of Service policies to users.


Class of Service (CoS) is accomplished by restricting users to specific meeting types. For example, administrators can grant users the ability to schedule meetings with only standard-rate video service (384 kbps video streams). Likewise, administrators can grant another set of users the ability to schedule meetings with high-definition video service (up to 2 Mbps).



Note

Cisco Unified Videoconferencing Manager can provide call detail records (CDRs) for all videoconferencing calls.

[Figure 8-6](#) illustrates Cisco Unified Videoconferencing Manager together with a Cisco IOS Gatekeeper. Using Microsoft Outlook, a participant in this example has scheduled the video terminals into a conference. When the meeting starts, Cisco Unified Videoconferencing Manager automatically out-dials all terminals into the call. All of this occurs without any user intervention aside from using Outlook to schedule the video conference. Optionally, Cisco Unified Videoconferencing Manager can also dynamically cascade all terminals. Thus, bandwidth consumption on the network is minimized to a single call.

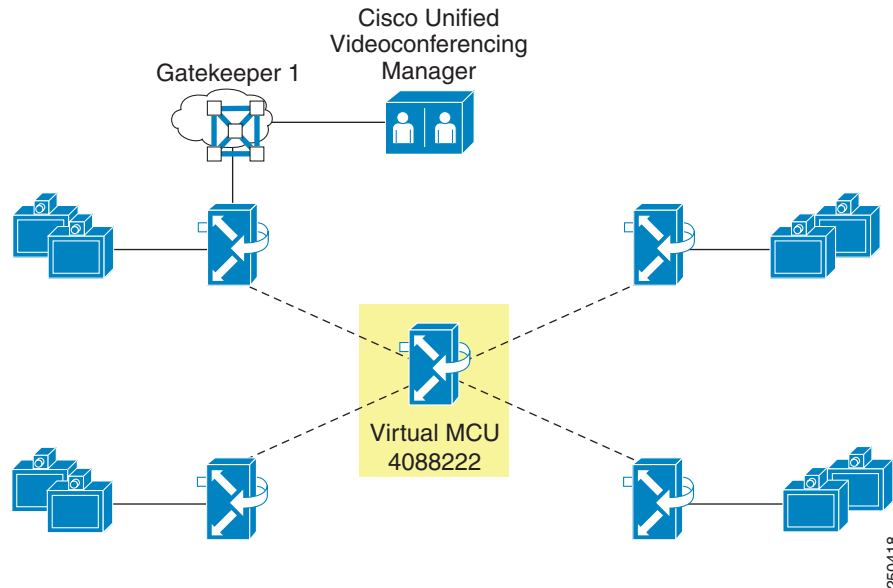
Figure 8-6 Cisco Unified Videoconferencing Manager**Note**

Video media streams do not flow through Cisco Unified Videoconferencing Manager. Video media terminates on the MCUs.

Virtual MCU

Cisco Unified Videoconferencing Manager has the capability to allow all MCU resources to appear as a single MCU. This virtual MCU concept provides a single conference ID across multiple MCUs. Therefore, this reduces the complexity of scheduling because participants can dial a single number and automatically be associated to the closest MCU. Likewise, the system can dial out to the participants. The virtual MCU provides the logic for dynamic cascading.

In [Figure 8-7](#), the same conference call illustrated in [Figure 8-6](#) is now depicted as a purely logical entity. The four distributed MCUs become a single virtual MCU when using Cisco Unified Videoconferencing Manager for scheduling the video conference. Participants can dial the number 4088222 located on the meeting invitation, and Cisco Unified Videoconferencing Manager will assign video terminals to local MCU resources. The virtual MCU has intelligent topology awareness of local resources, so it can avoid unnecessary bandwidth congestion over the WAN.

Figure 8-7 Virtual MCU Created by Cisco Unified Videoconferencing Manager**Note**

To enable dynamic cascading, Video terminals and MCUs must be associated to a location in the Resource Manager component of Cisco Unified Videoconferencing Manager.

Cisco Unified Videoconferencing Manager Components

Cisco Unified Videoconferencing Manager contains two major software components:

- Resource Manager — A scheduling component
- Network Manager — A network management component

Thus, Cisco Unified Videoconferencing Manager combines a scheduling application with a network management application.

Resource Manager Component of Cisco Unified Videoconferencing Manager

The Resource Manager component of Cisco Unified Videoconferencing Manager controls all aspects of videoconferencing and interfaces with MCUs and the gatekeeper. Resource Manager is the main control point for all incoming and outgoing video calls, and it makes all decisions regarding resource utilization and cascading of resources. In addition to basic MCU cascading to increase conference size, Resource Manager can intelligently select which MCU to use based on internally defined locations for each participant. Multisite video meetings can result in cascading MCUs local to user groups, thus creating one link across a WAN between cascaded MCUs.

The Resource Manager component of Cisco Unified Videoconferencing Manager has the following characteristics:

- Resource Manager resides on a separate Windows-based server and cannot reside on the same server with other Cisco Unified Communications components.
- Resource Manager sits between the MCUs and all other components.

- Resource Manager contains a special integrated gatekeeper to which only the MCUs register. The MCUs must register to Resource Manager.
- All routing decisions for MCU selection and cascading are made by Resource Manager.
- Resource Manager terminates and authenticates all inbound calls, and only media is sent to the MCUs.
- Resource Manager originates and maintains all outbound calls, and only media is sent to the MCUs.
- Resource Manager requires the use of an external gatekeeper for outbound calls.
- Resource Manager does not register to the gatekeeper it uses for outbound calls.
- Resource Manager does not support a direct SIP trunk connection from Cisco Unified Communications Manager.
- Resource Manager does not support a direct H.323 gatekeeper-controlled trunk connection from Cisco Unified Communications Manager.
- All scheduled video meetings are replicated in Resource Manager, with appropriate resources reserved.
- Participant information and real-time status are relayed from the MCUs to Resource Manager.
- Other than media stream termination, only Resource Manager communicates with the MCUs.
- Video Terminals can be defined in Resource Manager and reserved within Cisco Unified Videoconferencing Manager when users create a meeting.
- Any endpoint (H.323, SCCP, or SIP) can be defined in Resource Manager as a Video Terminal.
- MCU selection can be impacted by assigning locations to MCUs and Video Terminals defined in Resource Manager.

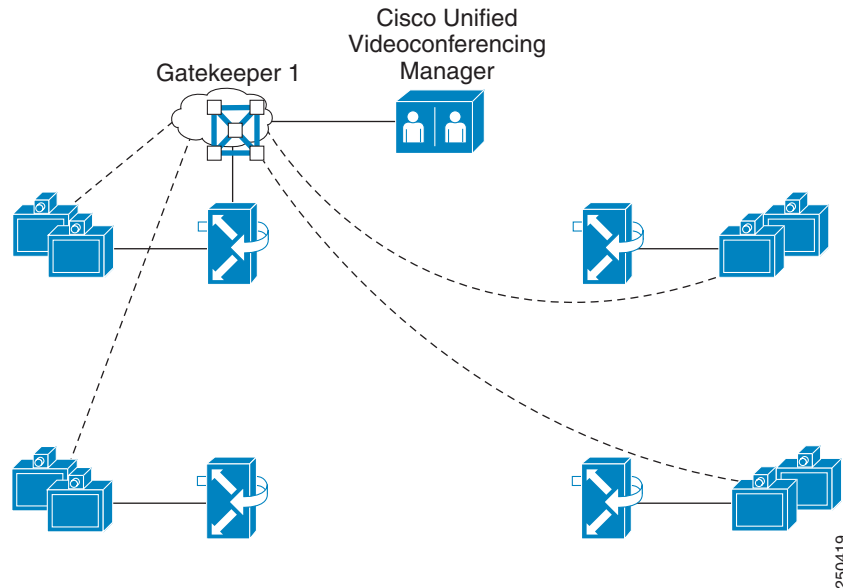
**Note**

The MCUs register to the Resource Manager internal gatekeeper as an MCU instead of as a gateway. The parameter is located on the MCU under Advanced H.323 Settings.

Video Endpoint and MCU Registration

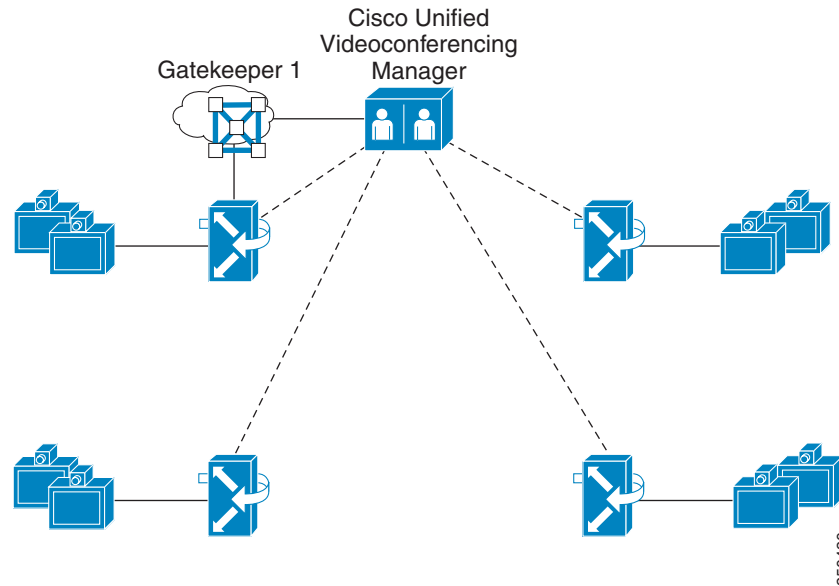
When using Cisco Unified Videoconferencing Manager, all H.323 video endpoints should be configured to register to the Cisco IOS gatekeeper. [Figure 8-8](#) shows all H.323 video endpoints configured with the Cisco IOS gatekeeper as their designated gatekeeper.

Figure 8-8 Video Endpoint H.323 Registration to a Cisco IOS Gatekeeper



You must register Cisco Unified Videoconferencing MCUs and Cisco Unified Videoconferencing Gateways with the Cisco Unified Videoconferencing Manager internal gatekeeper, which is included with the product. This preserves the virtual MCU features and conference setup, and it allows for conference control via a web interface during a meeting.

[Figure 8-9](#) shows all Cisco Unified Videoconferencing MCUs configured with the Resource Manager internal gatekeeper as their designated H.323 gatekeeper.

Figure 8-9 MCU H.323 Registration to Cisco Unified Videoconferencing Manager

For integration of video endpoints or Cisco Unified Videoconferencing MCUs with Cisco Unified Communications Manager, refer to the *IP Video Telephony* section in the *Cisco Unified Communications Solution Reference Network Design (SRND)*, available at

<http://www.cisco.com/go/ucsrnd>

Gatekeeper Integration with Resource Manager

The Resource Manager internal gatekeeper and the Cisco IOS Gatekeeper are configured as H.323 gatekeeper neighbors.

Example 8-1 illustrates the minimal configuration necessary for the Cisco IOS gatekeeper to integrate with the Cisco Unified Videoconferencing Manager. The Resource Manager internal gatekeeper is a remote zone from the perspective of the Cisco IOS gatekeeper.

Example 8-1 Cisco IOS Gatekeeper Minimal Configuration with Cisco Unified Videoconferencing Manager

```
gatekeeper
  zone local HQ cisco.com 10.1.1.1
  zone remote CUVM cisco.com 10.2.2.1 1719
  zone prefix CUVM 83* (This entry matches service prefixes configured
on the MCUs)
  gw-type-prefix 1#* default-technology
  lrq forward-queries add-hop-count
  no use-proxy HQ remote-zone CUVM inbound-to terminal
  no use-proxy HQ remote-zone CUVM outbound-from terminal
  no use-proxy HQ default inbound-to terminal
  no use-proxy HQ default outbound-from terminal
  no shutdown
```

When integrating Resource Manager into an infrastructure gatekeeper environment with Cisco Unified Communications Manager (Unified CM), the deployment options discussed in this section will bypass Unified CM when one H.323 endpoint calls another H.323 endpoint. Having the endpoints controlled by Unified CM is a deployment option not covered here due to the focus on environments that have yet to

deploy Cisco Unified CM. For information on integrating Cisco Unified CM with video endpoints, refer to the *IP Video Telephony* section in the *Cisco Unified Communications Solution Reference Network Design (SRND)*, available at

<http://www.cisco.com/go/ucsrnd>

Resource Manager Redundancy Considerations

Resource Manager is limited to server component redundancy, and it does not have any software-level redundancy. The server on which Video Administration is deployed should contain redundant components to minimize the risk of downtime. Cisco Unified Videoconferencing Manager is deployed on an MCS-7825 server regardless of the total number of endpoints, MCUs, or ISDN gateways under control.

Redundancy of the Cisco IOS Gatekeeper via Hot Standby Router Protocol (HSRP) provides for outbound calling redundancy from Resource Manager. If the primary gatekeeper fails, call resolution requests are automatically sent to the backup gatekeeper. Resource Manager does not need to know about the backup gatekeeper or primary gatekeeper state to accomplish this failover.

Gatekeeper Clustering and Alternate Gatekeeper

With the gatekeeper clustering and alternate gatekeeper methods, an issue arises with outbound calls from Resource Manager. Outbound calls from Resource Manager are sent to the Cisco IOS Gatekeeper without any registration to that gatekeeper. This prevents Resource Manager from being aware of the gatekeeper state and the existence of an alternate gatekeeper.

Network Manager Component of Cisco Unified Videoconferencing Manager

Cisco Unified Videoconferencing Manager can provide network management of MCUs and video endpoints. The Network Manager component of Cisco Unified Videoconferencing Manager queries the Cisco IOS Gatekeeper to discover all video endpoints on the network. This allows organizations to have a global view of MCUs, gatekeepers, and video endpoints. This video infrastructure can change dynamically yet be discovered automatically. Network Manager can also provide software upgrades to MCUs and many third-party endpoints.

In addition, the Network Manager component can provide real-time status of MCU health, alarms, events, active calls, and active conferences. Moreover, recipients can be alerted to events through email, and SNMP traps can be forwarded to other management systems.

Cisco Unified Videoconferencing Manager Deployment Considerations

There are three general types of deployments for Cisco Unified Videoconferencing Manager:

- Scheduling mode

This mode allows scheduling integration with Microsoft Outlook and simplifies all aspects of video conferencing setup. This mode can be further categorized as with or without dynamic cascading.

- Network Management mode

The network management features of Cisco Unified Videoconferencing Manager can be used solely for endpoint and MCU management without any participation in scheduling or dynamic cascading.

- Scheduling mode with Network Management mode

This method utilizes all of the components of the Cisco Unified Videoconferencing Manager.

Video Gateways

The Cisco Unified Videoconferencing 3545 Dual PRI, 3527 Single PRI, and 3522 4-Port BRI Gateways give enterprises the ability to connect ISDN-based H.320 systems with IP-based H.323 videoconference endpoints. These gateways provide translation services between H.320 and H.323 networks to convert multimedia information between circuit-switched ISDN and IP networks. The gateway also supports G.711, G.722, G.722.1, G.723.1, and G.728 audio codecs. In addition, voice transcoding between IP and the Public Switched Telephone Network (PSTN) is supported using G.711, G.723, or G.728. These systems enable users to videoconference with others users via the LAN or the PSTN, regardless of location. The Cisco 3500 Series Gateways also provide in-band DTMF conversion to out-of-band DTMF.

The three video gateway models provide the following features:

- **Cisco Unified Videoconferencing 3522 BRI Gateway**
This gateway can be configured with two or four BRI ports. When the gateway is equipped with BRI ports, it can support calls up to 384 kbps on aggregated channels.
- **Cisco Unified Videoconferencing 3527 PRI Gateway**
This gateway is a self-contained system that supports a high volume of calls over a single high-speed ISDN PRI connection, allowing dynamic allocation of its 23 B channels.
- **Cisco Unified Videoconferencing 3545 Gateway Modules**
These gateway modules are either a two-port PRI T1/E1 module that supports a high volume of calls over multiple high-speed ISDN PRI connections, or a four-port serial module for IP connectivity to older ISDN H.320 videoconferencing endpoints.

[Table 8-2](#), [Table 8-3](#), and [Table 8-4](#) list the maximum number of calls supported per platform.

Table 8-2 Cisco Unified Videoconferencing 3522 BRI Gateway Call Handling Capacity

Number of Calls	Capacity
1	384 kbps or 412 kbps
2	256 kbps
4	128 kbps
8	64 kbps

Table 8-3 Cisco Unified Videoconferencing 3527 PRI Gateway Call Handling Capacity

Call Type	Maximum Number of Calls Using One E1 PRI Line	Maximum Number of Calls Using One T1 PRI Line
Voice (64 kbps)	30	23
2B video (128 kbps)	15	11
6B video (384 kbps)	5	3
12B video (768 kbps)	2	1

Table 8-4 Cisco Unified Videoconferencing 3545 PRI Gateway Call Handling Capacity

Call Type	Maximum Number of Calls Using One E1 PRI Line	Maximum Number of Calls Using One T1 PRI Line	Maximum Number of Calls Using Two E1 PRI Lines	Maximum Number of Calls Using Two T1 PRI Lines
Voice (64 kbps)	30	23	60	46
2B video (128 kbps)	15	11	30	23
6B video (384 kbps)	5	3	10	7
12B video (768 kbps)	2	1	5	3

Outbound Dialing Service Prefixes

Video gateways must be configured with service prefixes to define the speed of outgoing calls and calling routing to the video gateway. In telephony systems, dialing 9 to access an outside line is very common. In order to keep dialing strings consistent with existing voice dial plans, Cisco recommends using 9# for video gateway service prefixes. Using the # in the service prefix ensures that ISDN users do not access the IVR and hairpin the call back out the ISDN network. The # is used as a delimiter by the gateway and prevents hair pinning from the ISDN network.

From the users' perspective, they will have to dial the service prefix, which in this case is equivalent to an access code, followed by the ISDN number of the H.320 video unit. For this configuration, a local (intra-zone) gateway is used whenever one is present. In zones that do not contain a gateway, the administrator should assign a gateway in another zone as the primary gateway for the local zone. Configure location request (LRQ) forwarding or a static hopoff statement to route all calls to a zone with a gateway for PSTN access.

Network Load Balancing

The Cisco Unified Videoconferencing Gateway supports the network load balancing (LAN to PSTN) feature. This feature allows users to build a pool of gateways for PSTN access. Network load balancing creates a larger number of access lines serviced by a single set of service prefixes.

Gatekeepers can perform load balancing on the network using feedback from the gateway in the form of Resource Availability Indication (RAI) messages that inform the gatekeeper of gateway resource availability. If the gateway is unavailable, the gatekeeper performs line hunting operations to route the call to an alternate gateway. When you set the gateway for RAI and Resource Availability Confirmation (RAC), it sends periodic RAI messages that inform the gatekeeper of the current resource availability in the gateway. The gatekeeper responds with Resource Available Confirmation (RAC) messages to acknowledge receipt of the RAI messages.

To implement network load balancing, you configure multiple gateways with identical service prefixes and register them with the same gatekeeper. Outbound PSTN calls are sent to the gateways based on resource availability, using RAI and RAC. In the gateway configuration, you set utilization parameters based on gateway resource percentages.

The main gateway configuration parameters for line hunting are:

- Utilization (percent load) for sending RAI ON message
- Utilization (percent load) for sending RAI OFF message

The RAI ON message tells the gatekeeper that resources are running low on the gateway that sent the message, and the gatekeeper should not forward any more calls to that gateway. (The default for sending a RAI ON is 80% load.) The RAI OFF message tells the gatekeeper that there are enough available resources on the gateway, and calls can be forwarded to the gateway again. (The default for sending a RAI off is 60% load.) Periodic RAI messages are sent from the gateway to the gatekeeper when one of the above thresholds is not achieved in a specified period of time (The default period for these messages is 30 seconds.)

Cisco Unified Border Element

The Cisco Unified Border Element is a Cisco IOS software component that incorporates the gatekeeper and border element functions for an H.323 video network. The Cisco IOS gatekeeper provides endpoint registration and call resolution for large H.323 video networks. The border element provides topology hiding, network isolation, call admission control and Quality of Service (QoS). The border element is increasingly used for dynamic call admission control with RSVP devices or on behalf of devices that do not support RSVP. Deployments with a border element simplify Network Address Translation (NAT) and firewall integration.

Gatekeeper

The Cisco gatekeeper performs all call routing and address registration (RAS) for all H.323 video components. The gatekeeper is one of the most important components in an H.323 network because it is the central management device for the H.323 video network and it performs functions required for a successful H.323 video deployment. Some of the most commonly used functions of the Cisco IOS gatekeeper include:

- H.323 component registration and call routing

The gatekeeper registers the IP address, E.164 address, H.323-ID, device type, and signaling ports for all the video infrastructure components. This registration allows the gatekeeper to provide call routing for all devices that are registered with the it.

- Bandwidth management

Managing video bandwidth on IP networks is an essential feature of any gatekeeper. By setting the following bandwidth parameters, you can configure the Cisco gatekeeper to manage the bandwidth in a zone, between zones, or per call.

- Inter-zone: Total bandwidth allowed from a local or default zone to and from all other zones.
- Remote: Total bandwidth allowed from all local zones to and from all remote zones.
- Session: Bandwidth allowed per session in a zone.
- Total: Total bandwidth allowed in a zone.
- Resource Availability Indicator (RAI): Endpoints supporting RAI update the gatekeeper of availability or unavailability.
- Circuit-id: Calls with this circuit-id call identifier are limited by the gatekeeper to max-calls per endpoint.
- Max-calls: Maximum number of calls that the gatekeeper will permit per endpoint.

- Authentication, authorization, and accounting (AAA) support

The Cisco gatekeeper works in conjunction with Remote Access Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System (TACACS) servers to provide authentication of devices and accounting via call detail recording (CDR).

- Via-zone

The gatekeeper can include the border element in the call flow by using the via-zone with **invia** and **outvia** configurations to include the border element for incoming calls or outgoing calls.

The Cisco gatekeeper also supports features such as the following, which enable users to build reliable and scalable H.323 networks:

- Hot Standby Router Protocol (HSRP) enables administrators to build a standby gatekeeper that becomes active if the primary gatekeeper fails.
- Gatekeeper Update Protocol (GUP) enables multiple gatekeeper entities to behave as a single gatekeeper cluster, providing greater resilience for endpoint registrations and calls. (Limits on scalability must be taken into consideration.)
- Directory Gatekeeper, or Location Request (LRQ) forwarding, enables administrators to build large multi-tier networks, minimizing the configuration required in the lower-tier gatekeepers. When a call is made in a lower-tier zone and a match is not found, the call is automatically forwarded up to the directory gatekeeper for resolution. (For more information on directory gatekeepers, refer to the [Call Routing](#) chapter.) [Figure 8-10 on page 8-20](#) illustrates a network configured with two regional directory gatekeepers, one at Site E and another at Site F.

Redundancy

A video network needs redundancy for the endpoints to register and for call resolution. Failure in call control devices can impact the ability to do videoconferencing. Any of the following methods can provide redundancy:

- Hot Standby Router Protocol (HSRP)

HSRP can be used to provide redundancy at the IP layer for the gatekeeper. Thus, endpoints can use one IP address for the gatekeeper address. Failures of devices providing hot standby for the gatekeeper IP address do not impact the endpoints drastically. The blackout period after the failure of the primary device and until the endpoints are fully registered and functional with the standby device, depends on the device time-outs and registration retries. This method is recommended for the directory gatekeeper because its main function is call resolution and it might not have any endpoints registered. (For more information on directory gatekeepers, see [Routing Inter-Zone Calls Using a Directory Gatekeeper](#), page 7-9.)

- Alternate gatekeeper

One of the simplest methods of providing redundancy is to configure the endpoints with a list of gatekeepers in order of preference, so that if the first gatekeeper in the list fails, the devices register with the next gatekeeper in the list. However, not all endpoints have this support. Cisco Unified Communications Manager, Cisco Unified Border Element, and Cisco Gateways support alternate gatekeeper functionality.

- Gatekeeper Update Protocol (Cisco GUP)

The Cisco Gatekeeper Update Protocol (GUP) enables gatekeepers in the cluster to update information on registrations and calls in the cluster. Alternate gatekeeper support is needed in the endpoints. When an endpoint registers or unregisters with its primary gatekeeper, that gatekeeper updates the remaining gatekeepers in the cluster about the change. Similar information is shared

when the gatekeeper resolves a call. If any gatekeeper in the cluster fails, the endpoint registers to the alternate gatekeeper that the primary gatekeeper had informed when the endpoint initially registered. With GUP, if a gatekeeper failure occurs, active calls are not disconnected. GUP is recommended where the number of endpoints is not very large. However, most video endpoints do not support alternate gatekeeper, therefore they cannot support the use of GUP.

Scalability

Call routing scalability can be achieved by using the hierarchical model. Common regions or areas are grouped together to be serviced by access gatekeepers which then point to a directory gatekeeper. The directory gatekeeper may not have any endpoints registered to it and may have a role of routing calls between access gatekeepers or to hop-off zones. The access gatekeeper may have endpoints registered to it or, in larger deployments, may have an additional level of gatekeepers to which it can route calls.

HSRP

Hot Standby Router Protocol (HSRP) enables a set of routers with the Cisco Unified Border Element to work together as a single virtual gatekeeper or border element. You can implement this feature by creating a *phantom* router that has its own IP and MAC addresses.

Based on the priority given by the network administrator, one of the HSRP gatekeepers in each group is selected to be active and the other to be standby. The gatekeeper with the highest priority serves as the active gatekeeper. The active gatekeeper does the work for the HSRP phantom. If an end node sends a packet to the phantom's MAC address, the active gatekeeper receives that packet and processes it. If an end node sends an Address Resolution Protocol (ARP) request for the phantom's IP address, the active gatekeeper replies with the phantom's MAC address.

The HSRP gatekeepers (both active and standby) watch for *hello* packets to monitor the status of each other. The gatekeeper group learns the hello and hold timers, as well as the standby address to be shared, from the active gatekeeper. If the active gatekeeper becomes unavailable for any reasons (such as power failure, scheduled maintenance, or failure to respond to three successive hello packets), the standby gatekeeper assumes the active role transparently within a few seconds. Because the new active gatekeeper assumes both the IP and MAC addresses of the phantom, video terminal registrations time out, and the terminals re-register with their same IP address to the newly active gatekeeper.



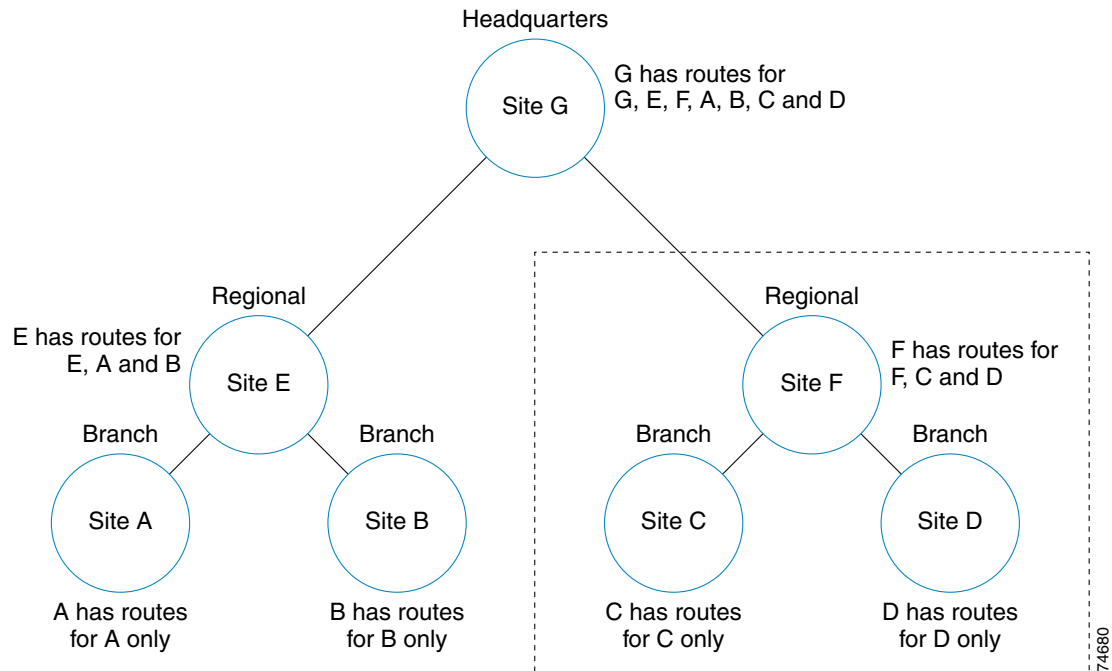
Note

When configuring gatekeepers and border elements on routers supporting HSRP, configure the border elements to register with the virtual, or phantom, IP address of the gatekeeper pair. This configuration enables both border elements to register with the active gatekeeper so that video calls are load-balanced between the two devices. If the primary router fails, the border element on the standby router registers with the now active gatekeeper, and calls are forward through it. The border element configured on the primary router will not re-register with the standby router if the primary router fails.



Note

Gatekeeper clustering is not supported in a videoconferencing environment. For clustering to work, video endpoints would have to support alternate gatekeepers, but currently there are no video terminals with this support.

Figure 8-10 Network with Two Directory Gatekeepers

Border Element

The Cisco Unified Border Element has the ability to proxy call signaling and media by termination and re-origination. The border element can register with the gatekeeper and be inserted in the call flows. The Cisco Unified Border Element has the following functionality:

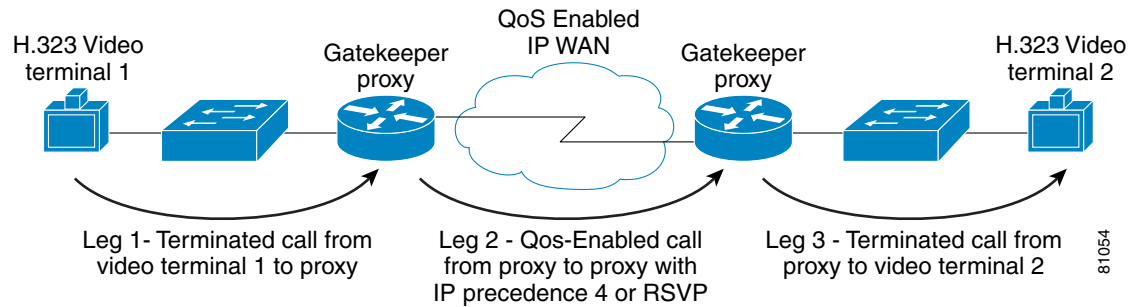
- Network isolation is achieved by re-origination of signaling and media for the call. You can use multiple interfaces to provide enhanced isolation.
- Classification of signaling, audio, and video traffic can be done with IP Precedence, DSCP, and Resource Reservation Protocol (RSVP).
- Call admission control with the help of the Resource Availability Indicator (RAI) as a gateway function is widely used. In addition, call treatment based on CPU, memory, total calls, circuit ID capacity, and gatekeeper bandwidth can be used in various combinations to provide efficient call admission control. RSVP can be used for dynamic call admission control.
- Topology hiding is used for inter-working various types of signaling and protocol.
- The border element can provide ease of deployment with firewalls and NAT because it can be a single trusted entity in either the inside network or the DMZ. With its protocol awareness for NAT and firewalls, the border element can also reduce management overhead.



Note

Because the gatekeeper is an optional device, the border element can also be used as a standalone device. The border element will have call routing management overhead in this case.

Figure 8-11 illustrates a call across a WAN link through border elements.

Figure 8-11 Call Across a WAN Through Border Elements**Note**

Single-legged Cisco Unified Border Element is supported in Cisco IOS Release 12.3(15)T or later.

Collaboration for Desktop Sharing

Deployments that want to leverage collaboration in their video networks can use any of the following methods to share user desktop screens.

H.239-Based

With this method, an additional video channel is established by the endpoints. The desktop screen is sent using this additional video channel during the call. In this method, the endpoint capabilities determine the screen sharing resolution and characteristics. All participants must support this mechanism. Audio and video for the conference is through regular channels by the endpoints, as in point-to-point calls.

The following design considerations apply to this method:

- Endpoints must support H.239.
- Intermediate devices such as call agents, border elements, NAT devices, or firewalls must support H.239.
- Additional WAN bandwidth might be needed and must be provisioned and accounted for.

Server and Client Communication

This method is commonly used by collaboration systems. Cisco Unified MeetingPlace, WebEx, and other systems use a similar method. In this method, a user-side client is used that can run on a PC or laptop and that makes desktop screen sharing possible. The server is the application server providing the services to the clients. Users that do not have access to a PC cannot see the screen being shared. Audio and video for the conference is through regular channels by the endpoints, as in point-to-point calls.

The following design considerations apply to this method:

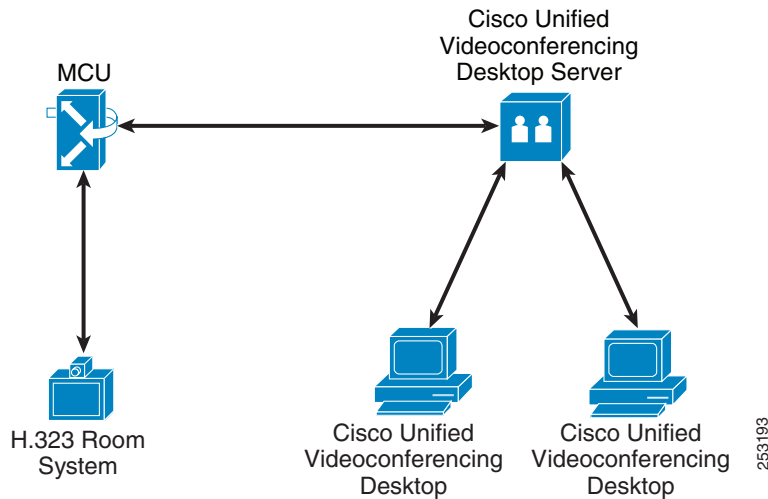
- Deployments need to consider the operating system and browsers for the user client on the desktop or the laptops.
- Consider the way the client software can be deployed for internal and external devices. Most systems will check installations or upgrades before users join the collaboration sessions through browsers.
- Intermediate devices such as NAT devices or firewalls must support traversal for the user-side client, if applicable.

- Quality of Service classification and call admission control are required for calls.
- Additional WAN bandwidth might be needed and should be accounted for and provisioned to accommodate the collaboration traffic.

Cisco Unified Videoconferencing Desktop

The Cisco Unified Videoconferencing Desktop solution is similar to the server and client communication. (See Figure 8-12.) The client component takes care of the video, audio, and desktop sharing capability. The server component is the application server for collaboration. In addition to this, it interacts with the MCU that hosts the video conference and enables traditional endpoints in the conference to view the content shared through the desktop. On the client side it uses a web protocol such as HTTP/HTTPS to tunnel the data, and on the endpoint side it uses H.239 for screen sharing. In this way, the Cisco Unified Videoconferencing Desktop provides a combination of H.239 and server/client communication for the collaboration solution.

Figure 8-12 Cisco Unified Videoconferencing Desktop



The following design considerations apply to this method:

- Video endpoints must support H.239.
- Intermediate devices such as call agents, border elements, NAT devices, or firewalls must support H.239.
- Deployments need to consider the operating system and browsers for the user client on the desktop or the laptops.
- Consider the way the client software can be deployed for internal and external devices. Most systems will check installations or upgrades before users join the collaboration sessions through browsers.
- Quality of Service classification and call admission control are required for calls.
- Intermediate devices such as NAT devices or firewalls must support traversal for the Cisco Unified Videoconferencing Desktop client, as applicable.
- Additional WAN bandwidth might be needed and should be accounted for and provisioned to accommodate the collaboration traffic.


- Cisco recommends co-locating the desktop server and the conference MCU because the Cisco Unified Videoconferencing Desktop server initiates a video channel to the MCU for every desktop client.

Conference Streaming

Some conferences such as training sessions, large team meetings, or panel discussions have the key people such as the trainer, team leader, or moderators and panel numbers participate interactively in the conference, while the large majority of users are viewing participants only. Rather than have every participant join the conference with their video endpoints, streaming such conferences is a more efficient and scalable way to present the conference to viewers. (See [Figure 8-13](#).) This not only optimizes the conferencing resources, but it also allows users to view the conference from simple devices such as media players that support streaming. For such solutions Real Time Streaming Protocol (RTSP) is commonly used.

The Cisco Unified Videoconferencing Desktop server supports streaming of live conferences. The participants who will be transmitting in the conference can join the conference through conventional endpoints that support interactive calls. The server uses the active conference from the MCU and converts it into an RTSP stream. The viewers can also view the conference using common players such as Quicktime player, VLC media player, Windows Media player, Real media player, or any other players that support RTSP.

Figure 8-13 Conference Streaming



The following design considerations apply to conference streaming:

- The streaming media might lag behind actual conference.
- Quality of Service classification and policing traffic are required.
- Unicast streaming may be used with a small number of viewers.
- Multicast streaming can be used if there is a need to conserve WAN bandwidth or if there is a large number of users viewing conferences.
- Additional WAN bandwidth may be needed and should be accounted for and provisioned to accommodate the streaming traffic.

Conference Recording

Recording the conference provides users with the capability to view the conference at a later time. Recording servers are separate application servers that store the recording and make it available for later viewing. The Cisco Unified Videoconferencing Desktop server supports recording of conferences. Conferences can be recorded by selecting that option when scheduling the conferences or by the moderator through the user desktop.

The following design considerations apply to conference recording:

- The recording server should be deployed on a separate server. Co-hosting with other Cisco Unified Videoconferencing server should be done only after carefully considering the server capacity and scalability.
- Recordings must be stored and managed on the desktop recording server. Access can be through the Recording access webpage. Security of these recordings depends on the operating system of the recording server.
- The recording server capacity should be planned based on the type of recordings, the frequency of recording, and how long the recordings will be available to users before the system archives them.
- Enterprises need to define a conference recording storage and retention policy.

Conference recording is dependent on the recording bitrate configured on the recording server. [Table 8-5](#) provides brief guidelines on the storage needed for recording meetings.

Table 8-5 **Storage Guidelines for Recording Servers**

Recording Bitrate (kbps)	Amount of Data Recorded (MBytes/min) ¹
256	2.3
384	3.5
512	4.7
768	7.0
1024	9.2

1. These values include a factor of 20% for overhead.

Use the following general formula to calculate the minimum server storage capacity needed for recording conferences:

Minimum Storage Capacity = [Recorded Mbytes/min (User corresponding values from [Table 8-5](#) based on the recording bitrate)] * [Recording time per day in minutes] * [Number of days] * [Number of simultaneous conferences]

For example, if the recording rate is 256 kbps and you want to be able to record 15 full days (24 hours) of 5 simultaneous conferences, the minimum required storage capacity would be:

Minimum Storage Capacity = 2.3 MB/min * [60 min/hr * 24 hr/day] * [15 days] * [5 conferences]

Minimum Storage Capacity = 248.4 GBytes

Integration and Interoperability

The Cisco Unified Videoconferencing MCU can be used for various system integration needs and to fulfill interoperability requirements with various systems, as described in the following sections.

Telepresence Integration

The Cisco Unified Videoconferencing 5000 Series or 3500 Series MCUs provide Cisco Telepresence Multipoint Switch integration to traditional videoconferencing networks. This allows H.323 terminals, desktop telephony, and other videoconferencing devices to participate in a conference with Cisco Telepresence systems. The Cisco Telepresence Multipoint Switch extends the conference to the Unified Videoconferencing MCU. This enables traditional and video telephony devices to join into the MCU conference, while the Cisco Telepresence system is conferenced through the Cisco Telepresence Multipoint Switch as a single conference.

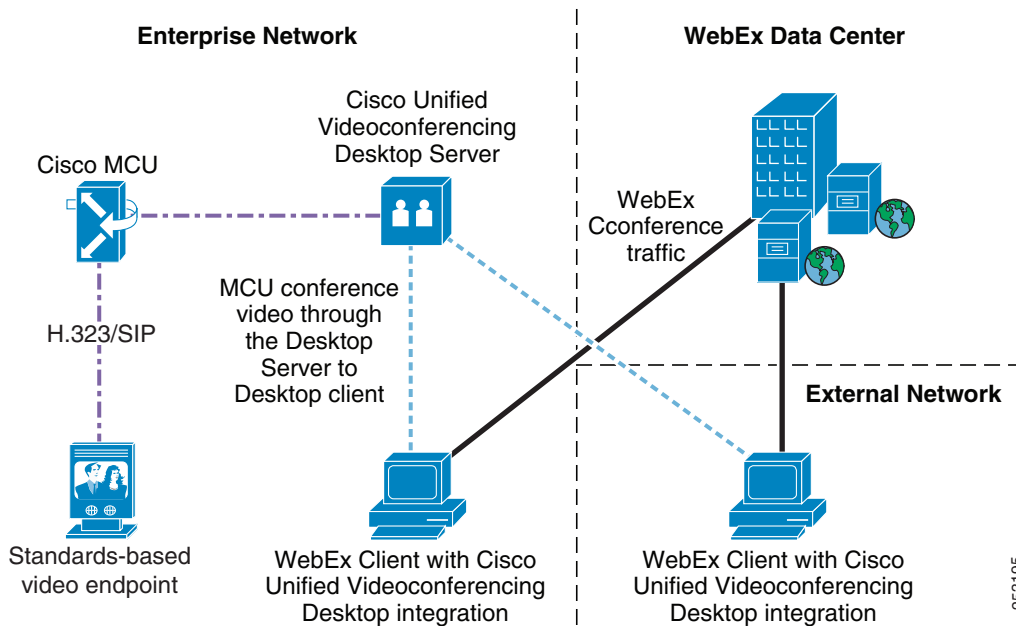
For additional details, refer to the Cisco Telepresence Multipoint Switch documentation at

http://www.cisco.com/en/US/products/ps7315/tsd_products_support_series_home.html

WebEx Integration

Through the deployment of the entire Cisco Unified Videoconferencing portfolio, especially the Desktop Streaming Server, WebEx clients on WebEx site versions T27 or later can be configured to use the Advanced Video plug-in. Advanced Video enables on-premise room systems, video telephony, and even telepresence systems to participate in a video conference with webcams both on-premise and over the Internet.

Integrating the Cisco Unified Videoconferencing MCU of the enterprise with WebEx leverages the Cisco Unified Videoconferencing Desktop solution to bring in the video of traditional desktop systems and traditional videoconferencing systems and presenting it in a WebEx collaboration session. In this way, the Native Video support in WebEx gets replaced with the one from the enterprise MCU. (See [Figure 8-14](#).)

Figure 8-14 WebEx Integration**Microsoft Office Communicator**

Enterprises that deploy Microsoft Office Communicator need integration for desktop video. The desktop client uses Cisco Unified Videoconferencing Desktop server integration with Microsoft Office Communicator to display video from the MCU in Microsoft Office Communicator, while using video from the desktop cameras in the conference as user video. Scheduling conferences is done through Outlook Plug-in integration.

Lotus Sametime

Enterprises that deploy Lotus Sametime need integration for desktop video. The desktop client uses Cisco Unified Videoconferencing Desktop server integration with Sametime to display video from the MCU, while using video from the desktop cameras in the conference as user video. Scheduling conferences is done through Lotus Notes Plug-in integration.

For more details on the integrations, refer to the latest Install and Upgrade Guides for the Cisco Unified Videoconferencing Desktop, available at

http://www.cisco.com/en/US/products/ps7088/prod_installation_guides_list.html

Connecting the Video Network with External Untrusted Networks

Enterprises want to leverage the benefits of videoconferencing not only for their internal users but also with external enterprises and users. The following design considerations can help identify the best approach.

External Connectivity

External connectivity to the enterprise network is a key requirement. The external devices can also be devices at the enterprise partners, vendors, or suppliers within a closed user group (CUG), but they might be untrusted by the enterprise. In some cases the external devices may be devices on the networks of other enterprises or on the untrusted Internet.



Note

Quality of Service for calls might require special consideration because QoS on the Internet is best-effort and cannot be guaranteed. Enterprise QoS can address quality within the enterprise network. For additional details, refer to the *Enterprise QoS Solution Reference Network Design Guide*, available at <http://www.cisco.com/go/designzone>.

Dial Plan

A key consideration for external connectivity is how the enterprise routes the calls. Does the enterprise act as a peer with other enterprises to route calls between themselves through call agents, gatekeepers, or SIP proxies, based on a protocol of choice?

Enterprises can also assign E.164 addresses to external endpoints so that calls can be routed only within the enterprise network to those external endpoints, thus controlling the dial plan to them. In this case, you must ensure that external networks cannot use the enterprise-assigned dial plan to reach these external devices. Enterprises can enforce rules or registration to prevent unauthorized access.

Enterprises that wish to support external devices that dial by IP address for conferences can use mechanisms that direct callers to a fixed enterprise number and present them to an auto-attendant that provides the endpoint information to select and join conferences using DTMF digits. Such deployments can leverage the Cisco Unified Border Element, which can service calls that do not have a called number in the call setup and can present that call with an auto-attendant number as the called number to an MCU with auto-attendant service.

External Endpoints or Devices

Identifying the type of endpoint can significantly change the way connectivity to external networks can be achieved.

Traditional endpoints or desktop endpoints support standards-based protocol such as H.323 or SIP and can register with the gatekeeper or SIP proxy as applicable. These endpoints then can register with internal trusted devices. Cisco recommends firewalls that support application-level gateway inspection. Topology hiding devices can provide a more efficient method for deployment and management.

Desktops or laptops with no clients might not support standards-based protocols but might use tunneling through HTTP/HTTPS to the communication servers. To service such clients, the Cisco Unified Videoconferencing Server can be deployed in the DMZ network. Users can access the server using the Microsoft Internet Explorer (IE) browser and can authenticate through credentials and the conference personal identification number (PIN) to participate in the conference. The lightweight software for the desktop or the PC can be downloaded and installed if needed.


Firewalls and Network Address Translation (NAT)

Firewalls provide security for unauthorized access to internal trusted networks. Cisco recommends that the most secure way to prevent unauthorized access for VoIP traffic is by inspection of the VoIP registrations and calls. Enterprises may need to deploy application-aware firewalls. These firewalls can then inspect on VoIP protocols too.

H.323 uses multiple signaling ports. An application-aware firewall can inspect the H.323 messages and open respective ports for the call to proceed. For an H.323 call to take place, it must first open an H.225 connection on TCP port 1720 using Q.931 signaling. Next, the H.245 management session is established. While this session can take place on a separate channel from the H.225 setup, it can also be done using H.245 tunneling, which takes the H.245 messages and embeds them in the Q.931 messages in the previously established H.225 channel. Next, the H.245 session opens dynamically assigned ports for the UDP-based RTP and RTCP video and audio data streams. The port numbers can range from 1024 to 65535. Because the port numbers are not known in advance, and because it would defeat the purpose of a firewall to open all these ports, a firewall must be able to snoop the H.323 data stream in order to open the additional ports needed for the call. This snooping is also known as *stateful inspection*. Firewalls that support H.323 message inspection in order to open just the needed ports per call are termed *application-aware firewalls*.

An additional problem encountered with most firewalls is the use of Network Address Translation (NAT). Within H.323, the H.225 and H.245 signaling channels make heavy use of the embedded IP address. For example, assume a terminal has a private address of 10.1.1.125, which gets translated to 206.165.202.125 when it tries to place a call to an H.323 terminal with an IP address of 206.165.201.78 on the outside network. The terminal on the outside still receives the private address within the H.225 signaling stream. Because this is a non-routable address, an attempt to make a connection back will fail. One way to work around this problem is to use an H.323-aware NAT firewall, which can rewrite the addresses in the signaling payload.

Using the border element in NAT or firewall environments allows administrators to target a single IP address to terminate all H.323 video calls. All incoming and outgoing video calls that access the public network will use the border element. With the Cisco ASA Firewall, administrators can enable H.323 fix-up and allow UDP port 1720 traffic to access the IP address of the border element. Without the border element, administrators would have to configure UDP port 1720 to all videoconferencing devices and have static NAT for each device, which may not be scalable with a large number of video endpoints in the network. If you use a Cisco IP/VC 352x or 3540 gateway, port 1820 must be configured for the videoconferencing devices. [Figure 8-15](#) illustrates the call flow in a network with NAT and a firewall.

Figure 8-15 Call Flow with NAT and a Firewall

When using the gatekeeper on the border element, you can have external devices on the outside untrusted network register to the gatekeeper through the UDP 1718 or 1719 port pinhole to the gatekeeper for RAS registrations. The application-aware firewall inspects RAS signaling and opens pinholes for H.225 call signaling based on the embedded ports in the RAS messages. Additional inspection of H.225 can open ports dynamically for H.245 communication and then RTP media, respectively, through the application-aware firewall. Organizations can use external gatekeepers to ensure E.164 number resolution.

