



CHAPTER 12

IP Video Telephony

Revised: April 30, 2012; OL-21733-18

Enterprises with Unified Communication systems may have widely deployed voice services such as point-to-point calls and conferencing within the enterprise as well as calls to the PSTN for external connectivity. When adding video to such networks, the following approaches can be used:

- [Enable Voice Devices to Support Video Calls, page 12-1](#)
- [Integrating Voice Network with Existing Video Network, page 12-1](#)

Enable Voice Devices to Support Video Calls

An enterprise may enable existing IP phones to support cameras or use software on PC-based systems to provide video capability in conjunction with IP phones. Where possible, newer devices that support video can be deployed, thus keeping the existing call control infrastructure intact.

This approach has the following advantages:

- Existing dial plan — Enterprises can use the existing dial plan and existing call agents to support video calls in the network.
- Call admission control — A single call admission control entity in the enterprise provides a way to optimize use of the network bandwidth.
- Existing network — The enterprise can use the existing network topology by adding the needed bandwidth to accommodate video on its network.
- Users — Enterprises can enable all existing users to make video calls.
- Video codec — A standardizing video codec enables the enterprise to optimize bandwidth usage for video calls, thus reducing the need for transcoding or transrating resources.
- Existing IP phones — The enterprise can leverage existing IP phones to add video by using soft clients or by connecting cameras IP phones.

Integrating Voice Network with Existing Video Network

An enterprise that has an existing video network used for conference room video calls or for video devices for executives, can integrate the video network with its voice network, thus enabling enterprise users to call video devices.

This approach has the following advantages:

- Separate dial plans — A separate dial plan for the video network enables enterprises to plan for video resources such as videoconferencing bridges and video PSTN gateways. Users can dial a prefix to access resources from the other network.

- Existing network — The enterprise can use the existing overlapped network topology as separate video and voice networks.
- Management and monitoring — Separate management and monitoring for voice and video make troubleshooting and problem isolation easier.
- Trunk protocol — Enterprises can choose the desired protocol to interwork between the voice and video call agents. This can enable enterprises to use similar methods for call transfers, DTMF, MWI, or other functionality.
- Independence from call agent features — Enterprises can leverage capabilities that video call agents provide for video endpoints, such as optimizing the bit rate of video traffic for calls, and those capabilities may be independent of the features of voice call agents.

An enterprise can choose either of the above approaches or can combine them to achieve its objective of enabling its users to make video calls.

Video is now fully integrated into Cisco Unified Communications Manager (Unified CM), and there are also many video endpoints available from Cisco and its strategic partners. For example, Cisco Unified Video Advantage is just as easy to deploy, manage, and use as a Cisco Unified IP Phone.

What's New in This Chapter

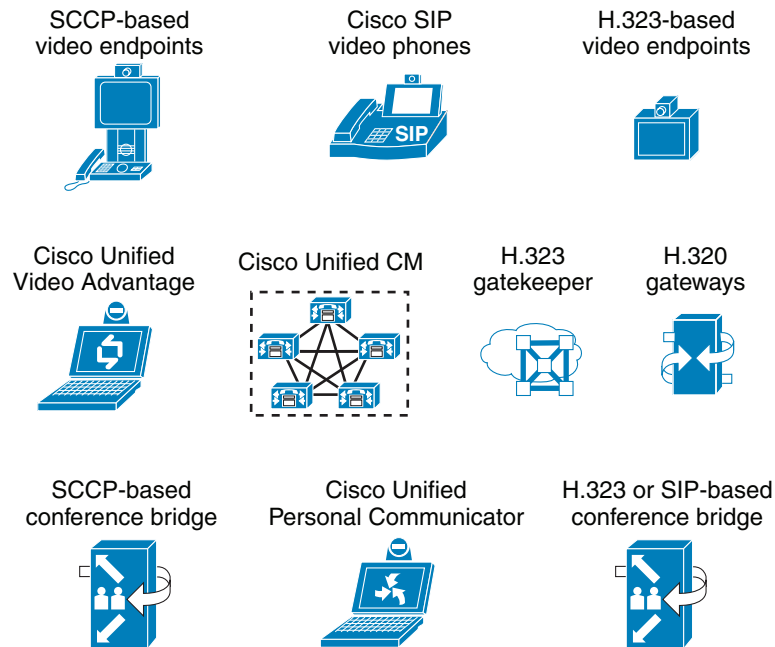
[Table 12-1](#) lists the topics that are new in this chapter or that have changed significantly from previous releases of this document.

Table 12-1 *New or Changed Information Since the Previous Release of This Document*

New or Revised Topic	Described in:	Revision Date
Video interoperability	Video Interoperability, page 12-44	April 30, 2012
Minor corrections and changes	Various sections throughout this chapter	September 30, 2011
Minor corrections and changes	Various sections throughout this chapter	June 2, 2011
SIP clients	Endpoints, page 12-5	November 15, 2010
Trunks for video calls	Trunks, page 12-15	November 15, 2010
Video endpoints	Endpoints, page 12-5	November 15, 2010
Cisco Unified IP Phone 9900 Series	Various sections throughout this chapter	April 2, 2010

IP Video Telephony Solution Components

The Cisco IP Video Telephony solution consists of Cisco Unified Communications Manager (Unified CM); Cisco Unified Videoconferencing 3500 and 5000 Series Multipoint Control Units (MCUs) for H.323, Session Initiation Protocol (SIP), and Skinny Client Control Protocol (SCCP) conference calls; Cisco Unified Videoconferencing 3500 Series H.320 Gateways; Cisco IOS H.323 Gatekeeper; Cisco Unified IP Phones 9900 Series; Cisco IP Video Phone E20; Cisco Unified Personal Communicator; Cisco Unified Video Advantage; Cisco Unified IP Phone 7985; third-party SCCP video endpoint solutions; and the existing range of H.323 or SIP-compliant products from partners such as Polycom, Lifesize, Sony, and others. (See [Figure 12-1](#).)

Figure 12-1 IP Video Telephony Components

Administration Considerations

This section discusses the following configuration elements in Unified CM Administration that pertain to Video Telephony:

- [Protocols](#), page 12-3
- [Endpoints](#), page 12-5
- [Regions](#), page 12-5
- [Topology-Aware Locations](#), page 12-8
- [Retry Video Call as Audio](#), page 12-10
- [Wait for Far-End to Send TCS](#), page 12-13
- [Trunks](#), page 12-15

Protocols

Unified CM supports a large number of protocols. Any device can call any other device, but video is supported only on SCCP, H.323, and SIP devices. Specifically, video is not supported in the following protocols in Cisco Unified CM Release 8.x:

- Computer Telephony Integration (CTI) applications (TAPI and JTAPI)
- Media Gateway Control Protocol (MGCP)

Therefore, Unified CM currently supports the types of calls listed in [Table 12-2](#).

Table 12-2 Types of Calls Supported in Unified CM Release 8.x

Calling Device Type	Called Device Type				
	SCCP	H.323	MGCP	TAPI/JTAPI	SIP
SCCP	Audio and video	Audio and video	Audio only	Audio only	Audio and video
H.323	Audio and video	Audio and video	Audio only	Audio only	Audio and video
MGCP	Audio only	Audio only	Audio only	Audio only	Audio only
TAPI/JTAPI	Audio only	Audio only	Audio only	Audio only	Audio only
SIP	Audio and video	Audio and video	Audio only	Audio only	Audio and video

Table 12-3 lists the audio and video algorithms and protocols currently supported in Unified CM.

Table 12-3 Capabilities Supported in Unified CM Release 8.x

H.323	SCCP	SIP
H.261	H.261	H.261
H.263, H.263+	H.263, H.263+	H.263, H.263+
H.264	H.264	H.264
G.711 A-law and mu-law	G.711 A-law and mu-law	G.711 A-law and mu-law
G.723.1	G.723.1	G.723.1
G.728	G.728	G.728
G.729, G.729a, G.729b, and G.729ab	G.729, G.729a, G.729b, and G.729ab	G.729, G.729a, G.729b, and G.729ab
G.722	G.722	G.722
G.722.1		
		iLBC
		iSAC
		AAC-LD
H.224 far-end camera control (supported by Unified CM but not by all endpoints); No protocol interworking	H.224 far-end camera control (supported by Unified CM but not by all endpoints); No protocol interworking	H.224 far-end camera control (supported by Unified CM but not by all endpoints); No protocol interworking
Out-of-band DTMF (H.245 alphanumeric) RFC2833 AVT Tones (only for H.323 intercluster trunk to SIP calls)	Out-of-band DTMF RFC2833 AVT Tones	RFC2833 AVT Tones Unsolicited SIP Notify KPML

Endpoints

IP phones are the most common video endpoints within the Cisco Unified Communications System. To add video to the IP phones for users, enterprises can use a camera for the Cisco Unified IP Phone 9971, use endpoints such as the Cisco E20 Video Phone, or connect the IP phone to a PC running Cisco Unified Video Advantage. Enterprises can also deploy soft clients such as Cisco Unified Personal Communicator.

Endpoints that support protocols such as H.323 and SIP can also be deployed with Cisco Unified CM. The decision to use H.323 is primarily determined by the use of H.239 for data sharing (for example, sharing PC screens during video calls) or if H.235 is used to pass secure tokens between endpoints for a video call with secure media between the endpoints. The type of user features, such as presence, also govern the type of protocol used for the endpoints, and the protocol choice is primarily dependent on the support for the features needed on the endpoints.

SIP video devices supported by Cisco Unified CM include the Cisco 9900 Series IP Phones, Cisco E20 Video Phone, Cisco Cius, third-party SIP devices (advanced), or the generic desktop and room system video device. Video can be enabled on the Cisco 9900 Series IP Phones by means of the video capabilities configuration for the devices. Configuration for the Cisco E20 Video Phone is on the phone itself. Cius supports video for calls with its front facing camera. The third-party SIP device (advanced) phone type or the generic video devices are additional options for endpoints from Polycom, Lifesize, Sony, and other manufacturers. While the configuration on Unified CM for these endpoints has not changed from earlier versions, the operation of Unified CM has been optimized to support the Cisco E20 Video Phone, Cisco Cius, Tandberg endpoints, and third-party endpoints more efficiently. Features such as the ability to process Early Offer from the endpoints and process them across SIP trunks without the use of MTP resources provide call signaling optimization and reduce the time to establish media for the call. Unified CM can also now support HD calls from these endpoints so that additional signaling (such as RTCP or parameters passed between the endpoints) is processed and sent across to achieve an optimal video call between the two devices. While Cius needs additional consideration due to its flexible use as a video phone when docked and as a Wi-Fi tablet when undocked, a Wi-Fi network is recommended for audio, and appropriate bandwidth should be used for calls. For additional information on wireless deployments, see the chapter on [Mobile Unified Communications, page 25-1](#).

For detailed information on the capabilities of the various IP phones and Cisco Unified Video Advantage, see the chapter on [Unified Communications Endpoints, page 18-1](#). For additional information on soft clients such as Cisco Unified Personal Communicator and the Client Services Framework, see the chapter on [Cisco Collaboration Clients and Applications, page 24-1](#).

Selecting the appropriate IP phones and endpoints for users to make video calls depends on the features desired and capabilities needed for video calls. The available options provide flexibility for designing and deploying clients for different types of users.

Regions

When configuring a region, you set two fields in Unified CM Administration: the Audio Codec and the Video Bandwidth. The audio setting specifies a codec type, while the video setting specifies the amount of bandwidth per call. However, even though the notation is different, the Audio Codec and Video Bandwidth fields actually perform similar functions. The Audio Codec field defines the maximum bit-rate allowed for audio-only calls as well as for the audio channel in video calls. For instance, if you set the Audio Codec for a region to G.711, Unified CM allocates 64 kbps as the maximum bandwidth allowed for the audio channel for that region. In this case, Unified CM will permit calls using either G.711, G.722, G.728, iLBC, or G.729. However, if you set the Audio Codec to G.729, Unified CM allocates only 8 kbps as the maximum amount of bandwidth allowed for the audio channel, and it will permit calls using only G.729 because iLBC, G.728, G.711, and G.722 all take more than 8 kbps.

**Note**

If both endpoints support G.711 and G.722, then G.722 will be negotiated because it is a wideband codec.

The Video Bandwidth field defines the maximum bit-rate allowed for the video channel of the call. However, for historical continuity with the practices used in traditional videoconferencing products, the value used in this field also includes the bandwidth of the audio channel. For instance, if you want to allow calls at 384 kbps using G.711 audio, you would set the Video Bandwidth field to 384 kbps and not 320 kbps.

**Note**

The Audio Codec setting also applies to the audio channel of video calls.

In summary, the Audio Codec field defines the maximum bit-rate used for audio-only calls and for the audio channel of video calls, while the Video Bandwidth field defines the maximum bit-rate allowed for video calls and should include the audio portion of the call.

Choosing the correct audio codec bandwidth limit is very important because each device supports only certain audio codecs, as shown in [Table 12-4](#). (For the most recent list of codecs supported by a particular endpoint, refer to the product documentation for that endpoint.)

Table 12-4 Types of Audio Codecs Supported by Endpoint Devices

Codec Type	Cisco 7900 and 9900 Series IP Phones	SCCP Third-Party Video Endpoints	Typical H.323 or SIP Endpoints	Cisco Unified Videoconferencing 3500 Series Gateways	Cisco Unified Videoconferencing 3500 and 5000 Series MCUs
G.729	Yes	Yes, depending on the model	No	No	Yes, depending on the model
G.728	No	Yes, depending on the model	Yes	Yes	Yes, depending on the model
G.711	Yes	Yes	Yes	Yes	Yes
G.722	Yes, depending on the model	Yes	Yes	Yes	Yes, depending on the model

As [Table 12-4](#) indicates, if you set the region to G.729, not all videoconferencing devices are able to support this type of codec. For example, calls between a Cisco Unified Video Advantage endpoint and a Tandberg T1000 endpoint would fail, or Unified CM would allocate an audio transcoding resource for the call.

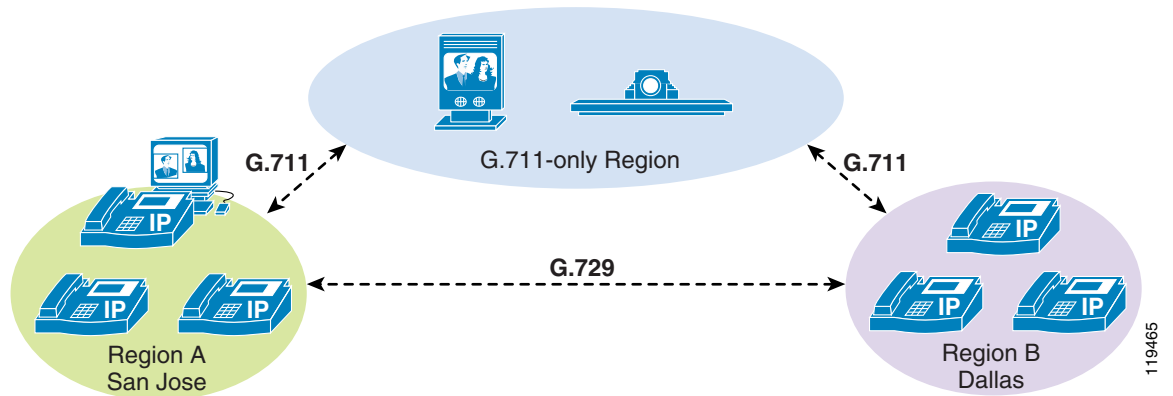
Cisco Unified CM Release 5.0 introduced audio transcoding resources based on Cisco IOS Enhanced Media Termination Point, which are able to support transcoding the audio stream of a video call while still supporting the video stream via a pass-through codec. The pass-through codec is used only for the video stream because the pass-through codec cannot be used for a stream that requires transcoding. The following three conditions must all be true for the pass-through codec to be used:

- The two endpoint devices have a matching CODEC capability.
- **MTP Required** is *not* checked for either endpoint.
- All intermediate resource devices (MTPs and transcoders) support the pass-through codec.

Traditional transcoders do not currently support the pass-through capability, so the call would connect as audio-only and would be transcoded between G.729 and G.711. To avoid this situation without using Cisco IOS Enhanced Transcoders, you would have to set the region to use G.711 instead. However, a region set for G.711 would also use G.711 for audio calls between two IP phones, which you might not want due to the increased consumption of bandwidth over the WAN.

If you want to use G.729 for audio-only calls to conserve bandwidth and to use G.711 for video calls, then you should configure one region to use G.711 for video endpoints that do not support G.729 and a separate region (or regions) to use G.729 for IP phones. (See Figure 12-2.) This method increases the number of regions needed but provides the desired codec and bandwidth allocations.

Figure 12-2 Using G.711 for Video Calls and G.729 for Audio-Only Calls



Note

It is possible to configure a pair of regions to prohibit video. If two video-capable devices in that region pair try to call each other, they will connect as audio-only unless Retry Video Call as Audio is not checked, in which case AAR rerouting logic will take over.

Table 12-5 lists some example configurations and their outcomes.

Table 12-5 Scenarios for Various Region Settings

Region Setting	Setting of Retry Video as Audio	Result
Region allows video	Enabled	Video calls allowed
Region allows video	Disabled	Video calls allowed
Region does not allow video	Enabled	Video calls will proceed as audio
Region does not allow video	Disabled	If AAR is not configured, video calls fail (with busy tone and "Bandwidth Unavailable" message displayed)

The Video Call Bandwidth field accepts values in the range of 1 to 32,256 kbps. However, to allow for compatibility with H.323 and H.320 videoconferencing devices, Cisco recommends that you always enter values for this field in increments of either 56 or 64 kbps. Therefore, valid values for this field include 112 kbps, 128 kbps, 224 kbps, 256 kbps, 336 kbps, 384 kbps, and so forth.

When the call speed requested by the endpoint exceeds the bandwidth value configured for the region, Unified CM automatically negotiates the call down to match the value allowed in the region setting. For instance, assume that an H.323 endpoint calls another H.323 endpoint at 768 kbps, but the region is set to allow a maximum of 384 kbps. The incoming H.225 setup request from the calling party would indicate that the call speed is 768 kbps, but Unified CM would change that value to 384 kbps in the outgoing H.225 setup message to the called party. Thus, the called endpoint would think that it was a 384-kbps call to begin with, and the call would be negotiated at that rate. The calling endpoint would show the requested bandwidth as 768 kbps, but the negotiated bandwidth would be 384 kbps.

However, if you set the Video Bandwidth to "None" in the region, Unified CM will either terminate the call (and send an H.225 Release Complete message back to the calling party) or will allow the call to pass as an audio-only call instead, depending on whether or not the called device has the Retry Video Call as Audio option enabled. (See [Retry Video Call as Audio, page 12-10](#).)

As the video resolution for the calls increases, so does the need for bandwidth. For video bandwidth in the region settings, the suggested values are 384 kbps for calls where CIF video resolution is desired, 768 kbps where VGA resolution is desired, and 1.5 Mbps for 720p resolution video calls. While most video endpoints have variable bit-rate encoders, video phones such as the Cisco Unified IP Phone 9900 Series have a constant bit-rate encoder for video. The constant bit-rate encoder provides better motion video and error resiliency.

Some endpoints such as the Cisco Unified Video Advantage use CIF resolution for calls. To restrict CIF resolution for video calls from such devices, the region configuration for video call bandwidth can be configured to 384 kbps and the devices can be associated to this region. Thus, calls to endpoints with higher resolution or conferences to MCUs negotiate CIF resolution for video.

Some Unified Communications endpoints with wireless capabilities (for example, Cisco Unified IP Phones 9900 Series and Cisco Cius) allow configuration of different bandwidth settings for the wireless media. For further details and design considerations, see [Wireless LAN Infrastructure, page 3-57](#).

Topology-Aware Locations

There are two methods to limit the amount of bandwidth available for calls between locations. Cisco Unified CM 4.0 introduced support for video calls in locations. More specifically, the location option in Cisco Unified CM defines the aggregate bandwidth allowed for all calls between that location and all other locations. This aggregate bandwidth value maps well to traditional hub-and-spoke network topologies. Cisco Unified CM Release 5.x introduced the option to use topology-aware locations based on Resource Reservation Protocol (RSVP) to determine if enough bandwidth is available along the path between two sites. RSVP enables per-hop checking that accommodates complex topologies while supporting the separation of audio bandwidth and video bandwidth via the use of RSVP Application IDs.

**Note**

Static locations and RSVP-based locations use different models to separate voice and video calls. For more details, see [Call Admission Control, page 11-1](#).

RSVP-based locations introduce the concept of an RSVP policy. While there are numerous policy options, there are two main categories:

- The RSVP reservation for the video stream is mandatory for the call to complete. The call will either fail (with busy tone played and the "Bandwidth Unavailable" message displayed to the user) or Automated Alternate Routing (AAR) will try to reroute the call.



Note While the call is being rerouted through the PSTN, no video will be available.

- The RSVP reservation for the video stream is desired.

First, the audio codec and the video bandwidth set in the region define the maximum speed (bit rate) of the video call. An RSVP reservation is sent from the Cisco RSVP Agent for the audio and video stream of the call using the maximum bit rate as the reservation requirement. If the RSVP reservation for the video stream fails, Unified CM checks the setting of RSVP policy to determine how to handle that call. If the policy for the audio stream is optional, the call will continue as audio-only. Otherwise, if the RSVP policy for the audio stream is mandatory, the call will continue as audio-only unless the audio stream also fails to get an RSVP Reservation. In that case, the call will either fail (with busy tone played and the "Bandwidth Unavailable" message displayed to the user) or Automated Alternate Routing (AAR) will try to reroute the call. (For more details on topology-aware locations, see [Call Admission Control, page 11-1.](#))



Note

If the video reservation fails while using a video-desired policy, the call will complete as audio-only. However, the user will not have any visual or audible feedback to indicate why the video failed.

When configuring static locations, you also set two fields in Unified CM Administration: the Audio Bandwidth and the Video Bandwidth. Unlike regions however, the Audio Bandwidth for static locations applies only to audio-only calls, while the Video Bandwidth applies to both the audio and video channels of video calls. The audio and video bandwidth are kept separate because, if both types of calls shared a single allocation of bandwidth, then it is very likely that audio calls would take all the available bandwidth and leave no room for any video calls, or vice versa. Also, separate bandwidth pools for audio and video correspond to the way queues are configured in the switches and routers in the network, which typically have a priority queue for voice traffic and a separate priority queue or a Class-Based Weighted Fair Queue for video traffic. (See [WAN Quality of Service \(QoS\), page 3-40](#), for more details.)

Both the Audio Bandwidth and the Video Bandwidth fields offer three options: None, Unlimited, or a field that accepts numeric values. However, the values entered in these fields use two different calculation models. For the Audio Bandwidth field, the values entered should include the Layers 3 to 7 overhead required for the call. For instance, if you want to permit a single G.729 call to or from a location, you would enter the value of 24 kbps. For a G.711 call, you would enter the value of 80 kbps. The Video Bandwidth field, by contrast, should be entered without the overhead included. For instance, for a 128-kbps call you would enter 128 kbps, and for a 384-kbps call you would enter 384 kbps. As with the values used in the Video Bandwidth field for regions, Cisco recommends that you always use increments of 56 kbps or 64 kbps for the Video Bandwidth field for locations as well.

For example, assume that a company has a three-site network. The San Francisco location has a 1.544-Mbps T1 circuit connecting it to the San Jose main campus. The system administrator wants to allow four G.729 voice calls and one 384-kbps (or two 128-kbps) video calls to/from that location. The Dallas location has two 1.544-Mbps T1 circuits connecting it to the San Jose main campus, and the

administrator wants to allow eight G.711 voice calls and two 384-kbps video calls to/from that location. For this example, the administrator would set the San Francisco and Dallas locations to the following values:

Location	Number of Audio Calls Desired	Audio Bandwidth Field Value	Number of Video Calls Desired	Video Bandwidth Field Value
San Francisco	4 using G.729	96 kbps (4 * 24 kbps)	1 at 384 kbps	384 kbps
Dallas	8 using G.711	640 kbps (8 * 80 kbps)	2 at 384 kbps	768 kbps

When the call speed requested by the endpoint exceeds the value configured for the location, Unified CM will *not* automatically negotiate the call speed (as it does for regions) to match the value allowed in the location setting. Instead, the call will be rejected or will be retried as an audio-only call (if the Retry Video as Audio setting is enabled on the called device). Therefore, you should always set the region's video bandwidth to a value lower than the location's video bandwidth value. For example, if you have two regions (region A and region B) and you set the video bandwidth between those two regions to 768 kbps but the devices in region A are in a location that has a video bandwidth set to 384 kbps, then all calls between those two regions will fail or will result in audio-only calls (depending on the Retry Video Call as Audio setting).

Retry Video Call as Audio

This check-box is available on all SCCP endpoint types that support video, including Cisco Unified IP Phones 7940, 7941, 7942, 7945, 7960, 7961, 7962, 7965, 7970, 7971, 7975, and Cisco IP Video Phone 7985, as well as third-party SCCP video endpoints, and all H.323 and SIP devices (clients, gateways and all types of H.323 trunks). When this option is activated (checked), if there is not enough bandwidth to reach the device (for example, if the Unified CM regions or locations do not allow video for that call), then Unified CM will retry the call as an audio-only call. When this option is deactivated (unchecked), Unified CM will not retry the call as audio-only but instead will either fail the call or reroute the call by whatever automated alternate routing (AAR) path is configured. By default, this retry option is enabled (checked).

This feature applies to the following scenarios only:


- The region is configured not to allow video.
- The location is configured not to allow video, or the requested video speed exceeds the available video bandwidth for that location when locations are not using an RSVP Policy.
- For calls between Unified CM clusters, the requested video speed exceeds the gatekeeper's zone bandwidth limits.

The Retry Video Call as Audio option takes effect only on the terminating (called) device, thus allowing the flexibility for the calling device to have different options (retry or AAR) for different destinations.

If the video call fails due to bandwidth limitations but automated alternate routing (AAR) is enabled, Unified CM will attempt to reroute the failed call as a video call to the AAR destination. If AAR is not enabled, the failed call will result in a busy tone and an error message being sent to the caller. (See [Figure 12-3](#).) Depending on the type of device that is calling, the failed call will result in one of the following conditions:

- If the calling device is an SCCP endpoint with an LCD screen, the caller will hear busy tone and will see the message "Bandwidth Unavailable" displayed on the device.
- If the calling device is an SCCP endpoint without an LCD screen (such as a Cisco Unified IP Phone 7902), the caller will hear busy tone.
- If the calling device is an H.323 or SIP device or a PSTN device connected through a gateway, the caller will hear busy tone and Unified CM will send the appropriate error message (such as a Q.931 Network Congestion cause code) back to the H.323, SIP, or MGCP device.

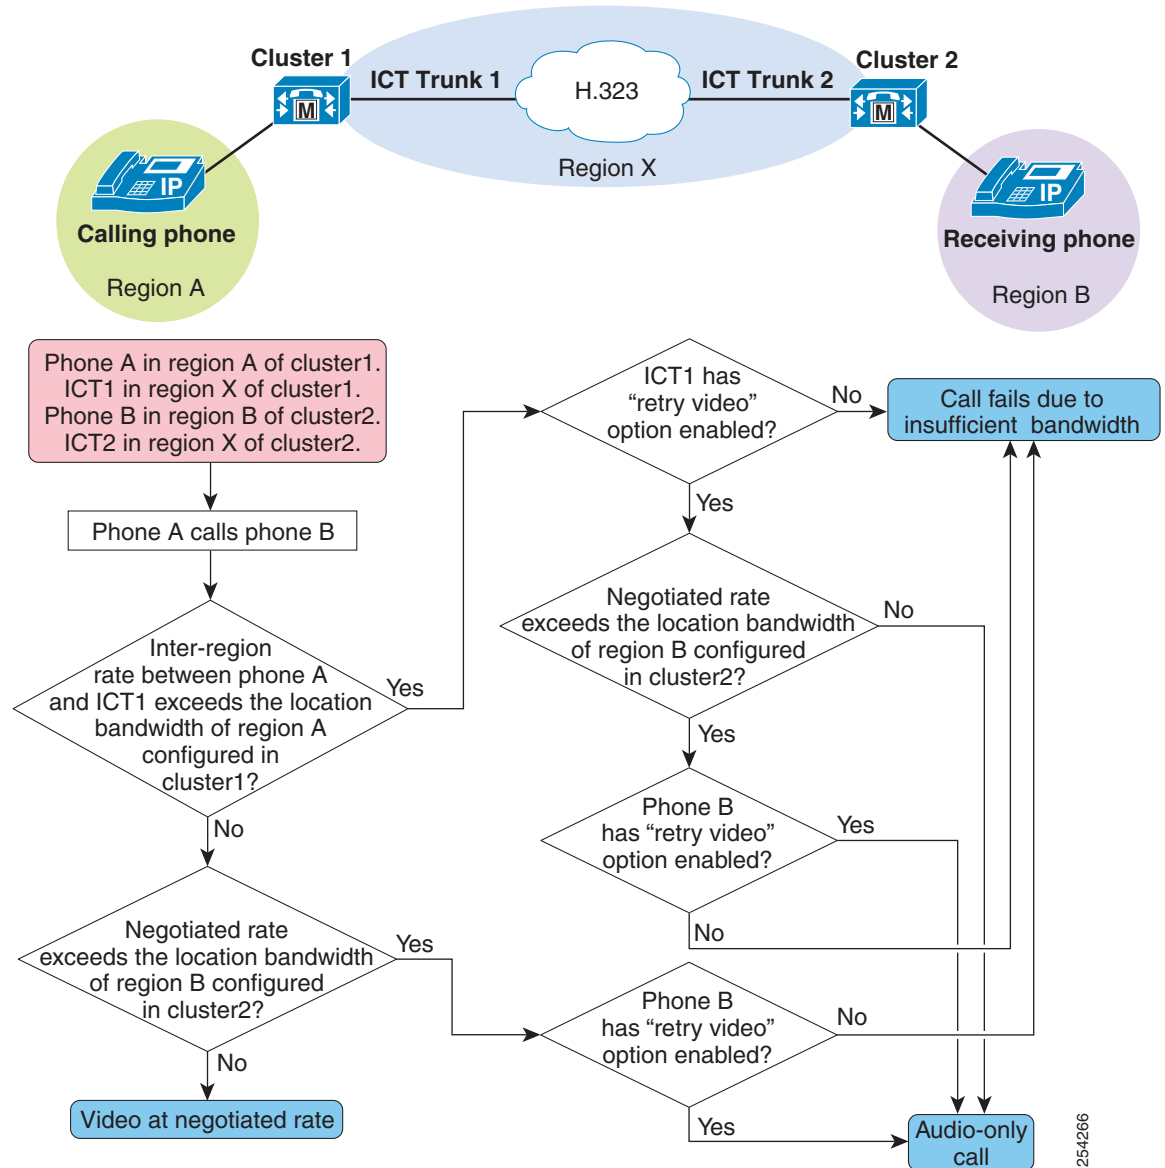
Figure 12-3 Possible Scenarios for a Video Call



See the chapter on [Call Admission Control](#), page 11-1, for further details on the use of AAR.

Figure 12-4 illustrates the steps of a call between two clusters using non-gatekeeper controlled intercluster trunks.

Figure 12-4 Call Flow Between Two Clusters Using Non-Gatekeeper Controlled Intercluster Trunks



Wait for Far-End to Send TCS

This check-box is available on all H.323 devices, including H.323 clients, H.323 gateways, and H.225 gatekeeper-controlled trunks. This feature pertains to the H.245 capabilities-exchange phase of H.323 calls. When this feature is enabled, Unified CM waits for the remote H.323 device to send its Terminal Capabilities Set (TCS) to Unified CM before Unified CM will send its TCS to the H.323 device. When the option is disabled, Unified CM does not wait but sends its TCS to the remote H.323 device immediately.

By default, the Wait for Far-End to Send TCS option is enabled (checked). However, you must uncheck (disable) it in the following circumstances:

- When the H.323 device communicating with Unified CM is also waiting for the far-end to send its TCS

In this case, a stalemate occurs because neither side sends its TCS, and the H.245 connection times out after a few seconds. Examples of devices that also wait for the far-end to send TCS are some H.323 routed-mode gatekeepers, H.320 gateways, H.323 proxies (or IP-to-IP gateways) and some H.323 multipoint conference bridges. These devices wait for the far-end to send TCS for the same reason Unified CM does: because they are waiting for both sides of the connection to send their TCSs before forwarding them to the other side.

- When communicating with another Unified CM cluster over an intercluster trunk

**Note**

For intercluster trunks and gatekeeper-controlled intercluster trunks, the Wait for Far-End to Send TCS option is always disabled and cannot be enabled.


In many scenarios, Unified CM performs the role of a software switch connecting two endpoint devices (such as two H.323 clients that are trying to call each other). In such cases, it is best if Unified CM waits until both devices have sent their TCS messages so that it knows the capabilities of each device and can therefore make the most intelligent decision about what TCS to send back to each party (depending on region and location configurations, among other things). In these cases, the Wait for Far-End to Send TCS feature should be enabled.

However, some other H.323 devices (such as an H.320 gateway, which connects an H.323 device to an H.320 device) also perform the function of connecting two or more parties together. The gateway also prefers to wait until both ends send their TCS messages, so that it can make the most intelligent choice about how to set up the call. A stalemate could result if Unified CM and the gateway both wait for the other side to send their TCSs. To avoid this stalemate situation, disable (uncheck) the Wait for Far-End to Send TCS feature.

For instance, consider the following call scenarios illustrated in [Figure 12-5](#):

- Scenario 1 — Cisco Unified Video Advantage calls an H.320 endpoint
- Scenario 2 — An H.323 client calls an H.320 endpoint


In both of these scenarios, Wait for Far-End to Send TCS feature can be left at its default setting of enabled (checked)

Figure 12-5 Scenarios with the Wait for Far-End to Send TCS Feature Enabled (Checked)

In Scenario 1 from Figure 12-5, Unified CM already knows the capabilities of the Cisco Unified Video Advantage client because SCCP devices provide Unified CM with their media capabilities during registration. But Unified CM does not know the capabilities of the H.320 gateway until the gateway sends its TCS to Unified CM during the H.245 phase of the call. Likewise, the H.320 gateway does not know what TCS to send to Unified CM until the H.320 endpoint sends its TCS to the gateway. In this case it is better to leave the Wait for Far-End to Send TCS feature enabled because the H.320 endpoint will send its TCS to the gateway, the gateway will send its TCS to Unified CM, and Unified CM will then have the TCSs from both endpoints with which to make a decision.

Figure 12-6 shows the following call scenarios, in which the call will fail unless the Wait for Far-End to Send TCS feature is disabled:

- Scenario 1 — Cisco Unified Video Advantage calls another Cisco Unified Video Advantage in a remote cluster via the ISDN network
- Scenario 2 — An H.323 client calls another H.323 client in the remote cluster via the ISDN network

Figure 12-6 Scenarios with the Wait for Far-End to Send TCS Feature Disabled (Unchecked)

In both scenarios from Figure 12-6, there will be a stalemate because both Unified CMs will wait to receive the TCSs from the gateways, and the gateways will both wait to receive the TCS from the ISDN side as well. The call will time-out after a few seconds and fail. From the perspective of the users, the

caller will hear ringback tone indicating that the call is progressing and the called party will hear a ring indicating an incoming call. When the called party attempts the answer the call, the H.245 phase will fail due to the stalemate, and the call will then fail by hanging up on both parties.

To work around this issue in scenarios such as these, Cisco recommends that you disable (uncheck) the Wait for Far-End to Send TCS option on the device representing the H.320 gateway in Unified CM. This device could be an H.225 gatekeeper-controlled trunk or an H.323 gateway device, depending on how you have configured Unified CM to reach the H.320 gateway.

However, if the Wait for Far-End to Send TCS option is disabled, there is a possibility that the initial capabilities exchanged will not work for the remote device. For instance, the Unified CM region might be set to 768-kbps video but the H.320 device might only support 384 kbps, or the selected audio codec might not work for the remote party. In such cases, the initially negotiated logical channels might have to be torn down and reopened with the correct speed and codec. Many legacy H.323 and H.320 devices do not handle this situation well and will disconnect the call when Unified CM sends a CloseLogicalChannel message to renegotiate the channel to a different value. Thus, you must be careful where and when you disable the Wait for Far-End to Send TCS option.

Trunks

Cisco Unified CM supports various types of trunks. H.323 trunks have been most popular for video deployments due to the large number of H.323 endpoint deployments, and it can be used to interwork with H.323 gatekeepers that route calls to video endpoints and gateways. H.323 trunks also provide a pass-through functionality for a number of video features used by the H.323 video endpoints, such as H.239 and H.235. The RASAggregator trunk enables Unified CM to provide advanced features such as call restrictions and bandwidth enforcement for calls for the endpoints registered to a Cisco IOS gatekeeper.

SIP trunks can provide interconnection to SIP networks. These trunks support video and SRTP across the trunks. Unified CM can provide a much tighter integration for video communication servers such as the Cisco TelePresence Video Communication Server (VCS). This capability expands the support for high-definition calls so that advanced signaling needed by Cisco VCS, Cisco Video devices, and third-party video endpoints can also work through Unified CM. In addition, Cisco Unified CM SIP trunks support Early Offer for video calls without the need of a media termination point (MTP). This can be important for deployments where calls go through multiple call control servers or where call cut-through time for establishing a bidirectional media path is important. For additional information, see the chapter on [Cisco Unified CM Trunks](#), page 14-1.

Some deployments may use DNS SRV in their networks. For Cisco TelePresence VCS deployments that use DNS SRV, Unified CM SIP trunks can also use DNS SRV. In such deployments, you need to consider the DNS server scalability and redundancy and also note that the load balancing and redundancy ability is dependent on the DNS server servicing the requests. Thus, the Unified CM trunk load balancing and redundancy will be in addition to the DNS server load balancing and redundancy.

Multipoint Conferencing

Whenever three or more parties want to engage in the same video call together, a Multipoint Control Unit (MCU) is required. An MCU consists of the following main components:

- Multipoint Controller (MC)
- Multipoint Processor (MP)

The MC handles all aspects of call setup and teardown for the conference, including media negotiation, call signaling, and choosing which MP to use for the call. The MP processes all the audio and video packets. The MC controls the MP, and one MC can control multiple MPs. The MP can be either software-based or hardware-based. Software-based MPs are typically not capable of advanced transcoding, transrating (multiple speeds), or composition features.

Cisco MCUs also support the Skinny Client Control Protocol (SCCP). This enables the MCU to be integrated with Unified CM as a video conference bridge similar to how the Cisco IOS conference bridges can be integrated for audio conferencing.

Cisco Unified Videoconferencing products have the Multipoint Controller (MC) and the Multipoint Processor (MP) functions built into the Multipoint Conference Units (MCUs) as an integrated device. Chassis-based MCUs provide flexibility and scalability by using an MCU module that manages and controls the Enhanced Media Processor (EMP) modules that do the video conferencing.



Note

An EMP is required for the Cisco 3545 and 5200 MCUs because they do not have a software-based MP. An MCU module is needed to control the EMPs in the MCU chassis.

Cisco Unified CM supports the Cisco Unified Videoconferencing MCUs in SCCP, H.323, and SIP modes. Each protocol offers different features and is used for different reasons, so each of these MCUs is equipped to run all three protocols. The Cisco Unified Videoconferencing MCUs can be configured to run all three protocols simultaneously and divide the total number of available MP resources between the three.

Regardless of signaling protocol, the MCU provides the same basic function of receiving the audio and video streams from each participant and sending those streams back out to all other participants in some sort of combined view. There are two types of views in a multipoint video conference:

- Voice-activated (switched)
- Continuous presence

Voice Activation

Voice-activated conferences take in the audio and video streams of all the participants, decide which participant is the dominant speaker, and send only the dominant speaker's video stream back out to all other participants. The participants then see a full-screen image of the dominant speaker (and the current speaker sees the previous dominant speaker). The audio streams from all participants are mixed together, so everyone hears everyone else, but only the dominant speaker's video is displayed.

You can use any of the following methods to select the dominant speaker:

- Voice activation mode

Using this mode, the MCU automatically selects the dominant speaker by determining which conference participant is speaking the loudest and the longest. To determine loudness, the MCU calculates the strength of the voice signal for each participant. As conditions change throughout the conversation, the MCU automatically selects a new dominant speaker and switches the video to

display that participant. A hold timer prevents the video from switching too hastily. To become the dominant speaker, a participant has to speak for a specified number of seconds and be more dominant than all other participants.

- Manual selection of the dominant speaker through the MCU's web-based conference control user interface

The conference controller (or chairperson) can log onto the MCU's web page, highlight a participant, and select that person as the dominant speaker. This action disables voice activity detection, and the dominant speaker remains constant until the chairperson either selects a new dominant speaker or re-enables voice activation mode.

- Configuring the MCU to cycle through the participant list automatically, one participant at a time
With this method, the MCU stays on each participant for a configured period of time and then switches to the next participant in the list. The conference controller (or chairperson) can turn this feature on and off (re-enable voice activation mode) via the web interface.

Continuous Presence

Continuous-presence conferences display some or all of the participants together in a composite view. The view can display from 2 to 16 squares (participants) in a variety of different layouts. Each layout offers the ability to make one of the squares voice-activated, which is useful if there are more participants in the conference than there are squares to display them all in the composite view. For instance, if you are using a four-way view but there are five participants in the call, only four of them will be displayed at any given time. You can make one of the squares in this case voice-activated so that participants 4 and 5 will switch in and out of that square, depending on who is the dominant speaker. The participants displayed in the other three squares would be fixed, and all of the squares can be manipulated via the conference control web-based user interface.



Note

Continuous presence requires the use of an Enhanced Media Processor (EMP) in the Cisco Unified Videoconferencing MCU.

MP Resources

For both types of conferences, the type of MP resource determines which video formats, transrating, and transcoding capabilities the MCU can support. If endpoints connect to the conference at different speeds, a transrating-capable MP is required. The RM and EMP modules are both capable of transrating between speeds. If a transrating-capable MP is not available, the MCU will send out flow-control messages to all the endpoints, instructing them to lower their transmit speed to match the maximum receive rate of the slowest endpoint. For example, if three participants are connected in a 384-kbps conference and a fourth participant joins at 128 kbps, the MCU will send flow-control messages to the other three participants instructing them to lower their transmit rates to match the 128-kbps participant. This method causes all participants to suffer degraded quality because one participant is less capable. A transrating-capable MP would, instead, convert the 128-kbps stream to 384 kbps, and vice versa, so that each participant can enjoy the best quality their connection allows.

For continuous-presence conferences, a transrating-capable MP is also very important. The software-based MP built into the MCU combines all the input streams and sends the resulting combination back out to each participant. For instance, if four participants are connected in a continuous-presence conference at 384 kbps using G.711 audio, each participant will transmit 320 kbps of video and 64 kbps of audio into the MCU. The MCU will take these four input video streams and combine them into the four-way composite view. The MCU will then transmit 1280 kbps of video back to each endpoint, along with the mixed 64 kbps of audio, for a total of 1344 kbps per endpoint. This method is known as Asynchronous Continuous Presence, and it can have a negative impact on bandwidth requirements, call admission control mechanisms, and interoperability with certain devices.

**Note**

Cisco strongly advises against the use of Asynchronous Continuous Presence.

With the RM or EMP modules, the MCU is capable of transrating each input stream before combining them, so that the total output bandwidth matches the input bandwidth. For instance, if the MCU is using the four-quadrant layout and each participant transmits 320 kbps of video and 64 kbps of audio into the MCU, the MCU would essentially transrate each input stream down to 80 kbps, combine them so that the resulting four-quadrant view is 320 kbps of video (4 * 80 kbps), combine this video with the mixed 64 kbps audio, and transmit the final combination back to each participant. This method is known as Synchronous Continuous Presence. Cisco strongly recommends that all continuous-presence conferences use the Synchronous Continuous Presence mode. However, use of this mode means that each MCU must have a transrating-capable MP (such as an RM or EMP) available, which does increase the cost of the MCU.

**Note**

For H.323 and SIP clients with built-in MCUs, Unified CM does not allow an H.323 client to generate a second call, thereby negating the functionality of the built-in MCU.

MCU Resources for Ad-Hoc Conferences

As stated previously, the Cisco Unified Videoconferencing MCUs offer support for SCCP beginning with software version 3.2+ for those models and with Cisco Unified CM Release 4.0. When configured in SCCP mode, Unified CM provides the MC function while the MCU provides the MP function. The SCCP MCU is controlled completely by Unified CM. Cisco Unified CM 8.6 and later releases support SIP-based TelePresence MCUs through conference bridges, thus providing an additional method of MCU integration for ad-hoc conferences that can support higher resolution.

Only the following events invoke ad-hoc MCU resources:

- The user of an SCCP or SIP endpoint (such as an IP Phone or a third-party SCCP video endpoint) presses the Conf, Join, or cBarge softkeys to invoke an ad-hoc conference
- The user of an SCCP or SIP endpoint (such as an IP Phone or a third-party SCCP video endpoint) presses the MeetMe softkey to invoke a reservationless meet-me conference.
- The user of Cisco Cius or Cisco Unified Personal Communicator in softphone mode uses the Join or Conference feature to join multiple calls into a conference.

Participants in either of these types of conferences can include any type of endpoint (that is, video and non-video devices using any signaling protocol that Unified CM supports via any supported gateway type); however, only SCCP endpoints, SIP Cisco IP Phones, or Cisco Unified Personal Communicator can invoke the ad-hoc MCU resources. In other words, an H.323 video endpoint cannot invoke an ad-hoc MCU resource, but an SCCP video endpoint can invoke the resource and then join an H.323 video participant to the call. For example, the user at the SCCP endpoint could press the Conf softkey, dial the directory number of an H.323 client, and then press the Conf softkey again to complete the transaction. The H.323 client will be joined as a participant on the SCCP MCU conference.


**Note**

Earlier versions of this document described alternative configurations for H.323 devices that do not support supplementary services (such as placing a call on hold). For details, refer to the section on *SCCP MCU Resources* in the *Cisco Unified Communications SRND Based on Cisco Unified Communications Manager 7.x*, available at http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/7x/uc7_0.html.

For reservationless conferences via the MeetMe softkey, the signaling protocol used by the other endpoints does not have to support being placed on hold and transferred. For these types of conferences, each endpoint dials the MeetMe dial-in number arranged by the SCCP client who initiated the conference.

Figure 12-7 illustrates how H.323 endpoints and Cisco IP Phones can participate in the same ad-hoc conference. In this example, the conference was initiated by an SCCP endpoint using the Conf softkey to invite the three members.

Figure 12-7 Ad-Hoc Conference Between SCCP, SIP, and H.323 Endpoints



Ad-hoc conferences support voice-activated mode as well as continuous presence, depending on the conferencing bridge used.

Media Resource Groups and Lists

When a user of an SCCP or SIP phone activates the Conf, Join, or MeetMe softkey, Unified CM uses the Media Resource Manager to select conference bridges. Conference bridges or MCUs resources are configured in the media resource groups (MRGs). The media resource group lists (MRGLs) specify a prioritized list of MRGs and can be associated with the endpoints. The Media Resource Manager uses MRGLs of the endpoints for selecting the conference bridge. How you group the resources is completely at your discretion, but it is typically done either by geographical placement (so that all endpoints at a given site use the conference bridges closest to them) or by endpoint type (so that video-capable endpoints use a video-capable MCU while audio-only endpoints use a different conference bridge resource).

Cisco Unified CM has the Intelligent Bridge Selection feature, which provides a method for selecting conference resources based on the capabilities of the endpoints in the conference. If there are two or more video endpoints when the conference is invoked and a videoconferencing resource is available, Intelligent Bridge Selection chooses that resource for the conference. On the other hand, if no videoconferencing resource is available or if there are no video-capable endpoints in the conference, Intelligent Bridge Selection chooses an available audio resource for the conference. Intelligent Bridge Selection provides an added functionality to select secure conference bridges for secure conferences. However, secure conference bridge selection is dependent on device capabilities. Unified CM may decide to allocate secure conference bridges in lieu of video or audio conference bridges. Flexibility to change the behavior of the Intelligent Bridge Selection functionality is provided through service parameter configurations in Unified CM.

For the purposes of Intelligent Bridge Selection, the following types of endpoints are considered to be video-capable:

- Cisco Unified Video Advantage (Should be running on a PC connected to the PC port of an IP phone)
- Cisco Unified IP Phone 7985G and 9900 Series
- Cisco Unified Personal Communicator or the Client Services Framework clients
- H.323 client (Third-party video endpoints)
- SIP Advanced endpoint (Third-party video endpoints)
- SIP trunk with no media termination point (MTP) for video calls
- H.323 trunk with no MTP for video calls

Intelligent Bridge Selection has the following advantages over other methods of conference bridge selection:

- Conference bridge selection by conference type – either secure, video, or audio conferences
- Simplified media resource configuration
- Optimized use of MCU video ports that potentially would have been used for audio-only conferences with other methods of bridge selection

All the conference bridge resources and MCUs can be in one MRGL, and Intelligent Bridge Selection will then select the conference bridge based on the need to do just an audio conference or a video conference.

Unified CM also supports an alternate way of selecting conference bridges, which can be specified by service parameter configurations. In this mode, Unified CM applies the following criteria to select the conference bridge resource to use, in the order listed here:

1. The priority order in which the media resource groups (MRGs) are listed in the media resource group list (MRGL)
2. Within the selected MRG, the resource that has been used the least

If the MCU is placed at the top of the MRGL for the phone, the MCU will always be chosen even for audio-only conferences that do not involve any video-capable participants. In this scenario, the MCU resources might be wasted on audio-only conferences and not be available to satisfy the request for a video conference when it occurs.


Note

Meet-me conferences do not use the Intelligent Bridge Selection feature.

Intelligent Bridge Selection

Cisco Unified CM includes the Intelligent Bridge Selection feature, which provides a method for selecting conference resources based on the capabilities of the endpoints in the conference. If there are two or more video endpoints when the conference is invoked and a videoconferencing resource is available, Intelligent Bridge Selection chooses that resource for the conference. On the other hand, if no videoconferencing resource is available or if there are no video-capable endpoints in the conference, Intelligent Bridge Selection chooses an available audio resource for the conference.

Intelligent Bridge Selection provides an added functionality to select secure conference bridges for secure conferences. However, secure conference bridge selection is dependent on device capabilities. Unified CM may decide to allocate secure conference bridges in lieu of video or audio conference bridges. Flexibility to change the behavior of the Intelligent Bridge Selection functionality is provided through service parameter configurations.

For the purposes of Intelligent Bridge Selection, the following types of endpoints are considered to be video-capable:

- Cisco Unified Video Advantage (Should be running on a PC connected to the PC port of an IP phone)
- Cisco Unified IP Phone 7985G
- Cisco Unified Personal Communicator
- H.323 client (Third-party video endpoints)
- SIP Advanced endpoint (Third-party video endpoints)
- SIP trunk with no media termination point (MTP) for video calls
- H.323 trunk with no MTP for video calls

Intelligent Bridge Selection has the following advantages over other methods of conference bridge selection:

- Conference bridge selection by conference type – either secure, video, or audio conferences
- Simplified media resource configuration
- Optimized use of MCU video ports that potentially would have been used for audio-only conferences with other methods of bridge selection

**Note**

Meet-me conferences do not use the Intelligent Bridge Selection feature.

H.323 and SIP MCU Resources

When configured in H.323 or SIP mode, the MCU provides the MC function and behaves like an H.323 or SIP peer to Unified CM. H.323 and SIP MCU conferences can be invoked in a number of different ways, but they all fall into two major categories:

- Scheduled
- Reservationless


A scheduled conference uses a scheduling application to reserve the MCU resources in advance of the call. The scheduling function typically is provided through a web-based user interface such as Cisco Unified MeetingPlace, or Cisco Unified Video Conferencing Manager. The scheduling application usually generates an invitation that provides the user with the date and time of the conference, the number of ports reserved for the conference, and the dial-in information. Alternatively, the scheduling system can be configured to dial out to some or all of the participants at the beginning of the conference.

For a reservationless conference, the MCU has a certain number of resources that are available on demand. To create a conference, a user simply dials into the MCU at any time. If that user is the first participant to dial in, the MCU dynamically creates a new conference using the settings defined in the service template. (For more information on service templates, see [Service Templates and Prefixes](#), page 12-22.) Subsequent users who dial into the same conference number are joined to that conference.

Any type of endpoint can create and participate in scheduled and reservationless H.323 or SIP conferences. For instance, an SCCP endpoint can dial into the H.323 MCU to create a reservationless conference just as well as an H.323 endpoint can.

Figure 12-8 illustrates how H.323 and SCCP endpoints can participate in the same H.323 conference. In this example, the conference was initiated by an SCCP endpoint that dialed into the H.323 MCU to create a new reservationless conference, and the other two parties subsequently dialed into that conference.

Figure 12-8 *SCCP and H.323 Endpoints in a Reservationless Conference*



H.323 and SIP conferences support both voice-activated and continuous-presence modes. Furthermore, H.323 conferences support all of the MP types, including the integrated software-based MP of the MCU, the Rate Matching (RM) module, and the Enhanced Media Processor (EMP) module.

Service Templates and Prefixes

A service on the MCU defines the settings that pertain to each conference. You can have different services defined for the different types of conferences. Each service defines, at a minimum, the following settings:

- Speed of the conference (the video bit-rate)
This setting might include multiple speeds if a transrating-capable MP is used.
- Minimum and/or maximum number of participants
The minimum number defines how many ports will be reserved at the beginning of the conference. The maximum number defines the maximum number of participants that the MCU will allow to join this conference.
- Video codec type (H.261, H.263, or H.264)
- Frame rate (15 or 30 fps)
- Resolution (QCIF or CIF)
- MP resource (Auto, MP, RM, or EMP)
- Video layout to be displayed (voice activated or continuous presence)

A conference can include multiple layouts or even dynamic layouts that change as the number of participants in the conference increases or decreases.

- H.323 and SIP, or SCCP

When the "SCCP service" check-box is enabled (checked), the service is used for SCCP conferences. When this box is disabled (unchecked), the service is used for H.323 and SIP conferences.

For H.323 and SIP services, each service is assigned a service prefix that is dialed by the endpoints to reach that particular service. The service prefix forms the leading digits of the conference number, and the trailing digits define the conference ID. This format allows multiple conferences to run simultaneously using the same service prefix. For instance, you can have a service prefix of 555, while the full dial string of the conference is seven digits. This scheme would allow for conference numbers in the range of 5550000 through 5559999, with four digits for the conference ID. The user must dial the full string to access the conference. When the call is received by the MCU, the MCU parses the dialed digits to try to match a service prefix. Once it determines which service prefix is being dialed, the MCU then uses the remaining digits as the conference ID. If the conference ID does not exist yet, the MCU creates a new reservationless conference with that ID. If the conference already exists, the user is added to that conference.

It is important to note that, if both H.323 and SIP are enabled on the MCU at the same time, the dial plan must be the same for both protocols. There is no separation between H.323 and SIP as there is with SCCP. If a conference is created using SIP, the MCU will register that conference via H.323. If the gatekeeper or SIP proxy rejects the registration, the conference will fail.

SCCP services must also have a service prefix defined, but the digits themselves do not mean anything because users do not "dial" into an SCCP service. The prefix is used only in the SCCP registration messages between Unified CM and the SCCP MCU resource. Users either use the Conf, Join, or cBarge softkeys to access the SCCP MCU conference (ad hoc), or they dial a MeetMe number assigned by Unified CM to join the conference (reservationless). Therefore, the digits you specify for the SCCP service prefix do not matter and can be any digits you wish, such as 999999 for instance. This prefix is not exposed outside of the SCCP signaling between the MCU and Unified CM (that is, it cannot be dialed, it is not included in the MCU's registration with the gatekeeper, and so forth).

**Note**

The Cisco MeetingPlace Express media server has support for SCCP. Using this integration, the Cisco MeetingPlace Express media server can provide ad-hoc conference bridges with Unified CM, similar to the way the Cisco Unified Videoconferencing MCUs do.

Sizing the MCU

There are several factors involved in determining the types and number of conferences that an MCU can support. These sizing factors are different for different models of MCUs. MCUs can also make available more ports when using desktop conferencing mode as compared to the high definition (HD) mode.

Calculating the size of MCUs depends on the following factors:

- The type of resolution for the video conference
- The total number of ports that the MCU can support
- The number of ports that the MCU can dedicate to each protocol
- Whether cascading conferences are needed between MCUs or between EMP cards

**Note**

A single SCCP conference cannot span multiple EMPs. Each SCCP conference can support up to 24 participants.

For specific information about the number of ports supported, refer to the product documentation for your MCU hardware, available on Cisco.com. Due to the almost infinite number of possible variations, it is very difficult to provide any concrete design guidance in this document. Many customers end up with a mixture of SCCP ad-hoc conferences, H.323 and SIP reservationless conferences, and H.323 and SIP scheduled conferences. The MCUs must be sized to accommodate all of those types of conferences at the correct speeds and video layouts. Needless to say, this can become quite complex to determine. Please consult with your Cisco sales representative for assistance on sizing the MCUs for your particular environment.

IVR for Dial-In Conference

Dial-in conferences typically use an interactive voice response (IVR) system to prompt users to enter the conference ID and the password (if one is configured) of the conference they want to join. You can use either of the following types of IVRs with the Cisco Unified Videoconferencing 3500 Series MCUs:

- The IVR built into the MCU
- Cisco Unified IP IVR

The built-in IVR of the MCU has the following characteristics:

- Can prompt to create a conference or join by conference ID
- Can prompt for the password of the conference
- Supports both in-band and out-of-band (H.245 alpha-numeric) DTMF
- Cannot be customized to provide more flexible menus or functionality

The only thing that can be customized is the recorded audio file that is played to the user.

If you want to have a single dial-in number and then prompt the user for the conference ID, you can use Cisco Unified IP IVR in conjunction with the MCU.

Cisco Unified IP IVR has the following characteristics:

- Can prompt for the conference ID and the password (among other things)
- Supports only out-of-band DTMF

That is, the calling device must support an out-of-band DTMF method, such as H.245 alpha-numeric on H.323 devices. These out-of-band DTMF messages are then relayed by Unified CM to the Cisco IP IVR server. If the calling device supports only in-band DTMF tones, the Cisco IP IVR server will not recognize them and the calling device will be unable to enter the conference.

- Can be highly customized to provide more flexible menus and other advanced functionality

Customizations can include such things as verifying the user's account against a back-end database before permitting that user to enter into the conference, or queuing the participants until the chairperson joins.



Note

Because Cisco Unified IP IVR supports only out-of-band signaling, it will not work with H.323 endpoints that use in-band DTMF tones.

With Cisco Unified IP IVR, users dial a CTI route point that routes the call to the Cisco Unified IP IVR server instead of dialing a route pattern that routes directly to the MCU. After collecting the DTMF digits of the conference ID, the Cisco Unified IP IVR then transfers the call to the route pattern that routes the call to the MCU. This transfer operation requires that the calling device supports having its media channels closed and reopened to a new destination. For example, an H.323 video device that calls the


Cisco Unified IP IVR will initially negotiate an audio channel to the Cisco Unified IP IVR server and then, after entering the appropriate DTMF digits, it will be transferred to the MCU, at which point Unified CM will invoke the Empty Capabilities Set (ECS) procedure described earlier in this chapter to close the audio channel between the endpoint and the Cisco Unified IP IVR server and open new logical channels between the endpoint and the MCU. If the H.323 video endpoint does not support receiving an ECS from Unified CM, it will react by disconnecting the call or, worse, crashing and/or rebooting.

Gatekeepers

Prior to the introduction of video support in Unified CM, H.323 videoconferencing networks relied on gatekeepers to perform device registration management, call routing, and bandwidth control. The Cisco IOS Gatekeeper, formerly known as the Multimedia Conference Manager (MCM), provides these functions. However, most gatekeepers, including Cisco's, offer only rudimentary call routing capabilities compared to what might be expected in a typical enterprise-class PBX. When used to route H.323 video calls, Unified CM supplements the basic gatekeeper features and provides full enterprise-class PBX functionality to H.323 video calls.

Unified CM and the gatekeeper work as a team to manage H.323 video endpoints. The gatekeeper handles all Registration, Admission, and Status (RAS) signaling, while Unified CM handles all of the H.225 call signaling and H.245 media negotiations. Therefore, you have to deploy gatekeepers along with the Unified CM servers if RAS signaling procedures are required for the H.323 endpoints in your network, as illustrated in [Figure 12-9](#).

Figure 12-9 Unified CM and Cisco IOS Gatekeeper Provide RAS Signaling for H.323 Endpoints



RAS signaling is required any time either of the following conditions exists:

- The endpoint does not use a fixed IP address.

If the endpoint uses a static IP address, Unified CM does not require RAS procedures to locate the endpoint. Instead, the endpoint is provisioned in Unified CM Administration with its static IP address, and calls to that H.323 client's directory number are routed directly to that static IP address. If the endpoint does not use a static IP address, then Unified CM must query the gatekeeper to obtain the endpoint's current IP address each time Unified CM extends a call to the endpoint.

- The endpoint requires RAS procedures to place calls to E.164 addresses.

Most H.323 videoconferencing endpoints are capable of dialing another endpoint directly only when dialing by IP address (that is, the user enters the IP address of the destination endpoint in dotted-decimal format and then pushes the call button). However, if the user dials an

E.164-formatted number (a numeric value not in the dotted-decimal format of an IP address) or an H.323-ID (in the format of *username* or *username@domain*), most endpoints today provide only one way to resolve these types of destinations – by a RAS query to their gatekeeper. A growing number of endpoints, however, can be configured so that, for any call to an E.164 address, they skip any RAS procedures and instead send an H.225 SETUP message directly to a specified IP address. This method of operation is known as peer-to-peer mode. Tandberg H.323 endpoints are one example that use this mode, in which you can either configure a gatekeeper address for them to register with, or configure the IP address of the Unified CM server they should use. In the latter case, the endpoint sends all calls directly to the specified IP address, bypassing the need for RAS procedures with any gatekeeper.

In addition to managing RAS procedures for H.323 video endpoints, gatekeepers also continue to play an important role in managing dial plan resolution and bandwidth restrictions between Unified CM clusters in large multisite distributed call processing environments. A gatekeeper can also integrate with large numbers of H.323 VoIP gateways within the organization, or it can act as a session border controller between an enterprise IP Telephony network and a service provider VoIP transport network.

Therefore, as it pertains to Cisco IP Video Telephony deployments, the Cisco IOS Gatekeeper can perform one or both of the following roles:

- Endpoint gatekeeper

An endpoint gatekeeper is configured to manage all RAS procedures for calls to, from, and between H.323 clients, MCUs, and H.320 video gateways. The endpoint gatekeeper directs all such calls to the appropriate Unified CM cluster so that Unified CM can perform all of the H.225 call routing and H.245 media negotiations.

- Infrastructure gatekeeper

An infrastructure gatekeeper is configured to manage all dial plan resolution and bandwidth restrictions (call admission control) between Unified CM clusters, between a Unified CM cluster and a network of H.323 VoIP gateways, or between a Unified CM cluster and a service provider's H.323 VoIP transport network.

In previous Cisco Unified CM releases, the endpoint gatekeeper and the infrastructure gatekeeper had to run on separate routers, and each endpoint gatekeeper could service only a single Unified CM cluster. If multiple Unified CM clusters existed within the enterprise, a separate endpoint gatekeeper had to be deployed for each Unified CM cluster. With the current Cisco Unified CM release, it is possible to combine these roles on a single gatekeeper, using it as an endpoint gatekeeper for one or more Unified CM clusters and as the infrastructure gatekeeper for managing calls between clusters or between a cluster and other H.323 VoIP networks. However, for the following reasons (among others), Cisco recommends that you still separate these roles onto two or more gatekeepers:

- Scalability

Depending on the Cisco IOS router platform you choose to deploy and your estimated busy hour call volume, you might need several gatekeepers to handle the load.

- Geographical resiliency

Putting all of your eggs into one basket may not be wise in a large, multi-national VoIP network. Having gatekeepers placed throughout your network (typically by geography) can provide better fault isolation in the event of a gatekeeper failure.

- Incompatibilities


Some configuration aspects of the gatekeeper are global in nature (they pertain to all endpoints registered with that gatekeeper). For example, the command **arq reject-unknown-prefix**, which may be useful in some H.323 VoIP transport environments, conflicts with the use of the **gw-type-prefix <prefix> default-technology** command, which is used in endpoint gatekeepers to route calls to Unified CM. While Cisco IOS does not stop you from configuring both commands on

the same gatekeeper, the **arq reject-unknown-prefix** command takes precedence and, therefore, calls to unknown numbers will be rejected instead of being routed to Unified CM. In this case, you would have to use one gatekeeper for the H.323 VoIP transport network and another gatekeeper for the Unified CM cluster(s).

Another example of incompatibility can occur in the way you configure the gatekeeper for redundancy. Most Cisco H.323 voice devices, including Cisco Voice Gateways and Unified CM, support the H.323v3 Alternate Gatekeeper feature, which would allow you to configure the gatekeepers as a gatekeeper cluster using the Gatekeeper Update Protocol (GUP) to keep in sync with each other. However, many H.323 video endpoints do not support Alternate Gatekeeper, so the gatekeepers must be configured to use Hot Standby Routing Protocol (HSRP) for redundancy. You cannot mix and match these two redundancy methods on the same gatekeeper. In this case, you might decide to use a gatekeeper cluster for those endpoints that support Alternate Gatekeeper and an HSRP pair of gatekeepers for those that do not.

Figure 12-10 illustrates a network scenario with two Unified CM clusters. Each cluster consists of SCCP and H.323 clients, H.323 MCUs, and H.320 gateways. To manage the RAS aspects of the H.323 clients, MCUs, and H.320 gateways, an endpoint gatekeeper is deployed with each cluster. A separate infrastructure gatekeeper manages dial plan resolution and bandwidth between the clusters. Gatekeeper redundancy is not shown in the figure, although each of these gatekeepers may actually be multiple gatekeepers configured for either Alternate Gatekeeper or HSRP-based redundancy.

Figure 12-10 Two Unified CM Clusters with Required Gatekeepers



Endpoint Gatekeepers

An endpoint gatekeeper is required any time both of the following conditions are met:

- The cluster contains H.323 clients, H.323 MCUs, or H.320 gateways (collectively referred to as H.323 endpoints). If none of these types of endpoints exists (for example, if all clients are SCCP endpoints and there are no MCUs or H.320 gateways), then an endpoint gatekeeper is not needed.
- And either of the following conditions is true:
 - The H.323 endpoints require RAS procedures to initiate calls to E.164 addresses. As mentioned earlier, a growing number of devices are capable of peer-to-peer call signaling, in which case there is no need for those devices to register with a gatekeeper.
 - The H.323 endpoints do not use static IP addresses.

The role of the endpoint gatekeeper is simply to handle the RAS aspects of communications with the endpoints, providing a place for these H.323 endpoints to register. The endpoint gatekeeper responds to all call requests made to, from, or between these endpoints by directing the call to the appropriate Unified CM server(s) so that Unified CM can perform all of the call routing and bandwidth control functions. To accomplish this call routing and bandwidth control, you configure Unified CM to register H.323 trunk(s) with the gatekeeper and configure the gatekeeper to route calls to those trunks for all calls to, from, or within that zone.

Cisco Unified CM should register to endpoint gatekeepers using a type of H.323 trunk called the RASAggregator trunk. This type of trunk is used for all H.323 client, H.323 MCU, or H.320 gateway zones, while the gatekeeper-controlled intercluster trunk and gatekeeper-controlled H.225 trunk are used to integrate with infrastructure gatekeepers.

Provisioning H.323 Clients

H.323 clients are provisioned much the same way as other phones are, in that you create a new phone (model type = H.323 Client), assign a directory number to it, and assign it a calling search space, device pool, and so forth. You configure the H.323 clients in Unified CM in one of the following ways. The method you use depends on whether or not the client uses a static IP address and whether or not the client requires RAS procedures to dial E.164 addresses.

- Gatekeeper controlled

This type of configuration is used for clients that do not have a static IP address assigned to them (they use a DHCP-assigned address) and that require RAS procedures to dial E.164 addresses. A RASAggregator trunk is used to communicate to and from these clients. (See [Figure 12-11](#) and [Figure 12-12](#).)
- Non-gatekeeper controlled, asynchronous

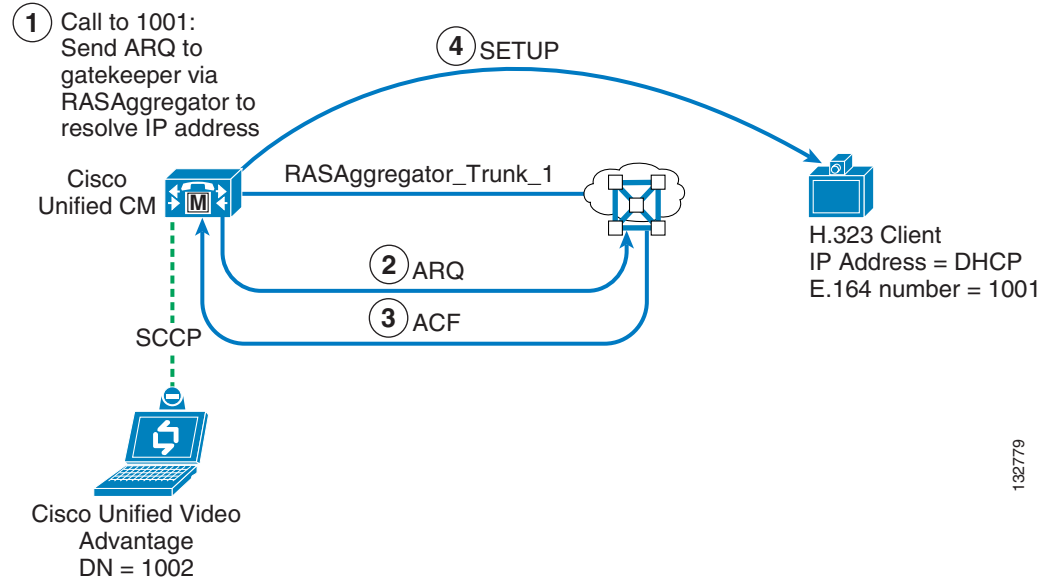
This type of configuration is used for clients that have a static IP address assigned to them but that require RAS procedures to dial E.164 addresses. While Unified CM can signal directly to them without the need of a gatekeeper to resolve their IP addresses, they are not able to signal directly to Unified CM but instead must query the gatekeeper to resolve the E.164 address they are trying to dial (thus, asynchronous communications). To support these types of clients, you must have at least one gatekeeper-controlled client defined in Unified CM for each zone on the gatekeeper, even if all the clients actually use static IP addresses. In this case, the non-gatekeeper controlled client may be a "dummy" client that does not actually exist. Its purpose is merely to create the RASAggregator trunk so that the gatekeeper will be able to route calls from the clients to Unified CM. (See [Figure 12-13](#) and [Figure 12-14](#).)

- Non-gatekeeper controlled, synchronous

This type of configuration is used for clients that have a static IP address and are also capable of peer-to-peer signaling (that is, they do not require RAS procedures to dial E.164 numbers). Unified CM signals directly to them, and they signal directly to Unified CM (thus, synchronous communications). No gatekeeper or RASAggregator trunk is needed for this type of client. (See Figure 12-15 and Figure 12-16.)

Figure 12-11 through Figure 12-16 illustrate the call signaling flows used in these three scenarios.

Figure 12-11 Call to Gatekeeper-Controlled Client from Unified CM



132779

Figure 12-12 Call from Gatekeeper-Controlled Client to Unified CM

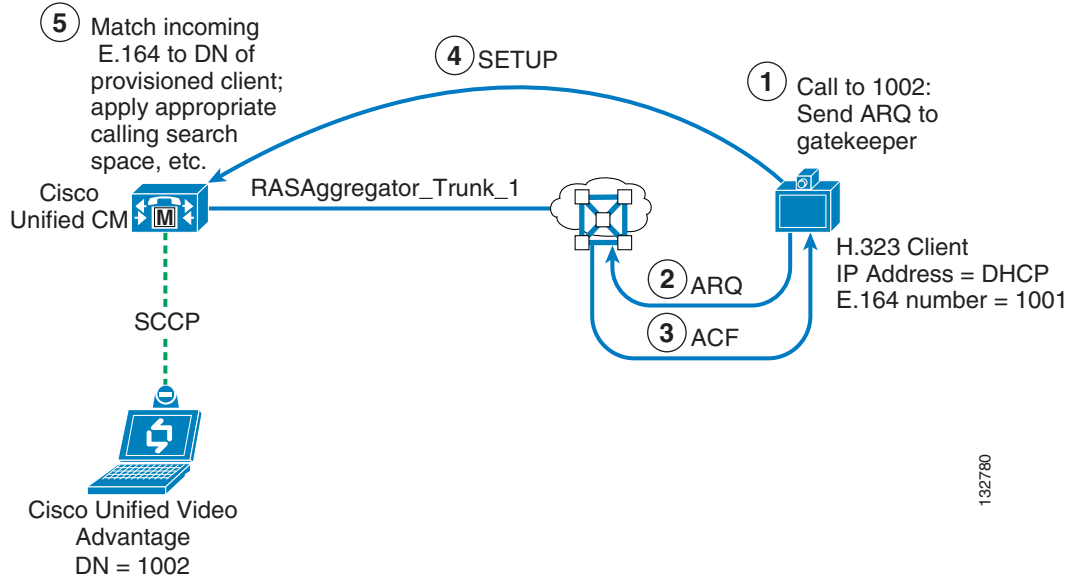


Figure 12-13 Call to Non-Gatekeeper Controlled Client from Unified CM (Asynchronous)

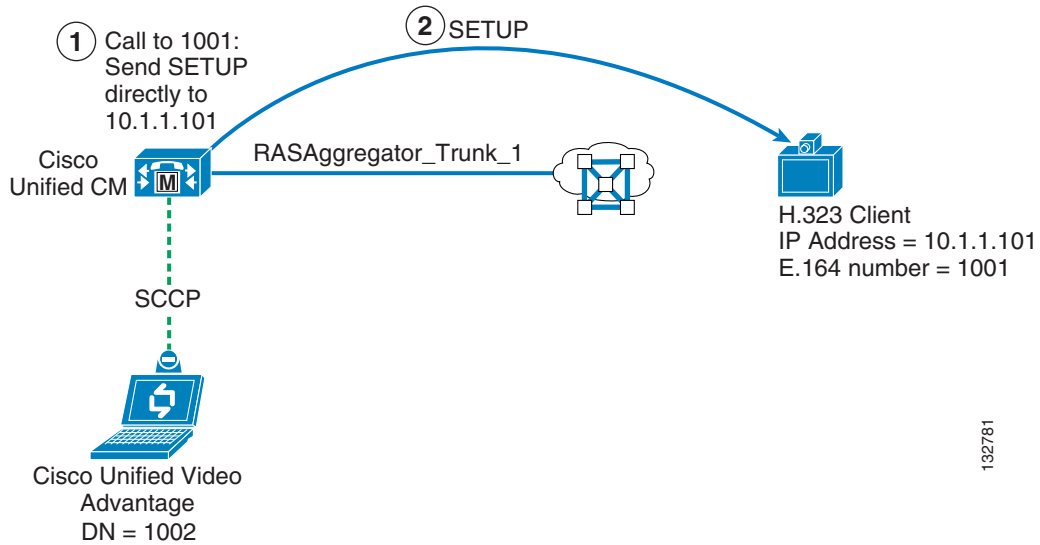


Figure 12-14 Call from Non-Gatekeeper Controlled Client to Unified CM (Asynchronous)

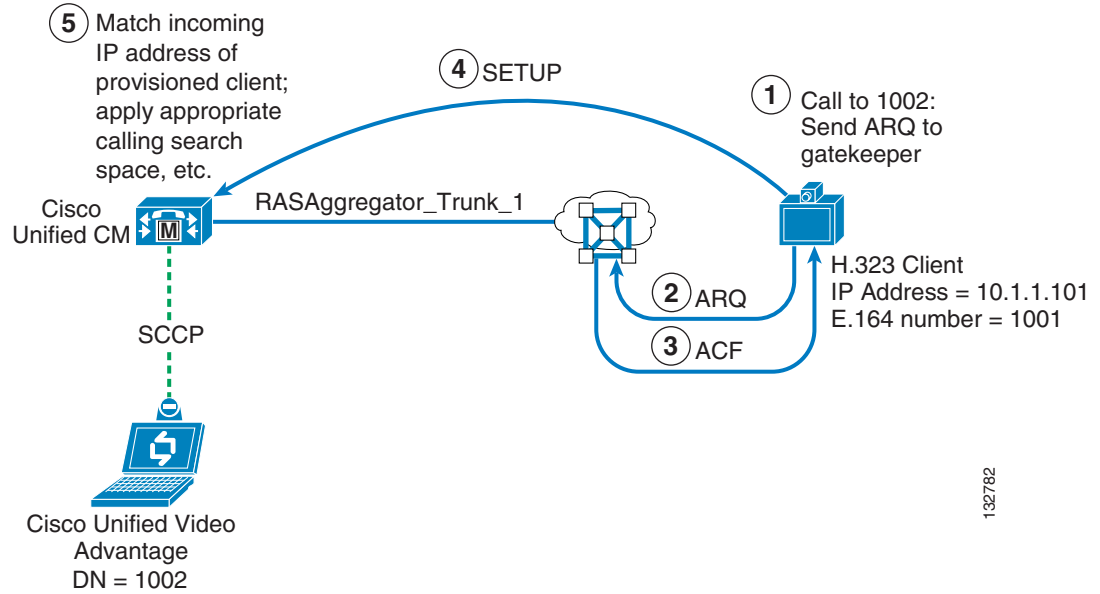


Figure 12-15 Call to Non-Gatekeeper Controlled Client from Unified CM (Synchronous)

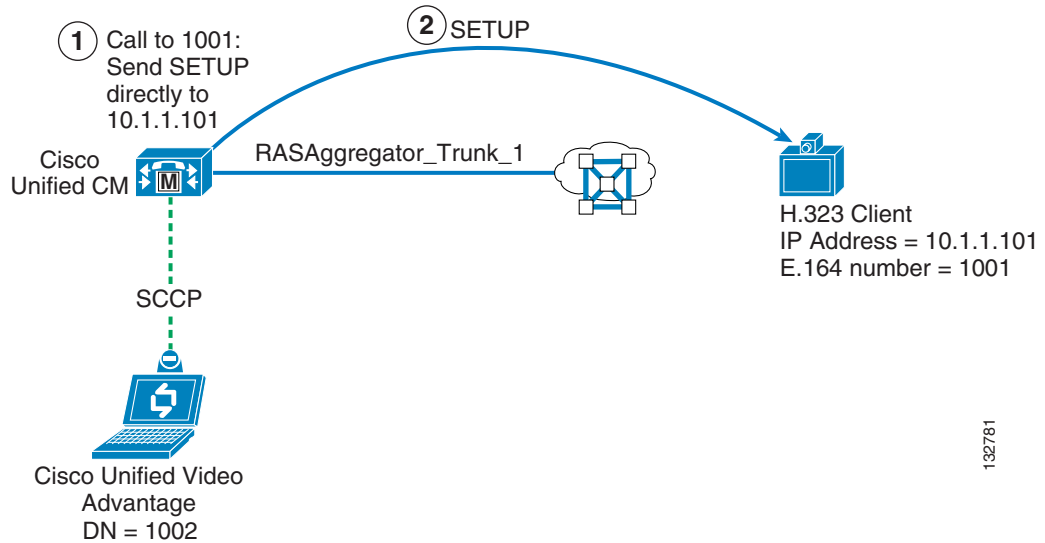
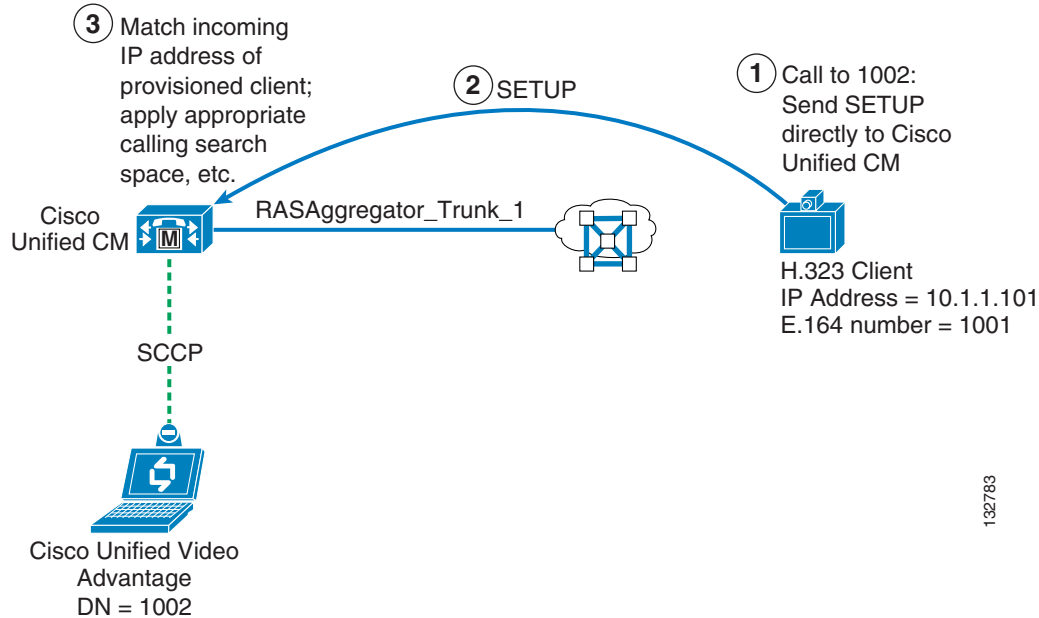


Figure 12-16 Call from Non-Gatekeeper Controlled Client to Unified CM (Synchronous)



132783

Gatekeeper-Controlled Clients

When you configure an H.323 client as gatekeeper-controlled, you may enter any alpha-numeric string (such as a descriptive name) in the Device Name field, check the **Gatekeeper-controlled** box, and fill in the following fields:

- **Device Pool**
The device pool you want the client to use. All H.323 clients (whether gatekeeper-controlled or non-gatekeeper controlled) that are registered in the same zone must use the same device pool. If you accidentally assign different device pools across the endpoints, Unified CM will register multiple RASAggregator trunks within the zone, and an inbound call might be rejected by Unified CM if the call is directed to the wrong RASAggregator trunk.
- **Gatekeeper**
A drop-down list of gatekeeper IP addresses. You must define the gatekeeper in Unified CM before configuring any gatekeeper-controlled H.323 clients.
- **Technology Prefix**
The technology prefix used by the RASAggregator trunk to register in the client zone on the gatekeeper. This technology prefix must match what is configured as the default technology prefix on the gatekeeper. All gatekeeper-controlled H.323 clients that are registered in the same zone must use the same technology prefix. If you accidentally assign different technology prefixes across the endpoints, Unified CM will register multiple RASAggregator trunks within the zone, and an inbound call might be rejected by Unified CM if the call is directed to the wrong RASAggregator trunk. Cisco recommends that you use **1#** for this prefix.

- Zone Name

The (case-sensitive) name of the client zone as configured in the gatekeeper. All gatekeeper-controlled H.323 clients that are registered in the same zone must use the same zone name. If you accidentally assign different zone names (remember, the field is case sensitive) across the endpoints, Unified CM will attempt to register multiple RASAggregator trunks with the gatekeeper (but the one with the incorrect zone name will fail to register), and an inbound call might be rejected by Unified CM if the call is directed to the wrong RASAggregator trunk.

Also, you must set the Unified CM service parameter **Send Product ID and Version ID** to **True**. This parameter allows the RASAggregator trunk to register with the gatekeeper as an H323-GW, so that the gatekeeper can direct all H.323 calls to, from, or within the client zone to the RASAggregator trunk.

Non-Gatekeeper Controlled Clients

When provisioning an H.323 client as non-gatekeeper controlled, you must enter the static IP address of the client into the Device Name field and leave all of the settings under the Gatekeeper-controlled section blank (unchecked). Unified CM then uses the static IP address to reach the client any time a call is extended to its directory number.

If the client is configured to use peer-to-peer mode, then no further configuration is required. If the client requires RAS procedures to place calls to E.164 addresses, then you must also configure a dummy gatekeeper-controlled H.323 client in order to create the RASAggregator trunk, by filling in the following fields:

- Device Name

A descriptive name that identifies this client as a dummy client used for the purpose of creating the RASAggregator trunk for the client zone.

- Device Pool

The device pool you chose when configuring the non-gatekeeper controlled H.323 client(s). If the device pool assigned to the dummy client is different than that assigned to the real clients, inbound calls from the real clients might be rejected by Unified CM.

- Gatekeeper

A drop-down list of gatekeeper IP addresses. You must define the gatekeeper in Unified CM before configuring the dummy gatekeeper-controlled H.323 client.

- Technology Prefix

The technology prefix used by the RASAggregator trunk to register in the client zone on the gatekeeper. This technology prefix must match what is configured as the default technology prefix on the gatekeeper. Cisco recommends that you use **1#** for this prefix.

- Zone Name

The (case-sensitive) name of the client zone as configured in the gatekeeper.

Also, you must set the Unified CM service parameter **Send Product ID and Version ID** to **True**. This parameter allows the RASAggregator trunk to register with the gatekeeper as an H323-GW, so that the gatekeeper can direct all H.323 calls to, from, or within the client zone to the RASAggregator trunk.

Provisioning H.323 MCUs

H.323 MCUs are provisioned in Unified CM as H.323 gateways, and then route patterns are configured to extend calls to these devices. When provisioning an H.323 gateway, you must enter the static IP address and TCP signaling port of the MCU into the Device Name field. Unified CM then uses the static IP address and TCP port to reach the MCU any time a call matches the route pattern(s) associated with it.

**Note**

The Cisco Unified Videoconferencing 3500 and 5000 Series MCUs do not listen on TCP port 1720 by default. (The Cisco Unified Videoconferencing 3500 and 5000 Series MCUs listen on port 2720 by default.) You must verify which TCP port they are listening on, and either change it to 1720 or provision the correct port in Unified CM.

If the MCU is configured to use peer-to-peer mode, then no further configuration is required. (Cisco Unified Videoconferencing MCUs do not currently support peer-to-peer mode, but some third-party MCUs do.) If the MCU requires RAS procedures to place calls to E.164 addresses, then you must also configure a dummy gatekeeper-controlled H.323 client in order to create the RASAggregator trunk, by filling in the fields for Device Name, Device Pool, Gatekeeper, Technology Prefix, and Zone Name as discussed in the section on [Non-Gatekeeper Controlled Clients](#), page 12-33.

MCU Service Prefixes

H.323 MCUs can use either E.164 addresses or technology prefixes (also referred to as service prefixes in the MCU) as the dial-in number(s) to reach reservationless or scheduled H.323 conferences running on them. Cisco recommends that you configure the MCUs to use E.164 addresses by setting the MCU Mode to **MCU** instead of **Gateway** in the MCU administration screens. If the **MCU** setting is not available on the model of MCU you are using, then you must use the following special configuration to properly route calls placed from other H.323 endpoints to the MCU:

If the MCU is configured in **Gateway** mode or is another vendor's MCU that (for whatever reason) requires its conference IDs to register as technology prefixes instead of as E.164 addresses, then the service prefix(s) of the MCU must begin with a # character. For example, if the MCU's service prefix is 8005551212, then you must provision the service prefix on the MCU as #8005551212. Thus, when other H.323 endpoints dial 8005551212, the gatekeeper will not find a matching technology prefix registered and will instead route the call to the RASAggregator trunk that is registered with the default technology prefix in the zone of the endpoint that is placing the call. Unified CM must then prepend the # character to the beginning of the called number before extending the call to the MCU. This character is prepended on the route pattern(s) associated with the H.323 gateway representing the MCU. Calls to the MCU from SCCP clients will therefore also have this # character prepended to the calling number.

If the MCU is configured in **MCU** mode or is another vendor's MCU that uses E.164 addresses for its conference IDs, then you do not have to prepend the # character. Also note that, if the MCU uses peer-to-peer mode and hence does not need to register its technology prefixes with any gatekeeper, then this situation does not apply and you do not have to prepend a # character.

Provisioning H.320 Gateways

As with H.323 MCUs, H.320 gateways are provisioned in Unified CM as H.323 gateways, and then route patterns are configured to extend calls to these devices. When provisioning an H.323 gateway, you must enter the static IP address and TCP signaling port of the H.320 gateway into the Device Name field. Unified CM then uses the static IP address and TCP port to reach the gateway any time a call matches the route pattern(s) associated with it.

**Note**

The Cisco Unified Videoconferencing 3500 and 5000 Series Gateways do not listen on TCP port 1720 by default. (The Cisco Unified Videoconferencing 3500 and 5000 Series Gateways listen on port 1820 by default.) You must verify which TCP port they are listening on, and either change it to 1720 or provision the correct port in Unified CM.

If the gateway is configured to use peer-to-peer mode, then no further configuration is required. If the gateway requires RAS procedures to place calls to E.164 addresses, then you must also configure a dummy gatekeeper-controlled H.323 client in order to create the RASAggregator trunk, by filling in the fields for Device Name, Device Pool, Gatekeeper, Technology Prefix, and Zone Name as discussed in the section on [Non-Gatekeeper Controlled Clients](#), page 12-33.

Gateway Service Prefixes

H.320 gateways use technology prefixes (also referred to as service prefixes in the gateway) as the prefix that users should dial to reach an ISDN destination. For calls to route correctly, you must configure the service prefix(s) of the gateway to begin with a # character. For example, if the gateway's service prefix that clients dial to reach an ISDN number is 9, then you must provision the service prefix on the gateway as #9. In this way, when H.323 clients dial 9 plus the PSTN number (such as 918005551212), the gatekeeper will not find a matching technology prefix registered and will instead route the call to the Unified CM trunk that is registered with the default technology prefix. Unified CM must then prepend the # character to the beginning of the called number before extending the call to the gateway. Note that, if the gateway uses peer-to-peer mode and hence does not need to register its technology prefixes with any gatekeeper, then this situation does not apply and you do not have to prepend a # character.

Gatekeeper Zone Configuration

The preceding sections discuss how to provision the endpoints in Unified CM Administration. You must also configure the endpoint gatekeeper(s) with the appropriate zone definitions. You must configure a zone for each type of endpoint (client, MCU, or gateway) and, optionally, for each device pool associated with these endpoints in Unified CM.

Each zone is configured to route all calls placed to, from, or within the zone to the RASAggregator trunk registered in that zone. You configure the zones on the endpoint gatekeeper by using the following command syntax:

```
zone local <zone_name> <domain_name> <ip_address> invia <zone_name>
outvia <zone_name> enable-intrazone
```

The command argument **invia** applies to calls placed to the zone from any other zone, **outvia** applies to calls placed from the zone to any other zone, and **enable-intrazone** applies to calls placed within the zone. The following sections illustrate the use of these commands.

Client Zones

The number of client zones you have to configure within each endpoint gatekeeper depends on the following factors:

- The device pools to which the H.323 clients are associated
 - The device pool determines which Unified CM servers are primary, secondary, and tertiary servers for each H.323 client. If you assign all H.323 clients to the same device pool, then you need to define only a single client zone in the endpoint gatekeeper. In other words, for each device pool used by H.323 clients, you must configure a separate client zone in the gatekeeper.
- Whether the endpoint gatekeeper provides services for a single Unified CM cluster or multiple Unified CM clusters

Each client zone is configured to route calls to a particular RASAggregator trunk. Therefore, if one endpoint gatekeeper is used to service multiple Unified CM clusters, then you must define a separate client zone for each cluster that the gatekeeper services.

To illustrate, the following examples show how client zones may be configured. [Example 12-1](#) shows a single client zone defined for a single Unified CM cluster in which all H.323 clients are associated with the same device pool. [Example 12-2](#) shows a single Unified CM cluster in which the H.323 clients are divided between two different device pools.

**Note**

Some of the commands shown in the following examples are the default values applied in the Cisco IOS Gatekeeper and, therefore, would not have to be configured explicitly, nor would they appear in the running configuration. They are included here for thoroughness but are marked by a ! at the beginning of the command line.

Example 12-1 Client Zone for a Single Unified CM Cluster and Single Device Pool

```
gatekeeper
zone local clients domain.com invia clients outvia clients enable-intrazone
gw-type-prefix 1# default-technology
no use-proxy clients default inbound-to terminal
no use-proxy clients default outbound-from terminal
! no arg reject-unknown-prefix
endpoint ttl 60
no shutdown
```

Example 12-2 Client Zones for a Single Unified CM Cluster and Two Device Pools

```
gatekeeper
zone local dp1-clients domain.com invia dp1-clients outvia dp1-clients enable-intrazone
zone local dp2-clients domain.com invia dp2-clients outvia dp2-clients enable-intrazone
gw-type-prefix 1# default-technology
no use-proxy dp1-clients default inbound-to terminal
no use-proxy dp1-clients default outbound-from terminal
no use-proxy dp2-clients default inbound-to terminal
no use-proxy dp2-clients default outbound-from terminal
! no arg reject-unknown-prefix
endpoint ttl 60
no shutdown
```

Disabling The Use of Proxy

The Cisco IOS Gatekeeper, formerly known as the Cisco Multimedia Conference Manager (MCM), previously offered an H.323 proxy function that has been at End of Life (EOL) for some time and is not compatible with Unified CM, but the commands in the gatekeeper to use a proxy for all calls to and from terminals (clients) are still enabled by default. You must disable this function for each client zone by using the following command syntax:

```
gatekeeper
no use-proxy <zone_name> default [inbound-to | outbound-from] terminals
```

The Cisco MCM proxy was replaced by a solution called the Cisco IOS Multiservice IP-to-IP Gateway and the associated via-zone-enabled Cisco IOS Gatekeeper. This document does not discuss the IP-to-IP Gateway, but Cisco Unified CM leverages the via-zone and IP-to-IP gateway constructs by registering its RASAggregator trunks with the gatekeeper, effectively mimicking an IP-to-IP gateway so that the gatekeeper will route all invia, outvia, and enable-intrazone calls to the RASAggregator trunk as if it were an IP-to-IP gateway.

Client Zone Prefixes

For H.323 client zones, there is no need to configure zone prefixes or technology prefixes of any kind, except for the default technology prefix. Instead, the **invia**, **outvia**, **enable-intrazone**, and **gw-type-prefix <I#> default-technology** commands ensure that all calls placed are routed to the RASAggregator trunk associated with the zone in which the call originated.

MCU Zones

The number of MCU zones you have to configure within each endpoint gatekeeper depends on the following factors:

- The device pools to which the MCUs are associated

The device pool determines which Unified CM servers are primary, secondary, and tertiary servers for each MCU. If you assign all MCUs to the same device pool, then you need to define only a single MCU zone in the endpoint gatekeeper. In other words, for each device pool used by MCUs, you must configure a separate MCU zone in the gatekeeper.
- Whether the endpoint gatekeeper provides services for a single Unified CM cluster or multiple Unified CM clusters

Each MCU zone is configured to route calls to a particular RASAggregator trunk. Therefore, if one endpoint gatekeeper is used to service multiple Unified CM clusters, then you must define a separate MCU zone for each cluster that the gatekeeper services.

Gatekeeper configuration for MCU zones is similar to the configurations shown in [Example 12-1](#) and [Example 12-2](#) for MCUs.

Disabling The Use of Proxy

By default, the Cisco IOS Gatekeeper is set to not use a proxy for calls to and from MCUs or gateways. However, if you have enabled the use of proxy for those types of endpoints, you must disable it for each MCU zone by using the following command syntax:

```
gatekeeper
no use-proxy <zone_name> default [inbound-to | outbound-from] [MCU | gateway]
```

If your MCU is registering as an MCU, then use the **MCU** argument at the end of the **no use-proxy** command; if your MCU is registering as a gateway, then use the **gateway** argument instead.

MCU Zone Prefixes

For H.323 MCU zones, there is no need to configure zone prefixes or technology prefixes of any kind, except for the default technology prefix. Instead, the **invia**, **outvia**, **enable-intrazone**, and **gw-type-prefix <I#> default-technology** commands ensure that all calls placed are routed to the RASAggregator trunk associated with the zone in which the call originated.

If your MCUs are registering their service prefixes as technology prefixes instead of E.164 addresses, use the special configuration described previously for prepending a # character to the MCU's service prefixes (see [MCU Service Prefixes, page 12-34](#)). Due to the way the Cisco IOS Gatekeeper selects a via-zone for calls to a technology prefix, when the endpoint dials the service prefix of the MCU, the call will fail if the gatekeeper finds a matching technology prefix registered. You must ensure that the client does not dial the # character, so that the gatekeeper will not find a matching technology prefix and will instead route the call to the RASAggregator trunk associated with the zone in which the call originated.

H.320 Gateway Zones

The number of H.320 gateway zones you have to configure within each endpoint gatekeeper depends on the following factors:

- The device pools to which the H.320 gateways are associated

The device pool determines which Unified CM servers are primary, secondary, and tertiary servers for each H.320 gateway. If you assign all gateways to the same device pool, then you need to define only a single gateway zone in the endpoint gatekeeper. In other words, for each device pool used by H.320 gateways, you must configure a separate gateway zone in the gatekeeper.

- Whether the endpoint gatekeeper provides services for a single Unified CM cluster or multiple Unified CM clusters

Each gateway zone is configured to route calls to a particular RASAggregator trunk. Therefore, if one endpoint gatekeeper is used to service multiple Unified CM clusters, then you must define a separate gateway zone for each cluster that the gatekeeper services.

Gatekeeper configuration for gateway zones is similar to the configurations shown in [Example 12-1](#) and [Example 12-2](#) for gateways.

Disabling The Use of Proxy

By default, the Cisco IOS Gatekeeper is set to not use a proxy for calls to and from gateways. However, if you have enabled the use of proxy for those types of endpoints, you must disable it for each H.320 gateway zone by using the following command syntax:

```
gatekeeper
no use-proxy <zone_name> default [inbound-to | outbound-from] gateway
```

Gateway Zone Prefixes

There is no need to configure zone prefixes of any kind for H.320 gateway zones. Instead, the **invia**, **outvia**, **enable-intrazone**, and **gw-type-prefix <1#> default-technology** commands ensure that all calls placed are routed to the RASAggregator trunk associated with the zone in which the call originated.

You must also use the special configuration described previously for prepending a # character to the gateway's service prefixes (see [Gateway Service Prefixes, page 12-35](#)). Due to the way the Cisco IOS Gatekeeper selects a via-zone for calls to a technology prefix, when the endpoint dials the service prefix of the gateway, the call will fail if the gatekeeper finds a matching technology prefix registered. You must ensure that the client does not dial the # character, so that the gatekeeper will not find a matching technology prefix and will instead route the call to the RASAggregator trunk associated with the zone in which the call originated.

Zone Subnets

As mentioned previously, the H.323 specification permits a single gatekeeper to manage multiple zones. However, the gatekeeper needs a way to decide which zone an endpoint should be placed in when it receives a Registration Request (RRQ) from that device. The RRQ message contains a Gatekeeper Identifier field that enables the endpoint to indicate the zone in which it would like to register. However, many H.323 video endpoints do not populate this field, and if the gatekeeper has multiple zones defined, it will not know which zone to place the endpoint into. Therefore, you must use of the **zone subnet** command to tell the gatekeeper which zone to associate with the endpoint. This command defines which IP addresses or IP address ranges are permitted to register in each zone. The command syntax requires that you enter a network mask. Therefore, you can specify either a particular host address by entering a 32-bit (/32) network mask or a range of addresses by specifying a smaller network mask.

Because MCUs, H.320 gateways, and Unified CM servers typically use fixed IP addresses but H.323 clients can use DHCP addresses, Cisco recommends that you define **zone subnet** commands only for the MCU and gateway zones but leave the client zones open so that any IP address is permitted in them. Note that you must also permit the Unified CM servers to register in the MCU and gateway zones, as illustrated in [Example 12-3](#).

**Note**

Some of the commands shown in the following example are the default values applied in the Cisco IOS Gatekeeper and, therefore, would not have to be configured explicitly, nor would they appear in the running configuration. They are included here for thoroughness but are marked by a ! at the beginning of the command line.

Example 12-3 Defining Zone Subnets

```
gatekeeper
no zone subnet MCUs default enable
zone subnet MCUs [MCUs_IP_addr]/32 enable
zone subnet MCUs [RASAggregators_IP_addr]/32 enable
no zone subnet gateways default enable
zone subnet gateways [gateways_IP_addr]/32 enable
zone subnet gateways [RASAggregators_IP_addr]/32 enable
! zone subnet clients default enable
no zone subnet clients [MCUs_IP_addr]/32 enable
no zone subnet clients [gateways_IP_addr]/32 enable
```

The configuration in [Example 12-3](#) explicitly permits the MCU and the RASAggregator for the MCU zone to register in the MCU zone, and it explicitly permits the gateway and RASAggregator for the gateway zone to register in the gateway zone. It also explicitly denies the MCU and gateway from registering in the client zone, while implicitly permitting all other IP addresses (including the RASAggregator for the client zone) to register in the client zone.

Endpoint Time to Live

Endpoints send lightweight Registration Requests (RRQs) to their gatekeeper periodically to maintain their registration status. The frequency with which they send these RRQs is referred to as the Time to Live (TTL) value. The endpoint may specify the TTL it wishes to use in the body of its RRQs. The gatekeeper may then honor the endpoint's requested TTL value by echoing it in the Registration Confirm (RCF) response or, alternatively, may override the endpoint's request by specifying a different TTL value in the RCF.

If the TTL value is not specified in the RRQ, the gatekeeper should specify one in its RCF response. The endpoint should then honor the TTL specified by the gatekeeper. The Cisco IOS Gatekeeper honors all TTL values specified by the endpoints. However, many H.323 video endpoints do not specify a TTL value in their RRQs. In such cases, the Cisco IOS Gatekeeper defaults to specifying a TTL value of 1800 seconds (30 minutes). The Cisco IOS Gatekeeper will flush the endpoint's registration after three TTL intervals have passed without receiving any messages from the endpoint (3 * 30 minutes = 90 minutes).

A large TTL value can cause problems with H.323 clients that do not use static IP addresses. For example, with the default TTL value of 1800 seconds, if you disconnect the client from the network and move it to another location in which it receives a different DHCP address, it will fail to register with the gatekeeper (Registration Reject (RRJ) cause value "duplicate alias") until three TTL intervals have passed, and the gatekeeper will flush that endpoint's original registration.

Therefore, Cisco recommends that you consider reducing the TTL value to as low a number as possible without causing any negative effect on your network. The Cisco IOS Gatekeeper permits you to set the TTL value anywhere in the range of 60 seconds to 3600 seconds. In most cases, 60 seconds should work well. However, if your gatekeeper is already heavily utilized, adjusting the TTL from the default of 1800 seconds to 60 seconds might cause it to become overwhelmed.

Use the following command syntax to set the TTL value:

```
gatekeeper
endpoint ttl <seconds>
```

Supported Gatekeeper Platforms

To act as an endpoint gatekeeper with Cisco Unified CM, the Cisco IOS Gatekeeper must run Cisco IOS Release 12.3(11)T or greater. For minimum Cisco IOS release requirements on the infrastructure gatekeeper, refer to the latest *Cisco Unified Communications System Release Notes for IP Telephony* available at:

<http://www.cisco.com/go/unified-techinfo>

To determine which release and feature set you should use for your router platform, use the Cisco Feature Navigator (requires a Cisco.com login account), available at:

<http://tools.cisco.com/ITDIT/CFN/jsp/index.jsp>

For more information, also refer to the *Cisco IOS H323 Gatekeeper Data Sheet*, available at:

http://www.cisco.com/en/US/prod/collateral/voicesw/ps6788/vcallcon/ps4139/data_sheet_c78_561921.html

Summary of Endpoint Gatekeepers

This section summarizes some key points to remember about endpoint gatekeepers and provides some example configurations that combine techniques used in the previous examples.

- Configure a separate zone in the endpoint gatekeeper for each type of endpoint (clients, MCUs, and H.320 gateways). If the endpoints are associated with multiple device pools, configure multiple zones for each type of endpoint.
- Configure a RASAggregator trunk to register in each zone. This trunk is automatically created when you configure gatekeeper-controlled H.323 clients in Unified CM Administration. However, for non-gatekeeper controlled H.323 clients, H.323 MCUs, and H.320 gateways, you must configure a dummy gatekeeper-controlled H.323 client in order to create the RASAggregator trunk for that zone.
- Set the service parameter **Send Product ID and Version ID** to **True** in order for the RASAggregator trunk to register with the gatekeeper as an IP-to-IP gateway. This setting enables the RASAggregator to be selected by the gatekeeper for all calls placed to, from, or within each zone due to the use of the **invia**, **outvia**, **enable-intrazone**, and **gw-type-prefix <I#> default-technology** commands applied to each local zone definition.
- You do not have to associate any zone prefixes for any of the endpoint zones. No matter what the endpoint dials, the gatekeeper should not find a matching zone prefix or technology prefix but should instead route the call to the RASAggregator trunk associated with the zone from which the call originated. To avoid having the gatekeeper accidentally match the dialed number to the technology prefix of your MCUs or gateways, mask all MCU and gateway service prefixes with a # character, and then prepend the # character in the route pattern associated with that MCU or gateway.

- Configure zone subnets if any of the H.323 endpoints do not support the ability to specify the Gatekeeper Identifier (name of the zone) with which they wish to register.
- Disable the use of the old MCM Proxy for all zones.
- Set the endpoint registration Time to Live (TTL) to as low of a value as you can without creating undo stress on the gatekeeper. In extreme cases where the gatekeeper is serving hundreds of endpoint registrations, setting the TTL to 60 seconds might cause an unmanageable amount of RAS traffic. In smaller environments, setting it to 60 seconds should work well.

[Example 12-4](#) shows a configuration for an endpoint gatekeeper servicing a single Unified CM cluster in which a single device pool is used to service all H.323 video endpoint types.



Note

Some of the commands shown in the following examples are the default values applied in the Cisco IOS Gatekeeper and, therefore, would not have to be configured explicitly, nor would they appear in the running configuration. They are included here for thoroughness but are marked by a ! at the beginning of the command line.

Example 12-4 Endpoint Gatekeeper Configuration for a Single Cluster and a Single Device Pool

```
gatekeeper
zone local clients domain.com invia clients outvia clients enable-intrazone
zone local MCUs domain.com invia MCUs outvia MCUs enable-intrazone
zone local gateways domain.com invia gateways outvia gateways enable-intrazone
! zone subnet clients default enable
no zone subnet clients [MCUs_IP_addr]/32 enable
no zone subnet clients [gateways_IP_addr]/32 enable
no zone subnet MCUs default enable
zone subnet MCUs [MCUs_IP_addr]/32 enable
zone subnet MCUs [RASAggregators_IP_addr]/32 enable
no zone subnet gateways default enable
zone subnet gateways [gateways_IP_addr]/32 enable
zone subnet gateways [RASAggregators_IP_addr]/32 enable
no use-proxy clients inbound-to terminals
no use-proxy clients outbound-from terminals
! no use-proxy MCUs inbound-to [MCU | gateway]
! no use-proxy MCUs outbound-from [MCU | gateway]
! no use-proxy gateways inbound-to gateway
! no use-proxy gateways outbound-from gateway
gw-type-prefix 1# default-technology
! no arq reject-unknown-prefix
endpoint ttl 60
no shutdown
```

Applications

Cisco IP Communications provides an expanding portfolio of applications that extend the features of Unified CM and provide advanced capabilities and integration with other communication media. Many of these applications can be used in conjunction with IP Video Telephony devices, even if they do not specifically support video. For instance, Cisco Unified CM does not support the negotiation of video channels for CTI applications using the TAPI/JTAPI protocols, but that does not necessarily preclude using a CTI application in conjunction with a video call. This section reviews some of the Cisco and third-party applications and discusses whether or not they can be used to provide advanced call treatment for video calls.

CTI Applications

The following applications are based on the Computer Telephony Integration (CTI) interface.

Cisco Emergency Responder

Cisco Emergency Responder (ER) routes emergency (911) calls to the correct Public Safety Answering Point (PSAP). It also provides the PSAP with the correct calling line ID of the originating device so that the PSAP can respond to the correct physical location of the incident and call the party back in the event that the call is disconnected. Cisco ER uses JTAPI to integrate with Unified CM. Emergency calls are routed to Cisco ER via a CTI route point, then Cisco ER decides which PSAP to forward the call to and what calling line ID to display. Cisco ER tracks each endpoint on the network to determine its physical location by using Simple Network Management Protocol (SNMP) and Cisco Discovery Protocol (CDP) to discover the physical port and specific Cisco Catalyst Ethernet switch to which the endpoint is connected. If CDP is not available, Cisco ER can be configured to locate endpoints by their IP subnet instead. Cisco ER then correlates this information with the physical location of the switch and stores the information in its database.

Both Cisco Unified Video Advantage and the Cisco IP Video Phone 7985 support CDP for the purpose of Cisco ER discovery. Cisco Unified Video Advantage does not send CDP messages to the switch directly but relies on its associated Cisco Unified IP Phone for this support. Therefore, if a Video Telephony user dials 911, Cisco ER is able to route the call to the correct PSAP.

Because third-party SCCP video endpoints do not support CDP, Cisco ER must track these endpoints by their IP subnet. Cisco ER is therefore able to route the call to the correct PSAP.

Because H.323 videoconferencing clients do not support CDP, Cisco ER must track them by their IP subnet. Cisco ER is therefore able to route the call to the correct PSAP. However, the H.323 device must support the Empty Capabilities Set (ECS) procedure in order to have its call routed by Cisco ER. If the H.323 endpoint does not support receiving an ECS from Unified CM, calls to 911 that are handled by Cisco ER will fail.

When Cisco Cius is connected to Unified CM through a 3G or 4G connection, Cisco ER will not be able to locate the appropriate PSAP for the user's location. Cisco Cius users should be instructed to use alternate means to call 911 while roaming with 3G or 4G connections.

Cisco Unified Communications Manager Assistant

Cisco Unified Communications Manager Assistant enables administrative assistants to provide coverage for the managers with whom they are associated. Unified CM Assistant uses JTAPI to integrate with Unified CM. While Unified CM Assistant is not specifically video-capable, there is nothing stopping Unified CM Assistant from being used with phones that are video-enabled. Once Unified CM Assistant has handled the call and the call is transferred to the final destination device, the two devices in the call communicate with each other directly and could establish video channels at that point. For example, if a video-capable endpoint dialed the directory number of a Manager, and the Assistant covered the call using the Unified CM Assistant application, there might not be any video established during the initial handling of the call. But once the Assistant transferred the caller to the Manager, Unified CM could then negotiate video channels. However, H.323 devices must support the Empty Capabilities Set (ECS) procedure in order to interoperate with Unified CM Assistant. If the H.323 endpoint does not support receiving an ECS from Unified CM, calls that are intercepted by Unified CM Assistant will fail when the Assistant attempts to transfer the call to the Manager.

Cisco Unified IP Interactive Voice Response and Cisco Unified Contact Center

Cisco Unified IP Interactive Voice Response (Unified IP IVR) and Cisco Unified Contact Center (Unified CC) use JTAPI to integrate with Unified CM. If a video-capable device calls into an IVR application (such as a help desk), the communication is audio-only while the caller is connected to the application server (that is, while the caller browses the IVR menu or waits in queue for a help-desk member to take the call). However, once the IVR application transfers the call to its final destination, video channels can be negotiated at that time. H.323 devices must support the Empty Capabilities Set (ECS) procedure in order to interoperate with Cisco Unified IP IVR and Unified CC. If the H.323 endpoint does not support receiving an ECS from Unified CM, calls that are intercepted by Cisco Unified IP IVR or Unified CC will fail when the application attempts to transfer the caller to the final destination.

IVR applications often use DTMF tones to select options in the IVR menu. An alternative is speech recognition, which enables the caller to speak commands to the IVR server instead of pressing keys on the phone. Because Cisco Unified IP IVR and Unified CC both use JTAPI to integrate with Unified CM, they pass DTMF tones through out-of-band signaling messages. Many H.323 devices on the market today use in-band DTMF tones, and these H.323 clients would not be able to use DTMF to navigate an IP IVR or Unified CC menu. However, these H.323 clients could use speech recognition if the IVR server is enabled for it. Video-capable devices such as Cisco Unified Video Advantage, third-party SCCP video devices, and any H.323 endpoint that uses H.245 alphanumeric out-of-band signaling for DTMF, can navigate the IVR menus using DTMF tones.

Cisco Unified Enterprise Attendant Console

The Cisco Unified Enterprise Attendant Console uses JTAPI to integrate with Unified CM. The Unified Enterprise Attendant Console is used as a front-office device to handle incoming calls. Although the Unified Enterprise Attendant Console does not specifically support video, video channels can be negotiated once the call is transferred to its final destination. However, H.323 devices must support the Empty Capabilities Set (ECS) procedure in order to interoperate with the Unified Enterprise Attendant Console. If the H.323 endpoint does not support receiving an ECS from Unified CM, calls that are intercepted by a Unified Enterprise Attendant Console will fail when the attendant attempts to transfer the caller to the final destination.

Cisco IP Communicator

Cisco IP Communicator is an SCCP software phone, and it does not support video. Video can be provided by associating the Cisco Unified Video Advantage client with Cisco IP Communicator. As of version 2.0, Cisco IP Communicator can associate to Cisco Unified Video Advantage 2.0 when both applications reside on the same PC. For more information, refer to the chapter on [Unified Communications Endpoints](#), page 18-1.

Cisco IP Communicator 2.1 and later releases also provide a SIP device protocol option for creating or adding a Cisco IP Communicator device. However, Cisco Unified Video Advantage is not supported for Cisco IP Communicator using SIP protocol.

Collaboration Solutions

The following technologies are sometimes used to provide video communications between endpoints.

T.120 Application Sharing

Some videoconferencing endpoints use the T.120 protocol to share documents, whiteboards, and text among participants in a conference. Unified CM does not support negotiating a T.120 channel. Instead of T.120, Cisco recommends using web-based collaboration solutions such as Cisco MeetingPlace or other third-party collaboration solutions.

Cisco Unified MeetingPlace

Cisco Unified MeetingPlace combines a high-end audio and video conferencing solution with a web-based front end for scheduling and participating in conferences. For more information, see [Cisco Unified MeetingPlace, page 22-13](#).

Video Interoperability

Video interoperability is the audio and video support for point-to-point calls between Cisco TelePresence System (CTS) endpoints, other Cisco Unified Communications (UC) video endpoints, and third-party video endpoints. Prior to Cisco Unified CM 8.5, video interoperability between the different families of video endpoints was possible only with the insertion of a video component between endpoints, such as a video transcoder or a multipoint control unit (MCU).

Cisco Unified CM 8.5 and later releases not only offer native video interoperability between different video endpoint family types, point-to-point, but also provide better video interoperability in general with H.264 codec negotiation in SIP and H.323 protocols and enable the endpoints to negotiate high definition (HD) resolutions when available.

As stated earlier, video interoperability in Unified CM also enables Cisco TelePresence System (CTS) endpoints to communicate with non-CTS endpoints, provided that the installed CTS software supports such interoperability. For further information, refer to *Interoperability Between CTS Endpoints and Other Cisco Endpoints or Devices*, available at

http://www.cisco.com/en/US/docs/telepresence/interop/endpoint_interop.html

The primary motivation for this newly added support is to facilitate the interaction of a diverse set of video endpoints without the need for deploying an expensive DSP infrastructure that would otherwise be required.

The following sections present general considerations and recommendations for the use of video interoperability:

- [Video Interoperability Architecture, page 12-44](#)
- [Design Considerations for Video Interoperability, page 12-45](#)

Video Interoperability Architecture

The video interoperability architecture includes the following elements:

- Video interoperability support is available only in Cisco Unified CM 8.5 and later releases.
- Two different video endpoint family types (Cisco TelePresence endpoints, Cisco UC endpoints, or third-party endpoints) engaged in a video call.

The following sections offer further information about the scope of the video interoperability support:

- [Video Interoperability Test Cases, page 12-45](#)
- [Limitations of Video Interoperability, page 12-45](#)

Video Interoperability Test Cases

In most cases a video endpoint that supports SIP or H.323 without using proprietary signaling would be able to interoperate with a Cisco UC video endpoint that supports video interoperability. For specific information on the scope of the interoperability between common sets of deployed devices and general information about the testing that was conducted to validate these more common examples of interoperability, refer to the *Cisco Unified Communications System Test Results for IP Telephony*, available at

http://www.cisco.com/en/US/docs/voice_ip_comm/uc_system/unified/communications/system/uc_tart.htm

Limitations of Video Interoperability

While video interoperability support attempts to enable any-to-any point-to-point video call interoperability, it is important to note that not all features of an individual video endpoint can be supported when interoperating with another endpoint. There are many reasons for this. For example, incompatibilities between different call control protocols could render a feature unavailable or offer a different representation of that feature. H.264 video media parameters can be represented differently in H.323 than in SIP, as another example. H.323 also does not have support for presence, but presence is quite commonly supported in SIP. Skinny Client Control Protocol (SCCP) does not have any notion of application sharing, which is commonly available in SIP and H.323 endpoint implementations. For instance, an SCCP user trying to share his/her PC screen would be hampered because Binary Flow Control Protocol (BFCP) and H.239 are not available in SCCP.

Design Considerations for Video Interoperability

Because Unified CM video interoperability decreases the reliance on video transcoding to achieve any-to-any video calling, some call flows might change substantially. That is not to say that video transcoding is unnecessary for all call flows, but it is no longer needed for the ones where the video interoperability capabilities of Unified CM can be employed. Therefore, enabling video interoperability reduces the need for DSP resources.

Additionally, the following areas should be considered when implementing the video interoperability capabilities of Unified CM:

- [Guideline and Restrictions for Video Interoperability, page 12-45](#)
- [Quality of Service \(QoS\) and Call Admission Control Considerations for Video Interoperability, page 12-46](#)

Guideline and Restrictions for Video Interoperability

The following guidelines and restrictions apply with regard to video interoperability in a Unified CM deployment:

- If H.323 or SCCP protocols are used in conjunction with video interoperability, Unified CM will support only a single H.264 payload and the packetization mode is treated as 0. An example side effect (but not the only one) of this circumstance is the fact that 1080p resolution is not available with these protocols because 1080p requires packetization mode 1.
- If multiple payloads are presented by an H.323 or SCCP endpoint engaged in a video interoperability call, Unified CM will use only the payload with the lowest codec profile. This, in turn, could result in less than the highest supported resolution being selected for the call.

- If a SIP endpoint omits the **level-asymmetry-allowed** parameter in the Session Description Protocol (SDP), Cisco products will assume that the endpoint can support asymmetric resolution transmission. Therefore, different receiving and sending video resolutions could be negotiated during a call.
- If a call is processed with video interoperability while Unified CM is performing protocol interworking with SIP and H.323, the H.323 endpoint must honor the proposed dynamic payload number specified by the SIP side, which means that no re-negotiation to a different payload would be supported.
- Unified CM will not negotiate Real-Time Transport Control Protocol (RTCP) feedback if the video call invokes a media termination point (MTP) or transcoder.

Quality of Service (QoS) and Call Admission Control Considerations for Video Interoperability

There are no changes to the configuration of regions and locations in Unified CM as a result of video interoperability support. However, regions play a significant role in determining the resolution between groups of endpoints, and they can be used to maximize or minimize the resolution that these devices use when interoperating. The **Max Video Call Bit Rate** field in the regions settings is used to determine the amount of bandwidth and, thus, the resolution that endpoints are able to negotiate.

For further information about QoS and call admission control with native video interoperability, see the section on [Call Admission Control Design Recommendations for TelePresence Video Interoperability Architectures](#), page 11-95.

Wireless Networking Solutions

Because video is so bandwidth intensive, Cisco does *not* recommend using a shared wireless medium such as 802.11b/g for video endpoints.

Take care to ensure that video endpoints do not share the wireless bandwidth with any production IP Phones because video will consume much of the bandwidth and make it difficult to support video, audio, and data all on the same wireless medium.

Cisco Unified Video Advantage relies on a physical Ethernet connection to the physical IP Phone with which it is associated. It is not uncommon for a user to have both a physical Ethernet interface and an Aironet 802.11b Wireless Adapter installed in the same PC. This configuration could cause problems for Cisco Unified Video Advantage if the wireless interface happens to be the preferred path to the network because Cisco Unified Video Advantage will not associate over that interface. Cisco recommends that you always make the physical Ethernet interface the preferred path. Also, when users connect to the PC port on the back of the IP Phone, instruct them to disable their Aironet Adapter to keep it from accidentally taking precedence.

Cisco Unified Wireless IP Phones 7925G and 7921G

The Cisco Unified Wireless IP Phones 7925G and 7921G do not support video. Nothing prevents a video endpoint from calling a Cisco Unified Wireless IP Phone; however, it will negotiate as an audio-only call. The Wireless IP Phone user can then put the call on hold, transfer it, or conference it. If the caller is an H.323 video endpoint, it must support the Empty Capabilities Set (ECS) procedure in order for these supplementary services to work.

XML Services

Currently there are no XML applications specifically written for the Cisco Unified Video Advantage client solution, Cisco IP Video Phone 7985, or third-party SCCP video endpoints. However, with the exception of a few third-party endpoints, these endpoints do support XML applications. Cisco Unified Video Advantage uses a Cisco Unified IP Phone, so any XML application supported on those phone models should also work with Unified Video Advantage.

Most third-party SCCP video endpoints support XML, but not all XML applications currently work on those endpoints. For example, Cisco Extension Mobility and the Berbee InformaCast product are two popular XML applications that currently do not work on the third-party SCCP endpoints. Support for these applications will require firmware upgrades to the endpoints as well as changes in Unified CM Administration in some cases.

