



CHAPTER 7

Deployment Guidelines for Video Networks

Revised: March 30, 2012, OL-27011-01

When deploying a video network, it is imperative to design, plan, and implement the video network in a way that provides the best user experience possible. Video is an application subject to the perceptions of the various parties during a call. If one of the video users in a meeting or collaboration effort does not have a satisfying experience, that perception can easily be transferred onto the rest of the attendees.

The following sections offer general guidelines for designing video networks:

- [Planning a Deployment Topology for Video, page 7-1](#)
- [Allocation of Video Resources, page 7-10](#)
- [Creating a Video-Ready Network, page 7-14](#)
- [Integration with Standalone Video Networks, page 7-19](#)

Planning a Deployment Topology for Video

When deploying video applications, it is fundamental to identify and plan the topology or topologies that are deployed or that should be deployed to cover the needs of the organization. Current video applications use the following main video deployment topology models:

- [Intra-Campus, page 7-2](#)
- [Intra-Enterprise, page 7-3](#)
- [Inter-Enterprise \(Business-to-Business\), page 7-4](#)

Video applications also use the following main call processing models:

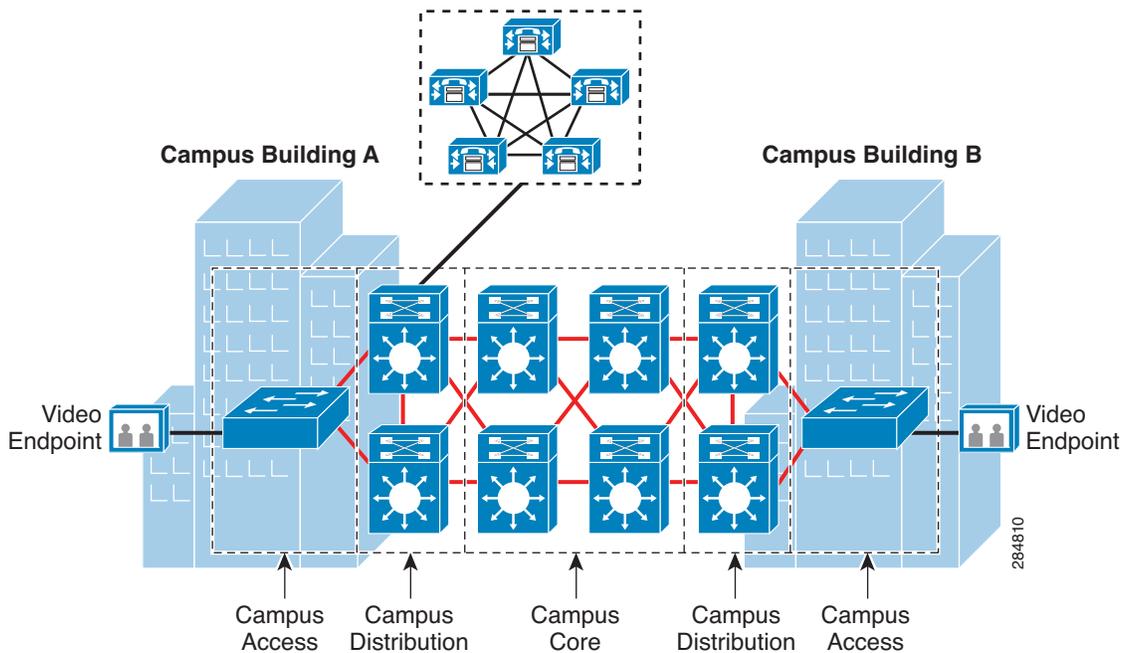
- [Single-Site Call Processing, page 7-5](#)
- [Multi-Site Call Processing, page 7-5](#)
- [Hosted Call Processing \(Video as a Service\), page 7-7](#)

The following sections provide an overview of these deployment models as well as guidelines for choosing a call processing model and topology.

Intra-Campus

An intra-campus topology for video is limited to offering video throughout a single company site or campus. This topology model is better suited for companies that require increased meeting effectiveness and productivity without requiring the users to move throughout the facilities. The intra-campus video deployment model can be used in conjunction with the intra-enterprise and inter-enterprise topology models to meet the needs of the organization. Figure 7-1 depicts an intra-campus topology using single-site call processing.

Figure 7-1 *Intra-Campus Deployment with Two Buildings (Everything in One Building Except Video Endpoints)*



With respect to call processing, the single-site and hosted call processing models are a better fit for intra-campus video topology models. Deciding between the two depends greatly on the endpoint density, planned growth, features required, and cost.

For instance, hosted video call processing deployment models offer a very feature-rich experience with a low cost investment. However, certain local multipoint call flows might require the media streams to travel outside of the premises when embedded video resources are not available or when their capacity is exceeded, and therefore the bandwidth-to-cost relationship becomes a crucial factor as user density increases.

On the other hand, a single-site call processing model has a higher initial investment because features require either hardware or software licensing for them to be available. Nevertheless, the single-site call processing model allows for user growth with a lower cost ratio after it has been deployed.

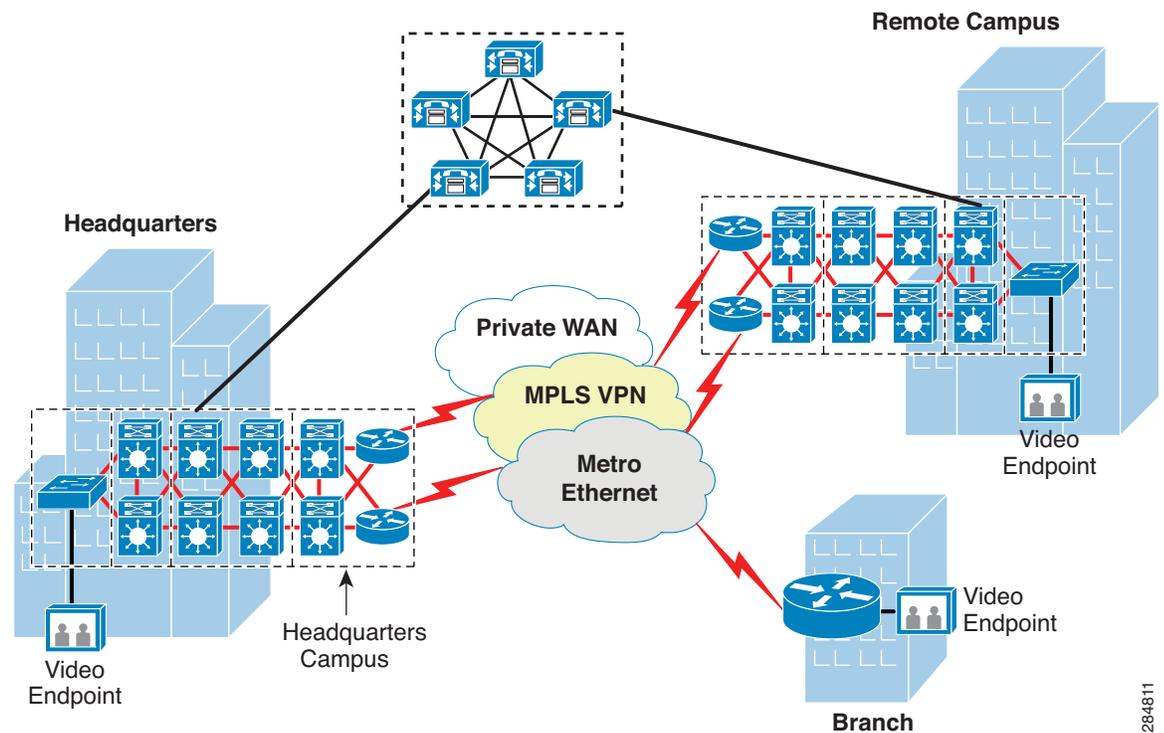
Intra-Enterprise

An intra-enterprise topology for video enables more than one site within the same company to use video applications through the WAN. The intra-enterprise deployment model is suitable for businesses that often require employees to travel extensively for internal meetings. Deploying video within the enterprise not only improves productivity by saving travel time and providing feature-rich collaboration, but it also reduces travel expenses. Furthermore, the overall quality of work and life is often improved when employees have to travel less.

To enable video across multiple sites, the intra-enterprise topology must make use of high-speed supporting WAN links to provide a rich video experience. For details about dimensioning the WAN links, see [Scalability and Performance](#), page 7-16.

Figure 7-2 illustrates an intra-enterprise topology that uses a multi-site call processing deployment model.

Figure 7-2 Multiple Sites with Distributed Multi-Site Cisco Unified Communications Managers



Intra-enterprise deployments can use either multi-site or hosted call processing deployment models. As with intra-campus deployments, the same considerations apply when deciding between the two call processing models.

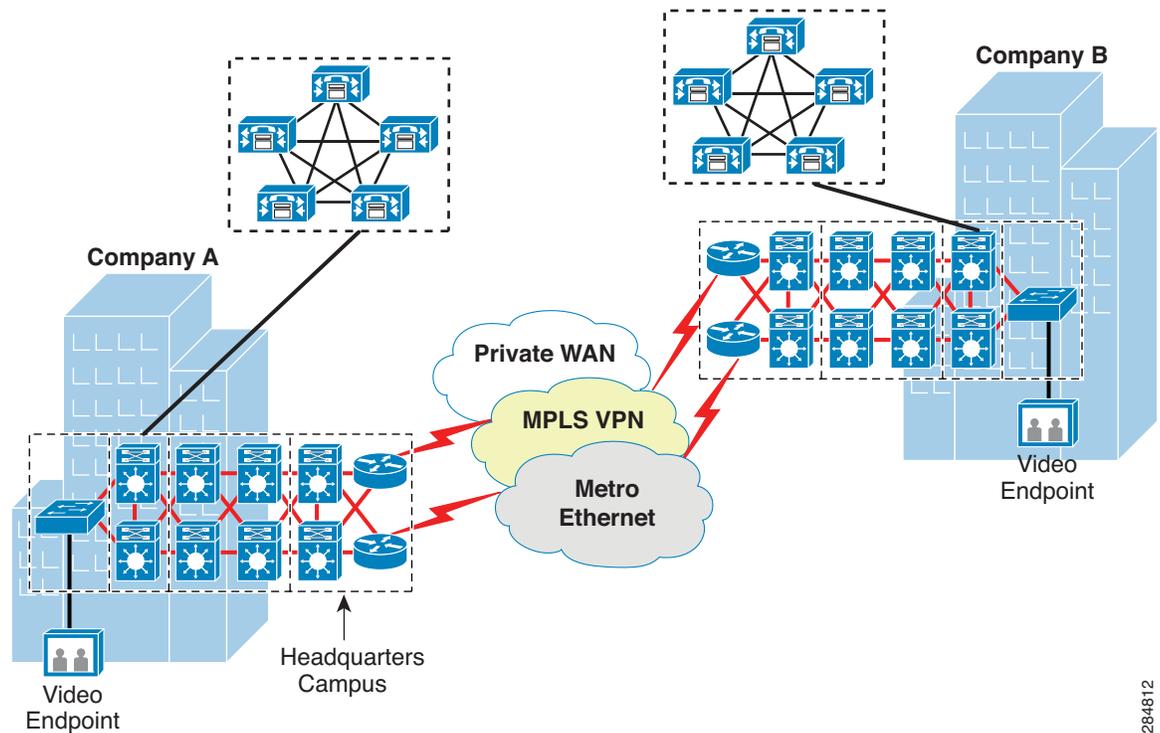
Inter-Enterprise (Business-to-Business)

The inter-enterprise network deployment model not only connects video endpoints within an enterprise, but also allows for video endpoints within one enterprise to call systems within another enterprise. The inter-enterprise model expands on the intra-campus and intra-enterprise models to include connectivity between different enterprises. This is also referred to as the business-to-business (B2B) video deployment model.

The inter-enterprise model offers the most flexibility and is suitable for businesses that often require employees to travel extensively for both internal and external meetings. In addition to the business advantages of the intra-enterprise model, the B2B topology deployment model lets employees maintain high-quality customer relations without the associated costs of travel time and expense.

Figure 7-3 depicts two companies communicating with each other through an inter-enterprise topology and single-site call processing.

Figure 7-3 *Inter-Enterprise (B2B) Deployment with Single-Site Call Processing at Each Company*

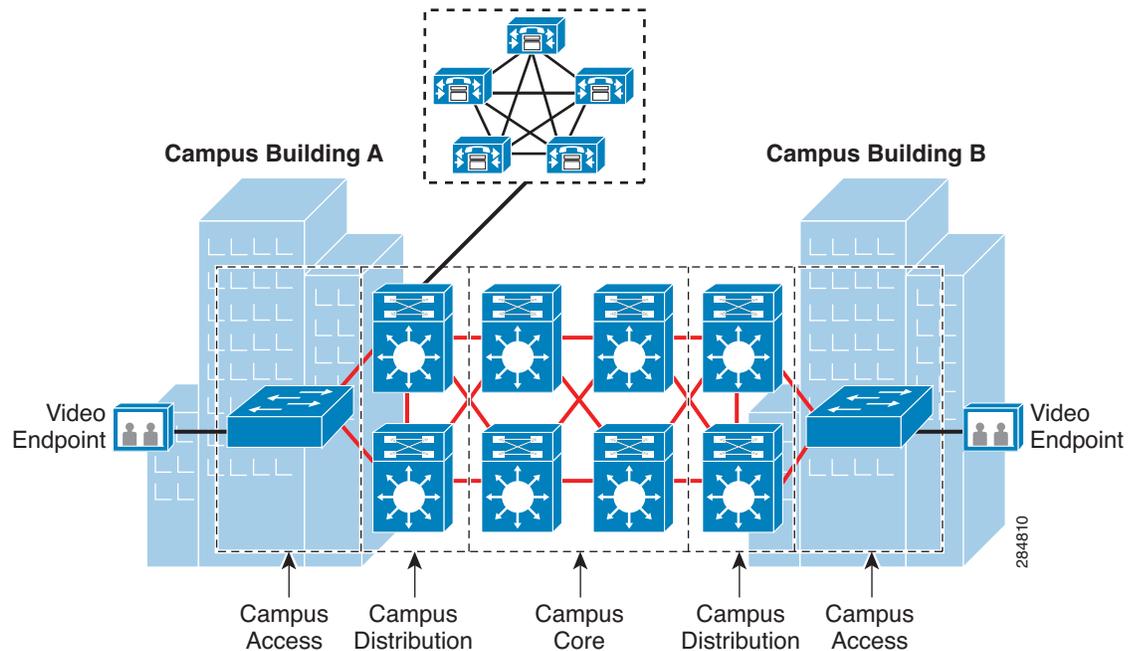


All three call processing models can work with the inter-enterprise deployment model, and the same considerations apply as with intra-campus and intra-enterprise deployments.

Single-Site Call Processing

A single-site call processing model confines call processing to service a single site, and the call processing agents are in the same location as the serviced endpoints. Whatever the distance is between the call processing agent and the endpoints, it should be serviced by LAN speed links. The single-site model is suitable for medium-sized businesses and government operations that reside at one site and that have basic video call processing needs, but where growth might be explosive or where the user density is very high, thus making the bandwidth-to-cost ratio of a hosted solution prohibitively expensive. [Figure 7-4](#) depicts an intra-campus topology using single-site call processing.

Figure 7-4 *Intra-Campus Deployment with Two Buildings (Everything in One Building Except Video Endpoints)*

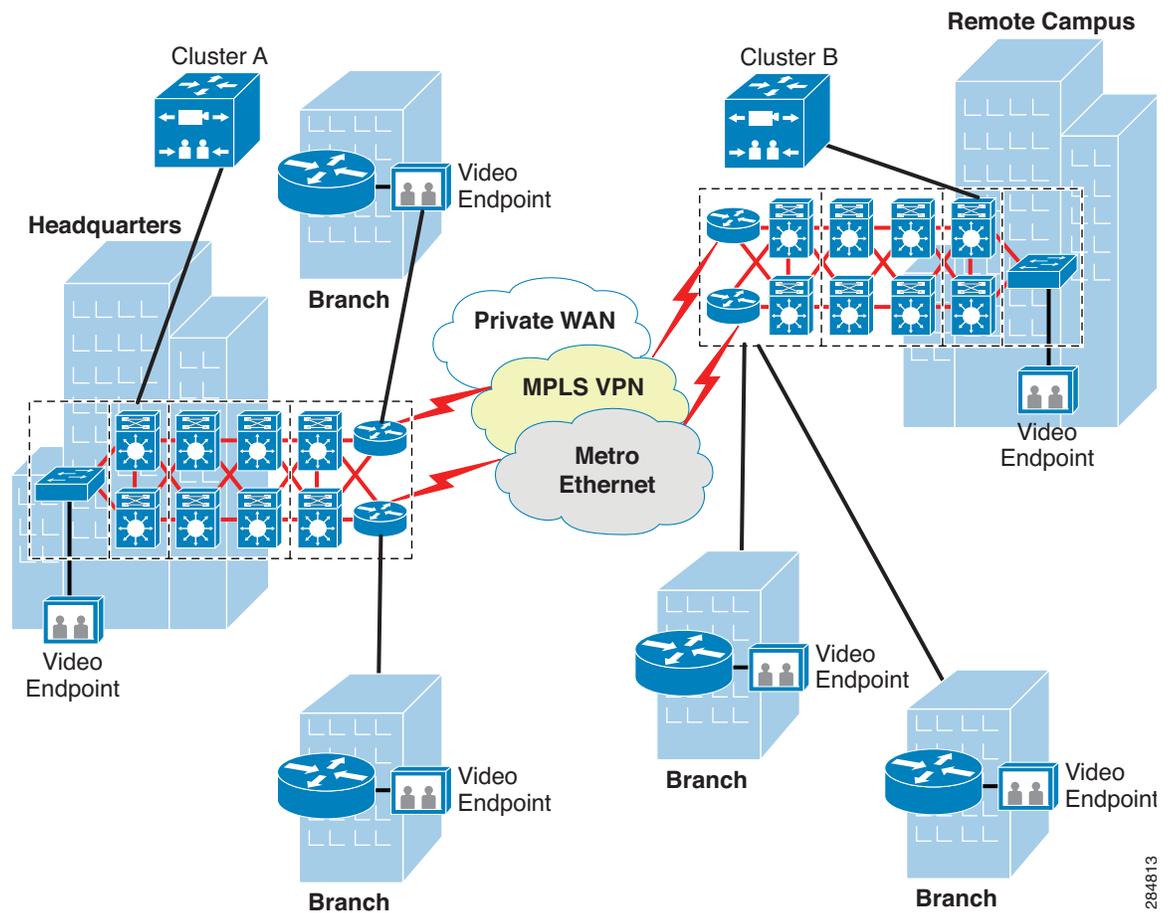


Multi-Site Call Processing

In a multi-site call processing model, the call processing agents can all be at the same location (centralized multi-site call processing) or distributed across various locations where high video user density exists or where service is critical and a backup is required.

Within the same call processing agent cluster, multi-site call processing can service a variety of topologies (for example, hub-and-spoke and multiple-hubs-to-spokes topologies) using either a centralized or distributed multi-site call processing model. [Figure 7-5](#) illustrates a distributed call processing model where a multi-site call processing deployment services a large central site and multiple remote or branch sites, with the home office sites or smaller branches depending upon the larger ones (multiple hubs to spokes) for call processing services.

Figure 7-5 Distributed Multi-Site Call Processing

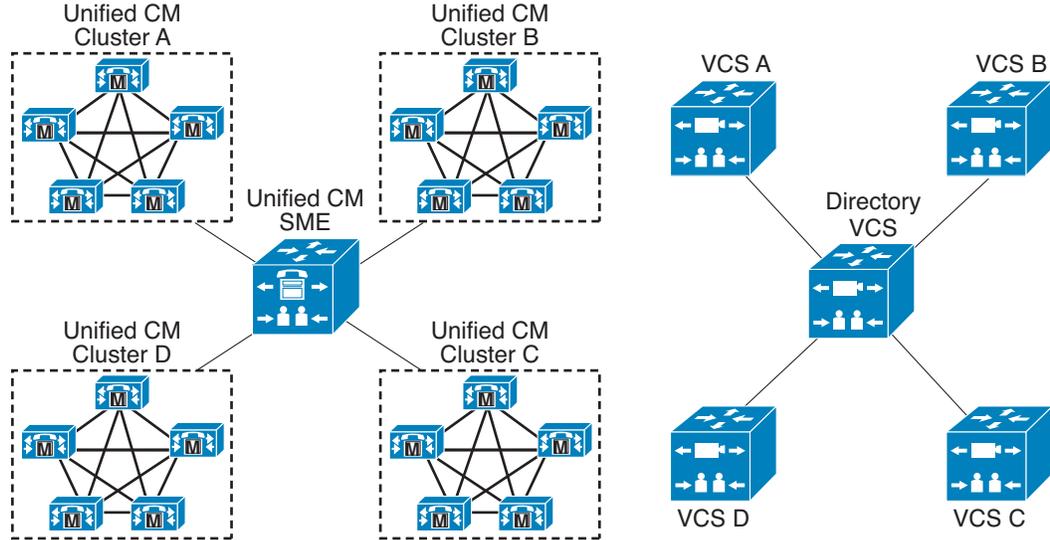


284813

The multi-site call processing model also includes deployments where clustered call processing agents interact with other clustered call processing agents through a call processing element deployed for the sole purpose of aggregating the call routing between the clusters. Deploying an aggregating call processing entity is advantageous because it eliminates the need to implement full-mesh connectivity between all the call processing clusters. Instead, the various leaf clusters engage with the aggregating call processing element when the leaf clusters communicate with each other. If the dial plan is implemented correctly to provide hierarchy and allow for increasing the capacity of the overall solution, no dial plan updates are needed in the leaf clusters when new leaf clusters are added.

Figure 7-6 illustrates two examples of multi-site aggregated call processing deployments, one using Cisco Unified Communications Manager Session Management Edition (SME) the other one using a Cisco TelePresence Video Communication Server (VCS) as a directory gatekeeper.

Figure 7-6 Two Examples of Multi-Site Aggregated Call Processing Deployments



Note

Although H.323 gatekeepers do allow the concept of a cluster (alternate gatekeeper), a single gatekeeper is also considered as a clustered call processing agent for purposes of the above discussion because H.323 gatekeepers are self-deterministic call routing entities.

Hosted Call Processing (Video as a Service)

Hosted deployment models refer to services provided and managed from the cloud. The offering is compelling because of the lower cost of ownership (low investment) and feature-rich experience it provides. Hosted video solutions are aimed at mid-sized and smaller businesses, providing them with an affordable entry point into the world of enterprise-grade video, and some hosted solutions also provide a migration path that protects the original investment as the business grows.

In this deployment model, more than just the call processing elements are often off-premises, thus requiring that the video streams travel off premises in some multipoint call flow scenarios. In these cases, higher user density in a service location will require higher bandwidth to service them.

Guidelines for Choose a Call Processing Topology and Video Endpoints

Choosing the right elements and deployment models to implement or expand a video network is instrumental for ensuring that the desired features, performance, and scalability are achieved. Moreover, choosing the wrong elements and models for the video network can result in costly changes to provide the functionality that the organization requires.

The following sections provide general guidelines for selecting the elements and deployment models for a video network:

- [Call Processing Model and Call Processing Agent Selection Guidelines, page 7-8](#)
- [Endpoint Selection Guidelines, page 7-9](#)
- [Design Considerations for Video Networks, page 7-9](#)

For information about call signaling protocol selection, see the chapter on [Call Control Protocols and IPv6 in IP Video Solutions, page 4-1](#).

Call Processing Model and Call Processing Agent Selection Guidelines

To choose the correct call processing agent and its deployment model and topology, it is important to consider the following points in addition to the requirements of the organization during the design phase of the deployment:

- What features are needed to fulfill the use cases for the success criteria of the deployment? For example, SIP, SRTP, BFCP, or IPv6 video support.
- Are any video endpoints already deployed in the network that will have to be serviced? Are there requirements for previous video endpoints that have to be serviced by the newly selected call agent, interoperability with a previous video network, and so forth?
- If there are any previous video elements to be included in the deployment, what protocols do they use or can they use? For example, multi-protocol endpoints (SIP and H.323) or single-protocol endpoints.
- Has the call control protocol been selected? If not, are there any features dependent on a particular call control protocol? For example, H.239 can be used only in conjunction with H.323.
- What are the locations that will need video service? What is their user density? How critical will redundancy be for each particular location? It is advisable to provide a call agent for redundancy purposes at each site that has a high user density (distributed call processing). A very high user density will also create challenges (internet bandwidth usage) under certain call flows for hosted call processing.
- Will protocol interworking be used? Interworking will significantly influence the placement of the call agents under certain circumstances because the media stream might have to traverse the call processing agent in order to reach its destination.
- What is the maximum number of video calls you expect to service? A single Cisco TelePresence Video Communication Server (VCS) 7.0 has a limit of 500 non-traversal calls. Beyond that, a second VCS is needed (in a cluster or standalone) to service the rest of the calls.

The outcome of these considerations, coupled with the customer requirements, product data sheets, and product release notes, will determine which call processing agent to select and which call processing model and topology to use. For example, if the requirements include the use of BFCP as an application sharing technology and service for a maximum of 2000 calls, then a VCS cluster would be the best choice.

Endpoint Selection Guidelines

Selecting the right endpoints for the job is just as important as selecting the call processing agent. In addition to the customer requirements, the following points help determine the selection:

- What call control protocol will be used, H.323 or SIP?
- Will embedded video resources be needed for video conferencing? If so, Cisco TelePresence System EX90 would be a suitable choice.
- What video resolution formats will be required? For example, HD 720p.
- What other endpoints will be engaged in a call with the given endpoint(s) being selected? For example, Cisco Unified IP Phone 9971.
- What application sharing technologies will be needed? For example, BFCP over UDP.
- What are the mobility requirements? Will this endpoint be a mobile endpoint (collaboration tablet)?

The outcome of these considerations, coupled with the customer requirements, product data sheets, and product release notes, will determine which endpoints to select.

Design Considerations for Video Networks

When designing a video network, it is important to consider the implications of interworking. Depending on the call processing agent selected, the media streams might have to traverse the call agent when interworking is used. Therefore, interworking might have a negative impact on bandwidth usage if the call processing agent is remote to the endpoints engaged in the call because the call would not flow point to point.

Additionally, although DNS SRV records are typically for scalability purposes (the number of SIP trunks required to integrate a system is reduced when using SRV records), call processing agents behave differently with regard to endpoint registration and call processing peering when it comes to the use of DNS SRV records. These differences can create unexpected conditions when integrating different call agents if the behavioral differences are not understood prior to the integration.

Allocation of Video Resources

Whether servicing one or many physical locations, there are a number of considerations that need to be weighed to determine the best network locations for the video resources. Video resources can be either dedicated or embedded. Embedded video resources lie inside of the endpoints and service calls only for the endpoints that contain those resources. Dedicated video resources, on the other hand, reside in appliances separate from the endpoints, and they service any endpoints that have access to those resources.

Correct distribution of the video resources is necessary to achieve the desired level of user experience and, more often than not, the right level of redundancy and availability. The more factors you take into consideration, the more reliable your determination will be for the locations of the video resources. The following factors should be integrated into the decision making process to determine the best allocation model and locations for the video resources:

- Branch available bandwidth
- Bandwidth cost
- Remote video resource cost
- Usage patterns at headquarters and remote sites
- Call agent bridge selection algorithm
- Type of video resources

The following basic models can be used to allocate dedicated video resources in a deployment:

- [Centralized Video Resource Allocation, page 7-10](#)
- [Distributed Video Resource Allocation, page 7-13](#)

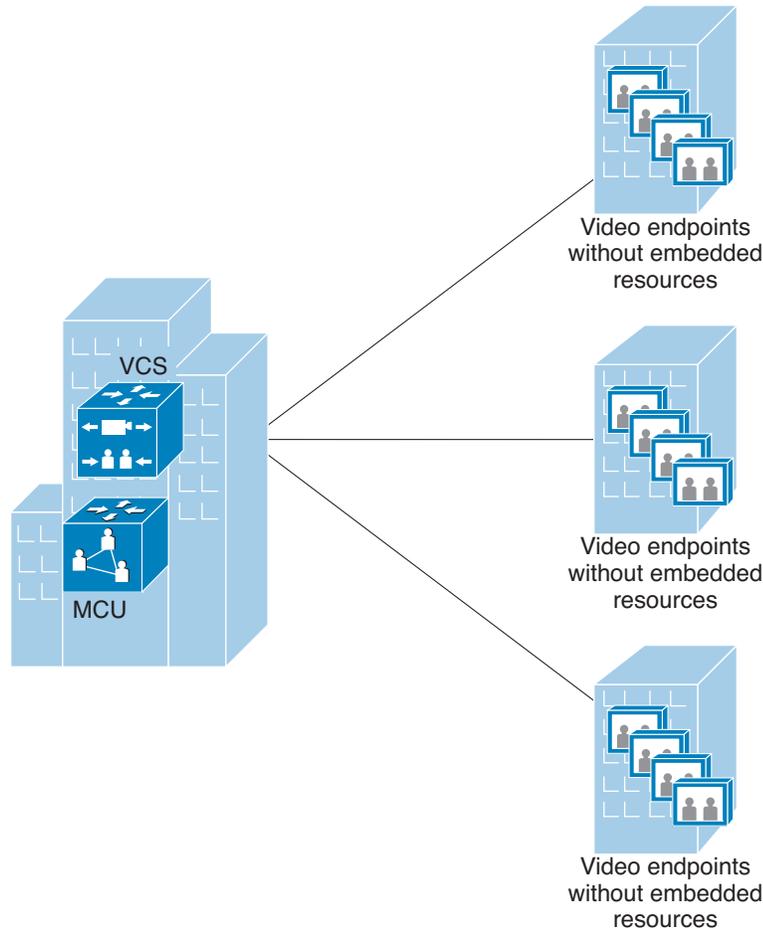
You can also combine embedded video resources with these models for dedicated video resources to form a hybrid model, if needed, to fit the necessities of your video solution more precisely.

Centralized Video Resource Allocation

Centralized resource allocation should be considered when the combined costs of placing the resources in the same location are less than distributing them. The feasibility of a centralized resource allocation should also be considered. For example, not all scenarios will be suitable for a centralized resource architecture if the concentration of the resources induces an undesirable condition on the endpoints (for example, unacceptable jitter).

As previously stated, the factors listed in the section on [Guidelines for Choose a Call Processing Topology and Video Endpoints, page 7-8](#), should be integrated into the decision making process to select the best approach to fit a given scenario. For instance, consider the scenario depicted in [Figure 7-7](#). Although at first it might seem less costly to concentrate all the video resources in the headquarters, such resource concentration would create the side effect of increasing the bandwidth requirements for the WAN links between the hub and spokes. Furthermore, multipoint conference capacity would be limited in the remote sites by the bandwidth provided in the WAN links.

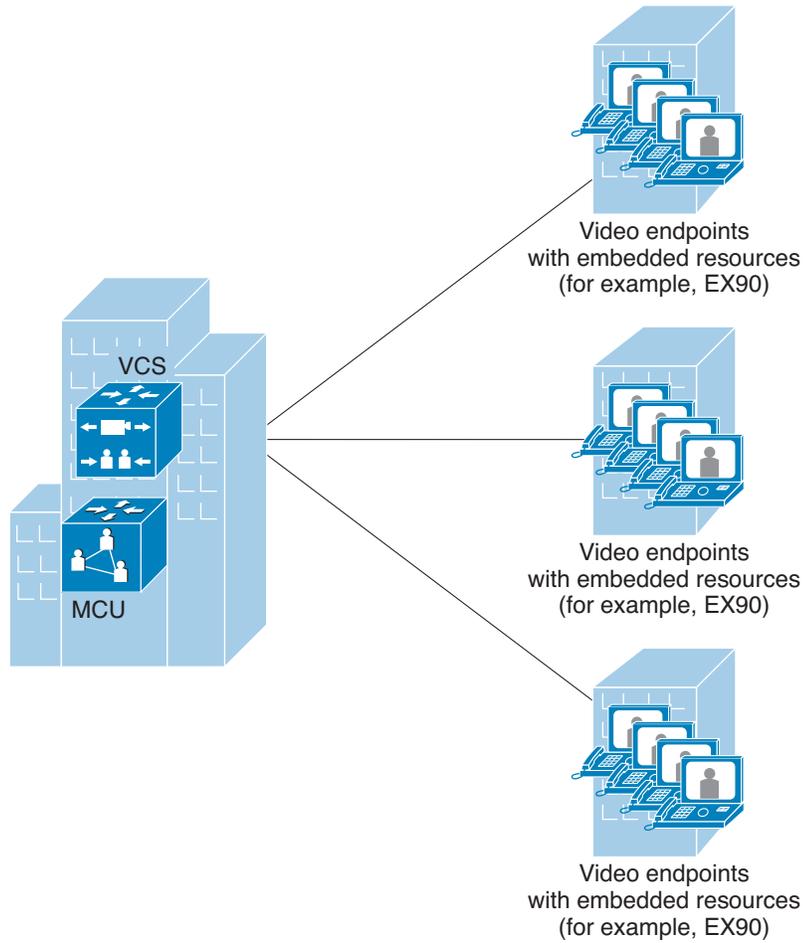
Figure 7-7 *Hub and Spoke Network with Centralized Dedicated Video Resources (No Embedded Resources)*



In general, centralized video resource allocation architectures are better suited for scenarios where not many remote endpoints exist and their remote location and bandwidth usage does not induce undesirable effects on the media streams to be transported in the WAN links.

A modification to the previous scenario is presented in the hybrid example in [Figure 7-8](#), in which embedded video resources are located at the remote sites. In this scenario, the centralized dedicated video resources would be utilized only when a video conference involves a video endpoint without embedded resources in the central location or when the number of participants in the video conference exceeds the capacity that the embedded video resources can handle.

Figure 7-8 *Hub and Spoke Network with Centralized Dedicated Video Resources and Embedded Video Resources at the Remote Sites*



Many other scenarios are possible, and therefore different strategies and/or restrictions can also apply to the centralized video resource approach.

Distributed Video Resource Allocation

Distributing the dedicated video resources throughout various locations has several advantages; chief among them are the WAN link bandwidth savings and less likelihood of inducing undesirable effects on the media streams (since many of them are locally terminated). However, a distributed allocation also has some limitations. For example, certain video calls still have to traverse the WAN links, and these streams are still limited by the WAN link properties.

Cost effectiveness of the solution can be maximized by considering the following points when determining whether or not to deploy distributed dedicated video resources:

- What is the expected video call utilization pattern at the location(s) of the video resources?
- Can the current bandwidth of the WAN link(s) support the expected usage pattern(s) of the remote location(s) without inducing undesirable effects in the video streams?
- Will the limitations or effects (if any) of the media transmissions over the WAN links be acceptable for the intended use cases?
- Would using distributed embedded video resources satisfy the planned use cases?
- Will the video solution be able to grow adequately without distributed dedicated video resources, given the current and planned network topologies?

Additionally, when using a distributed allocation for dedicated video resources, it is important to understand the bridge selection algorithm of the call control elements in order to decide where best to locate the video resources. For instance, the dedicated video resources might be reserved based on the time zones of the endpoints and the resources. Other alternatives involve video resource reservation based on video location or manual reservation. In any case, the importance of understanding the selection algorithm derives from the need to understand where the streams are ultimately terminated in order to distribute the video resources more efficiently.

Creating a Video-Ready Network

Video brings significant benefits for businesses, such as superior collaboration, lower travel costs, and personalized advertisements. However, video applications also introduce additional challenges for the underlying network infrastructure and IT departments. For instance, how does one configure the network for video? How should IT departments prioritize and scale video? How do they protect other applications from being swamped by high-bandwidth video streams? To support these enterprise video applications, a tightly controlled network foundation providing the following services is required:

- [Optimized Video Delivery, page 7-14](#)
- [Security of Video Applications, page 7-16](#)
- [Scalability and Performance, page 7-16](#)

Optimized Video Delivery

For video to be an efficient collaboration tool, the user experience must be of high quality. To ensure the user experience quality, the video delivery must be optimized to meet the organization's requirements. The following sections offer guidelines on how to optimize the video delivery:

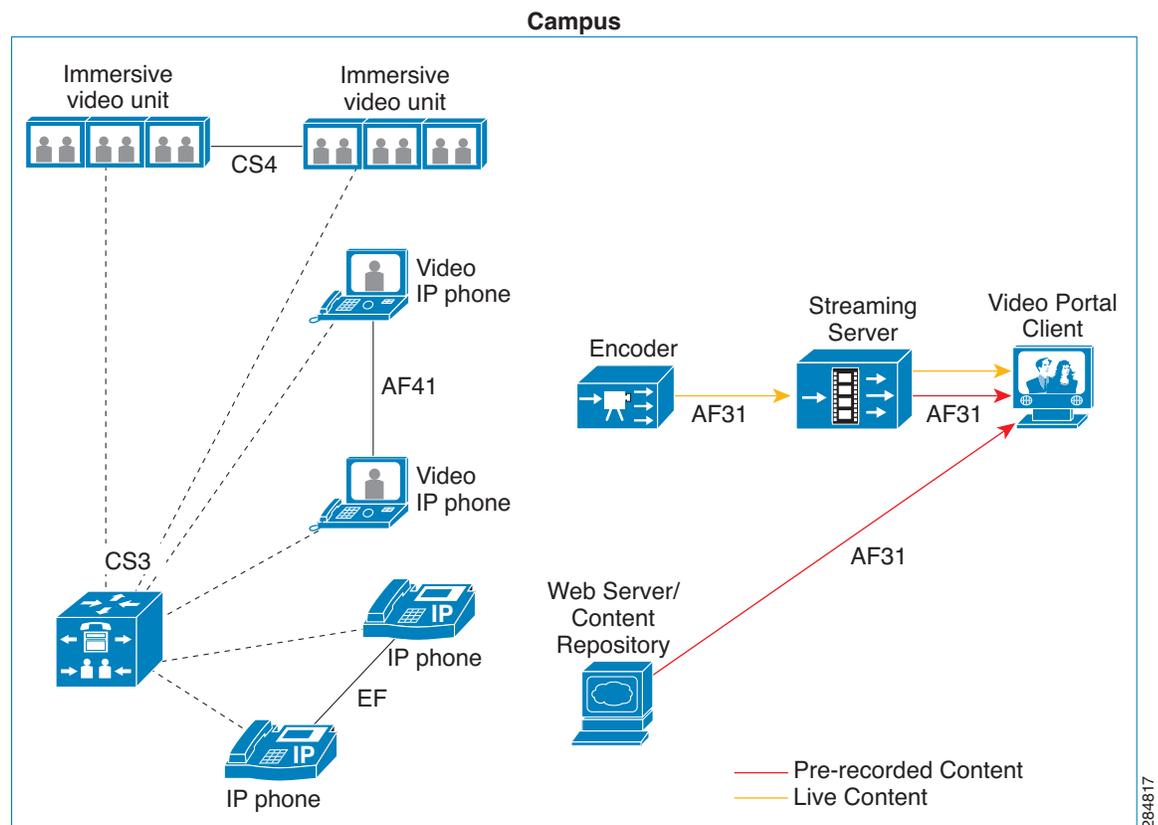
- [Quality of Service \(QoS\), page 7-14](#)
- [Content Sharing Technologies, page 7-15](#)
- [Reliability, page 7-16](#)

Quality of Service (QoS)

The first step in optimizing the delivery of video is to identify the traffic of interest and to apply a differentiated Quality of Service (QoS). QoS helps the organization to provide application intelligence to differentiate between business-critical and noncritical video streams, and to keep the latency, jitter, and loss for selected traffic types within an acceptable range. Furthermore, priority queuing should be used over other queuing mechanisms whenever possible to provide a better video experience. In the case of multiple video applications on converged networks (TelePresence and IP video telephony combined in the same network), Cisco recommends differentiation of QoS classes per application.

[Figure 7-9](#) depicts an example of multiple IP video applications and IP voice converging in the same network. In this example, immersive video, videoconferencing, video-on-demand, and voice over IP are identified and assigned the recommended QoS markings to provide the required service level, thus avoiding over-provisioning or overlap of applications in the queues.

Figure 7-9 Recommended QoS Traffic Markings in a Converged Network



For more information about QoS in video solutions, refer to the chapter on [Quality of Service and Call Admission Control](#), page 5-1.

Content Sharing Technologies

Depending on the endpoints being deployed as part of the video solution, it is important to consider the content sharing standards supported and required by the video endpoints and how they will converge or interoperate if necessary. There are currently three main content sharing technologies used by IP video solutions: Binary Floor Control Protocol (BFCP), H.239, and auto-collaborate. H.323 endpoints make use of the H.239 standard to provide the content sharing functionality, while SIP endpoints may use auto-collaborate or the newer standard BFCP.

For information about presentation sharing and content sharing technologies, refer to the *Cisco TelePresence Interoperability Deployment Guide*, available at

http://www.cisco.com/en/US/products/ps8332/products_device_support_tables_list.html

Reliability

The performance and availability of a video-ready campus must be monitored pro-actively and measured across the network. Alternate paths in the case of failure should also be offered to ensure the reliability of the video solution. For information about how to design highly available networks, refer to the *Campus Network for High Availability Design Guide*, available at

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns815/landing_cHi_availability.html

Security of Video Applications

Whenever possible, a video-ready network should integrate video security to protect against unauthorized access to video applications. Mitigation of attacks and protection of traffic from snooping and intrusion by malicious users is essential, and so is preventing malicious users from transmitting unauthorized video. A variety of techniques can be used to secure a video network, from network virtualization techniques to segregate video traffic, to the use of Trusted Relay Points (TRP) for software clients residing in the data VLAN or Session Border Controllers (SBC) for topology hiding from the exterior world. For information about how to secure video networks, see the chapter on [Security for Video Communications](#), page 9-1.

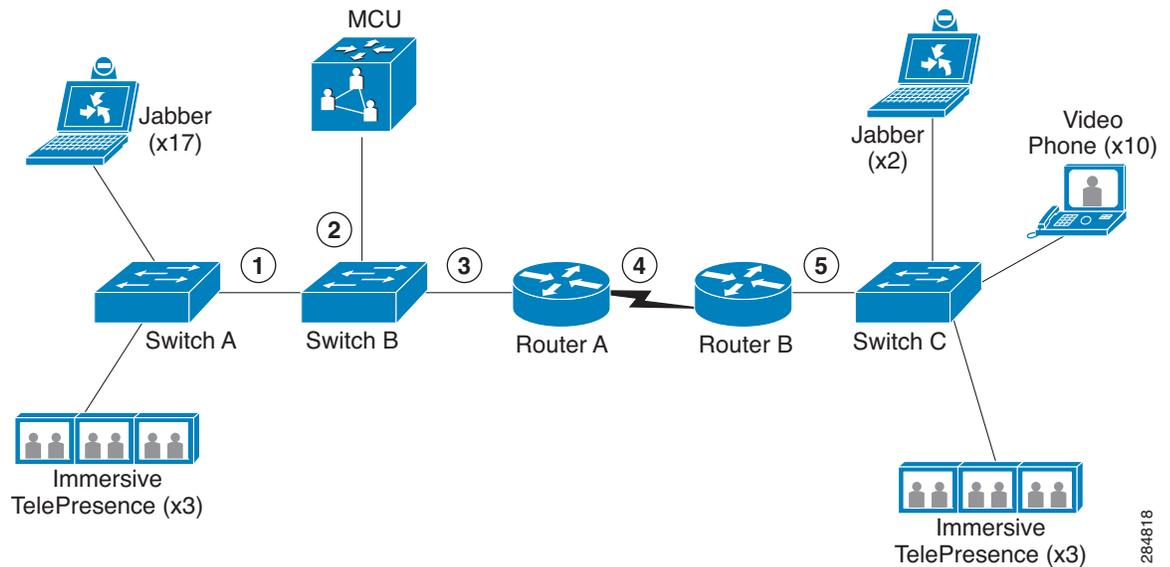
Scalability and Performance

Network scalability is critical to supporting increasing bandwidth demands as more video applications or video users are deployed. To maintain optimal performance, the network should easily accommodate higher bandwidths, scaling to support high-definition (HD) video streams and in some cases even multiple HD video streams simultaneously. Therefore, it is crucial to size the network adequately for the expected traffic generated by the video applications to be deployed.

The first step in sizing the network is to understand the endpoint and user requirements. Next, determine how the media flows will behave during conferences and point-to-point calls. Then add the considerations for voice and data traffic and backup measures required for reliability. [Figure 7-10](#) illustrates an example scenario where 20 immersive and desktop video endpoints are located in the headquarters campus while 15 miscellaneous endpoints are located in the branch. The expected usage pattern in this example is as follows:

- Headquarters desktop endpoint users and immersive endpoint users will generate peak usage of 7 calls among each other.
- Headquarters endpoints will generate and/or receive peak usage of 6 calls to/from the branch office.
- Calls on headquarters desktop endpoints use 1.3 Mbps while the immersive endpoints use 12 Mbps.
- Branch video IP phones use 1 Mbps, desktop endpoints use 1.3 Mbps, and immersive endpoints use 12 Mbps.
- At peak times, branch users will generate a maximum 4 calls among each other and generate or receive 6 calls (as listed above) to/from the headquarters.
- Headquarters and branch users need to access 10 Mbps (combined) of data applications between sites.

Figure 7-10 Determining Capacity Requirements for a Video-Ready Network



284818

The above requirements are obviously simplistic. A very complex network and deployment would have a longer set of requirements and applications to support. But with the list of requirements in the example above, we can determine the following:

- The maximum expected bandwidth usage is 49 Mbps for the video streams in the uplink between switch A and switch B (link 1 in [Figure 7-10](#)), assuming the worst-case scenario of:
 - Seven participants in a local multipoint call of the desktop video endpoints:

$$7 * (1.3 \text{ Mbps}) = 9.1 \text{ Mbps}$$
 - Six participants in a multipoint call of the immersive endpoints, 3 local and 3 remote:

$$3 * (12 \text{ Mbps}) = 36 \text{ Mbps}$$
 - Six participants in a multipoint call of the local desktop endpoints, 3 local and 3 remote video IP phones:

$$3 * (1.3 \text{ Mbps}) = 3.9 \text{ Mbps}$$



Note

This bandwidth calculation does not consider the extra bandwidth for data applications or any other extra bandwidth that is necessary (for call signaling, for example). Before dimensioning the LAN, you need to include these other bandwidth requirements.

- In a worst-case scenario, the branch WAN link (link 5 in [Figure 7-10](#), but the same considerations apply for links 3, 4, and 6) would use 43.6 Mbps of bandwidth for video streams, assuming:
 - Four participants in a multipoint call between 4 video phones. Because no local Multipoint Control Unit (MCU) is available, all the streams have to travel to the headquarters for the multipoint call to occur.

$$4 * (1 \text{ Mbps}) = 4 \text{ Mbps}$$
 - Six participants in a multipoint call between 5 desktop video endpoints and 1 video phone, 3 local endpoints (2 desktop video endpoints and 1 video phone) and 3 remote endpoints:

$$2 * (1.3 \text{ Mbps}) + 1 \text{ Mbps} = 3.6 \text{ Mbps}$$
 - Six participants in a multipoint call on the immersive endpoints, 3 local and 3 remote:

$$3 * (12 \text{ Mbps}) = 36 \text{ Mbps}$$



Note This bandwidth calculation does not consider the extra 10 Mbps requested for data applications or any other extra bandwidth that is necessary (for call signaling, for example). These requirements also need to be added as part of the sizing process.

- Finally, for link 2 servicing the MCU in [Figure 7-10](#), we can anticipate that 92.6 Mbps of video streams will traverse it at its peak if we consider the two previous bandwidth calculations:

$$49 \text{ Mbps} + 43.6 \text{ Mbps} = 92.6 \text{ Mbps}$$

**Note**

The general rule that has been thoroughly tested and widely used is to over-provision video bandwidth by 20% in order to accommodate a 10% burst and the Layer 2-to- Layer 4 network overhead. Furthermore, the above calculations are based on the worst-case scenario for the usage patterns provided in the example, but they do not consider the case where all video users want to make video calls at the same time (known as 100% call completion).

In summary, the network capacity and performance design in a video-ready network should allow for video forwarding without introducing significant latency of the call completion rate desired or the usage patterns expected. Refer to your endpoint documentation to determine the amount of bandwidth required per call.

Integration with Standalone Video Networks

Whether replacing a previous video network solution or trying to converge the video network solution under the same call processing platform elements, integration could pose some challenges for the IT department. Understanding the options and guidelines for integration will provide a better experience for the integrator and the user.

The integration approach differs depending upon the call signaling protocol used. The following sections outline the general guidelines for the two most widely used protocols in video networks today:

- [Integration with Standalone H.323 Video Networks, page 7-19](#)
- [Integration with Standalone SIP Video Networks, page 7-20](#)

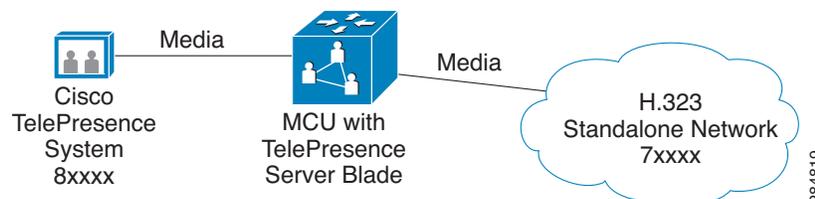
Integration with Standalone H.323 Video Networks

H.323 is a very well define protocol, which makes interoperability with H.323 call processing elements considerably easier than with multi-vendor SIP elements; however, H.323 is not as service-rich as its counterpart SIP. For example, Cisco has implemented the ability to switch screens depending on who the active speaker is (smart switching), but this feature is not natively available in H.323 networks.

Whenever possible, use the native interoperability of the video endpoints to connect them directly to the H.323 network, provided that this does not cause the loss of required features such as smart switching. Otherwise, if feature retention is critical or if interoperability cannot be obtained point-to-point natively, then connect the H.323 endpoints through a video transcoder or an interoperability-enabled conference bridge. A non-overlapping dial plan is also recommended, and different access codes can be used between the networks to indicate to the call processing agents that a hop to the next video system is required to complete the call.

[Figure 7-11](#) illustrates the use of a Multipoint Control Unit (MCU) to connect a Cisco TelePresence System to a third-party H.323 standalone video network.

Figure 7-11 Integrating Standalone H.323 Networks



Integration with Standalone SIP Video Networks

SIP video networks are more feature-rich than H.323 networks and can enable very useful features when all the endpoints support the. However, SIP is not as well defined as H.323, thus making interoperability with it more challenging.

If an endpoint conforms tightly to the SIP standard, then call agents can make use of the native video interoperability features available within the call processing agents. Otherwise, video networks can be bridged to each other with video transcoders or interoperability-ready multipoint control units.

Cisco recommends that you do not use an overlapping dial plan between the video networks, but do use an access code to indicate to the call processing agents to route between video networks and avoid inter-digit time-out. If uniform resource identifier (URI) dialing is used for the dial plan, Cisco recommends using different domains for ease of administration.

Figure 7-12 illustrates integration of a Cisco TelePresence System with a SIP standards-based third-party system. This example uses native interoperability along with different domains for dialing.

Figure 7-12 Integrating Standalone SIP Networks

