CISCO SYSTEMS

Cisco IP Videoconferencing Solution
Reference Network Design Guide

July 2002

CONTENTS

Preface

This preface describes the purpose, scope, intended audience, and general organization of this *Cisco IP Videoconferencing Solution Reference Network Design Guide*. It also provides information on how to order documentation from Cisco Systems.

Purpose

This document provides guidelines, recommendations, and best practices to help you design an IP videoconferencing solution for your enterprise using the Cisco Architecture for Voice, Video, and Integrated Data (AVVID).

Scope

This document describes the products and features used to build a Cisco IP Videoconferencing (IP/VC) system, and it gives recommendations on how to combine those elements into an effective solution for your enterprise. However, this document does not contain specific implementation or configuration details for the products and features. For details about a particular product or feature, refer to the technical documentation available online at Cisco.com. (See Obtaining Documentation, page ix.)

Note Unless stated otherwise, the solution designs presented in this document require the minimum software releases listed in Table 1, and the information presented here applies only to those releases.

Table 1   Cisco IP/VC Hardware Platforms and Minimum Software Releases

Platform	Minimum Required Software Release
IPVC 3510 Multipoint Conference Unit (MCU)	2.2.1
IPVC 3520 Gateway	2.2.3
IPVC 3525 Gateway	2.2.3
IPVC 3530 Video Terminal Adapter (VTA)	1.0
IPVC 3540 Gateway Module	1.0.9.1
IPVC 3540 Multipoint Conference Unit (MCU)	2.155
Multimedia Conference Manager (MCM)	Cisco IOS Release 12.2(8)T

Audience

This document is intended for Cisco customers, partners, and systems engineers who will be designing and implementing an IP videoconferencing solution in the enterprise environment.

Organization

This guide contains the chapters and information listed in the following table.

Note   Cisco strongly recommends that you carefully read chapters 1 and 2 before attempting to design an IP videoconferencing solution and before reading any other sections of this guide.

Chapter	Title	Description
1	Introduction	Presents basic concepts related to IP videoconferencing and the H.323 standard.
2	Deployment Models	Describes the primary models used to deploy an IP videoconferencing solution and explains when to use each model. **Note** This guide makes frequent references to these deployment models. Cisco recommends that you read this chapter carefully and understand the main characteristics of each model.
3	Campus Infrastructure	Lists considerations and guidelines for deploying IP videoconferencing with Quality of Service (QoS) in a campus environment (or LAN).
4	WAN Infrastructure	Presents considerations and guidelines for deploying videoconferencing across an IP WAN.
5	WAN QoS	Describes key Quality of Service (QoS) features of the Cisco AVVID network infrastructure and how they apply to IP videoconferencing over a WAN.
6	Dial Plan Architecture	Lists important considerations for designing an effective videoconferencing dial plan, and explains some of the implementation mechanisms available.
7	Call Routing	Describes the main call routing methods used with Cisco gatekeeper and Cisco IP/VC equipment in an H.323 video network, and lists guidelines for using each method.
8	Cisco Video Infrastructure Components	Describes the various components of the video network infrastructure, such as the Cisco Multimedia Conference Manager and the Multipoint Conference Units, and presents guidelines for their use in the enterprise environment.
9	Multi-Zone WAN Case Study	Presents an extended example of a multi-zone WAN implementation that employs many of the concepts and techniques discussed in this guide.
A	Resource Reservation Protocol (RSVP)	Gives a few brief recommendations about using RSVP for call admission control.

Obtaining Documentation

The following sections explain how to obtain documentation from Cisco Systems.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at the following URL:

http://www.cisco.com

Translated documentation is available at the following URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which is shipped with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual subscription.

Ordering Documentation

Cisco documentation is available in the following ways:

- Registered Cisco Direct Customers can order Cisco product documentation from the Networking Products MarketPlace:

  http://www.cisco.com/cgi-bin/order/order_root.pl

- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:

  http://www.cisco.com/go/subscription

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco corporate headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

If you are reading Cisco product documentation on Cisco.com, you can submit technical comments electronically. Click **Leave Feedback** at the bottom of the Cisco Documentation home page. After you complete the form, print it out and fax it to Cisco at 408 527-0730.

You can e-mail your comments to bug-doc@cisco.com.

To submit your comments by mail, use the response card behind the front cover of your document, or write to the following address:

Cisco Systems
Attn: Document Resource Connection
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain documentation, troubleshooting tips, and sample configurations from online tools by using the Cisco Technical Assistance Center (TAC) Web Site. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC Web Site.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com is a highly integrated Internet application and a powerful, easy-to-use tool that provides a broad range of features and services to help you to

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

You can self-register on Cisco.com to obtain customized information and service. To access Cisco.com, go to the following URL:

http://www.cisco.com

Technical Assistance Center

The Cisco TAC is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two types of support are available through the Cisco TAC: the Cisco TAC Web Site and the Cisco TAC Escalation Center.

Inquiries to Cisco TAC are categorized according to the urgency of the issue:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration.
- Priority level 3 (P3)—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- Priority level 2 (P2)—Your production network is severely degraded, affecting significant aspects of business operations. No workaround is available.
- Priority level 1 (P1)—Your production network is down, and a critical impact to business operations will occur if service is not restored quickly. No workaround is available.

Which Cisco TAC resource you choose is based on the priority of the problem and the conditions of service contracts, when applicable.

Cisco TAC Web Site

The Cisco TAC Web Site allows you to resolve P3 and P4 issues yourself, saving both cost and time. The site provides around-the-clock access to online tools, knowledge bases, and software. To access the Cisco TAC Web Site, go to the following URL:

http://www.cisco.com/tac

All customers, partners, and resellers who have a valid Cisco services contract have complete access to the technical support resources on the Cisco TAC Web Site. The Cisco TAC Web Site requires a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to the following URL to register:

http://www.cisco.com/register/

If you cannot resolve your technical issues by using the Cisco TAC Web Site, and you are a Cisco.com registered user, you can open a case online by using the TAC Case Open tool at the following URL:

http://www.cisco.com/tac/caseopen

If you have Internet access, it is recommended that you open P3 and P4 cases through the Cisco TAC Web Site.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses issues that are classified as priority level 1 or priority level 2; these classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer will automatically open a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to the following URL:

http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Before calling, please check with your network operations center to determine the level of Cisco support services to which your company is entitled; for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). In addition, please have available your service agreement number and your product serial number.

Introduction

This chapter provides an overview of the H.323 standard and the video infrastructure components used to build an H.323 videoconferencing network. It describes the basics of the H.323 video standard and infrastructure components used throughout this guide.

H.323 Basics

The H.323 standard provides a foundation for audio, video, and data communications across Internet Protocol (IP) networks. H.323 is an umbrella recommendation from the International Telecommunications Union (ITU) that sets standards for multimedia communications over local area networks (LANs). The H.323 standard is part of a larger range of videoconferencing standards (H.32x) for videoconferencing over various network media. For example, H.320 supports videoconferencing over Integrated Services Digital Network (ISDN), H.321 supports videoconferencing over Asynchronous Transfer Mode (ATM), H.324 supports videoconferencing over standard Plain Old Telephone Service (POTS) lines, and H.323 supports videoconferencing over IP LANs.

The H.323 specification consists of multiple protocols, including:

- H.245 — Provides control signaling used to exchange end-to-end control messages. These control messages carry information relating to:

  - Capabilities exchange

  - Opening and closing of logical channels used to carry media streams

  - Flow control messages

  - General commands and indications

- H.225 — Provides registration, admission, and status (RAS), which is the protocol used between H.323 devices and the gatekeeper for device registration. The RAS protocol is used to perform registration, admission control, bandwidth utilization updates, status, and disengagement procedures between H.323 devices and the gatekeeper. H.225 is also used during call setup to open a call signaling channel using standard Q.931 messaging protocol.

Table 1-1 lists some of the standards supported by the H.323 specification.

Table 1-1   Protocols Supported by the H.323 Standard

Standard	Supported Functions
H.225	RAS, Call Setup and Tear Down (Q.931 call establishment)
H.245	Call Control Messaging
H.261 H.263	Video Formats
G.711 G.722 G.723 G.728	Audio Formats

Videoconferencing with H.323

Historically, videoconferencing was done primarily over ISDN and time division multiplexed (TDM) networks using standard H.320. Running interactive video over data networks was not an option due to video's shared media characteristics, connection-less nature, and lack of guaranteed data flows.   With the introduction of switched 10/100 Mbps networks, high-end routers, and Layer 2 and Layer 3 quality of service (QoS), delivering interactive video over IP is now a reality. Today there is a large installed base of H.320 networks that incur large monthly access and switched usage charges.

With the current advances to the IP networks, it is now possible to run interactive video over an IP network, thus saving customers thousands of dollars a month by converging voice, video, and data traffic over a common path. Costs drop even further as videoconferencing terminals no longer need to support complex network aggregation devices such as Inverse Multiplexers (IMUXs) and can instead rely on simple Ethernet network interface cards (NICs) for network connectivity.

H.323 builds on top of existing IP data networks, ultimately saving money and scaling to larger deployments. The resulting drop in cost per seat is expected to cause an exponential increase in the number of H.323 terminals deployed as users move videoconferencing assets from shared areas, such as conference rooms, to the user desktop. For example, distance learning and business meetings are two common applications that can be deployed effectively with H.323 over IP networks.

H.323 Videoconferencing Components

Five components make up an H.323 videoconferencing network:

- Video Terminal, page 1-4
- Gatekeeper, page 1-5
- Gateway, page 1-6
- Multipoint Conference Unit (MCU), page 1-7
- Proxy, page 1-8

Cisco offers product solutions for all the above components except video terminals, which are covered in detail in *Chapter 8, Video Infrastructure*. Figure 1-1 illustrates a typical H.323 videoconferencing network.

Figure 1-1   H.323 Videoconferencing Infrastructure Components
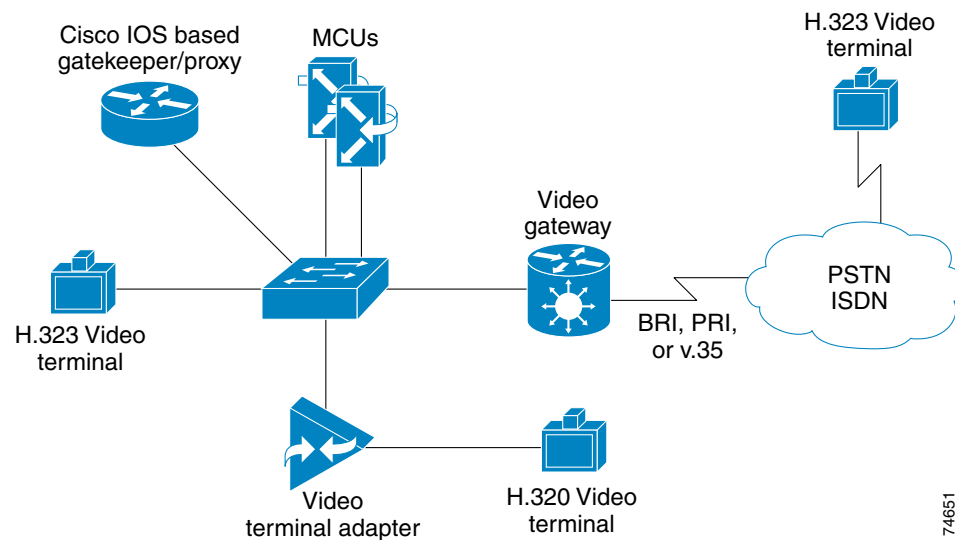
Video Terminal

Video terminals come in many forms, including video systems installed on PCs as standalone desktop terminals and group-focused shared conference room devices.   Figure 1-2 illustrates the functional components in an H.323 video terminal.

Figure 1-2   Functional Components of a Video Terminal
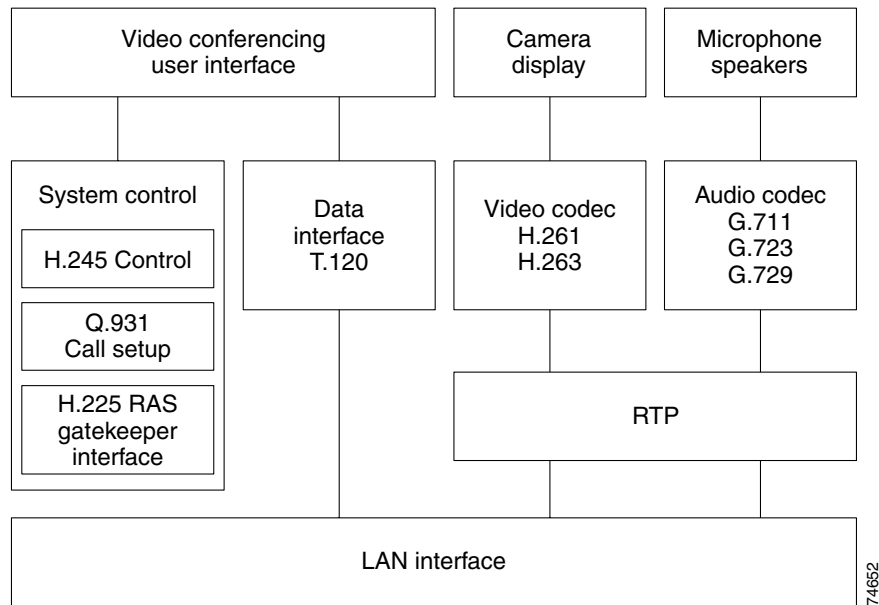
Gatekeeper

The gatekeeper is one of the most important components of an H.323 videoconferencing network. Although the H.323 standard lists the gatekeeper as an optional device, you cannot build a scalable video network without the application controls the gatekeeper provides. Each video infrastructure component registers with the gatekeeper. The gatekeeper performs all address resolution, bandwidth management, admission control, zone management, and intra-zone and inter-zone call routing.

A zone is a logical grouping of H.323 infrastructure components registered to, and managed by, a single gatekeeper. Zones are not dependent on physical network topology or IP subnets. Zones, which may span one or more network segments or IP subnets, are simply a logical grouping of devices. As such, zones can be defined based on geographical proximity, bandwidth availability, or other criteria.

The most fundamental function of a gatekeeper is to provide address resolution, thus allowing terminals, gateways, and MCUs to be addressed using the international E.164 address standard and/or an H.323 alias. Each endpoint that is registered to a gatekeeper must be assigned a unique E.164 address (numeric identifier). As a result, zone prefixes are used in the H.323 video network to identify zones, similar to the use of area codes in telephony systems.

Throughout this document are example topologies that are based on single-zone and multi-zone configurations. For example, Figure 1-3 illustrates a single zone.

Figure 1-3   Single H.323 Zone
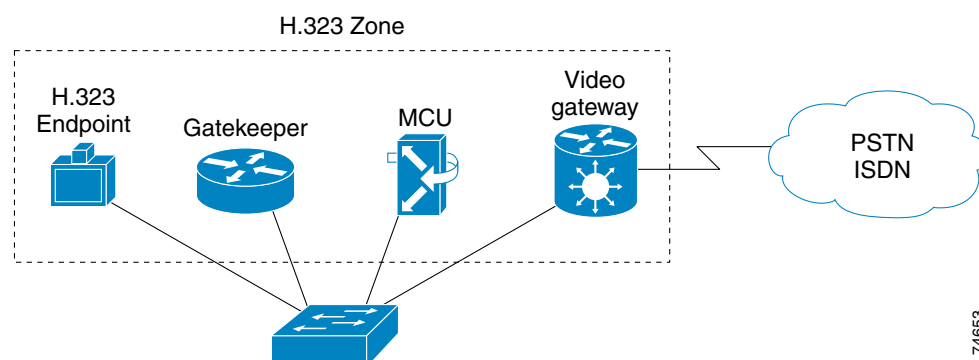
Gateway

Gateways provide interoperability between H.323 elements and an installed base of H.320 units. The H.323 gateway allows H.323 video terminals to communicate with other H.32*x* video terminals, such as H.320 and H.321 video terminals. Video gateways perform translation between different protocols, audio encoding formats, and video encoding formats that may be used by the various H.32*x* standards. For example, the ISDN H.320 standard uses the H.221 protocol for signaling, while the H.323 standard uses H.225. The gateway must translate between these two protocols to allow devices of different network media and protocols to communicate with each other. Figure 1-4 illustrates the role of a gateway in an H.323 video network.

Figure 1-4   Functional Components of an H.323 Video Gateway
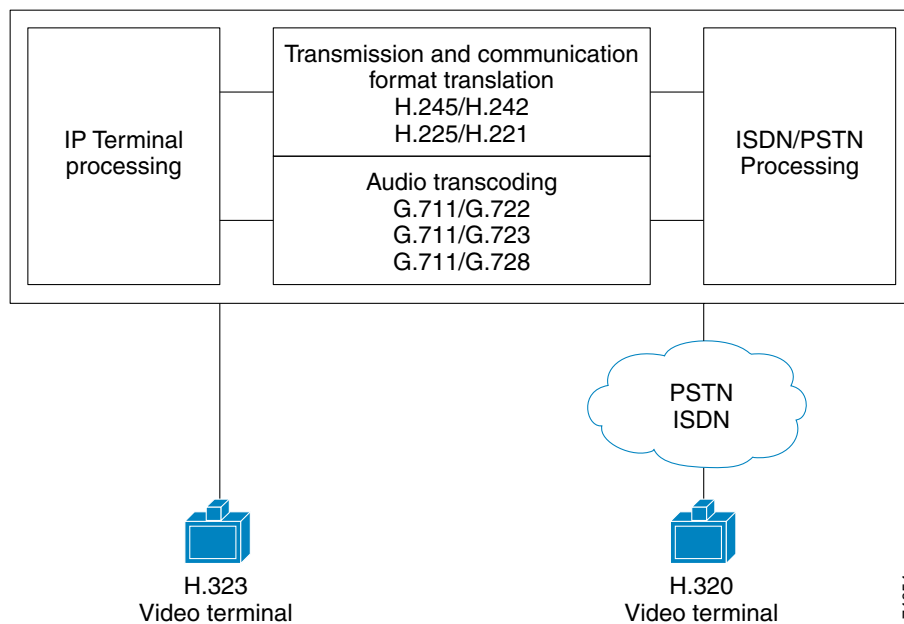
Multipoint Conference Unit (MCU)

Video terminals are generally point-to-point devices, allowing only two participants per conversation. A multipoint conference unit (MCU) allows video conferences to be extended to three or more participants. An MCU consists of a multipoint controller (MC) and a multipoint processor (MP). The MC manages all call setup control functions and conference resources as well as the opening and closing of media streams. The MP processes audio and video media streams only. Cisco MCUs can be stacked to create more conferences or cascaded to create larger conferences. Stacking and cascading are covered in detail in *Chapter 8, Video Infrastructure*. Figure 1-5 illustrates the function of an MCU.

Figure 1-5   Functional Components of an MCU
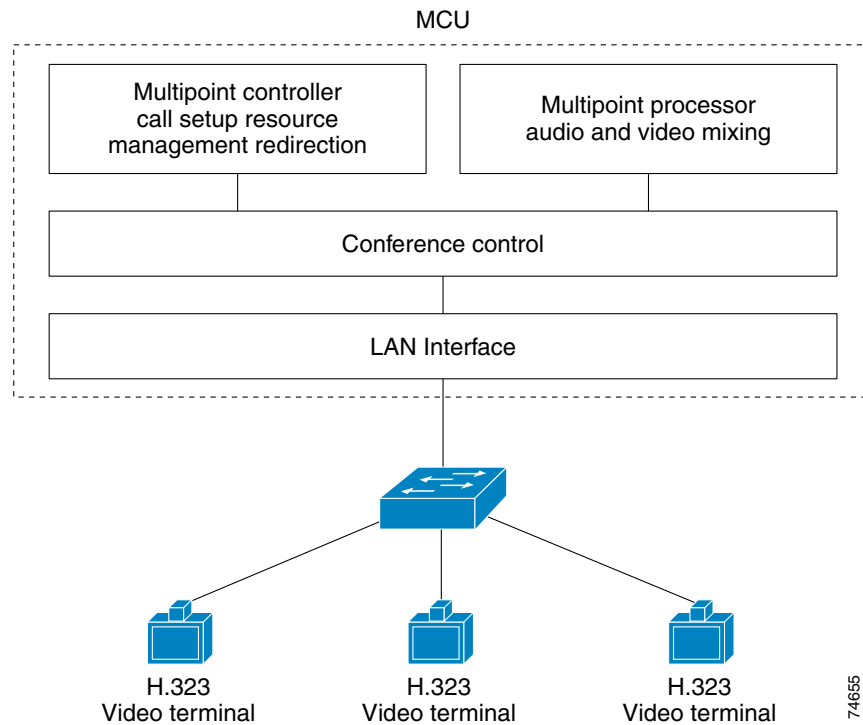
Proxy

A proxy is a call processing agent that terminates H.323 calls from a local LAN or zone and establishes sessions with H.323 endpoints located in other LANs or zones. In so doing, the proxy provides network administrators with the ability to set and enforce quality of service (QoS) on inter-zone segments. The proxy also provides a method of identifying H.323 videoconferencing connections for tunneling through firewalls and Network Address Translation (NAT) environments. Figure 1-6 illustrates a proxy call over a WAN link.

Figure 1-6   Proxy Call Over a WAN Link

CHAPTER

2
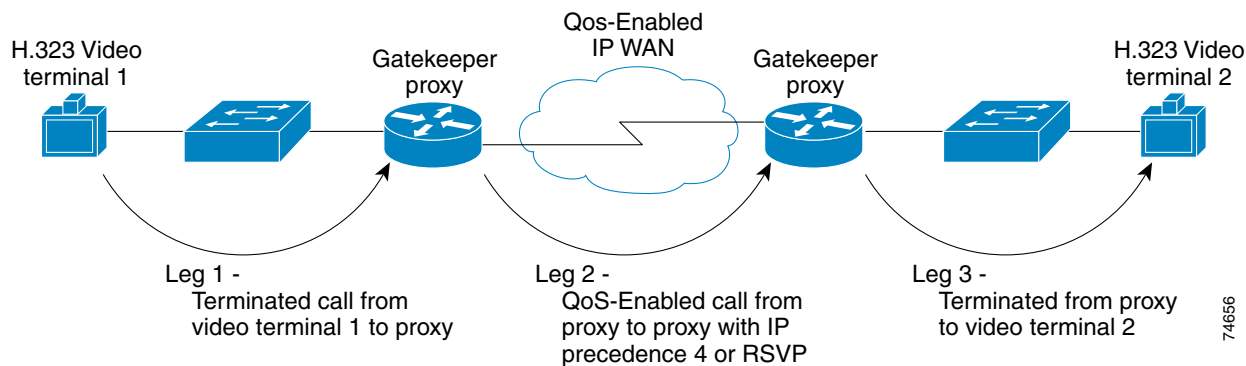
Deployment Models

This chapter introduces four basic design models used to deploy IP videoconferencing solutions:

This chapter provides basic design criteria and guidelines for selecting the correct deployment model. Subsequent chapters of this design guide describe in more detail each of the basic models introduced here.

Composite Deployment Model

Figure 2-1 illustrates a composite topology that encompasses all of the deployment models discussed in this guide. All designs discussed in this chapter are supported with currently shipping products.

The overall goals of a Cisco-based H.323 videoconferencing solution are as follows:

- Provide end-to-end IP video connectivity across the corporate infrastructure, with *business quality* transmission. Business quality video is defined as 30 frames per second operation with a minimum of Common Intermediate Format (CIF) resolution. Typically, this level of quality requires 384 kbps of application bandwidth for most video terminals.
- Provide quality of service (QoS) — high availability with low latency and jitter (delay variability).
- Reduce Integrated Services Digital Network (ISDN) costs by eliminating the need for ISDN attachments directly to video terminals.
- Allow Public Switched Telephone Network (PSTN) access to legacy H.320 systems through shared gateway resources.
- Support multipoint calling through Multipoint Conference Units (MCUs).
- Conserve WAN bandwidth by distributing MCU and gateway resources across the IP infrastructure.
- Lower total cost of ownership for the video network by utilizing the existing IP infrastructure.
- Support manageability of multiple H.323 elements in a distributed network topology.

Figure 2-1   Composite Deployment Model
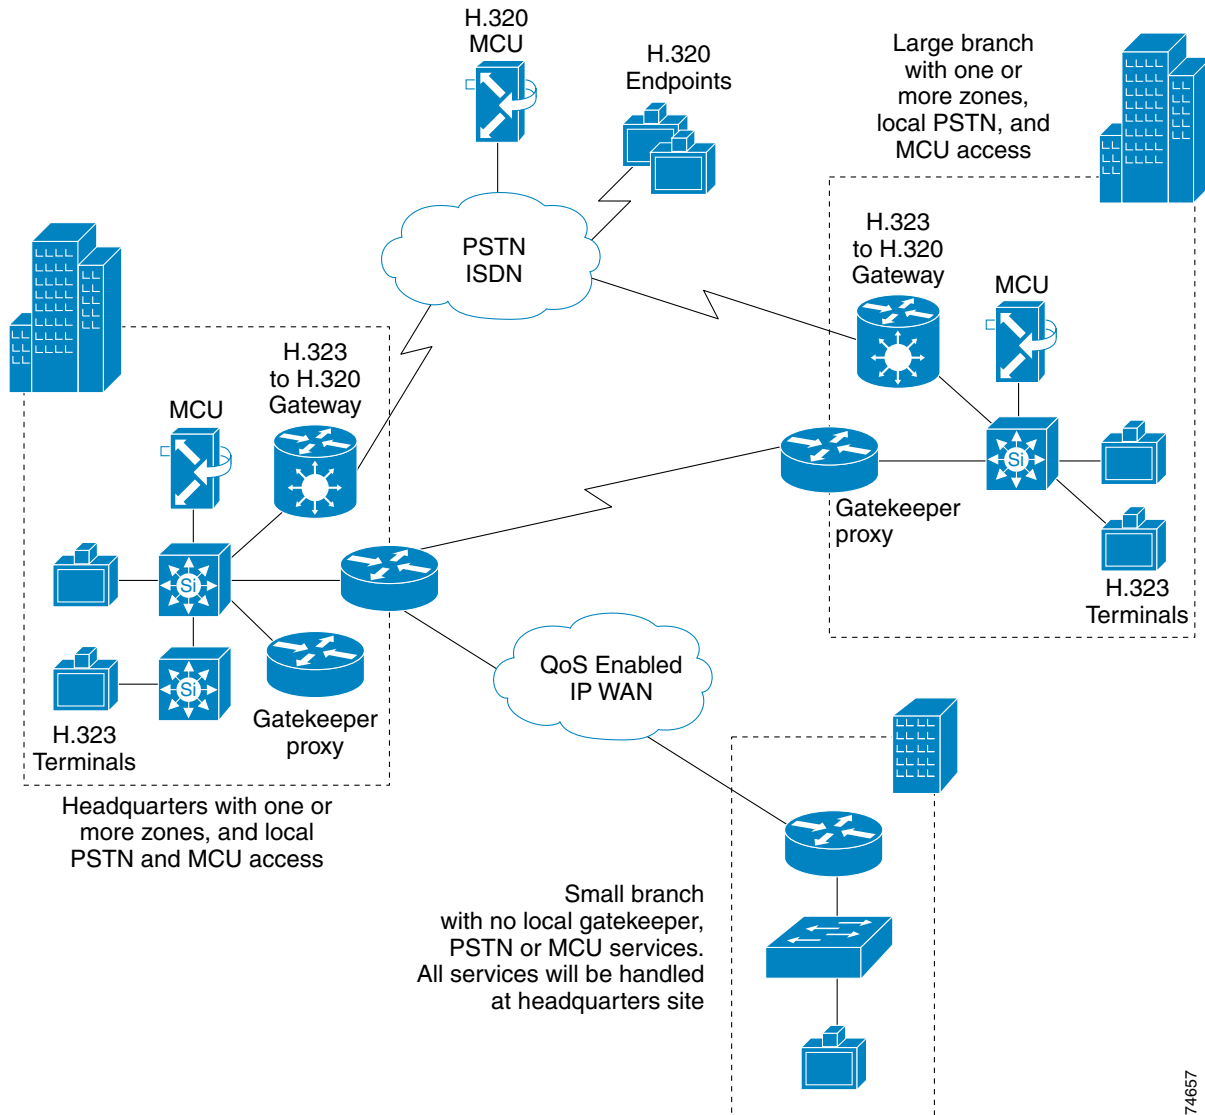
Campus Single Zone

Figure 2-2 illustrates an H.323 network in a campus environment configured with a single zone. This is the most basic design model to implement and is used in pilot installs and smaller video environments.

Figure 2-2   Campus Single Zone
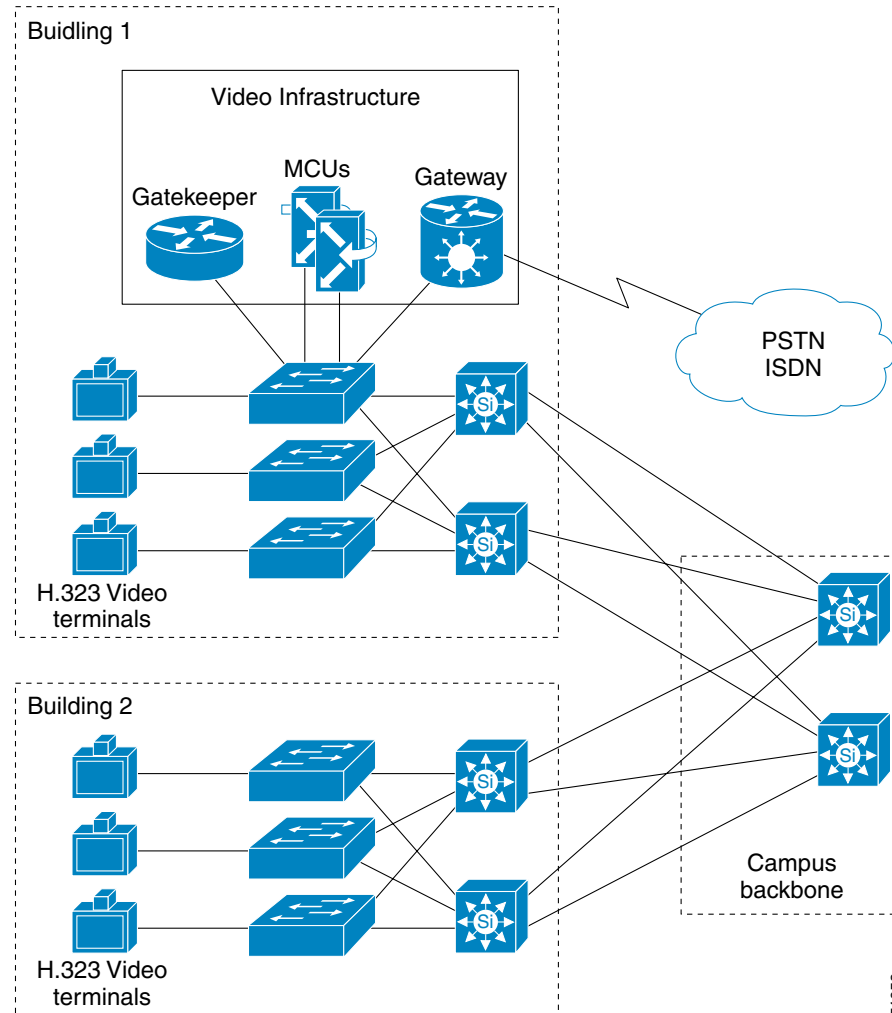
The campus single-zone deployment model has the following design characteristics:

- A single gatekeeper supporting a single zone for H.323 video.

- All H.323 video users registered with the single gatekeeper. (See *Chapter 8* for gatekeeper registration limits.)

- Optional PSTN access available through the Cisco IP/VC 352X gateway.

- Optional multipoint conferencing available through the Cisco IP/VC 3510 MCU.

- Zone bandwidth managed by the configured gatekeeper.

- All gateway and MCU services registered and managed by a single gatekeeper.

- Call routing between endpoints using fully qualified E.164 addresses or H.323-ID.

Campus Multi Zone

Figure 2-3 illustrates a multi-zone H.323 video network in a campus environment. This model is most often implemented in an enterprise campus network. Depending on business function, administrators may choose to create different zones for security reasons. For example, company executives may be registered in a single zone that is separate from other users to allow administrators to limit access to those video terminals. In addition, as a video network grows, a single zone may not be manageable because of the number of users or the ability to manage network resources.

Note   Multiple zones can be configured on a single router. If you configure multiple local zones on a single router, and MCUs and/or gateways are registered with the zones, you must add hopoff statements for each service prefix. If hopoffs are not added for each service prefix, the video terminal will not be able to access MCUs or gateways outside its local zone. See Routing Inter-Zone Calls Using Hopoff Statements, page 7-8, for more information.

Figure 2-3   Campus Multi Zone
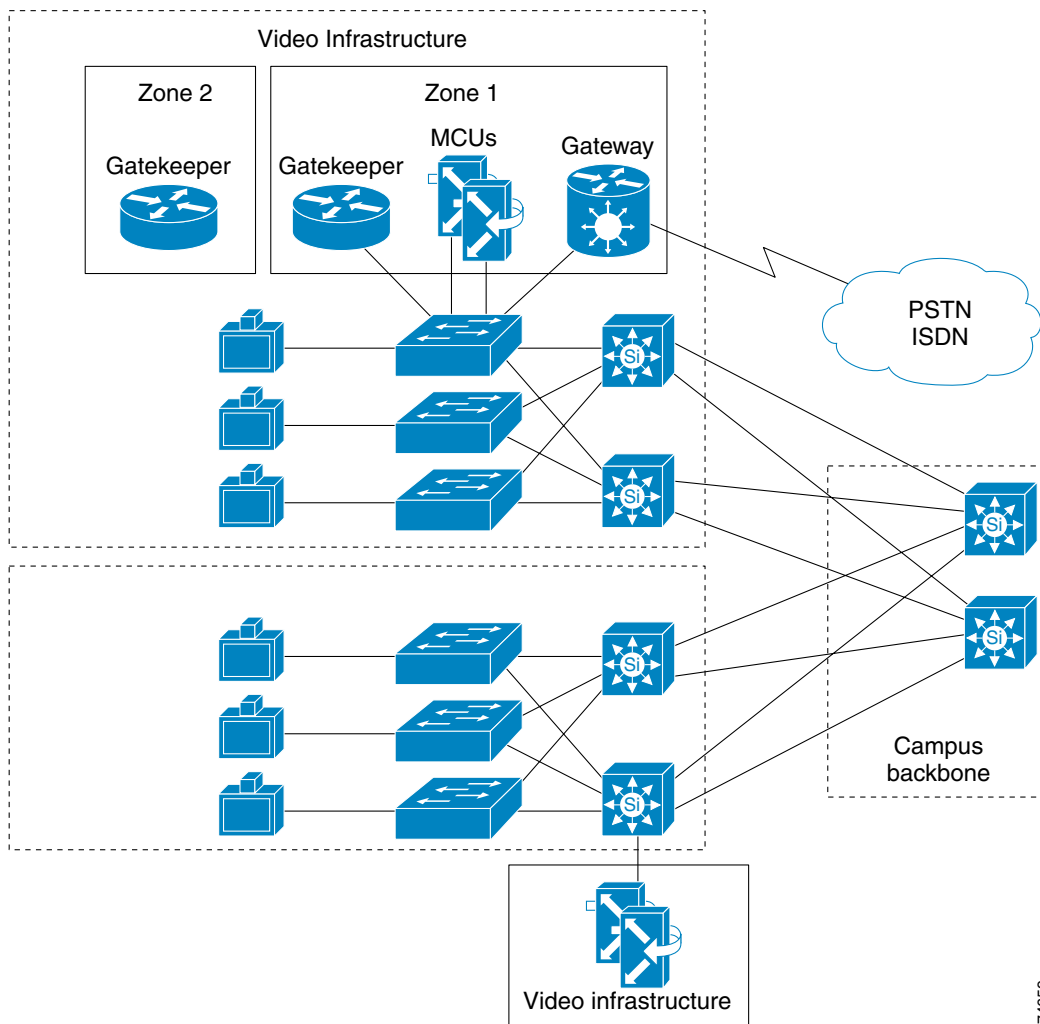
The campus multi-zone deployment model has the following design characteristics:

- Multiple gatekeepers supporting multiple zones for H.323 video.
- H.323 endpoints register with one of the multiple gatekeepers. (See Chapter 8 for gatekeeper registration limits.)
- Bandwidth management for each zone and between zones is controlled by configured gatekeepers.
- Optional PSTN access available through Cisco IP/VC 352X gateway.
- Gateway and MCU services are registered and managed across multiple gatekeepers.
- Gateway and MCU services may be distributed throughout the campus.
- H.323 users and services are segmented for security, bandwidth control, and resource allocation.
- Intra-zone and inter-zone call routing using fully qualified E.164 address or H.323-ID.

WAN Single Zone

Figure 2-4 illustrates a single-zone H.323 video network in a WAN environment. This deployment model is used when remote sites have a small number of video endpoints, usually no more than one or two at each remote site on a T1 WAN link. From a management or economic standpoint, it might not make sense to create a zone at each remote site for one or two video terminals. Call admission control (CAC) across the WAN is not usually an issue with only one or two video terminals at each remote site, but it is an issue when the number of remote endpoints exceeds the provisioned video bandwidth.

In the absence of a gatekeeper, implement quality of service on the WAN ports by using one of the following methods:

- Priority queuing on traffic classification IP Precedence 4, or Differentiated Services Code Point (DSCP) AF41
- Access control list (ACL) for each video terminal at the remote site, to direct the video streams to the appropriate priority queue

Figure 2-4   WAN Single Zone
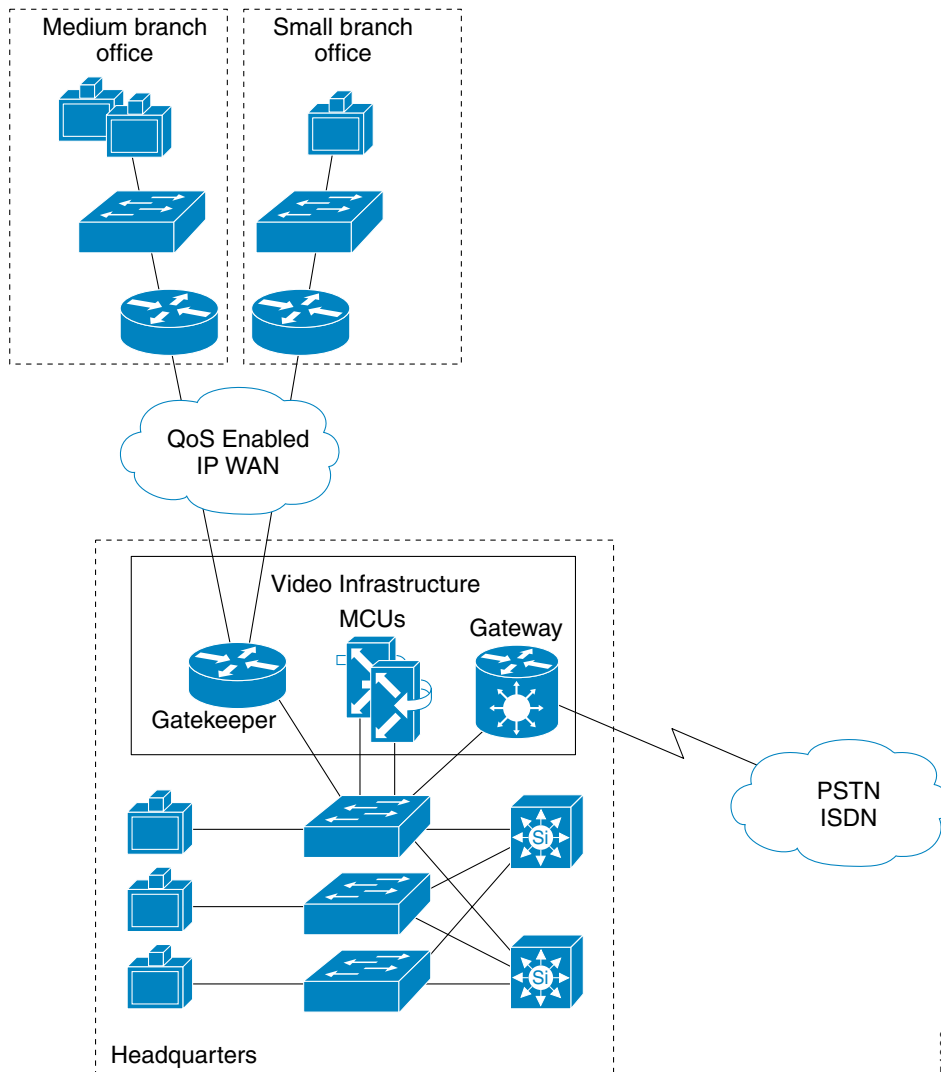
The WAN single-zone deployment model has the following design characteristics:

- A single gatekeeper supporting a single zone for H.323 video.
- All H.323 video users registered with the single gatekeeper. (See *Chapter 8* for gatekeeper registration limits.)
- Optional PSTN access available through Cisco IP/VC 352X gateway.
- Optional multipoint conferencing available through the Cisco IP/VC 3510 MCU.
- H.323 video bandwidth managed by a single gatekeeper.
- All gateway and MCU services registered and managed by a single gatekeeper.
- WAN QoS, with priority queuing by means of traffic classification or ACL entries.
- Call routing between endpoints using fully qualified E.164 addresses or H.323-ID.

WAN Multi Zone

Figure 2-5 illustrates a multi-zone H.323 network in a WAN environment. This deployment model is used in large enterprise, government, and educational networks. QoS can be implemented using either the proxy and priority queuing (PQ) features in Cisco IOS software, traffic classification by the video terminals, or Layer 3 switches in conjunction with priority queuing on the WAN ports of the routers.

Creating multiple zones in a WAN environment allows administrators to manage network resources and assure video quality across low-speed WAN links. Call admission control (CAC) is very important in a large WAN environment. With multiple zones enabled, the gatekeeper can manage the total amount of H.323 video bandwidth allowed across a particular network link. For example, you could limit the total H.323 video bandwidth across a T1 WAN link to 768 kbps, and the gatekeeper would then reject any call request that exceeds this limit of 768 kbps.

Figure 2-5   WAN Multi Zone
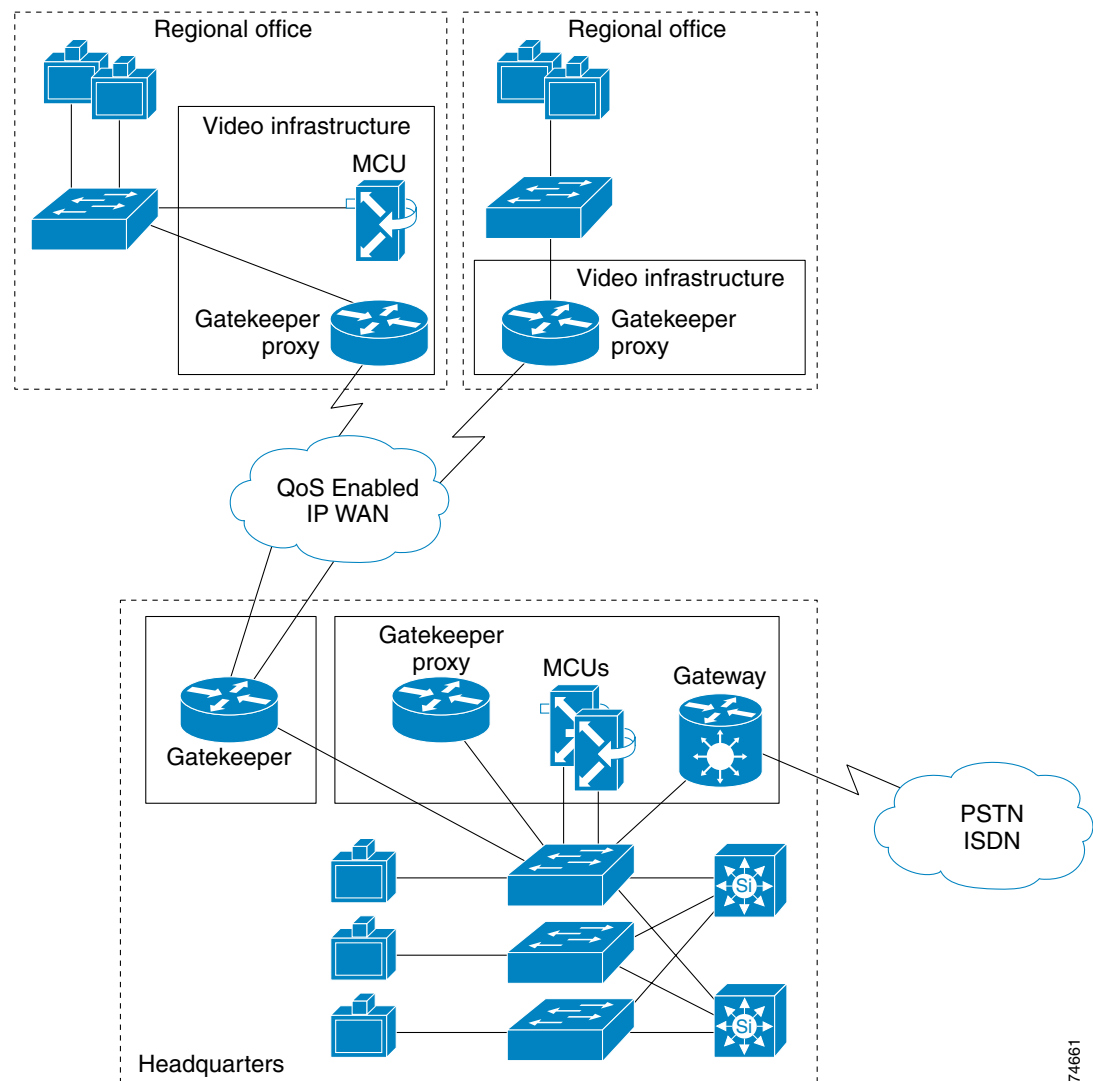
The WAN multi-zone deployment model has the following design characteristics:

- Multiple gatekeepers supporting multiple zones for H.323 video.

- H.323 endpoints and services registered with the assigned gatekeeper, usually at the local site.

- Optional PSTN access available through Cisco IP/VC 352X.

- Bandwidth management available in each zone and across the WAN, using the gatekeeper at each site.

- Distributed services available at larger branch sites to conserve bandwidth.

- Inter-zone and intra-zone call routing using fully qualified E.164 addresses or H.323-ID.

- Proxy enabled at each site with priority queuing (PQ) on the WAN, or PQ based on traffic classification implemented on the WAN ports.

Campus Infrastructure

This chapter provides guidelines for deploying H.323 videoconferencing with Quality of Service (QoS) on a campus network using one of the following basic H.323 video designs:

- Single-Zone Campus, page 3-2
- Multi-Zone Campus, page 3-3

Network Infrastructure

Building an end-to-end H.323 video network requires an infrastructure based on Layer 2 and Layer 3 switches and routers. It is important to have all H.323 video endpoints, gateways, and multipoint conference units (MCUs) connected to a dedicated 10/100 switched Ethernet port. Cisco recommends using a 100-Mbps full duplex connection to the Cisco gatekeeper to ensure adequate bandwidth on all router platforms. Some endpoints, however, do not support 100-Mbps full duplex. For example, older Polycom ViewStations and the Cisco IP/VC 3530 both support 10-Mbps half duplex only.

> **Note** There are known issues with some Cisco Catalyst switches and video endpoints negotiating half/full duplex. If the negotiation fails, the endpoint still functions but the system experiences video freezing every three to five seconds. Cisco recommends that you set all switch ports attached to H.323 video devices to 100-Mbps full duplex whenever possible. If the video unit supports only 10 Mbps, configure the switch port for 10-Mbps half duplex.

Single-Zone Campus

Figure 3-1 illustrates an H.323 single-zone campus network.

Figure 3-1   Single-Zone Campus
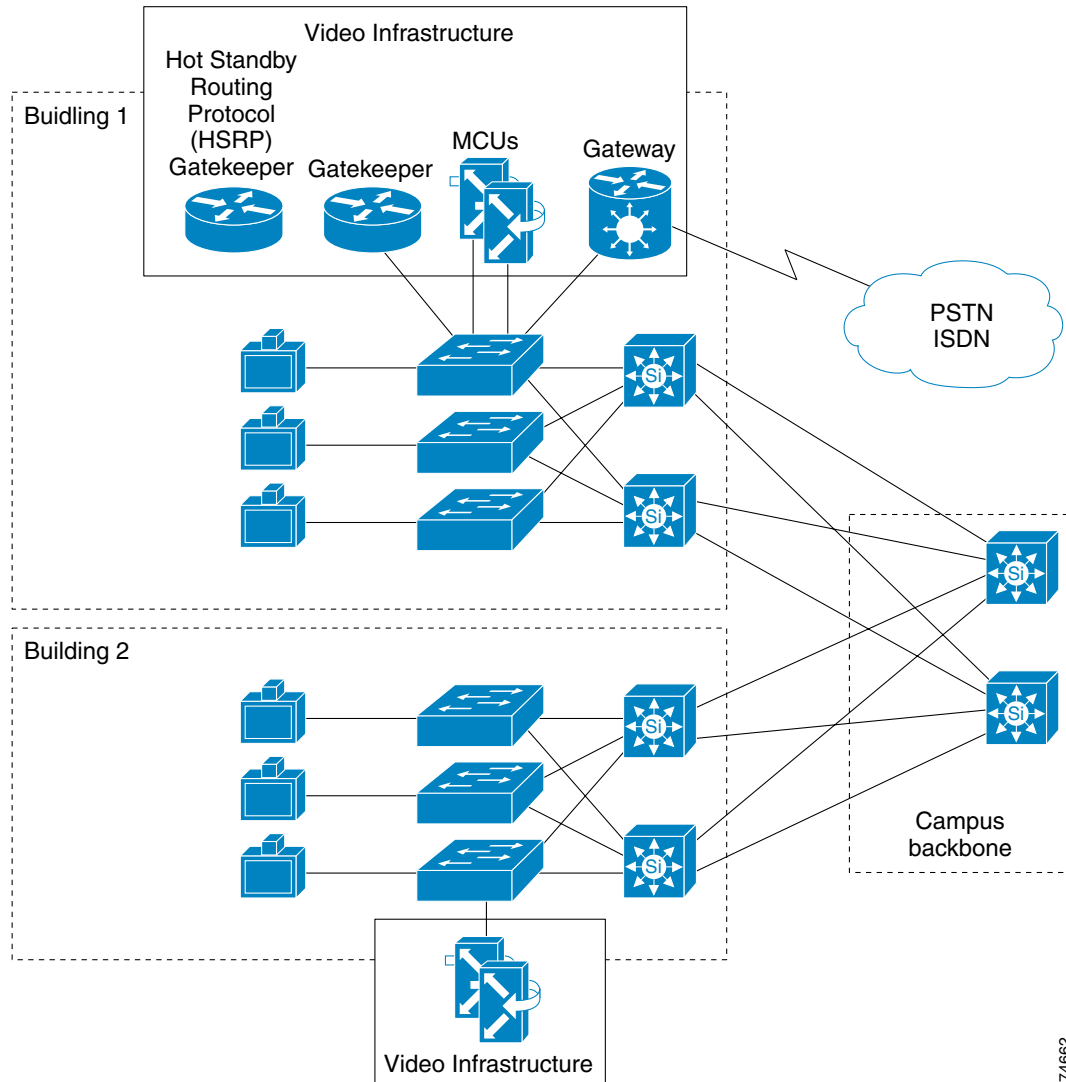
Single-zone campus networks are usually used in pilot deployments or in campuses with a small number of video terminals or endpoints. The single-zone campus deployment allows an administrator to deploy H.323 video on the campus while keeping management overhead to a minimum. There is only one gatekeeper to manage, and the dial plan is very simple with no inter-zone call routing.

It is important to consider multi-zone dial plans when deploying a single-zone model. If you deploy a single-zone dial plan but need to upgrade to a multi-zone model in the future, you will have to change the entire dial plan. Therefore, to simplify future network scaling, Cisco recommends that you use a multi-zone dial plan even for a single-zone campus.

In summary, a single-zone campus model consists of:

- Campus environment
- Pilot environments
- Small number of video endpoints
- No bandwidth limitations

Multi-Zone Campus

Figure 3-2 illustrates an H.323 multi-zone campus network.

Figure 3-2   Multi-Zone Campus
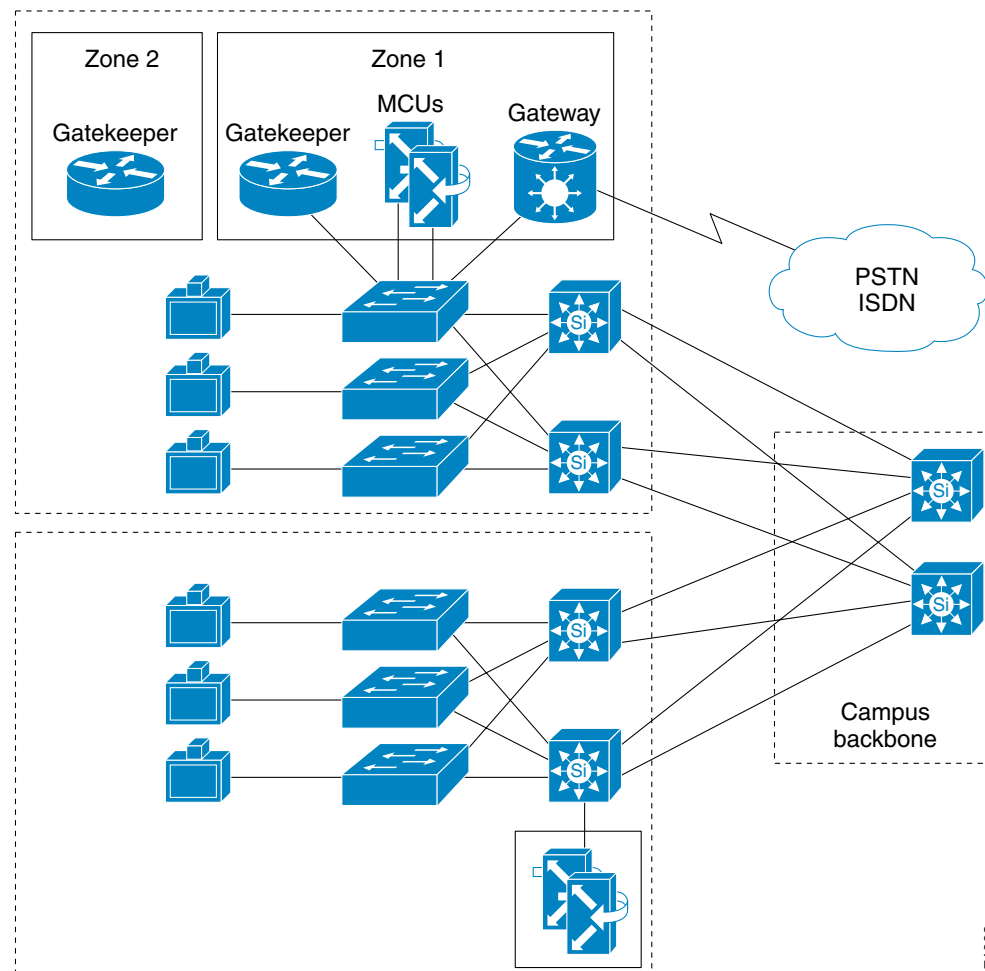
Multi-zone campus networks are common in large campus environments. Creating multiple zones allows administrators to segment user groups for security, better management of the H.323 video network, and bandwidth control in and between zones. For example, company executives may be registered in a single zone containing their own gateway and MCU resources.

In campuses with a large number of video terminals, it is important to control the amount of video bandwidth on the network. With a single zone, bandwidth management capabilities are very limited. Creating multiple logical zones on the campus allows an administrator to manage bandwidth within and between zones.

Physical placement of gatekeepers, MCUs, and gateways depends on customer preference and network configuration. Some deployments locate all of the gatekeepers, MCUs, and gateways in a single data center, while others may decide to distribute the equipment throughout the campus.

In summary, the multi-zone campus model consists of:

- Campus environment
- Large numbers of video terminals
- Users segmented into separate video domains
- Restricted access for some users

Note   Multiple zones can be configured on a single router. If you configure multiple local zones on a single router, you must add hopoff commands for each service prefix registered. If hopoffs are not added for each service prefix, the video terminal will not be able to access MCUs or gateways outside its local zone. See Routing Inter-Zone Calls Using Hopoff Statements, page 7-8 for more information.

Quality of Service

In a converged environment, voice, video and data traffic all travel over a single transport infrastructure. Not all traffic types should be treated equally. Data traffic is bursty, loss tolerant, and not sensitive to delay. Video traffic, on the other hand, is bursty, has very little tolerance for loss, and is latency sensitive. The challenge is to provide the required level of service for all three traffic types.

Running both video and data on a common network requires the proper QoS tools to ensure that the delay and loss parameters of video traffic are satisfied in the face of unpredictable data flows. Some of these tools may be available as a feature in some video terminals (for example, Polycom, VCON, and PictureTel), switches, and routers.

Traffic Classification Types

The first step in preserving video quality on a data network is to classify video traffic as high priority and allow it to travel through the network before lower priority traffic. Data traffic can be classified at a lower priority without adversely affecting its performance because of its characteristics as provided by the Transfer Control Protocol (TCP), which handles flow control and error correction. For video, classify traffic at Layer 2 and Layer 3 as follows:

- At Layer 2, use the three bits in the 802.1Qp field, referred to as class of service (CoS), which is part of the 802.1Q tag.
- At Layer 3, use the three bits of the Differentiated Services Code Point (DSCP) field in the type of service (ToS) byte of the IP header.

Traffic classification is the first step toward achieving QoS. Ideally, you should perform this step as close to the source as possible. However, you can also set this field within a router using the Cisco Multimedia Conference Manager (MCM), a Cisco IOS feature.

Table 3-1 lists the recommended traffic classifications for various applications.

Table 3-1   Recommended Traffic Classifications

Layer 2	Layer 3 Classification			
CoS	IP Precedence	Pre-Hop Behavior (PHB)	DSCP	Application
7	7	–	56-63	Reserved
6	6	–	48-55	Reserved
5	5	EF	46	Voice Bearer
4	4	AF41	34	Video Conferencing
3	3	AF31	26	Call Signaling
2	2	AF2y	18, 20, 22	High Priority Data
1	1	AF1y	10, 14, 16	Medium Priority Data
0	0	BE	0	Best Effort Data

Trust Boundaries

The concept of trust is an important and integral part of deploying QoS. Once the end devices have set ToS values, the switch has the option of trusting them or not. If the switch trusts the ToS values, it does not need to do any reclassification; if it does not trust the values, then it must reclassify the traffic for appropriate QoS.

The notion of trusting or not trusting forms the basis for the trust boundary. Ideally, traffic classification should be done as close to the source as possible. If the end device is capable of performing traffic classification, then the trust boundary for the network is at the access layer in the wiring closet. If the device is not capable of performing traffic classification, or if the wiring closet switch does not trust the classification done by the end device, the trust boundary should shift to other devices.

Shifting of the trust boundary depends on the capabilities of the switch in the wiring closet. If the switch can reclassify the packets, then the trust boundary remains in the wiring closet. If the switch cannot perform this function, then the task falls to other devices in the network going toward the backbone. In this case, reclassification occurs at the distribution layer, which means that the trust boundary has shifted to the distribution layer. For this shift to occur, there must be a high-end switch in the distribution layer with features to support traffic reclassification. If possible, try to avoid performing traffic reclassification in the core of the network.

In summary, try to maintain the trust boundary in the wiring closet. If necessary, move it down to the distribution layer on a case-by-case basis, but avoid moving it to the core of the network. This advice conforms to the general guidelines for keeping the trust boundary as close to the source as possible.

Note This discussion assumes a three-tier network model, which has proven to be a scalable architecture. If the network is small and the logical functions of the distribution layer and core layer happen to be in the same device, then the trust boundary can reside in the core layer if it has to move from the wiring closet. For detailed configuration information, refer to the *Cisco AVVID Network Infrastructure Enterprise Quality of Service Design* guide.

QoS Features Summary

Table 3-2 shows supported QoS features on each switch platform.

Table 3-2   Supported QoS Features by Switch Platform

Platform	Ability to Trust	Reclassify CoS	Reclassify ToS	Congestion Avoidance (WRED)[1]	Priority Queues	Multiple Queues	Congestion Management (WRR)[2]	Policing
Catalyst 2900XL	No	Yes	No	No	No	Yes	No	No
Catalyst 2950	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Catalyst 3500XL and 3524-PWR-XL	Yes	Yes	No	No	Yes	Yes	Yes	No
Catalyst 3550	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Catalyst 4000 with Supervisor Engine II	Yes	Yes (Switch-wide)	No	No	No	Yes	No	No
Catalyst 4006 with Supervisor Engine II	Yes	Yes (Switch-wide)	No	No	No	Yes	No	No
Catalyst 4006 with Supervisor Engine III	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Catalyst 5000	Yes	Yes	Yes	Yes (Does not work for VoIP on bottom threshold)	Yes	Yes	No	Yes
Catalyst 6000 with Policy Feature Card (PFC)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

1. Weighted random early detection (WRED).

2. Weighted round robin (WRR).

Note   Currently the only Cisco LAN switches that support a minimum of two queues and that can guarantee video quality are the Catalyst 8500, Catalyst 6000 family, Catalyst 4000 family, Catalyst 3500XL, and Catalyst 2900XL.

In summary, follow these recommendations for QoS deployment:

- Create a trust boundary at the network edge in the wiring closet. Enable the trust boundary on ports on the wiring closet switch where video terminals have the ability to set IP precedence. A rule of thumb is to trust the classification from conference room systems and *not* trust classification from desktop video units.

- Reclassify ToS at the edge if devices (both room systems and desktop units) cannot be trusted.

- Shift the trust boundary to the distribution layer and reclassify ToS there if reclassification is not possible at the edge.

- Use a priority queue for delay-sensitive video traffic.

WAN Infrastructure

This chapter provides guidelines for deploying H.323 video across an IP WAN, and it describes IP WAN infrastructure design considerations for:

Single-Zone WAN

Figure 4-1 illustrates a single-zone WAN network.

Figure 4-1   Single-Zone WAN
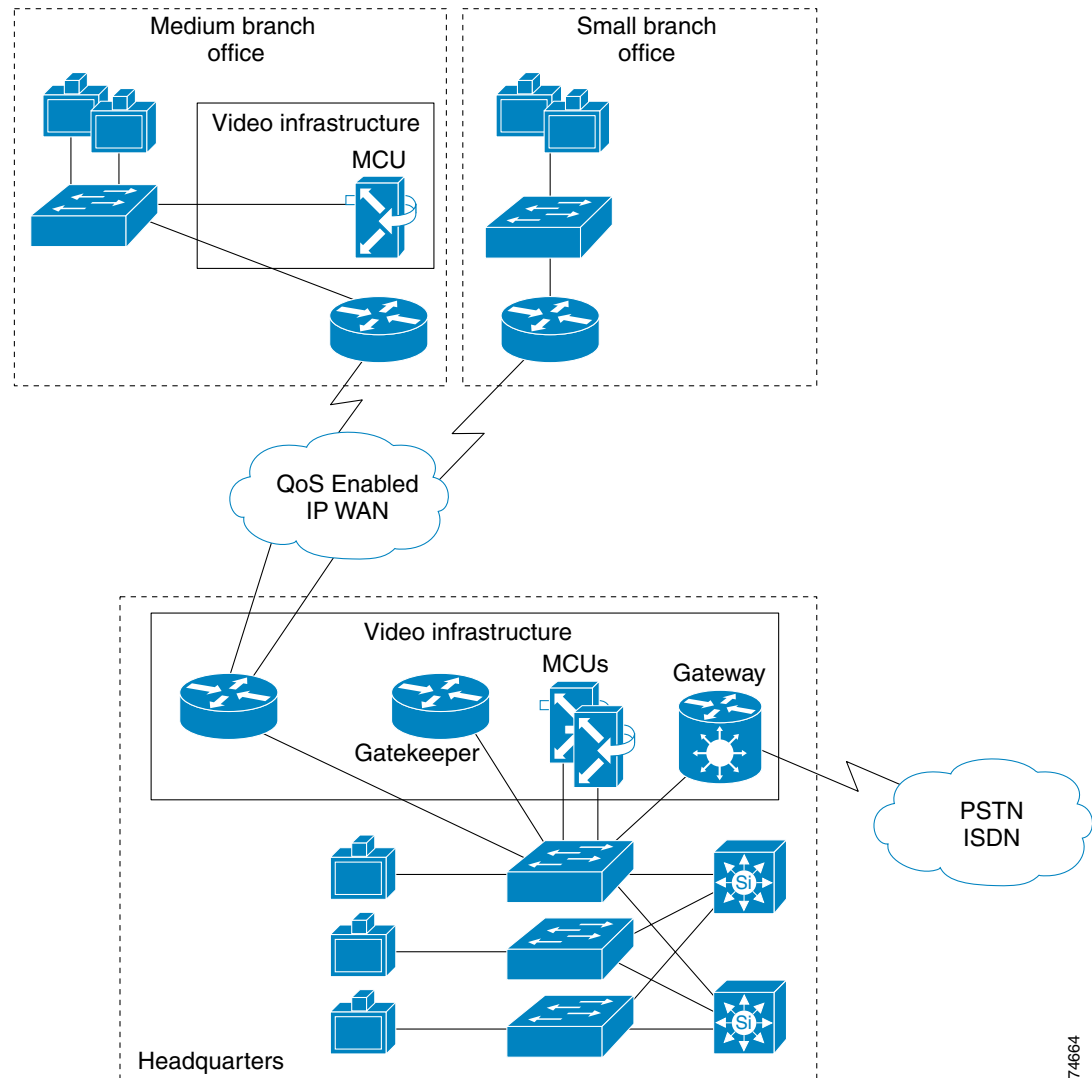
A single-zone WAN model consists of the WAN environment and less than three videoconferencing terminals per remote site. (This limit is based on a T1 WAN link.) Cisco recommends that you configure a gatekeeper and a zone for a remote site with one or two video terminals, but this configuration is not strictly required.

Due to the limited number of endpoints and traffic classification options, you can achieve quality of service (QoS) and call admission control (CAC) by following these basic rules:

- The total data rate of the video terminals plus 20% should not exceed 33% of the WAN link capacity.

- The priority queue must be provisioned for the maximum data rate of the video terminals plus 20%.

  For example, assume a site has a link capacity of 1.544 Mbps and contains two video terminals that support a maximum data rate of 256 kbps each. Therefore, the required queue size for the two video terminals is (256+256)x120% = 614 kbps. Provisioning the priority queue for 614 kbps allows both video terminals to be in a call across the WAN at the same time, without the possibility of overrunning the priority queue. If we add a third video terminal in this example, we would need to add a gatekeeper and create a zone to provide call admission control.

The key elements for successful deployment of videoconferencing in a single-zone WAN environment are:

- Traffic Classification, page 4-3
- Call Admission Control (CAC), page 4-4
- Provisioning, page 4-4
- Priority Queuing on the WAN, page 4-4
- Entrance Criteria, page 4-4

Traffic Classification

Classify traffic at one of the following places:

- Video endpoint (Polycom, PictureTel, Tandberg, and VCON); IP Precedence 4 or DSCP AF41
- Switch port (Layer 3 switch required); IP Precedence 4 or DSCP AF41 (recommended)
- Router (ACL entry); IP Precedence 4 or DSCP AF41

Figure 4-2 illustrates these three options for traffic classification.

Figure 4-2   Traffic Classification Options for Single-Zone WAN
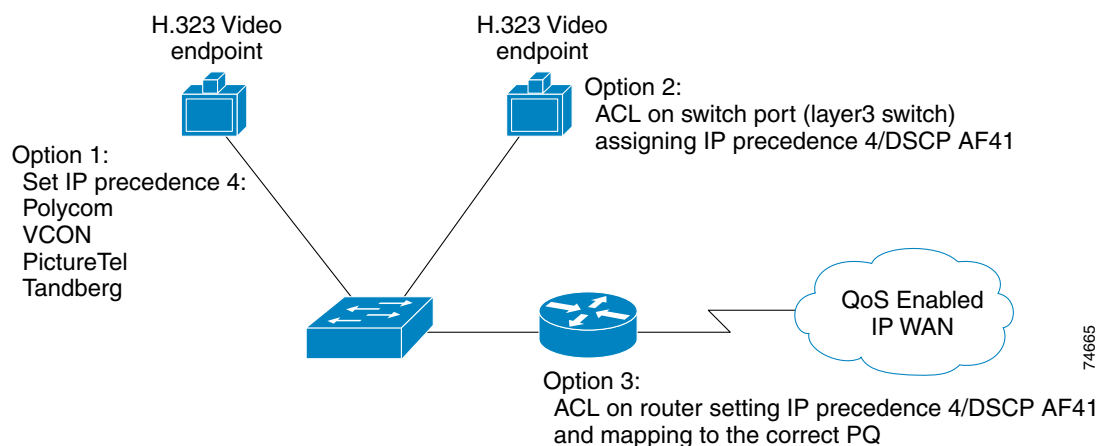
Call Admission Control (CAC)

For remote sites that do not have a gatekeeper to enforce CAC, provision the priority queue and limit the number of video terminals at each site. The number of video terminals multiplied by the maximum call data rate, must not exceed the capacity of the priority queue. Cisco recommends that you use a gatekeeper and zones for remote sites with more than two video terminals. You can install a gatekeeper at each remote site with more than two video terminals, or you can install one gatekeeper at the central site and define a separate zone for each remote site.

Note   This recommendation is based on a T1 WAN link.

Provisioning

Provision WAN queues according to the following equation:

Priority queue size = (Number of users) x (Maximum data rate of video terminals) x 120%

The priority queue must be provisioned to handle the maximum data rate used by any of the video terminals, otherwise the priority queue has the potential to become oversubscribed. Add 20% to the maximum data rate of the video terminals to allow for IP and transport overhead. Refer to the WAN QoS chapter for more information.

Priority Queuing on the WAN

Configure multiple queues for the WAN ports on routers. Videoconferencing traffic goes into a priority queue (PQ) that services IP Precedence 4 or DSCP AF41. Class-based weighted fair queuing (CBWFQ) is *not* recommended for interactive video.

Entrance Criteria

In the single-zone WAN model, use access control lists (ACLs) to access configured priority queues at remote sites. ACLs ensure that only traffic from the video terminals has access to the configure PQ. The small number of video terminals at remote sites makes ACL entries a viable option.

Central sites that have either Layer 3 switches or video terminals capable of setting IP Precedence, should set the entrance criteria for the PQ to any packets with IP Precedence 4 or DSCP AF41. This method, however, is not as secure as the ACL option but works properly if the trust boundaries are configured correctly.   This method can also be used at remote sites if ACLs are not acceptable.

Multi-Zone WAN

Figure 4-3 illustrates a multi-zone WAN network.

Figure 4-3   Multi-Zone WAN
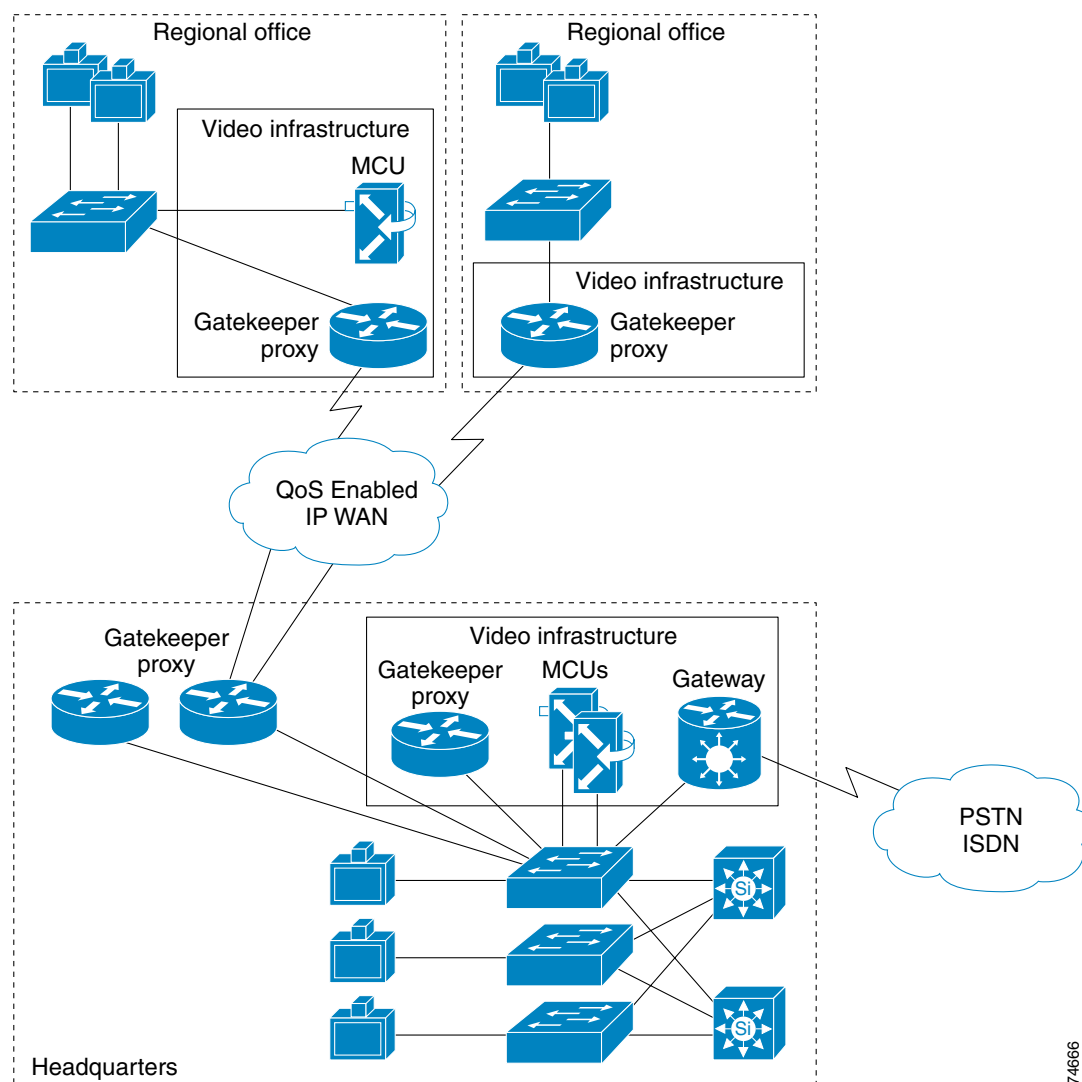
A multi-zone WAN model consists of the WAN environment and three or more videoconferencing terminals per remote site. (This model is based on a T1 WAN link.) Multi-zone WAN deployments are found in large enterprises and state-based distance-learning networks. Remote sites containing three or more video terminals are managed by either a centralized or local gatekeeper (local gatekeeper is recommended). The gatekeeper manages bandwidth within the local zone and across the WAN between zones.

Currently, it is possible to manage bandwidth only in a hub-and-spoke environment with gatekeeper bandwidth controls. An intermediate gatekeeper is not aware of a call passing through its zone. Only the originating zone gatekeeper and terminating zone gatekeeper are aware of the active call. Resource Reservation Protocol (RSVP) can be used in conjunction with Differentiated Services Code Point

(DSCP) to scale larger than hub-and-spoke environments. This configuration may, however, cause issues with other applications such as IP telephony. See the appendix on Resource Reservation Protocol (RSVP) for more information.

Figure 4-3 shows each remote site running the gatekeeper and proxy on the WAN router, and dedicated routers running Hot Standby Routing Protocol (HSRP) for the gatekeeper and proxy at the central site. Two factors determine whether to use a dedicated router or a shared router for the gatekeeper and proxy:

- Is the site currently running the appropriate router software for gatekeeper and proxy support? If not, either upgrade the router software or use a dedicated router for the gatekeeper and proxy.

- What is the number of registered endpoints and simultaneous calls being processed? If there are more than 20 registered endpoints at a given site, Cisco recommends using a dedicated router. For registration numbers and CPU utilization, refer to the chapter on Cisco Video Infrastructure Components.

The deployment guidelines for a multi-zone WAN environment are similar to those for a single-zone WAN. The biggest difference is the ability to control bandwidth in the multi-zone WAN through an added classification point (gatekeeper and zone). The key elements for successful deployment of videoconferencing in a multi-zone WAN environment are:

- Traffic Classification, page 4-7
- Bandwidth Control and Call Admission Control (CAC), page 4-7
- Provisioning, page 4-7
- Priority Queuing on the WAN, page 4-8
- Entrance Criteria, page 4-8

Traffic Classification

Classify traffic at one of the following places:

- Video endpoint (Polycom, PictureTel, Tandberg, and VCON); IP Precedence 4 or DSCP AF41

- Proxy classification; IP Precedence 4 or RSVP (recommended for traffic reclassification)

- Switch port (Layer 3 switch required); IP Precedence 4 or DSCP AF41 (recommended classification method for all video endpoints)

- Router (ACL entry); IP Precedence 4 or DSCP AF41 (Due to the larger number of video terminals at each site, this option is not typically used.)

Figure 4-4 illustrates classification options for a multi-zone WAN model.

Figure 4-4   Traffic Classification Options for Multi-Zone WAN
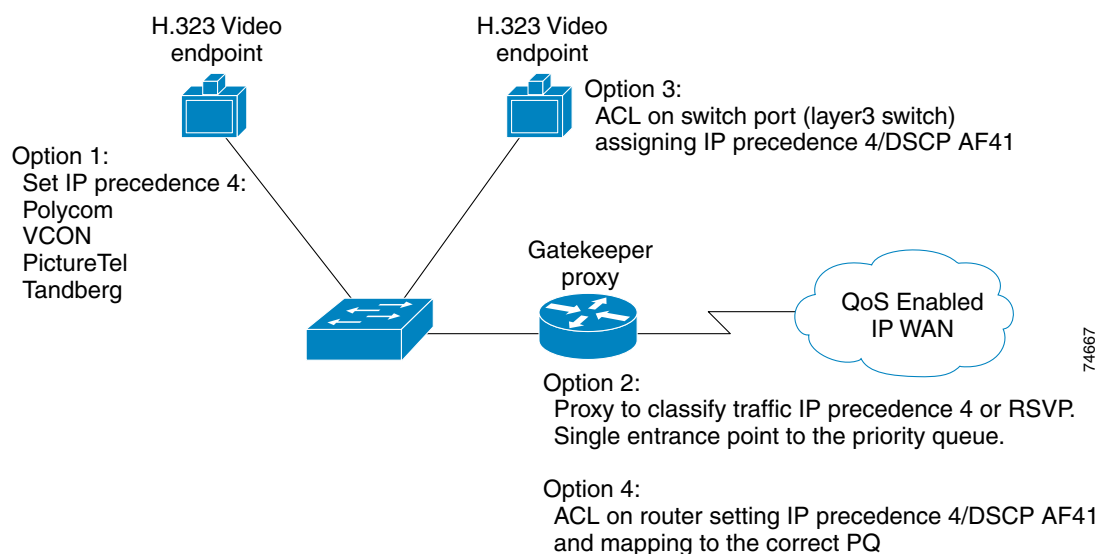
Bandwidth Control and Call Admission Control (CAC)

Because each remote site in a multi-zone WAN has its own gatekeeper and zone, bandwidth control between zones is possible. By configuring the *remote* bandwidth in each remote gatekeeper, administrators can limit the amount of available bandwidth for calls to and from the WAN. Use the global **bandwidth remote** command at remote sites to control video calls across WAN links. For more information on the gatekeeper and bandwidth control, refer to the chapter on Cisco Video Infrastructure Components.

Provisioning

Provision WAN queues based on the bandwidth limits set in the gatekeeper, and do not provision more than 33% of the link capacity for voice and video applications. Cisco recommends that voice and video traffic combined use no more than 33% of the link capacity.

Priority Queuing on the WAN

Configure multiple queues for WAN ports on routers. Videoconferencing traffic goes into a PQ that services the proxy only, or streams marked with IP Precedence 4 or DSCP AF41.

Entrance Criteria

Using the proxy allows administrators to limit access to the priority queue by configuring an ACL on the WAN router. Only video calls authenticated by the gatekeeper have access to the proxy. The ACL allows only packets received from the proxy to access the configured priority queue. The ACL prevents unauthorized users from installing a video terminal on their desk, making video calls using IP addresses, and accessing the priority queue. By restricting access to the priority queue, the configured ACL ensures that unauthorized users cannot oversubscribe the priority queue. Rouge users are serviced out of the default queue, thus ensuring video quality for authorized video terminals.

If the proxy is not used, the entrance criteria for the priority queue should be any packets with IP Precedence set to 4 or DSCP AF41. It is important to configure trust boundaries properly to prevent unauthorized traffic from accessing the priority queue.

5

WAN QoS

This chapter addresses quality of service (QoS) requirements for implementations of H.323 videoconferencing solutions over the enterprise WAN. By applying the prerequisite tools, you can achieve excellent video, voice, and data transmissions over an IP WAN, irrespective of media and even low data rates.

WAN QoS Model

Figure 5-1 illustrates the typical hub-and-spoke topology of the enterprise WAN model described in this chapter.

Figure 5-1   Enterprise WAN Model
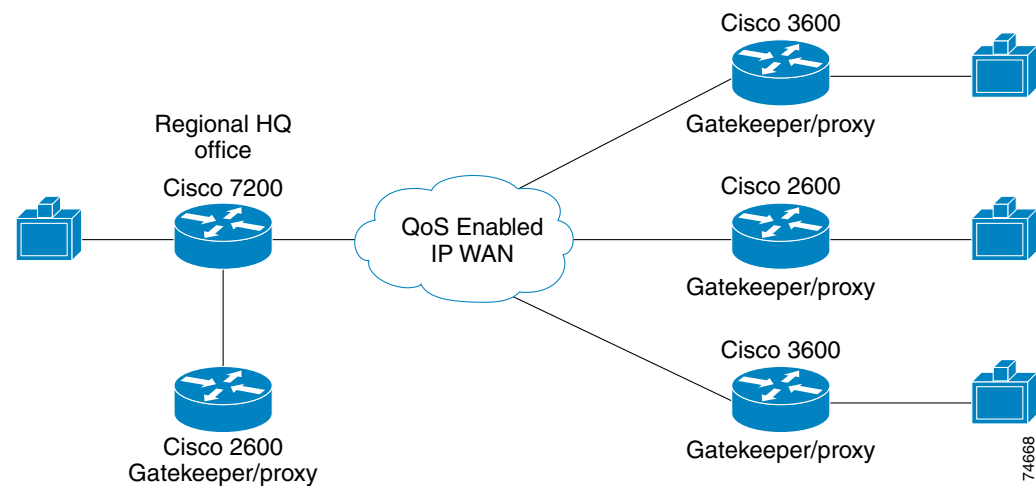
Capacity Planning

Before placing video traffic on a network, ensure that adequate bandwidth exists for all required applications. First, calculate the minimum bandwidth requirements for each major application (for example, voice, video, and data). This sum represents the minimum bandwidth requirement for any given link, and it should consume no more than 75% of the total bandwidth available on that link. This 75% rule assumes that some bandwidth is required for overhead traffic, such as routing and Layer 2 keepalives, as well as additional applications such as email and HyperText Transfer Protocol (HTTP) traffic. Figure 5-2 illustrates capacity planning on a converged network.

Figure 5-2   Capacity Planning on a Data, Voice, and Video Network

Video data rate + 20% = bandwidth required

Video   Voice   Data   Routing etc.

Video data rate
BW required

128kbps = 153kbps
384kbps = 460kbps
512kbps = 614kbps
768kbps = 921kbps
1.5Mbps = 1.8Mbps

< 33% of link capacity

< 75% of link capacity

Link capacity

74669
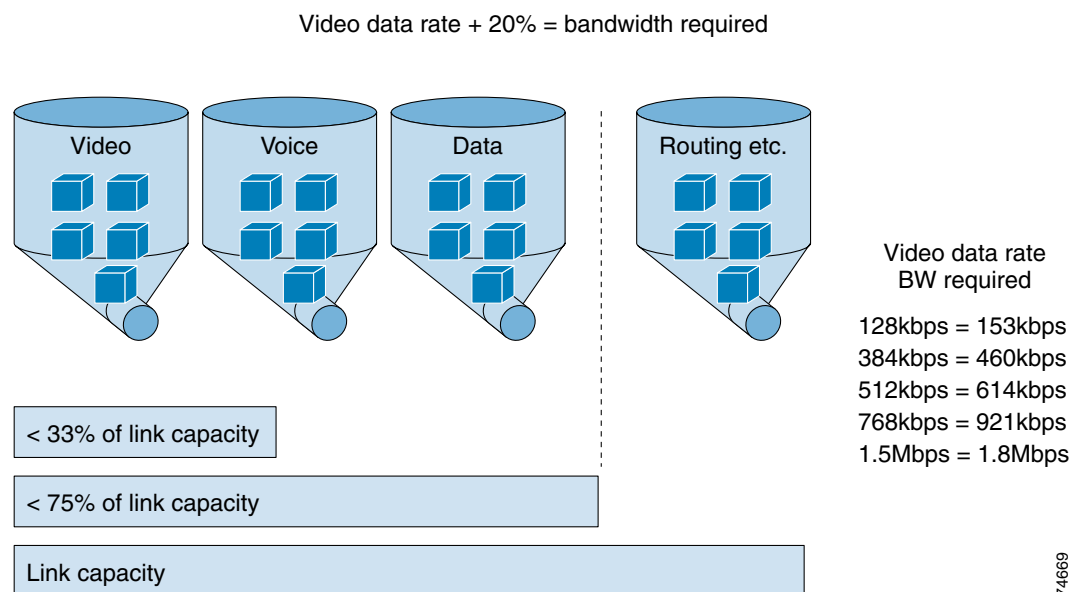
QoS Tools

This section discusses the tools used to implement QoS for H.323 videoconferencing over an enterprise WAN. These tools include:

- Traffic Classification, page 5-3
- Proxy Usage, page 5-3
- Traffic Prioritization, page 5-3

This section concludes with a summary of best practices for each of the applicable data link protocols.

Traffic Classification

Before traffic can be handled according to its unique requirements, it must be identified or labeled. There are numerous classification techniques, including Layer 3 schemes such as IP Precedence or Differentiated Services Code Point (DSCP).

In many cases, traffic classification is done at the edge of the network by the video terminal or an Ethernet switch such as the Catalyst 6000. In these cases, the trust boundary is extended to the edge of the enterprise network and resides in the access or distribution layer. For a more detailed discussion of trust boundaries, see Trust Boundaries, page 3-5.

In some cases, however, the ability to classify and define a trust boundary at the edge of the network might not exist, such as in a branch with Ethernet switches and video endpoints that cannot classify traffic. In this situation, you can implement the trust boundary and classification on the router itself by using ACL entries for small sites without a gatekeeper or by using the proxy in larger branch sites that contain a gatekeeper.

Proxy Usage

In the multi-zone WAN model, Cisco recommends that you use the proxy whenever possible. The proxy allows the classification or reclassification of video streams with IP Precedence or Resource Reservation Protocol (RSVP). The proxy also provides a single access point for the priority queue to keep unauthorized video streams from oversubscribing the priority queue. Video terminals must be registered with the gatekeeper to obtain access to the proxy. The gatekeeper is configured for a maximum video bandwidth allowed outside its local zone. This maximum bandwidth should match the amount of bandwidth provisioned for the priority queue to ensure proper queuing functionality.

Traffic Prioritization

In choosing from among the many available prioritization schemes, the major factors to consider include the type of traffic being put on the network and the wide area media being traversed. For multi-service traffic over an IP WAN, Cisco recommends low-latency queuing for the WAN. This allows up to 64 traffic classes, with the ability to use multiple queues for different traffic types, such as priority queuing behavior for videoconferencing and voice, a minimum bandwidth for Systems Network Architecture (SNA) data and market data feeds, and weighted fair queuing for other types of traffic.

Figure 5-3 shows this prioritization scheme as follows:

- Video traffic is placed into a queue with priority queuing (PQ) capabilities and is allocated a bandwidth of 460 kbps. The entrance criterion for this queue could be any video stream received from the specific IP address of a proxy or any traffic with IP Precedence set to 4. Traffic in excess of 460 kbps would be dropped if the interface becomes congested. Therefore, an admission control mechanism (such as gatekeeper bandwidth limits) must be used to ensure that this limit is not exceeded.

- SNA traffic is placed into a queue that has a specified bandwidth of 56 kbps. Queuing operation within this class is first-in-first-out (FIFO) with a maximum allocated bandwidth of 56 kbps. Traffic in this class that exceeds 56 kbps is placed in the default queue. The entrance criterion for this queue could be Transmission Control Protocol (TCP) port numbers, Layer 3 address, IP Precedence, or DSCP.

- All remaining traffic can be placed in a default queue. If you specify a bandwidth, the queuing operation is FIFO. Alternatively, if you specify the keyword **fair**, the queuing operation is weighted fair queuing (WFQ).

Figure 5-3 illustrates optimized queuing for videoconferencing on the WAN.

Figure 5-3   Optimized Queuing
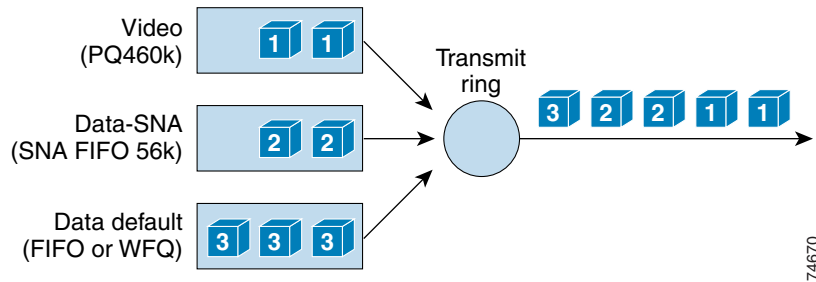
Keep in mind the following points when configuring low-latency queuing:

- For leased lines and Asynchronous Transfer Mode (ATM), the minimum system software is Cisco IOS Release 12.0(7) T.
- For Frame Relay, the minimum system software is Cisco IOS Release 12.1(2) T.

Table 5-1 gives the minimum bandwidth requirements for video and data networks. Note these values are minimum, and any network should be engineered with adequate capacity for all the applications that will use it.

Table 5-1   Minimum Bandwidth Requirements

Traffic Type	Leased Lines	Frame Relay	ATM	ATM Over Frame Relay
Video + Data Maximum video data rates up to 384 kbps	768 kbps	768 kbps	768 kbps	768 kbps
Video + Data Maximum video data rates > 384 kbps	1.544 Mbps	1.544 Mbps	1.544 Mbps	1.544 Mbps

Best Practices

Table 5-2 shows the minimum recommended software release for enterprise video over the WAN, and it includes recommended parameters for QoS tools.

Table 5-2   Minimum Recommended Software Releases

Data Link Type	Minimum Cisco IOS Software Release	Classification	Prioritization	LFI[1]	Traffic Shaping	cRTP[2]
Serial Lines	12.0(7)T	IP Precedence = 4, DSCP = AF41 for video; other classes of traffic have a unique classification	LLQ[3] with CBWFQ[4]	N/A	N/A	N/A
Frame Relay	12.1(2)T	IP Precedence = 4, DSCP = AF41 for video; other classes of traffic have a unique classification	LLQ with CBWFQ	N/A	Yes	N/A
ATM	12.0(7)T	IP Precedence = 4, DSCP = AF41 for video; other classes of traffic have a unique classification	LLQ with CBWFQ	N/A	Yes	N/A
ATM Over Frame Relay	12.1(2)T	IP Precedence = 4, DSCP = AF41 for video; other classes of traffic have a unique classification	LLQ with CBWFQ	N/A	Yes	N/A

1. Link Fragmentation And Interleaving (LFI)
2. Compressed Real-time Transport Protocol (cRPT)
3. Low Latency Queuing (LLQ)
4. Class-Based Weighted Fair Queuing (CBWFQ)

Note that cRTP is not recommended for use with IP videoconferencing. Best practices for cRTP are as follows:

- Use cRTP only with low bit rate voice codecs such as G.729. If G.711 is used as the audio codec for a voice or videoconferencing call, the statistical throughput gains achieved with cRTP are not significant enough to merit its use.

- Use cRTP only when low bit rate voice is a significant percentage of the offered load. In general, this feature is beneficial only when low bit rate voice is greater than 30% of the offered load to a circuit.

- cRTP can affect forwarding performance, and Cisco recommends that you monitor CPU utilization when this feature is enabled.

Call Admission Control

Call admission control (CAC), or bandwidth control, is required to ensure that the network resources are not oversubscribed. Calls that exceed the specified bandwidth limit are rejected to ensure video quality.

There are three schemes for providing CAC for video calls over the WAN:

- Limiting the number of video terminals

    Limiting the number of video terminals for CAC is necessary only in the single-zone WAN model. With no gatekeeper at the remote sites in this model, the only way to control the amount of bandwidth used for video across the WAN is to limit the number of video terminals at the remote sites. The priority queue at each site must then be provisioned for the maximum possible data rate of all the video endpoints at any given site. See Single-Zone WAN, page 4-2, for more information on this CAC scheme.

- Gatekeeper CAC

    This method of CAC is available only in the multi-zone WAN model.   The gatekeeper allows administrators to set bandwidth limits for inter-zone calls, intra-zone calls, or sessions. This scheme allows administrators to set an inter-zone or remote bandwidth limit, provision a priority queue for the same amount of bandwidth, and ensure the integrity of that queue. Currently, gatekeeper CAC is limited to hub-and-spoke configurations. See Multi-Zone WAN, page 4-5, for more information on this CAC scheme.

- Resource Reservation Protocol (RSVP)

    With the proxy RSVP, reservation requests can be made across the network on a per-call basis. This method allows a video network to scale larger than a hub-and-spoke environment.   However, the current implementation of the RSVP reservation request is not synchronized with the call setup. Therefore, if the call setup fails, the call will go through with no quality guarantees. This issue will be addressed in a future release of the proxy. There are several methods of implementing RSVP for CAC. See the appendix on Resource Reservation Protocol (RSVP) for more details.

C H A P T E R **6**

Dial Plan Architecture

This chapter defines and explains the key elements in designing a dial plan for an H.323 network. An H.323 video dial plan is a numbering scheme that allows H.323 video endpoints to dial other video endpoints or video services (multipoint conference unit or gateway). This chapter discusses each of these components in the context of a single-zone or multi-zone scenario.

This chapter contains the following sections:

Dial Plan Components

A well designed dial plan is a key component of a successful H.323 video network, and it is one of the first things you need to consider when designing an H.323 video network. Without a well constructed dial plan, it is impossible to scale the network.

H.323 dial plans consist of four key elements:

- E.164 address

    An E.164 address is a numeric identifier defined on each H.323 video endpoint, just as E.164 is used in telephony systems.

- H.323-ID

    An H.323-ID is an alphanumeric identifier defined on each H.323 video endpoint, and it can be used to dial the H.323 endpoint. An alias may also be used to refer to an H.323-ID. For example, email addresses are often used as H.323-IDs. H.323-IDs cannot be used to dial to the PSTN or to a Cisco IP/VC 3510 multipoint conference unit (MCU).

- Zone prefix

    A zone prefix is a numeric prefix that identifies a zone. Zone prefixes are used for inter-zone call routing, the same way an area code is used in a telephony system. Each zone in an H.323 network has one unique zone prefix. Area codes are often used as zone prefixes in H.323 networks.

- Service or technology prefix

    A service prefix is a numeric prefix used in an H.323 dialing string to access a defined service on an MCU or gateway. The Cisco gatekeeper refers to the service prefix as a technology prefix, which is also used by H.323 voice gateways. (This document refers to these prefixes as *service prefixes*.) Service prefixes are used on video gateways and MCUs to define parameter settings and to route calls for the device. On a Cisco IP/VC 352X video gateway, a service prefix defines the type of call being made (voice or video) and the data rate of the call. On a Cisco IPVC/3510 MCU, service prefixes define the data rate of the call, number of participants, and picture format. When an MCU or video gateway registers with the gatekeeper, it registers all defined service prefixes. When an H.323 endpoint uses a video gateway or MCU, the dial string must start with the service prefix followed by the PSTN number being dialed (in the case of a gateway call) or the conference ID being created or joined (in the case of an MCU call).

Table 6-1 shows the correlation between components of a video and IP telephony dial string.

Table 6-1   Correlations Between Video and IP Telephony Dial Stings

Video Dial String	IP Telephony Dial String
Service prefix	Technology prefix
Zone prefix	Area code
E.164 address	Local exchange
	Unit ID

Service Prefix Design

Service prefixes are a very important part of the dial plan. Inter-zone and intra-zone calls to an MCU or gateway are routed using the service prefix.   The single-zone and multi-zone models are very similar, and both are discussed in this section, with minor differences between them noted.

It is important to keep dial strings intuitive. For example, the models in this section use dial strings that are very similar to telephony dial strings. Dial strings are reviewed in the chapter on Call Routing.

In a single-zone network, Cisco recommends that you reserve a block of numbers for service prefixes, such as 8* for MCUs and 9* for gateways.

Note   The asterisk is a wildcard that represents any dialed digits. For example, the string 8* represents any dialed string beginning with the digit 8 followed by any number. Users do not dial the asterisk (*) when placing a call.

Cisco also recommends that you add the local area code to the service prefixes of MCUs. For example, a San Jose MCU might have a service prefix of 40880. Gateway prefixes should remain 9* to keep dial strings consistent with telephony dial plans. This service prefix structure also allows an easy migration to a multi-zone dial plan.

E.164 addresses must not overlap with service prefixes. For example, if an MCU registers with a service prefix of 40880* and a video terminal registers with 4088011212, all calls made to the video terminal would be routed to the MCU instead.

In a multi-zone network, service prefixes need to route between zones. Therefore, all service prefixes must be unique across all zones. All inter-zone or intra-zone calls are routed based on the service prefix. Cisco recommends that you design service prefixes in a multi-zone network to allow user dial strings to

be consistent. To achieve this consistency, use the different approaches outlined in the following sections for service prefixes on MCUs and gateways.   Service prefixes, E.164 addresses, and zone prefixes must not overlap, or call routing issues will arise.
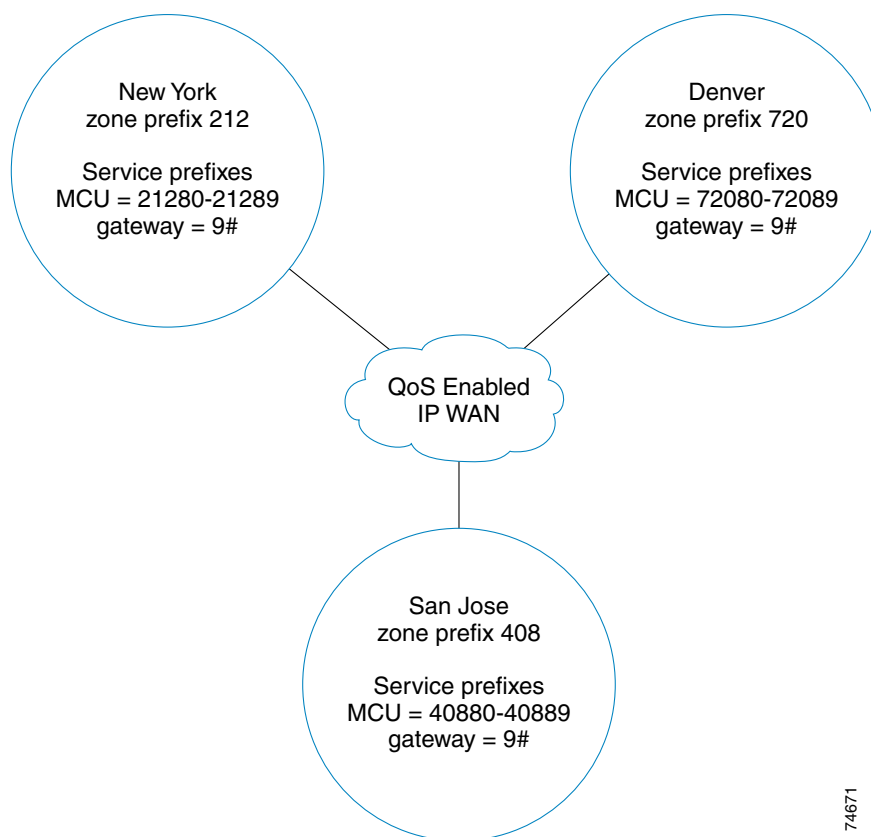
MCU Service Prefixes

MCUs must be accessible from any H.323 endpoint on the network, which means that all service prefixes in all zones must be unique. In order to accommodate unique service prefixes without reserving large blocks of numbers, Cisco recommend that you design the MCU service prefixes to be a combination of the zone prefix and a service number. This design allows all the service prefixes for MCUs to be consistent in all zones.

For example, if the reserved block of numbers is 8*, the service prefix for a 384 kbps call with five users could be 40880 in the 408 zone and 41580 in the 415 zone. The dial string for a 384 kbps conference call in zone 408 would be 40880<*conference ID*>. This design eliminates the need for hopoff entries, and it associates the service with the zone in which the service resides. (For more information on hopoffs, see the Call Routing chapter.)

Gateway Service Prefixes

Gateway services in a multi-zone network are similar to those in the single-zone model. Reserve a block of numbers for gateway services. In zones that contain gateways, off net calls always use the local gateway. For zones without a gateway, add a hopoff entry or use location request (LRQ) forwarding to route the call to a zone containing a gateway. (See the Call Routing chapter for more information regarding hopoffs, LRQ forwarding, and directory gatekeeper.)

For example, if the reserved block of numbers is 9*, the gateway service configured for all outbound calls could be 9#. Configure these service prefixes on all gateways in all zones. In zones that have a zone prefix starting with 9, ensure that the zone prefix and gateway service prefixes do not overlap. For example, if the zone prefix is 916, a gateway service prefix of 9 cannot be used in that zone, otherwise all calls in the zone would be routed to the gateway. To avoid this problem, Cisco recommends that you configure the gateway service prefixes to include a # sign, such as 9#. Figure 6-1 illustrates service prefix design in a multi-zone network.

Figure 6-1   Service Prefix Design in Multi-Zone Networks

Note   Use a # in the service prefix for gateways to ensure that calls coming in from the PSTN network do not have the ability to hair-pin back onto the PSTN through the gateway. When a user dials a # from the PSTN to the gateway, the # is treated as a delimiter and the call fails.

Single-Zone Dial Plan

Dial plans for single-zone networks are straightforward. There are a few rules that you must follow to ensure that call routing in a single zone works properly. When developing a dial plan for a single-zone network, use the following components and guidelines:

- Incoming PSTN call routing method

    As a general rule, the incoming PSTN routing method is a good place to start when designing a dial plan because it determines the number strings and the E.164 numbering structure used in the dial plan. If Direct Inward Dialing (DID) is used, each H.323 endpoint is assigned a valid E.164 directory number (DN). If interactive voice response (IVR) or TCS4 is used, the administrator can choose the E.164 number structure. Cisco recommends using 10-digit numbers for E.164 addresses because 10-digit numbers allow for an easy migration to a multi-zone dial plan. (Endpoints should be configured with a local extension, and that extension plus the zone prefix together should consist of 10 digits.) Incoming PSTN routing methods, DID, IVR, and TCS4 are covered in detail in the Call Routing chapter.

- Service prefixes

  Services prefixes must not overlap with E.164 addresses; therefore, it is a good idea to reserve a block of numbers for service prefixes. In Figure 6-2, the reserved block of numbers is 8* for MCUs, the zone prefix is 40856, and the two service prefixes for the MCU are 4085680 and 4085681. Gateway services are 9# and do not include the zone prefix.

- H.323-IDs

  H.323-IDs are alphanumeric strings used to identify an H.323 terminal. H.323-IDs are often email addresses of individual users or conference room names for room systems. Using H.323-IDs to place calls is intuitive, as long as the user-to-endpoint mapping is static. Some H.323 room systems are used in multiple conference rooms, and naming these units can be a challenge.

Figure 6-2 illustrates a single-zone design for a campus network.

Figure 6-2   Single-Zone Configuration for a Campus Network
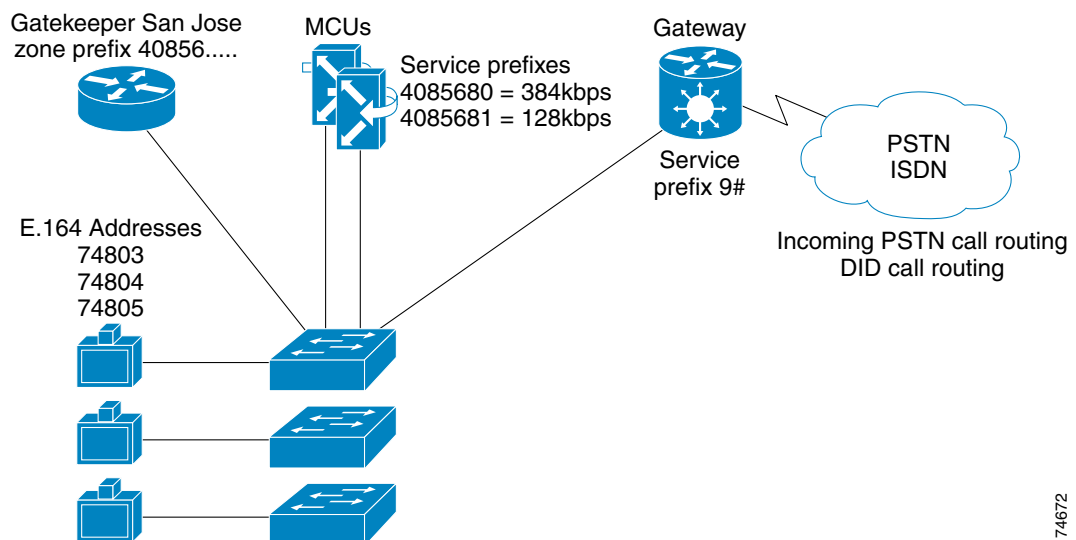
When creating a dial plan for a single zone in a WAN environment, it is always a good idea to use a numbering scheme that allows an easy migration to a multi-zone dial plan. For this purpose, Cisco recommends that you use fully qualified E.164 addresses for the video terminals.

Figure 6-3 illustrates a single-zone WAN dial plan. All video terminals, gateways, and MCUs register in one zone and are routed according to the E.164 address, H.323-ID, or service prefix registered by each device.

Figure 6-3   Single-Zone WAN Dial Plan
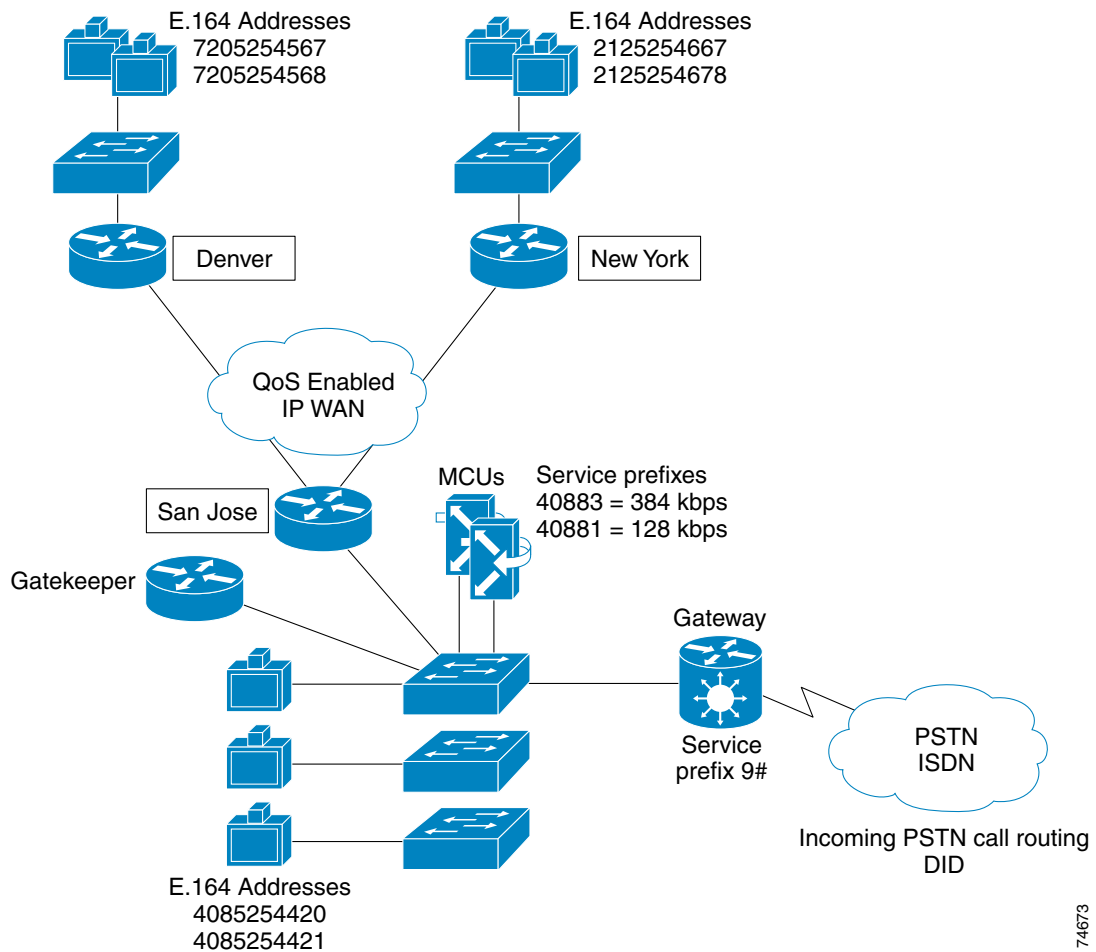

Zone Prefix Design

Zone prefixes are used in an H.323 video network to allow inter-zone call routing between H.323 endpoints, in the same way an area code is used in the PSTN. Each zone on the network must have a unique zone prefix that is used to identify the zone. Cisco recommends using the local area code for the zone prefix. For example, in Figure 6-4 there are three zones: San Jose campus zone 408*, New York 212*, and Denver 720*.

Zone prefixes can be configured with a wildcard (408*) or with dots (408.......). Cisco recommends that you use the dot method when configuring zone prefixes because this method lets you specify the exact number of digits to match, whereas the wildcard matches any number of digits. Zone prefixes can vary in length, and using more digits in the zone prefix reduces the number of available terminal addresses.

> **Note** Never use a wildcard for the zone prefix of a Directory gatekeeper zone prefix. Doing so would cause all calls, including local calls, to be forwarded to the Directory gatekeeper. Instead, use the dot (.) method to specify the zone prefix of the Directory gatekeeper. For more information on using the Directory gatekeeper, refer to the chapter on Call Routing.

Figure 6-4   Example Network with Unique Zone Prefixes
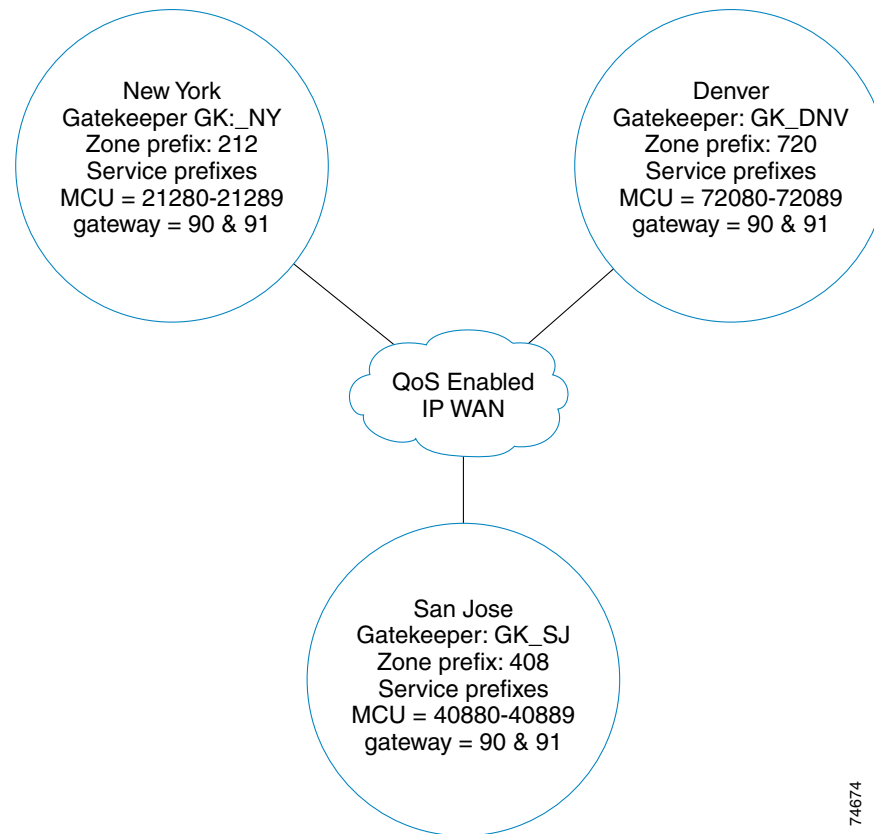
Large sites that need more than one zone can still use the local area code and expand the zone prefix to include some of the E.164 address. For example, the San Jose campus in Figure 6-5 has three zones: one configured as (40852…..), the second as (40856…..), and the third as (40857…..). Video terminals can then register with five-digit extensions, allowing extension-based dialing within the local zone, but 10-digit dialing is still required between zones.

Figure 6-5   Using a Single Area Code for Multiple Zones
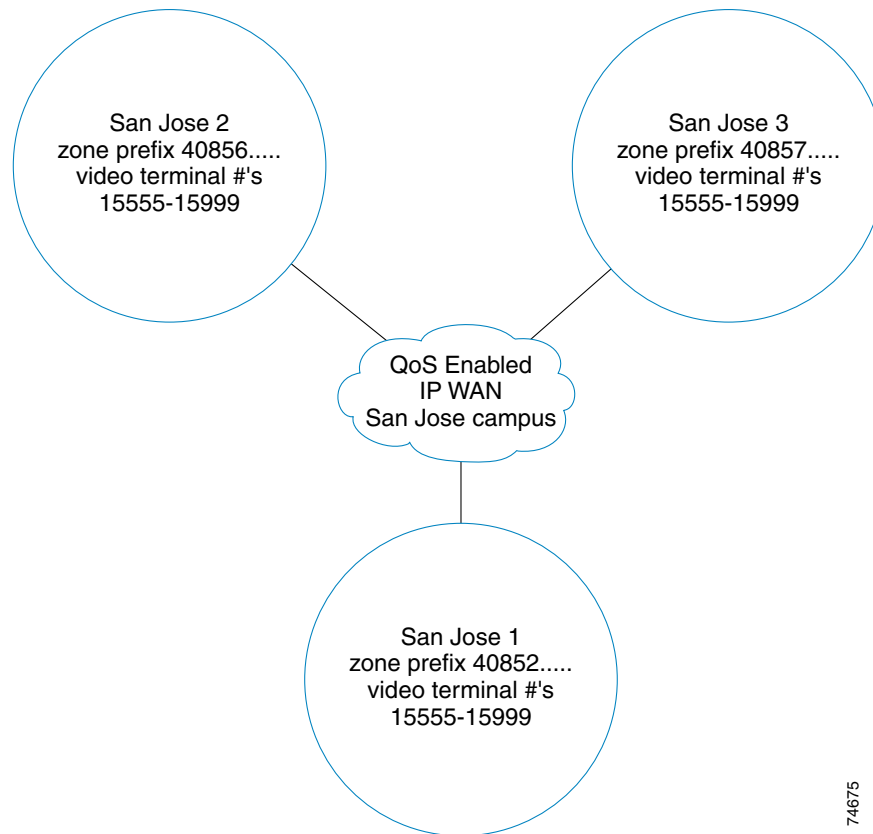
Multi-Zone Dial Plan

Dial plans for multi-zone networks have the added complexity of zone prefixes and inter-zone call routing. When developing a dial plan for a multi-zone network, consider the following components and guidelines:

- Incoming PSTN Call Routing

    Again, it is a good idea to start with the incoming PSTN routing method when developing the dial plan because the routing method determines which E.164 numbering structure is used in the dial plan. Unless there is at least one gateway in each PSTN area code, direct inward dialing (DID) is not recommended for use as the primary incoming PSTN routing method in a multi-zone network because the DID number is within one area code but the remote zone prefix may be in a different area code.

    Rather than configuring your remote zone prefixes to match the area code, which would confuse the dial plan, Cisco recommends that you place a gateway in each area code. It is important that you order enough DID numbers for all zones located in the area code serviced by the gateway. Because the Cisco gatekeeper does not support digit manipulation, it is very difficult to route incoming DID

calls between zones. There are cases in a multi-zone network where you might want to use a mix of incoming call routing methods; for example, you could use DID for endpoints but use IVR for MCU meet-me conferences.

If interactive voice response (IVR) or TCS4 is used, the administrator can choose the E.164 number structure. Cisco recommends using 10-digit numbers for E.164 addresses because 10-digit numbers allow for an easy migration to a multi-zone dial plan. (Endpoints should be configured with a local extension, and that extension plus the zone prefix together should consist of 10 digits.) Incoming PSTN routing methods, DID, IVR, and TCS4 are covered in detail in the Call Routing chapter.

- Service Prefixes

  Service prefixes must not overlap with E.164 addresses; therefore, it is a good idea to reserve a block of numbers for service prefixes. When a range of numbers is reserved for MCUs (for example, 8*), append the zone prefix to the reserved number to create a unique service prefix. For example, if the zone is 408 and the reserved block of numbers is 8*, the first service prefix might be 40880. In this case, an H.323 endpoint may not register with an E.164 address that starts with 40880-40889. If an MCU registers with a service prefix of 40880 in the zone and an H.323 endpoint registers with 4088012, all calls to 4088012 would be routed to the MCU.

- Zone Prefixes

  Zone prefixes are also very important in the development of the dial plan. Zone prefixes are much like area codes in a telephony system. Cisco recommends that you use local area codes for zone prefixes because area codes are unique, already defined, and people are familiar with them. Again, it is up to the administrator to choose the zone prefixes, but it is also important that the prefixes be intuitive and capable of growing with the network. Zone prefixes must not overlap with service prefixes, otherwise call routing issues will arise. (If you use a zone prefix plus a service prefix for MCUs, overlap with MCUs will not be an issue.) See Service Prefix Design, page 6-2, for details.

- H.323-IDs

  H.323-IDs are alphanumeric strings used to identify an H.323 terminal. H.323-IDs are often email addresses of individual users or conference room names for room systems. Using H.323-IDs to place calls is intuitive, as long as the user-to-endpoint mapping is static. Some H.323 room systems are used in multiple conference rooms, and naming these units can be a challenge.

If IVR is the chosen method for incoming PSTN call routing, observe the following guidelines:

- All systems dialing in from the PSTN must support Dual Tone Multi-Frequency (DTMF).
- Implement a private numbering plan.

If DID is the chosen method for incoming PSTN call routing, observe the following guidelines:

- Gateways must reside in each area code for zone prefix consistency.
- Use IVR to route MCU calls.

Figure 6-6 illustrates a multi-zone design using IVR, and Figure 6-7 illustrates a multi-zone design using DID.

Figure 6-6   Using IVR in a Multi-Zone Configuration
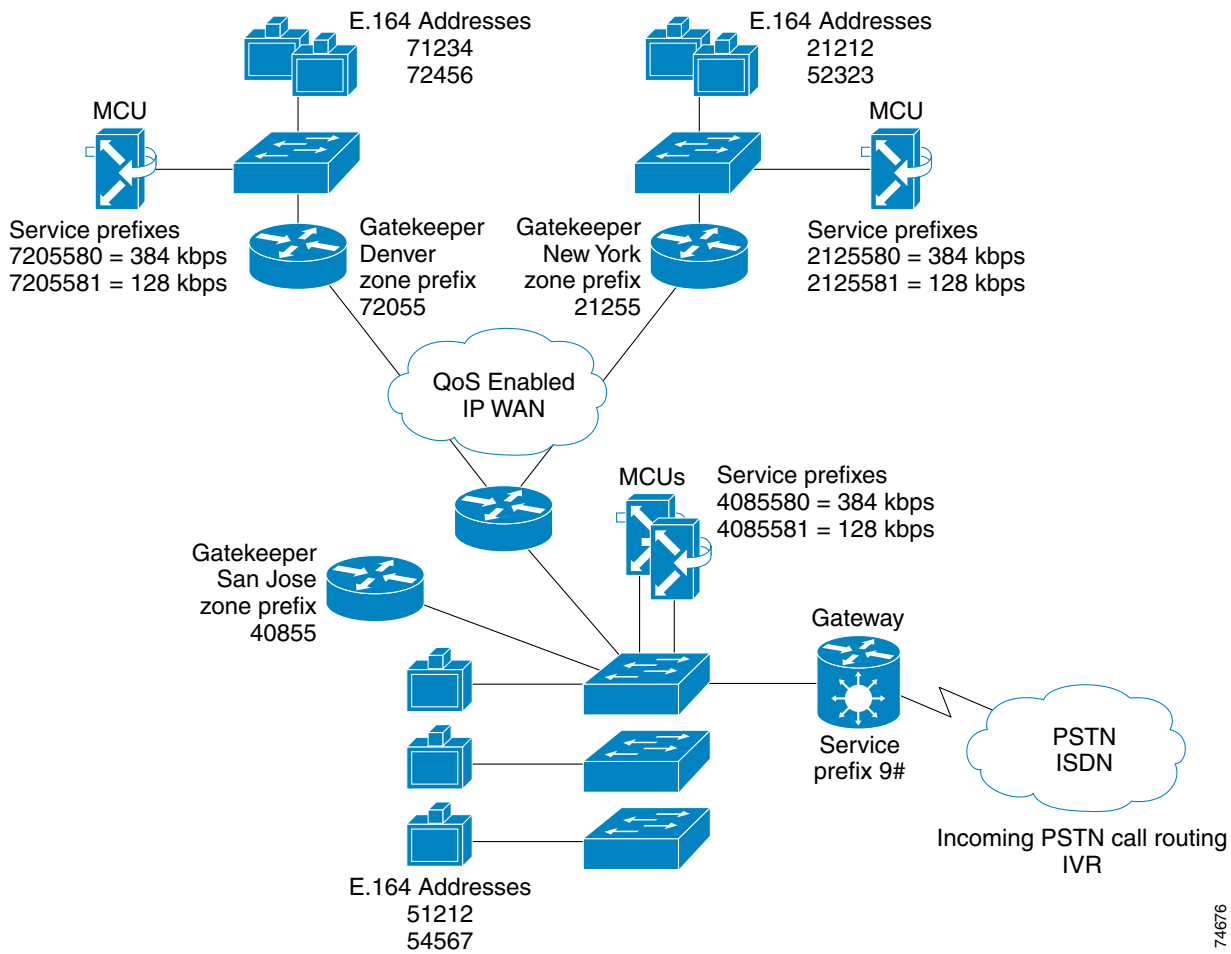
Figure 6-7   Using DID in a Multi-Zone Configuration
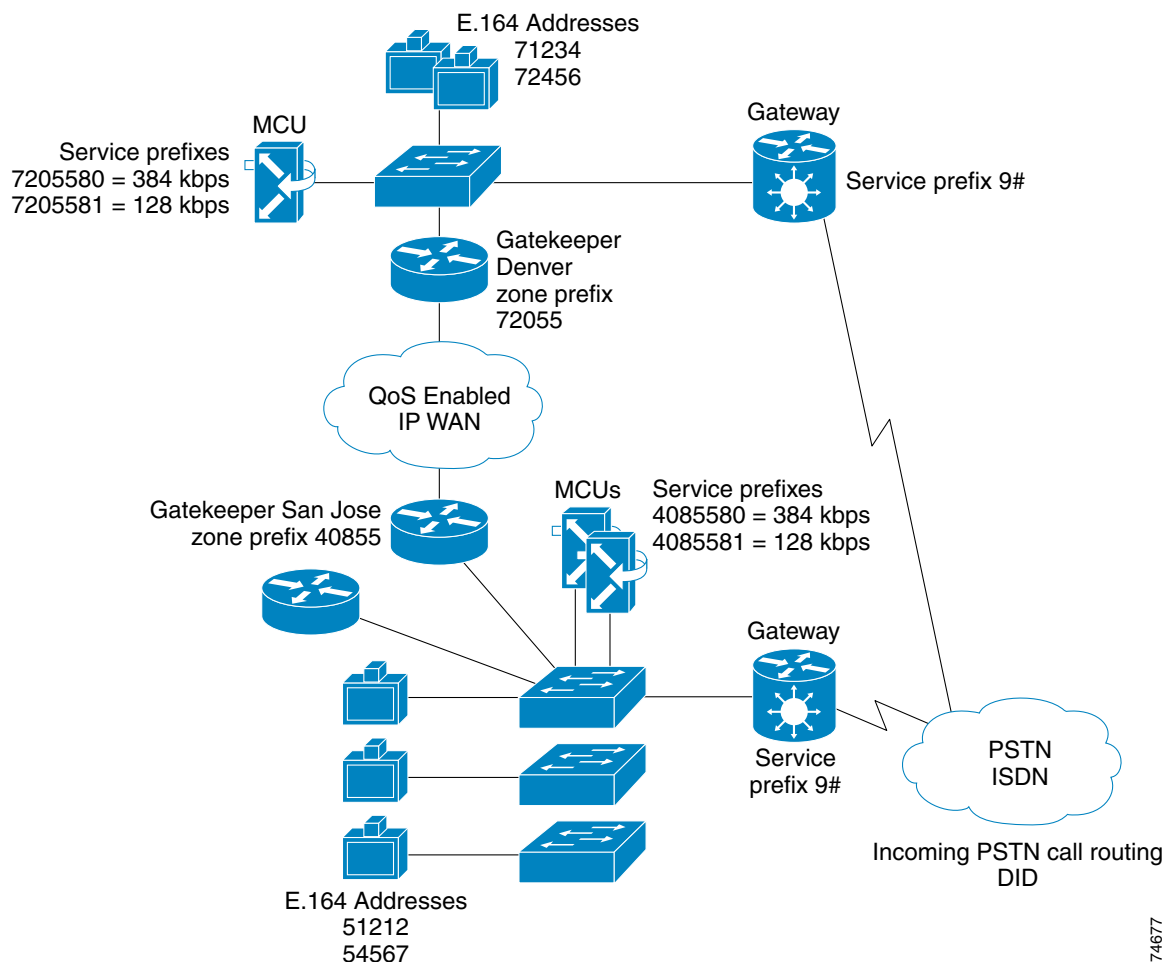
Note   In Figure 6-7, the IVR is still enabled for access from the PSTN to the MCUs.

CHAPTER 7

Call Routing

This chapter describes various call routing methods used by Cisco gatekeeper and Cisco IP/VC equipment in an H.323 video network. Calls can be routed to and from many types of devices in a variety of ways.

Call Routing Scenarios

There are four possible call routing scenarios in an H.323 network:

- H.323 endpoint to H.323 endpoint using the E.164 address

    Routing calls between H.323 endpoints is the simplest type of call routing in an H.323 network. To dial within a single zone, the endpoint initiating the call enters the E.164 address of the endpoint being called. (In most cases, the E.164 address is a video terminal extension). If the call is an inter-zone call, the initiator must enter the zone prefix and terminal extension. Using this type of dial string is similar to dialing outside the local area code in a telephony system. In multi-zone networks, service prefixes for Multipoint Conference Units (MCUs) should contain the zone prefix.

- H.323 endpoint to H.323 endpoint using H.323-ID

    To use the H.323-ID to route calls between H.323 endpoints, the calling station must dial the H.323-ID of the video terminal being called. H.323-IDs are supported only for calls from video terminal to video terminal or from video terminal to Video Terminal Adapter (VTA). When using a VTA, exercise care in addressing because some H.320 units cannot send alphanumeric strings. In these cases, E.164 addresses are the only usable route table mechanism. Between zones, Domain Name Service (DNS) may be used to reach the H.323-ID of an endpoint registered to a remote gatekeeper. To use DNS, the calling station dials *H.323-ID@Domain*, which allows the gatekeepers to resolve the remote zone destination using DNS.

- H.323 endpoint to an H.323 service (gateway or MCU)

    Routing calls from an H.323 endpoint to a service is also simple. In a single zone, an H.323 endpoint dials the service prefix followed by either the conference ID (for an MCU call) or the Integrated Services Digital Network (ISDN) telephone number of the H.320 endpoint. The service prefix can also route inter-zone calls to services, but in this case the service prefix contains the zone prefix for the MCUs, and the inter-zone calls use hopoffs for gateways.

- Incoming PSTN to H.323 endpoint or service

  You can use any of the following methods to route calls from the Public Switched Telephone Network (PSTN) to H.323 endpoints or services:

  - Multiple Subscriber Number (MSN) and Direct Inward Dialing (DID)
  - Interactive Voice Response (IVR)
  - TCS4
  - Default extension

  For more details on these routing methods, see Routing PSTN Calls to H.323, page 7-4.

Example

Figure 7-1 illustrates a multi-zone network with video terminals and services in each zone. The following dial strings apply to the scenarios in Figure 7-1:

- H.323 endpoint to H.323 endpoint:

  Intra-zone call, User1 to User2 — User1 dials 4085558072

  Inter-zone call, User1 to User3 — User1 dials 7207125543

- H.323 endpoint to H.323 endpoint using H323-ID:

  Intra-zone call, User1 to User2 — User1 dials User2@cisco.com

  Inter-zone call, User1 to User3 — User1 dials User3@cisco.com

- H.323 endpoint to service:

  Intra-zone call, User1 to H.320 system — User1 dials 9#12125551212

  Inter-zone call, User3 to H.320 system — User3 dials 9#12125551212

  Gateway calls always use the local gateway if one is present.

- PSTN endpoint to H.323 endpoint or service using IVR (see Routing PSTN Calls to H.323, page 7-4):

  Intra-zone call, H.320 system to User1 — H.320 system dials 4085552000, IVR answers, and the H.320 system enters 40855558071. Or, if DID is enabled to User1, H.320 system dials 4085558071 directly.

  Intra-zone call, H.320 system to 408, and MCU conferences 40880123 — H.320 system dials 4085552000, IVR answers, and the H.320 system enters 40880123. (For DID and IVR deployments, IVR should be used for routing calls to an MCU conference.)

  Intra-zone call, H.320 system to User3 — H.320 system dials 4085552000 followed by 7207125543. (If there is a gateway in the 720 area code, DID could be enabled to User3 and IVR could be used to reach User3 in the 720 zone instead of having to dial the 408 gateway.)

  Inter-zone call, H.320 system to MCU, with conference to 72080111 — H.320 system dials 408555200 followed by 72080111. (For DID and IVR deployments, IVR should be used for routing calls to an MCU conference.)
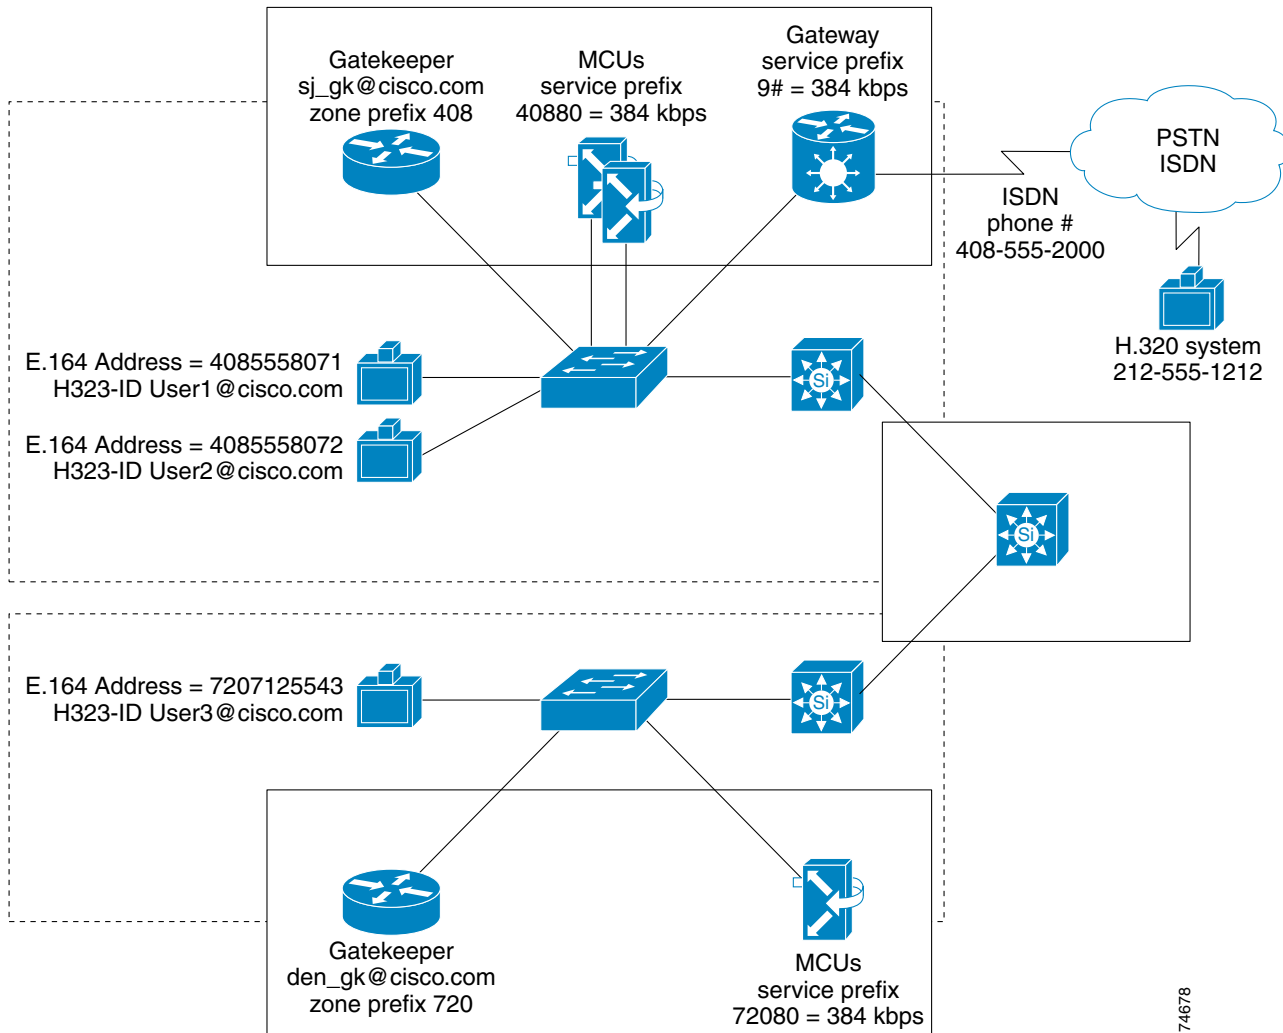
Figure 7-1   Call Scenarios in an Example Multi-Zone Network

Table 7-1 shows the dial strings for the intra-zone call types and Table 7-2 shows the dial strings for the inter-zone call types in Figure 7-1.

Table 7-1   Dial Strings for Intra-Zone Calls

Call from:	Call to:	Dial String
H.323 Endpoint	H.323 Endpoint	*<E.164 address>* or *<H.323-ID>*
H.323 Endpoint	Service	*<Service Prefix> <Conference ID or PSTN E.164 address>*
Service	H.323 Endpoint	*<E.164 address>*
Service	Service	*<Service Prefix> <Conference ID or PSTN E.164 address>*

Table 7-2   Dial Strings for Inter-Zone Calls

Call from:	Call to:	Dial Sting
H.323 Endpoint	H.323 Endpoint	*<Zone prefix + E.164 address> or <H.323-ID>*
H.323 Endpoint	Service	*<Zone Prefix and/or Service Prefix> <Conference ID or PSTN E.164 address>*
Service	H.323 Endpoint	*<Zone Prefix and/or E.164 address>*
Service	Service	*<Zone Prefix and/or Service Prefix> <Conference ID or PSTN E.164 address>*

Routing PSTN Calls to H.323

There are several methods for routing calls from the PSTN to H.323 endpoints and services:

An H.323 video network can use one or more of these available routing methods, and each routing method has advantages over the others in different situations.

Multiple Subscriber Numbering (MSN) with Direct Inward Dialing (DID)

Multiple Subscriber Numbering (MSN) is a group of phone numbers assigned to a single ISDN Basic Rate Interface (BRI) line. MSN is not available in most regions of the United States, Canada, or South America, but it is widespread in Europe.

Direct Inward Dialing (DID) is supported on Primary Rate Interface (PRI) lines. DID allows multiple directory numbers to be assigned to a single PRI line. DID is supported throughout the United States and Europe.

Interactive Voice Response (IVR)

IVR is a widely deployed automated call answering system that responds with a voice menu, allowing the H.320 endpoint to access H.323 endpoints by entering an extension from a keypad. When an incoming call arrives, the IVR answers the call and asks for the extension. The caller enters an E.164 address, and the call is transferred to the appropriate H.323 endpoint. Using IVR requires the calling H.320 endpoint to support Dual Tone Multi-Frequency (DTMF). Most legacy conference room systems support DTMF. (See Table 7-3 for DTMF support.)

TCS4

TCS4 is a special method for routing incoming H.320 video calls by using extensions. TCS4 allows direct extension dialing to an H.323 endpoint on the LAN, which register to the gatekeeper with an E.164 address. When an H.320 endpoint dials a gateway's phone number followed by a TCS4 delimiter and the E.164 address, the call is routed directly to the corresponding H.323 endpoint. TCS4 is new, and only some of the H.320 endpoints permit the user to enter a TCS4 extension when dialing. (See Table 7-3 for TCS4 support.) Due to the limited support for the TCS4 standard in H.320 devices, TCS4 is not frequently used for incoming call routing and, therefore, DID or IVR are typically better choices.

Default Extension

Specifying a default extension in the gateway forces all calls received by the video gateway to be routed directly to a default E.164 address. A default extension can also be used in conjunction with any of the routing methods mentioned previously. If the call cannot be routed by one of the previous methods, the call is then forwarded to the default E.164 address.

Routing Inbound PSTN Calls in a Single-Zone Network

You can use any of the routing methods to route calls from the PSTN to H.323 endpoints and services in a single-zone network. Each method offers the following functions and numbering structures:

- DID

    Using DID in a single-zone network allows administrators to order blocks of DID numbers and assign each H.323 endpoint a DID number to be used as its E.164 address. This method allows H.320 users and H.323 users to dial the same number to access an H.323 endpoint. (This method assumes that, in most cases, the carrier sends 10 digits.) DID can also be used for MCU conferences; however, in order to route calls to an MCU service in a zone, the zone prefix, the service prefix, and the conference ID combined must match one of the DID numbers associated with the ISDN line. This method disables the use of ad-hoc conference IDs created by the users on the MCU, but it may be preferable over using IVR to reach these conferences. This method does, however, require that the conference ID match the statically registered directory number. DID call routing is very desirable because the dial strings are exactly the same as those used in telephony systems, but routing H.323 service prefixes can become complex when using DID call routing.

- IVR

    IVR allows administrators to define the dial plan. E.164 addresses can be four-digit extensions or 10-digit directory numbers. (The video terminal extension and the zone prefix combined should be 10 digits). When routing incoming PSTN calls with IVR, the call initiator must dial the directory number of the PRI gateway and enter the E.164 address or service prefix dial string after the IVR has answered. IVR requires DTMF support on the dialing endpoint, but some older H.320 systems do not support DTMF. (See Table 7-3 for TCS4 and DTMF support.)

- TCS4

    When using TCS4 to route incoming calls, the administrator again defines the numbering plan. When using TCS4, the initiator dials the directory number of the gateway, a TCS4 delimiter, and the E.164 address or service. The delimiter must be configured in each video gateway, and the options are # or *. Using TCS4 requires the dialing endpoint to support TCS4. (See Table 7-3 for TCS4 and DTMF support.) TCS4 is not a commonly used routing method at the present time.

- Default Extension

    A default extension is usually used in special cases such as call center applications or to route calls to a single H.323 video terminal.

Note   All of these dial-in methods are mutually exclusive, and you can implement multiple incoming routing methods on the same gateway. If an incoming PSTN call arrives at a gateway supporting all of the routing methods, the gateway first tries to resolve the address using the routing methods in this order: DID, IVR, TCS4, and default extension. (A DID environment is a typical example that would use a gateway supporting multiple incoming call routing methods.) Cisco recommends that you do not assign a DID number for ad-hoc MCU calls; instead, use IVR to route incoming calls to an MCU and use DID to route incoming calls to video terminals.

Table 7-3 summarizes partner product capabilities as they relate to interoperability with the Cisco IP/VC gateways.

Note The information included in Table 7-3 is subject to change, and you should contact the product manufacturer directly for updated information.

Table 7-3   Partner Product Capabilities and Interoperability

Partner and Product	DTMF	Software Version	TCS4	Software Version.
PictureTel				
PictureTel I-Series	Yes	Any	Yes	Any
PictureTel Concorde 4500	Yes	6.1	Yes	6.3
PictureTel Venue 2000	Yes	1.3	No	
Proshare Video System 500	Yes	5	No	
Teamstation	Yes	4	No	
SwiftSet II	Yes	1.04	No	
VTEL				
Galaxy 725	Yes	1	Yes	1
Galaxy 755	Yes	1	Yes	1
Galaxy 2500	Yes	1	Yes	1
Galaxy 5500	Yes	1	Yes	1
Gateway	Yes	1.2	Yes	1.1
Smart Station	Yes	5	No	
WG500	Yes	5	No	
ESA TC1000	No		No	
ESA TC2000	No		No	
ESA TC5000	No		No	
Smart Link MCS	No		No	
Settop 250	No		No	
VCON				
Escort 25	Yes	4.01	Yes	4.01
Cruiser 75	Yes	4.01	Yes	4.01
Cruiser 150	Yes	4.01	Yes	4.01
Cruiser 384	No		Yes	4.01
Media Connect 8000	No		Yes	4.01
Media Connect 6000	No		Yes	2

Table 7-3   Partner Product Capabilities and Interoperability (continued)

Partner and Product	DTMF	Software Version	TCS4	Software Version.
Tandberg				
Vision Classic	Yes	K 2.8	No	
Vision 600	Yes	B 1.3	No	
Vision 770	Yes	B 1.3	No	
Vision 800	Yes	C 4.0	No	
Vision 1000	Yes	B 1.3	No	
Vision 2000	Yes	B 4.3	No	
Vision 2500	Yes	C 4.0	No	
Vision 5000	Yes	C 4.0	No	
Tandberg 500	Yes	B 4.3	Yes	B 4.3
Tandberg 550	Yes	B 4.3	Yes	B 4.3
Tandberg 800	Yes	B 4.3	Yes	B 4.3
Tandberg 880	Yes	B 4.3	Yes	B 4.3
Tandberg 1000	Yes	B 4.3	Yes	B 4.3
Tandberg 2500	Yes	B 4.3	Yes	B 4.3
Tandberg 6000	Yes	B 4.3	Yes	B 4.3
Tandberg 7000	Yes	B 4.3	Yes	B 4.3
Tandberg 8000	Yes	B 4.3	Yes	B 4.3
Zydacron				
Z350 Windows 98	Yes	2.2	Yes	2.2
Z350 for NT	Yes	2.3	Yes	2.3
OnWAN240/250 Win 95	Yes	2.04	Yes	2.04
OnWAN250 for OS/2	Yes	2	Yes	2
OnWAN240/250 Win NT	Yes	2.02	Yes	2.02
Z220 Plus for Win 95	Yes	2	Yes	2
Z360 for Win NT	Yes	1.1	Yes	1.1
Polycom				
ViewStation SP	Yes	5.X	Yes	5.X
ViewStation FX	Yes	6.X	Yes	6.X
ViaVideo	Yes	1.5X	Yes	1.5X

Routing Inbound PSTN Calls in a Multi-Zone Network

Call routing in a multi-zone network becomes more complicated due to the use of zone prefixes and inter-zone routing of service prefixes. For example, the executive staff of a company can be assigned to a single zone to keep the dial strings simple, and DID can be implemented in the executive zone. Other zones on the network might use IVR due to the lack of video gateway services in every zone. By using the dial plans outlined in this document, you can keep the dial strings consistent across all zones.

You can use any of the following routing methods to route calls from the PSTN to H.323 endpoints and services in a multi-zone network:

- DID

  If you use DID to route calls from the PSTN to the H.323 endpoints and services, each E.164 address and service is a valid DID number associated with a PRI line attached to a Cisco IP/VC gateway. In order to use DID in a multi-zone network where zones may reside in different geographic regions, PSTN area codes and "data" boundaries require a video gateway in each area code.

- IVR

  IVR allows administrators to define the number structure of the dial plan. E.164 addresses can be four-digit extensions or 10-digit directory numbers. (The video terminal extension and the zone prefix combined should be 10 digits). IVR is the easiest method for routing incoming PSTN calls in a multi-zone network. When using IVR, calls terminate at the gateway, and then the caller enters the E.164 address or service prefix. If the call is in the local zone, only the E.164 address or service prefix is needed. If the call is going to another zone, the caller enters the zone prefix followed by the E.164 address.   Services hosted on a remote MCU are dialed the same way (that is, zone prefix + service prefix + conference ID). IVR requires DTMF support from the dialing endpoint, but some older H.320 systems do not support DTMF. (See Table 7-3 for TCS4 and DTMF support).

- TCS4

  When using TCS4, the dial string from the H.320 endpoint contains the ISDN directory number for the gateway followed by a TCS4 delimiter and the E.164 address of the H.323 endpoint. If the incoming call is destined for an H.323 endpoint outside of the local zone, the zone prefix must be added to the dial string. Using TCS4 requires the dialing endpoint to support TCS4. (See Table 7-3 for TCS4 and DTMF support). Because TCS4 is not a commonly used routing method, Cisco recommends IVR instead.

- Default Extension

  A default extension is usually used in special cases such as call center applications or to route calls to a single H.323 endpoint.

Routing Inter-Zone Calls Using Hopoff Statements

You can add hopoff statements in the gatekeeper to route calls between zones without using a zone prefix. Hopoffs are used for routing gateway services because the service has no association with the zone where the gateway resides. To create a common dial plan, strategically deploy gateways in major sites and use hopoff statements in all smaller *stub* zones that do not contain a gateway. Then users, regardless of what zone they are in, can dial a common service prefix to access the outside world.

Use of the hopoff statement eliminates the need for users in a stub zone to dial the zone prefix of the zone that contains the gateway. Hopoffs override the gatekeeper parse order and direct calls with the defined service to a specific zone. Use the following command syntax to configure hopoff statements in the gatekeeper:

```
gw-type-prefix <prefix #> hopoff  <gatekeeper name>
```

MCUs do not require hopoff statements because the zone prefix is always embedded in the service prefix.

Note When creating multiple zones on a single router and registering MCUs or gateways in any of the zones, enter a hopoff command for each service prefix. Routing of service prefixes between local zones also requires a hopoff.

In Figure 7-2, District Site A and District Site B have hopoffs configured to forward all gateway calls (service prefix 9#) to the regional site. These hopoff statements forward calls matching 9#* to the regional site.

Figure 7-2   Inter-Zone Routing with Hopoff Statements Configured

Cisco IP Videoconferencing Solution Reference Network Design Guide
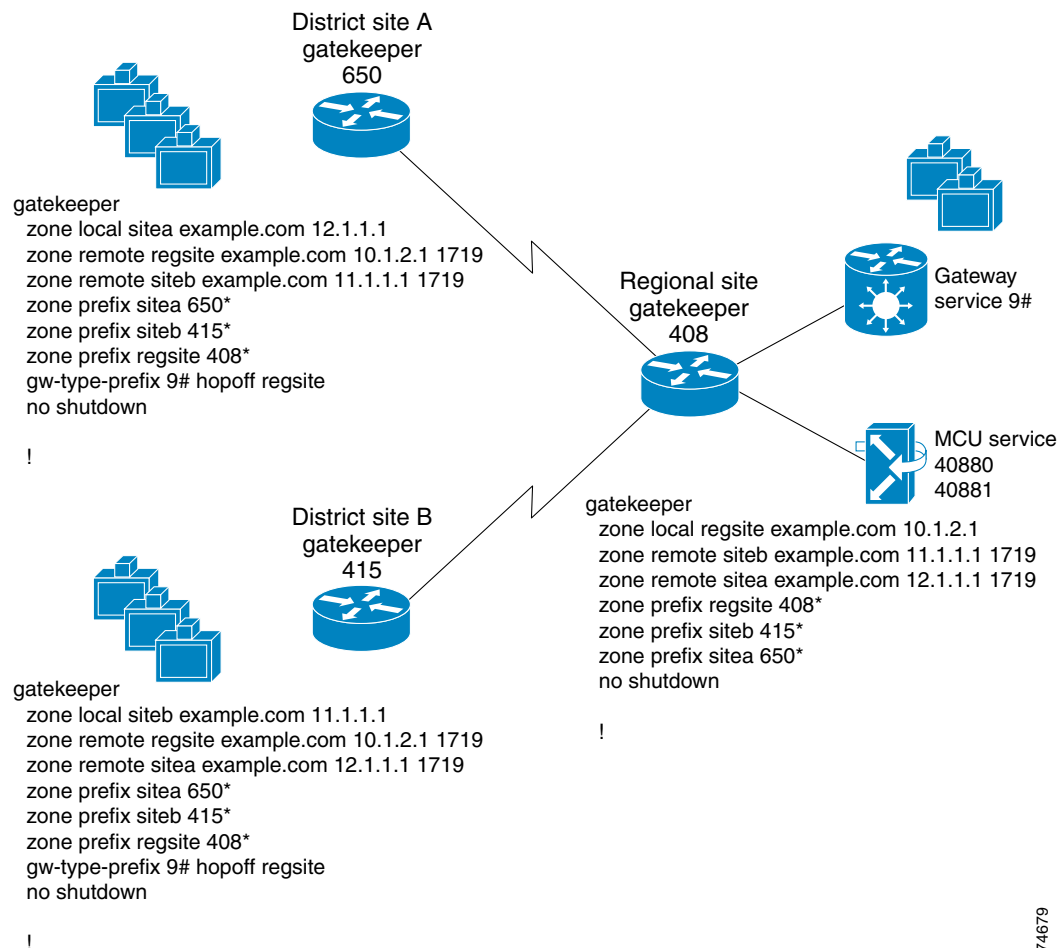
Routing Inter-Zone Calls Using a Directory Gatekeeper

Currently, there is no gatekeeper protocol that allows gatekeepers to update each other with routing information. This limitation implies a full-mesh topology, where every gatekeeper must be statically configured to know about every other gatekeeper to which it is going to route calls. In effect, all gatekeepers must be known to each other. This poses scalability problems when a new zone or service is added because the administrator must add an entry in every gatekeeper for the new zone or service.

By using a directory gatekeeper and Location Request (LRQ) forwarding, a hierarchical gatekeeper design can limit the administrative overhead in a large multi-zone network. LRQ forwarding allows an administrator to create a directory gatekeeper that maintains all zone prefixes for the network or subset of the network. In Figure 7-3, sites A, B, C, and D are configured to forward all LRQs that cannot be resolved locally to directory gatekeeper sites (E and F).

Figure 7-3   *Inter-Zone Routing with Directory Gatekeepers*
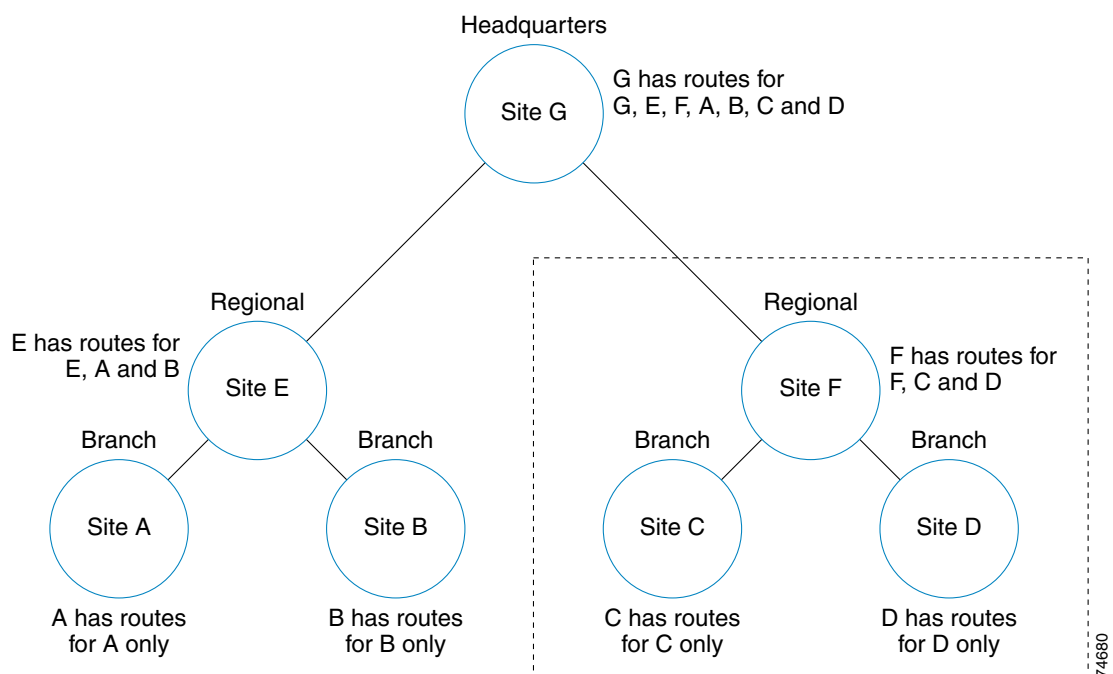
In Figure 7-4 there are three zones, two district zones and a regional zone that has a connection back to headquarters. Each district zone contains information about its local zone only. The command line **zone prefix regsite ……….** routes any call placed with 10 digits, but not matched in the local zone, to the regional site.

The regional site contains the routing information for its own zone as well as the two district zones associated with it. Zone prefix and hopoff statements are added to the regional site as the zones are added to the network. There is also a **zone prefix hqsite ……….** entry in the regional gatekeeper that forwards any 10-digit call with no match to the headquarters gatekeeper. If LRQs are going to be forwarded past the directory gatekeeper, an **lrq forward-queries** entry must be added to the gatekeeper, otherwise LRQs will not be forwarded past the directory gatekeeper. (LRQ forwarding has a maximum limit of seven hops.) This model can be expanded in a large network to make an H.323 network more manageable.

Figure 7-4   Directory Gatekeeper Example

District site A
MCM/gatekeeper
proxy

gatekeeper
  zone local sitea example.com 12.1.1.1
  zone remote regsite example.com 10.1.2.1 1719
  zone prefix sitea 650.......
  zone prefix regsite..........
  no shutdown

!

District site B
MCM/gatekeeper
proxy

gatekeeper
  zone local siteb example.com 11.1.1.1
  zone remote regsite example.com 10.1.2.1 1719
  zone prefix siteb 415.......
  zone prefix regsite..........
  no shutdown

!

To
headquarters

Regional
site
MCM/gatekeeper
proxy

gatekeeper
  zone local regsite example.com 10.1.2.1
  zone remote siteb example.com 11.1.1.1 1719
  zone remote sitea example.com 12.1.1.1 1719
  zone remote hqsite example.com 10.1.1.1 1719
  zone prefix regsite 408.......
  zone prefix siteb 415*
  zone prefix sitea 650*
  zone prefix hqsite..........
  lrq froward-queries
  no shutdown

!

74681

Note When configuring the directory gatekeeper, do not use the wildcard (*) as the directory gatekeeper zone prefix, otherwise calls will not be routed properly. For example, the command **zone prefix regsite *** will route all calls, even local ones, to the directory gatekeeper.

In Figure 7-4 the directory gatekeeper entry is **zone prefix regsite ....**, which allows any 10-digit dial string that is not matched locally to be forwarded to the directory gatekeeper. If there is a need for users to dial 11- or 12-digit dial strings, you can enter multiple zone prefix entries for the directory gatekeeper. Deployments that support international locations are more likely to require multiple zone prefix entries for the directory gatekeeper.

If a root zone contains a video gateway, and multiple directory gatekeeper zone prefixes are configured, you might have to add a hopoff to the configuration. If any of the directory gatekeeper zone prefix lengths match the dial string minus the service prefix, the call is forwarded to the directory gatekeeper. For example, if a local gateway service prefix is 9#, PSTN calls will be either nine digits (local calls) or 12 digits (long distance) including the service prefix.

When the gatekeeper starts to parse the dial string, it strips the service prefix and starts looking for a match. In the preceding example, local PSTN calls are parsed on seven digits and long distance PSTN calls are parsed on 11 digits. If the gatekeeper configuration contains a directory gatekeeper entry with

Cisco IP Videoconferencing Solution Reference Network Design Guide

seven dots or 11 dots, a hopoff is needed. The same rule applies to MCUs, but in most cases MCU calls are parsed on five digits or less, while most directory gatekeeper zone prefix entries are matched on 10 digits or more.

Example 7-1 illustrates the configuration of a root zone containing multiple directory gatekeeper zone prefix entries and a hopoff for 9#. The reason for the hopoff is to eliminate long distance calls (which are parsed on 11 digits) from matching the DGK zone prefix entry with 11 dots. Figure 7-5 and Figure 7-6 illustrate the parse order for Admission Requests (ARQs) and Location Requests (LRQs) in the Cisco gatekeeper.

Example 7-1   Configuration of Root Zone with Multiple Director Gatekeepers

```
gatekeeper
 zone local HKG cisco.com 10.1.3.1
 zone remote APAC_DGK cisco.com 10.1.2.1
 zone prefix HKG 852.......
 zone prefix APAC_DGK ..........
 zone prefix APAC_DGK ...........     (This entry matches long distance PSTN calls to a gateway)
 zone prefix APAC_DGK ............
 gw-type-prefix 9#* hopoff HKG     (This entry keeps all 9# dial strings in the HGK zone)
 no use-proxy HKG default inbound-to terminal
 no use-proxy HKG default outbound-from terminal
 bandwidth remote 1000
 no shutdown
```

Figure 7-5   Gatekeeper Address Resolution for ARQ
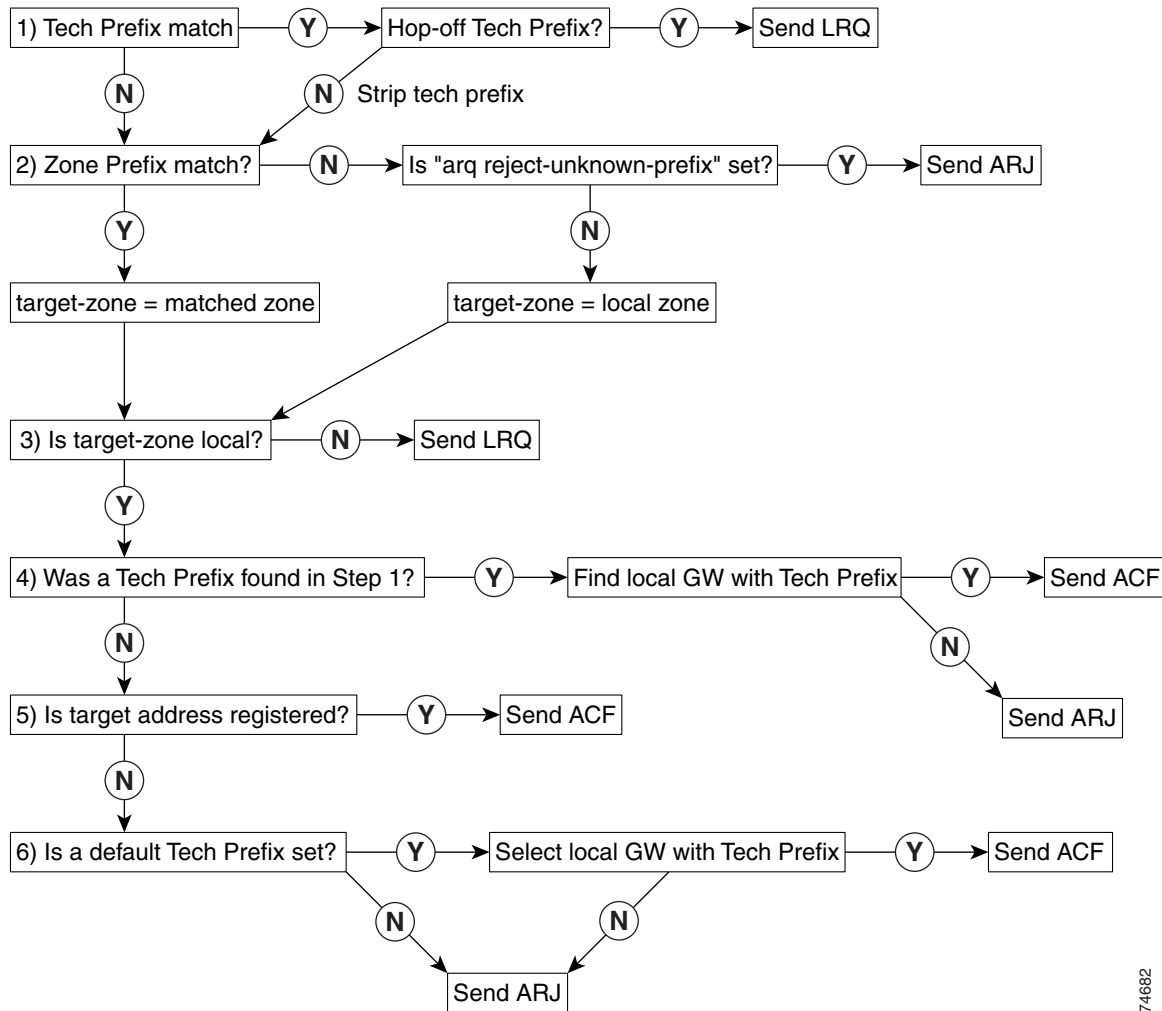
Figure 7-6   Gatekeeper Address Resolution for LRQ
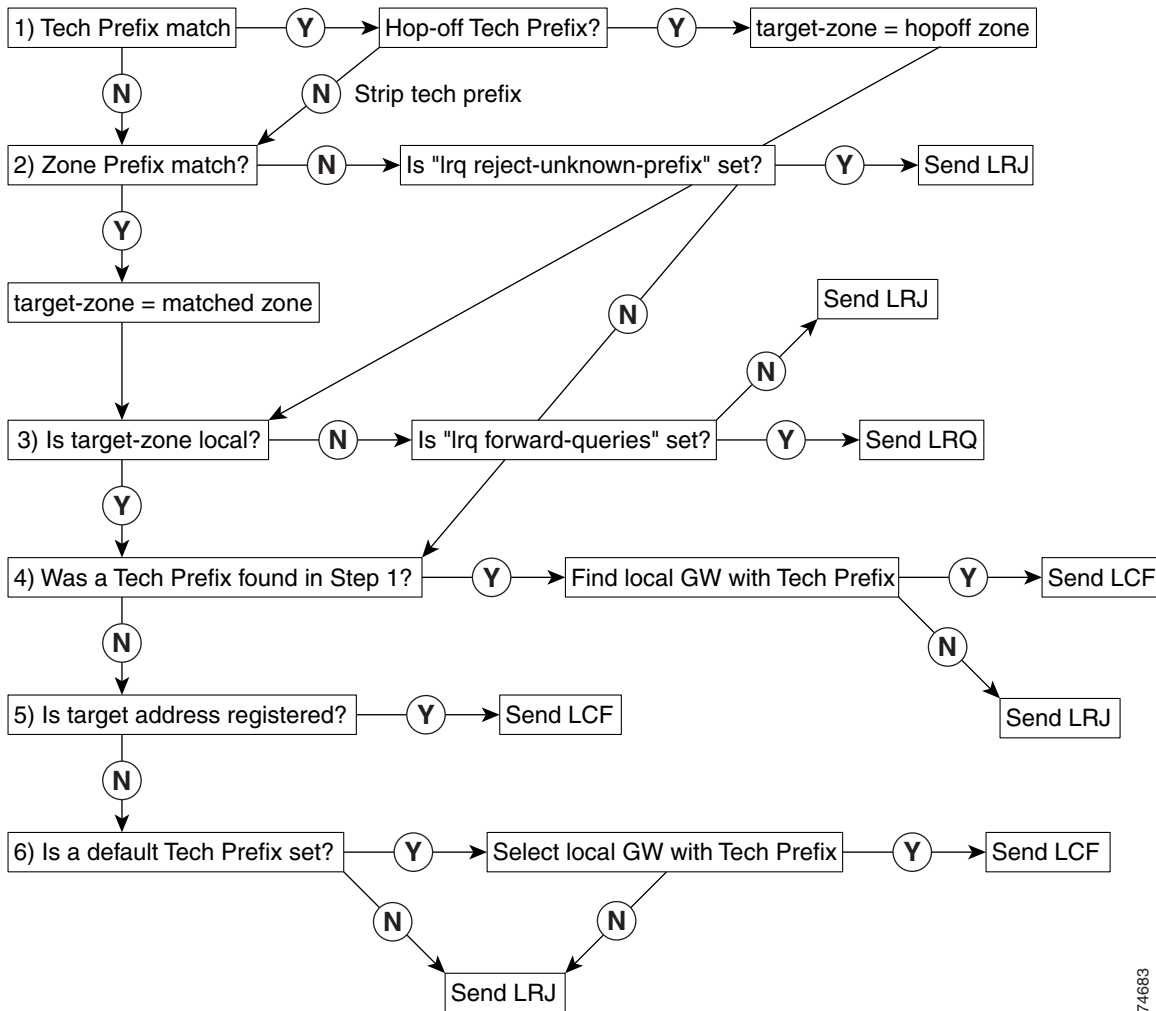
Cisco Video Infrastructure Components

This chapter describes the Cisco IP Videoconferencing (IP/VC) infrastructure components and network design considerations relating to those Cisco IP/VC components. The Cisco video infrastructure consists of the following products:

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

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

Cisco IP/VC 3540 MCU and Gateway

The Cisco IP/VC 3540 is a chassis-based unit supporting MCU modules with a capacity of up to 100 128-kbps video calls, dual PRI H.320 gateway modules, and T.120 application server modules. The Cisco IP/VC 3540 supports both voice activated and continuous-presence conferencing. Each Cisco IP/VC 3540 MCU blade supports a 10/100-Mbps Ethernet interface, H.261 and H.263 video codecs, G.711, G.722, and G.728 voice codecs, and video data rates from 128 kbps to 1.5 Mbps. Configuration of the MCU depends on the desired function and network layout.

The Cisco IP/VC 3540 is designed to support a large number of scheduled and ad hoc conferences. There are three MCU blades available for the Cisco IP/VC 3540 MCU. Table 8-1 lists the three modules and the supported number of calls at each data rate. The Cisco IP/VC 3540 supports cascaded calls and can be used in conjunction with the Cisco IP/VC 3510 to create a distributed MCU architecture.

Note   Continuous-presence bandwidths are asymmetrical. A 384-kbps continuous-presence call with four users actually consumes 1.344 Mbps.

Table 8-1   Data Rate and Maximum Users Supported by a Cisco IP/VC 3540

Module	Data Rate	Maximum Users
100 Session	128 kbps	100
	384 kbps	70
	768 kbps	25
	1.5/2.0 Mbps	10
	Voice only	150
60 Session	128 kbps	60
	384 kbps	42
	768 kbps	15
	1.5/2.0 Mbps	5
	Voice only	90
30 Session	128 kbps	30
	384 kbps	21
	768 kbps	9
	1.5/2.0 Mbps	3
	Voice only	45

Note   The continuous-presence feature decreases the total number of supported participants by approximately 35%.

The gateway card provides connectivity between ISDN-based H.320 participants and IP-based H.323 endpoints. The gateway module has a single 10/100 Ethernet interface, two PRI ports that are configurable for T1 or E1 speeds, and support for a wide range of switch protocols. The gateway module supports H.261 and H.263 video and G.711, G.722, G723.1, G.728, and G.729 audio for optimum videoconference quality. The T.120 features of the module allow multimedia data conferences to take place among IP and ISDN users. The gateway supports conferences at bandwidths up to 384 kbps, and is available with optional audio transcoding capabilities.

The T.120 Application Server is a Windows NT server platform that hosts applications critical to multimedia conferences, including the T.120 data conferencing server application. The combination of Application Server and Data Conferencing Server makes data sharing an integral part of multipoint conferences. PC-based H.323 endpoints can be equipped with a T.120 application that allows users to share views of an application such as spreadsheets or Web pages and gives users the ability to change numbers interactively in an analysis; point to a Web-page feature; view diagrams, graphic presentations, or slide lectures; or engage in text chats, whiteboard exchanges, or rapid file transfers. All of these capabilities can greatly enhance the videoconference.

Cisco IP/VC 3510 MCU

The Cisco IP/VC 3510 MCU enables conferences involving three or more endpoints. The Cisco IP/VC 3510 MCU has one 10/100-Mbps Ethernet connection and supports video data rates from 128 kbps to 1.5 Mbps as well as G.711-based voice. The Cisco IP/VC 3510 MCU supports both voice activated and continuous-presence calls.   Configuration of the Cisco IP/VC 3510 MCU depends on the desired function and network layout.

The Cisco IP/VC 3510 MCU is designed to support ad hoc conferences with an average of three to five users. Table 8-2 shows the maximum number of users for a single Cisco IP/VC 3510 MCU at each supported data rate. The maximum number of users on a single MCU can reside in one or more conferences. Conferences at different rates can also reside on the same Cisco IP/VC 3510 MCU, but if there are not enough resources at the time of the call, the call is rejected.

Note   Continuous-presence bandwidths are asymmetrical. A 384-kbps continuous-presence call with four users actually consumes 1.344 Mbps.

Table 8-2   Data Rate and Maximum Users Supported on a Cisco IP/VC 3510 MCU

Data Rate	Maximum Users
128 kbps	15
384 kbps	9
512 kbps	7
768 kbps	5
1.5 Mbps	3

Initiating a Call

To initiate a multipoint call using an IP/VC 3540 or 3510 MCU, the endpoint dials the appropriate service prefix followed by a conference ID. (The conference ID can be up to 9 digits long.) If a service on an MCU is 40880 for a 384-kbps call, the user might dial 4088011223, the call would be routed to the MCU using the 40880 service prefix, and the MCU would initiate an ad hoc conference with an ID of 4088011223. Users can also initiate a call and invite the other participants by dialing the conference ID, the invite string **, and the E.164 address of the another participant. Figure 8-1 shows the dial sequence of an MCU call with Video 1 initiating an MCU call to 40880123 and inviting Video 2 and Video 3.

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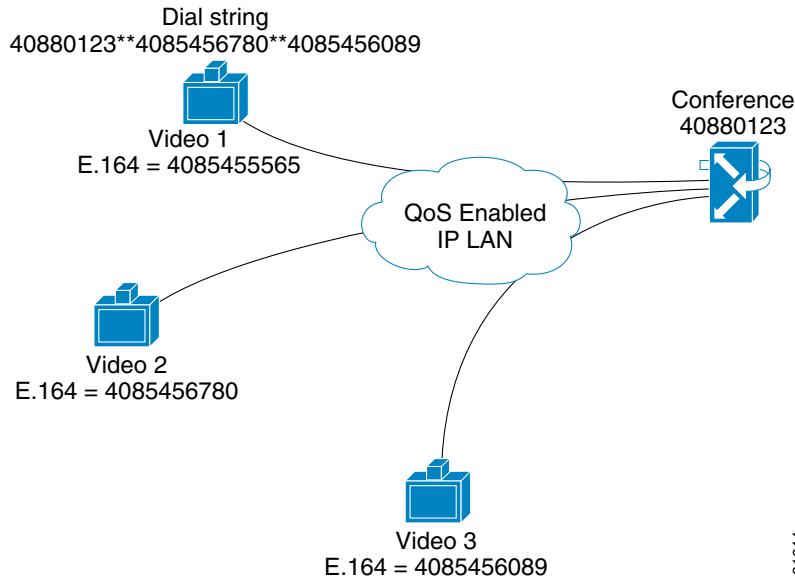
Cascading MCUs

Cascading MCUs allows larger conferences to be created by combining resources from multiple MCU blades or units. Cascading is also used in distributed MCU environments to save bandwidth on low-speed WAN links (see Distributed MCUs, page 8-5).   Both the IP/VC 3540 and 3510 MCUs support cascading.

An administrator can cascade MCUs by inviting conference calls on different MCUs to join in a single combined call. In Figure 8-2, a conference was started on each of the four MCUs. To cascade the MCUs in this example, an administrator accessed the web interface on MCU 1 and invited conferences 4088112 on MCU 2, 4088632 on MCU 3, and 4088552 on MCU 4.

Figure 8-2   Cascaded MCU Conference
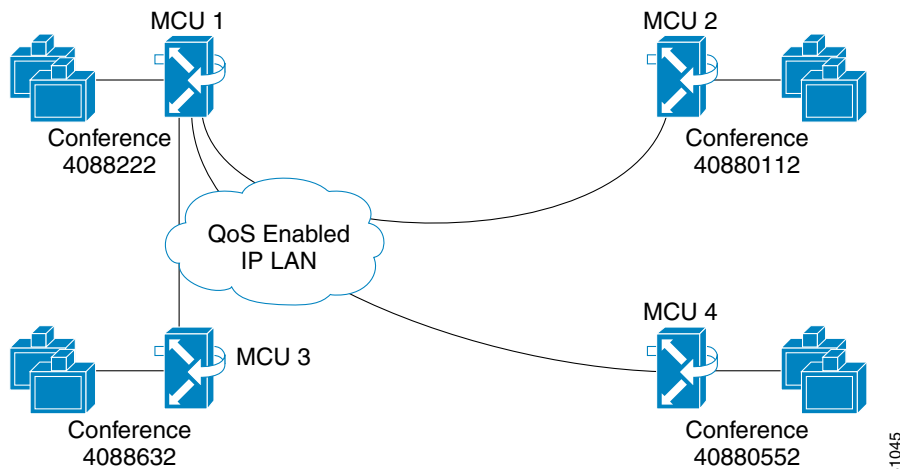
> **Note**   There is no physical connection made between the cascaded devices. Cascading occurs over the LAN or WAN, allowing MCUs to be distributed across a network. An MCU can invite H.323 endpoints, H.320 endpoints through a gateway, or other MCUs through the web interface on the MCU.

Distributed MCUs

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

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

Figure 8-3   Centralized MCU Services, with All Calls Traversing the WAN
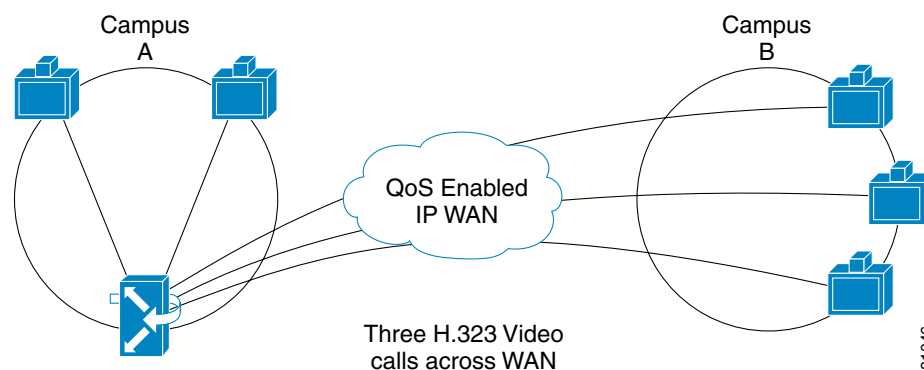
Figure 8-4 illustrates a distributed MCU model with two video terminals at Campus A calling into a local MCU, and three video terminals at Campus B calling into a local MCU. In this model, there is just a single call cascading the two conferences across the WAN. In most distributed networks, an IP/VC 3540 will be located at the headquarters site and an IP/VC 3510 will be located at each large remote site.

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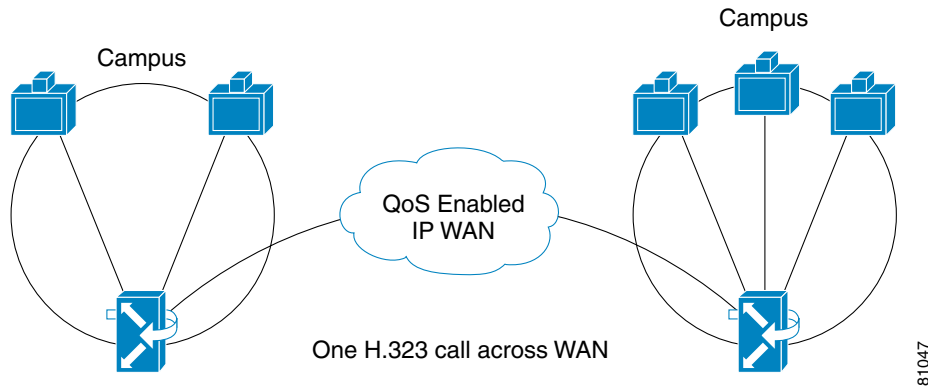
Video Gateways

The Cisco IP/VC 3540, 3525, and 3520 videoconferencing gateways give enterprises the ability to connect ISDN-based H.320 systems with IP-based H.323 videoconference endpoints. These gateways provide translation services between H.320 and H.323 networks to convert multimedia information between circuit-switched ISDN and IP networks. The gateway also supports G.711 and voice transcoding between IP and the Public Switched Telephone Network (PSTN). These systems enable users to videoconference with others users via the LAN or the PSTN, regardless of location.

The three video gateway models provide the following features:

- Cisco IP/VC 3520

   This gateway can be configured with two or four BRI ports, two or four V.35 ports, or two BRI and two V.35 ports. When equipped with V.35 ports, the Cisco IP/VC 3520 supports RS-366 or V.25bis signaling, allowing the gateway to set up circuit-switched connections through a data communications equipment (DCE) device such as an inverse multiplexer (IMUX) or access concentrator at speeds up to 768 kbps. When the gateway is equipped with BRI ports, it can support calls up to 384 kbps on aggregated channels.

- Cisco IP/VC 3525

   This gateway is a self-contained system that supports a high volume of calls over a single high-speed ISDN PRI connection, allowing dynamic allocation of its 23 B channels. With the Cisco IP/VC 3525, multiple H.323 endpoints can share this PRI T1/E1 system when communicating with ISDN-based endpoints. This gateway can support either eight sessions at 128 kbps, three sessions at 384 kbps with T1, four sessions at 384 kbps with E1, or a mixture of all three. Sessions at different speeds can take place simultaneously.

- Cisco IP/VC 3540 MCU

   This gateway is a two-port PRI T1/E1 module that supports a high volume of calls over multiple high-speed ISDN PRI connections. The IP/VC 3540 gateway has the ability to span a single call across the two PRI connections on a single blade, providing more efficient use of B channels.

Table 8-3 shows the maximum numbers of calls supported per platform.

Table 8-3   Maximum Number of Calls per Platform

Platform and Call Data Rate	Maximum Number of Calls
IP/VC 3520 4 X BRI	
• 128 kbps	• 4
• 384 kbps	• 1
IP/VC 3520 4 X V.35	
• 128 kbps[1]	• 12
• 384 kbps[1]	• 4
• 768 kbps[2]	• 4
IP/VC 3525	
• 128 kbps	• 8
• 384 kbps	• 3 for T1, 4 for E1
IP/VC 3540 2X PRI	Maximum number of calls per T1/E1
• 128 kbps[3]	• 23/30
• 256 kbps[3]	• 11/15
• 384 kbps	• 7/10

1. Numbers based on an IMUX with three BRI lines.

2. Requires an IMUX with PRI connectivity or an ISDN switch with a PRI connection.

3. Call capacity decreases if T.120 or audio transcoding features are used.

Service Prefixes

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

From the users' perspective, they will have to dial the service prefix, which in this case is equivalent to an access code, followed by the ISDN number of the H.320 video unit. For this configuration, a local (intra-zone) gateway is used whenever one is present. In zones that do not contain a gateway, the administrator should assign a gateway in another zone as the primary gateway for the local zone. Configure location request (LRQ) forwarding or a static hopoff statement to route all calls to a zone with a gateway for PSTN access. Figure 8-5 illustrates the service prefix configuration for a PSTN gateway.

Figure 8-5   Service Prefix for an IP/VC 3525 PSTN Gateway
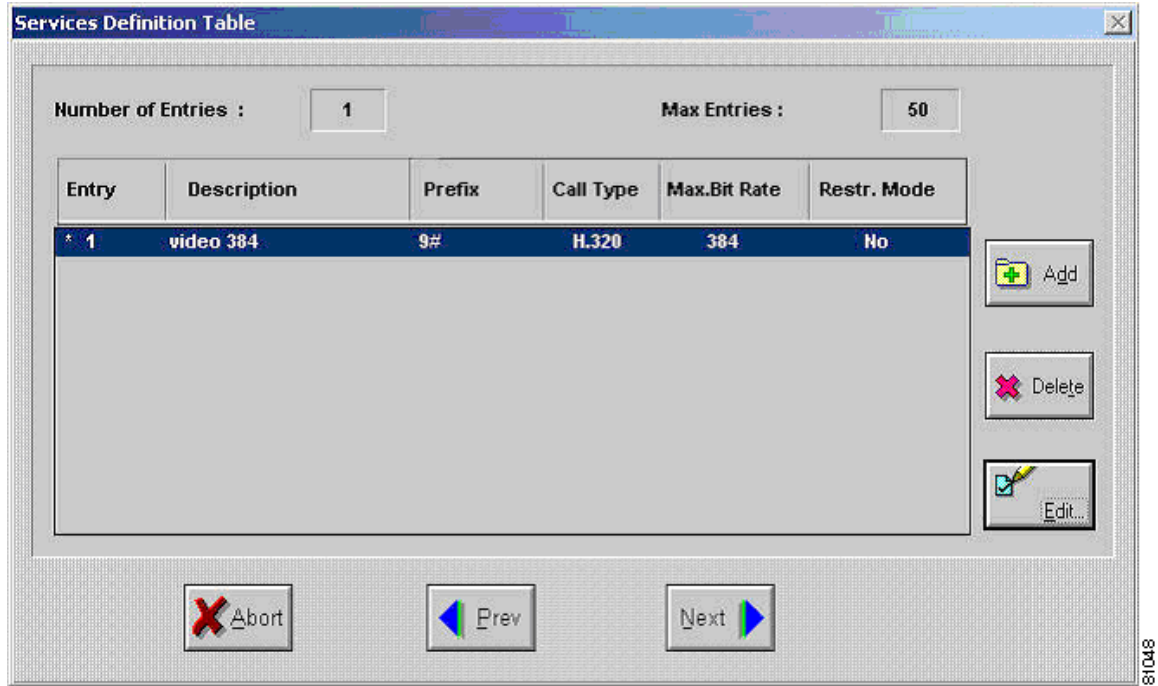
Line Hunting

The Cisco IP/VC 352X supports the Outbound Line Hunting (LAN to PSTN) feature. Line hunting allows users to build a pool of gateways for PSTN access. This feature creates a larger number of access lines serviced by a single set of service prefixes. With the current release of gateway code (version. 2.2) for the Cisco IP/VC 3520 and Cisco 3525, line hunting is support using Resource Availability Information (RAI) and Resource Availability Confirmation (RAC).

To implement the line hunting feature, you configure multiple gateways with identical service prefixes and register them with the same gatekeeper. Outbound PSTN calls are sent to the gateways based on resource availability, using RAI and RAC. In the gateway configuration, you set utilization parameters based on gateway resource percentages. Figure 8-6 shows the configuration screen from a Cisco IP/VC gateway.

The main gateway configuration parameters for line hunting are:

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

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

Figure 8-6   Line Hunting Configuration Screen for an IP/VC 3540 Gateway
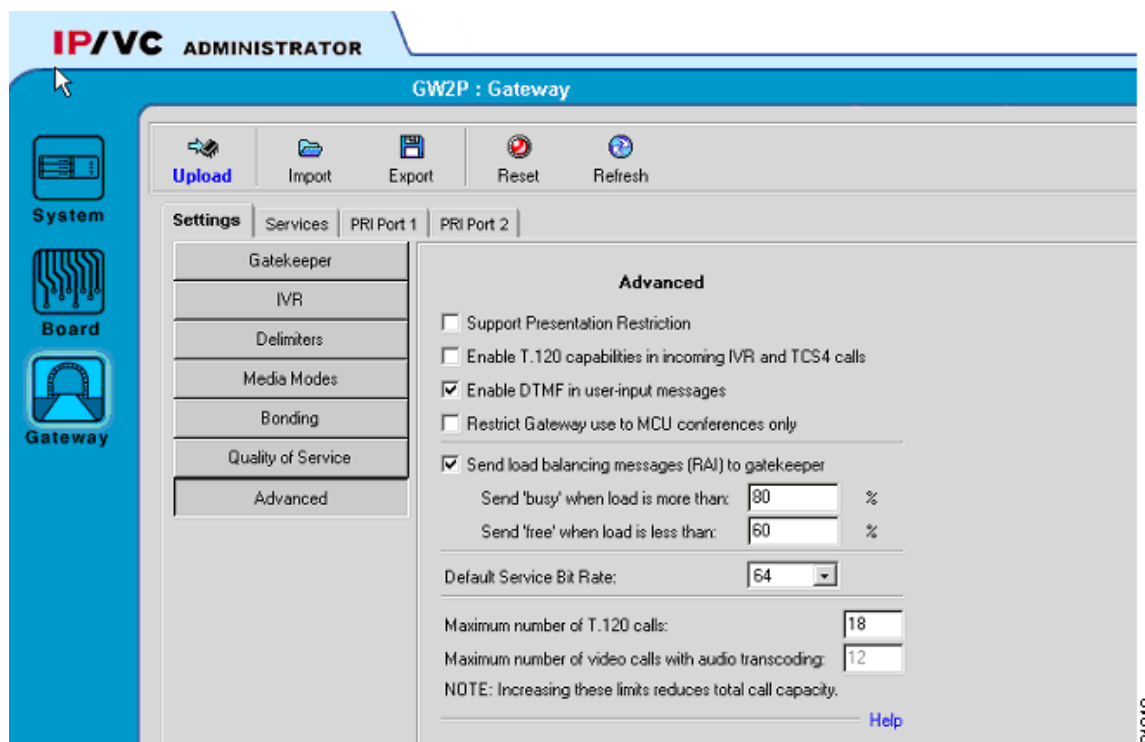
Inbound line hunting (PSTN to LAN) can be done but is not very eloquent. Telephone companies are able to build hunt groups across multiple PRI lines, allowing calls to be rolled to a second PRI if the first is busy. Note that all B channels on a PRI line must be busy for a call to roll to a second PRI. In a telephony system, this is not a problem because each voice call takes a single B channel, but a video call requires several B channels, depending on the data rate of the call. For example, a 384-kbps video call consumes six B channels of the PRI line. If the first of two PRI lines in a hunt group has two B channels available, the PSTN will send another 384-kbps video call to that PRI line because the PRI is not completely busy, but the call will fail at the gateway due to lack of resources (only two channels available instead of the required six).

The only way to make inbound line hunting work is to standardize on a data rate (for example, 384 kbps) and use the "busy out" feature on the gateway. The "busy out" feature allows the gateway to busy out remaining channels on the gateway when there are not enough channels available to connect a call at the standardized data rate (in our example, six channels for 384 kbps). With the additional channels busied out, the next inbound call will be forwarded to the next gateway in the trunk group.

Currently this method of using the "busy out" feature is the best option for incoming line hunting. Figure 8-7 illustrates the use of this method with two PRI gateways in a single hunt group.

Figure 8-7   Line Hunting Example
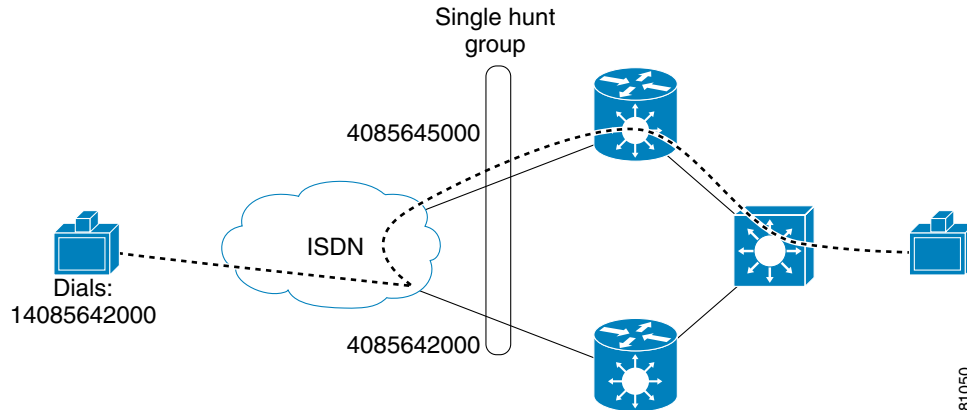
Cisco IP/VC 3530 VTA

The Cisco IP/VC 3530 Video Terminal Adapter (VTA) allows you to migrate existing H.320 video units from ISDN to an IP network. The IP/VC 3530 VTA is a single-port gateway that converts H.320 to H.323 for a single H.320 video terminal.

The IP/VC 3530 VTA has two physical ports that connect directly to an H.320 codec. The supported connection from the VTA is V.35/RS-366 only; if the H.320 codec does not have V.35/RS-366 ports, you must add them.   Most H.320 codecs support this connection, but some H.320 systems are installed with a different interface such as BRI or PRI. Even though there are two ports on the VTA, only one H.320 codec can be attached to each VTA. If all calls made from an H.320 codec through a VTA use aggregated channels, only one V.35/RS-366 connection is required. If 2 X 56-kbps or 2 X 64-kbps calls are going to be made, both ports must be connected. Figure 8-8 illustrates the connection from a VTA to an H.320 codec.

Figure 8-8   Connection Between Cisco IP/VC 3530 VTA and H.323 Codec

H.320 room system
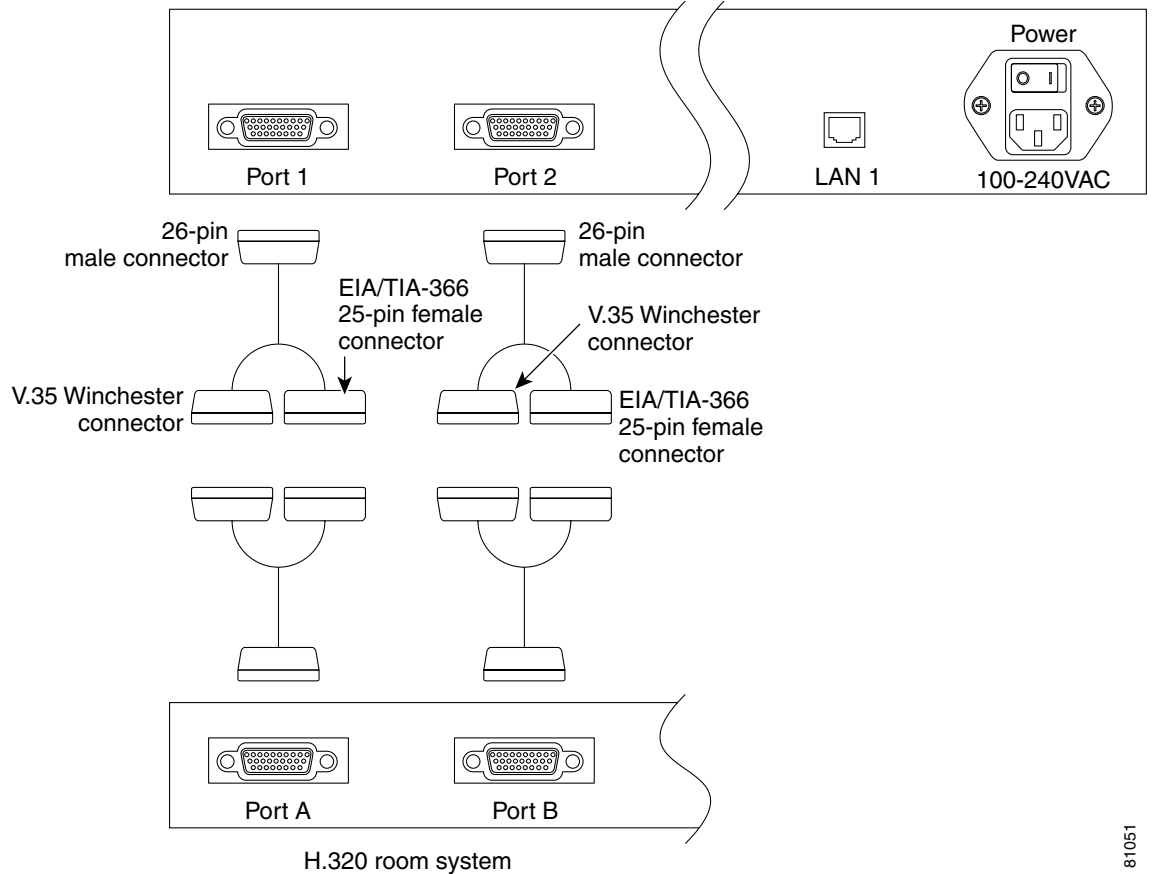
Because the VTA is a gateway device and does not negotiate the call data rate automatically, prefixes and suffixes must be added to calls with different data rates. When the VTA is initially configured, the administrator selects a default incoming and outgoing data rate. When calls are placed or received at the defined default data rate, no prefix or suffix is needed.   If a call is going to be placed or received at a speed other than the defined default data rate, the prefix or suffix must be added.

To place a call at a data rate other than the default outgoing data rate, the user must add a prefix to the dial string. Table 8-4 lists the prefix for each outgoing data rate. When placing a call at the selected default rate, the user does not need add the prefix to the dial string. When a #XX prefix is added to the dialed number, the VTA sees the prefix, strips it from the number, and places the call at the specified data rate.

Table 8-4   Data Rate Prefixes for Outbound Calls

Desired Data Rate for Call	Dialing Prefix
2 X 64 kbps (2B)	#00
128 kbps	#10
256 kbps	#20
384 kbps	#30
768 kbps	#70

Table 8-4   Data Rate Prefixes for Outbound Calls (continued)

Desired Data Rate for Call	Dialing Prefix
2 X 56 kbps (2B restricted)	#01
112 kbps (restricted)	#11
224 kbps (restricted)	#21
336 kbps (restricted)	#31
672 kbps (restricted)	#71

To receive calls at a data rate other than the specified default incoming data rate, a suffix must be added to the dial string. When the VTA is configured, a default incoming speed is set, and that data rate is used for calls received with no suffix. When the VTA registers with the gatekeeper, it registers with six E.164 addresses. If the E.164 address of a VTA were configured as 408565212, the VTA would register with the first six E.164 addresses listed in Table 8-5.

Table 8-5   Data Rate Suffixes for Incoming Calls

E.164 Address	Supported Data Rate
408565212	Default Data Rate
40856521200	2 X 64 kbps
40856521210	128 kbps
40856521220	256 kbps
40856521230	384 kbps
40856521270	768 kbps
40856521201	2 X 56 kbps (restricted)
40856521211	112 kbps (restricted)
40856521221	224 kbps (restricted)
40565621231	336 kbps (restricted)
40856521271	672 kbps (restricted)

Cisco Multimedia Conference Manager (MCM)

The Cisco Multimedia Conference Manager (MCM) is a Cisco IOS software component that supplies gatekeeper and proxy functions for an H.323 video network.   The Cisco IOS gatekeeper enables large H.323 video networks to be built and managed on Cisco hardware. The proxy supplies needed functions that are not currently provided by devices in some IP networks. For example, Quality of Service (QoS), access to Network Address Translation (NAT) networks, and firewall access are some of the functions that the proxy supplies.

Gatekeeper

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

- H.323 component registration and call routing

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

- Bandwidth management

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

   - Inter-zone: Total bandwidth allowed from a local or default zone to and from all other zones

   - Remote: Total bandwidth allowed from all local zones to and from all remote zones

   - Session: Bandwidth allowed per session in a zone

   - Total: Total bandwidth allowed in a zone

- Authentication, authorization, and accounting (AAA) support

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

Table 8-6 lists the various router platforms that support Cisco IOS gatekeeper functionality, along with their relevant performance data.

Table 8-6   Platforms that Support Cisco IOS Gatekeeper

Platform	IP Routing (Packets per Second)	Registered H.323 Endpoints	Simultaneous Video Calls	Video Proxy Sessions
Cisco 7200	50-100,000	3000	500	50 at 768 kbps 75 at 384 kbps 100 at 128 kbps
Cisco 3660	25-100,000	1800	250	25 at 768 kbps 35 at 384 kbps 50 at 128 kbps
Cisco 3640	10-40,000	1800	150	10 at 768 kbps 15 at 384 kbps 30 at 128 kbps
Cisco 3620	10-15,000	1800	75	10 at 768 kbps 15 at 384 kbps 30 at 128 kbps
Cisco 2620	5-10,000	900	60	2 at 768 kbps 4 at 384 kbps 8 at 128 kbps

Table 8-6   Platforms that Support Cisco IOS Gatekeeper (continued)

Platform	IP Routing (Packets per Second)	Registered H.323 Endpoints	Simultaneous Video Calls	Video Proxy Sessions
Cisco 2610	2-5,000	900	60	2 at 768 kbps 4 at 384 kbps 6 at 128 kbps
Cisco 3810	2-5,000	900	60	2 at 768 kbps 4 at 384 kbps 6 at 128 kbps
Cisco 37XX (New platform)	Numbers not available yet.			

Figure 8-9 illustrates CPU utilization for each router platform, based on a maximum of 32 calls at a data rate of 384 kbps.

Figure 8-9   Proxy CPU Usage by Platform
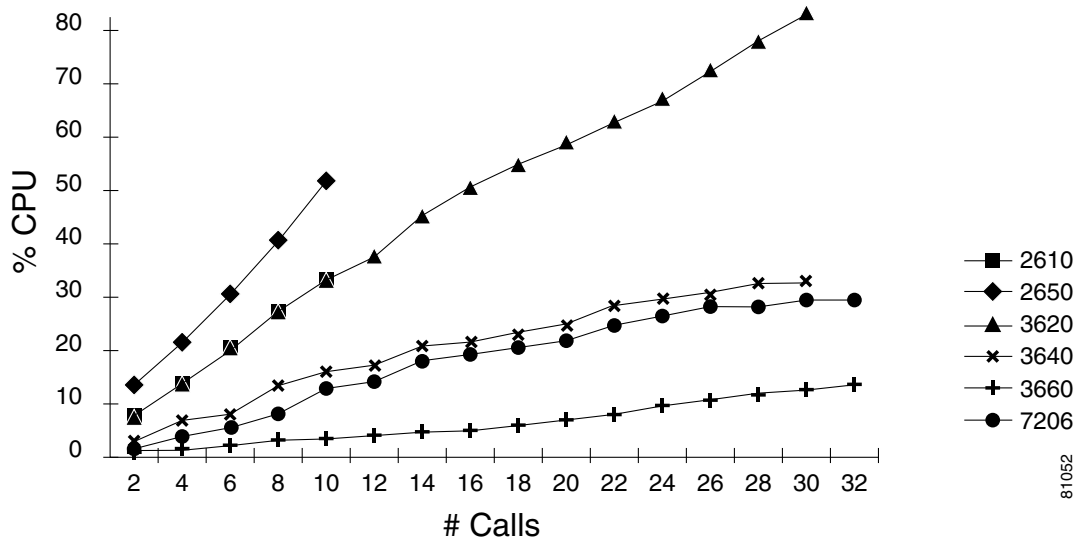
The following limits apply to the platforms listed in Figure 8-9:

- Cisco 2600 Series routers have a software limit of 10 proxy calls.

- Cisco 3620 and 3640 routers have a software limit of 30 proxy calls.

- Cisco 3660 and 7206 routers did not reveal any limitations when tested with a maximum of 32 proxy calls.

Note   The maximum number of proxy calls for the Cisco 2600 and Cisco 3600 series routers are higher than the numbers in Table 8-6. If the router is being used for a gatekeeper and proxy only (with no routing functions), the maximum number of proxy calls can be based on the higher software limits.

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

- Hot Standby Router Protocol (HSRP) enables administrators to build a standby gatekeeper that becomes active if the primary gatekeeper fails.

- Directory Gatekeeper, or Location Request (LRQ) forwarding, enables administrators to build large multi-tier networks, minimizing the configuration required in the lower-tier gatekeepers. When a call is made in a lower-tier zone and a match is not found, the call is automatically forwarded up to the directory gatekeeper for resolution. (For more information on directory gatekeepers, refer to the Call Routing chapter.) Figure 8-10 illustrates a network configured with two regional directory gatekeepers, one at Site E and another at Site F.

HSRP

Hot Standby Router Protocol (HSRP) enables a set of routers with Cisco Multimedia Conference Manager (MCM) to work together as a single virtual gatekeeper. You can implement this feature by creating a *phantom* router that has its own IP and MAC addresses.

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

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

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

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

Figure 8-10   Network with Two Directory Gatekeepers
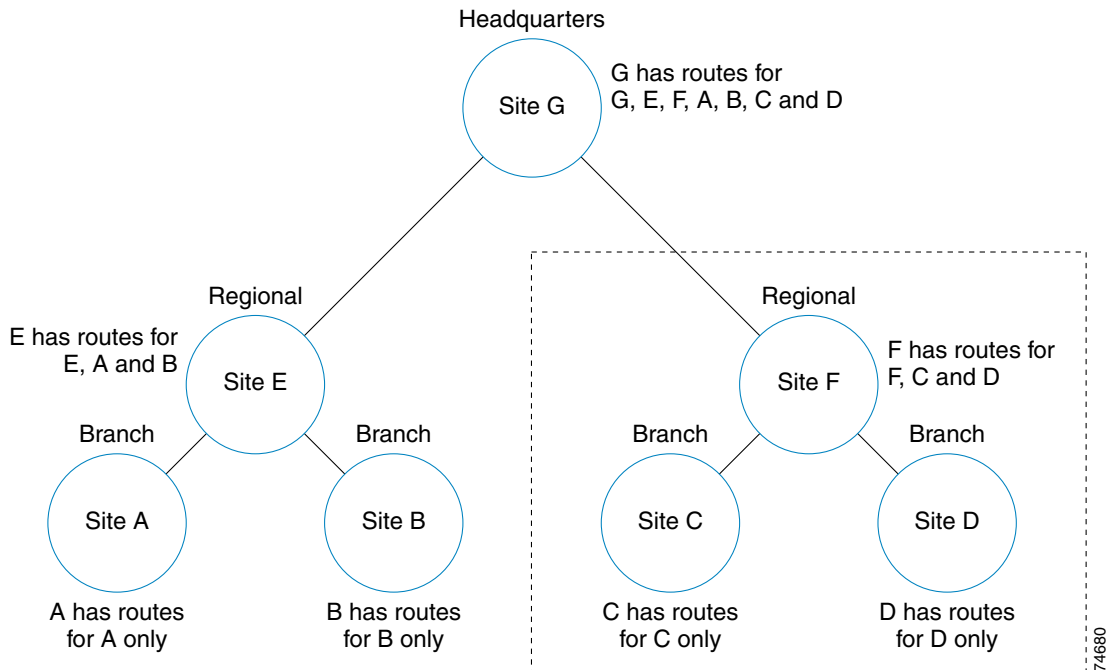
Proxy

The Multimedia Conference Manager (MCM) also provides a proxy, which functions as a call processing agent to terminate H.323 calls from one local LAN or zone and to establish sessions with H.323 terminals located in different LANs or zones. The proxy provides the following features:

- Classification of video and audio streams, with IP Precedence or Resource Reservation Protocol (RSVP), for QoS

- Access through firewalls and Network Address Translation (NAT) environments

- Secure WAN queue access

The Cisco proxy allows video terminals with no QoS capabilities to obtain traffic classification across IP WAN links. The proxy is configured to support RSVP or IP Precedence, and it registers with the local gatekeeper as an H.323 endpoint. The proxy is then used to classify traffic across low-speed WAN links with the configured traffic classification. Each proxy call contains three call legs: one from the calling video terminal to the proxy registered in its zone, one from proxy to proxy across the WAN, and one from the remote proxy to the receiving video terminal. Figure 8-11 illustrates a proxy call across a WAN link.

Figure 8-11   Proxy Call Across a WAN
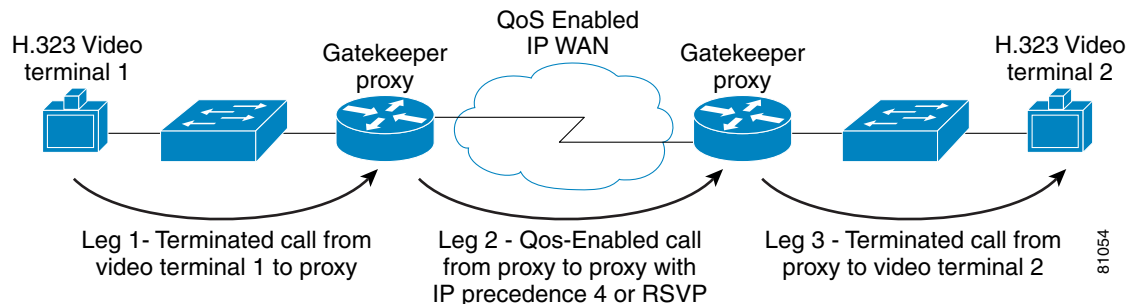
Note   Single-legged proxy is supported in Cisco IOS software release 12.2(2) or later.

Firewalls and Network Address Translation (NAT)

H.323 is cumbersome to run through a firewall because it uses multiple data ports for a single call. For an H.323 call to take place, it must first open an H.225 connection on TCP port 1720, using Q.931 signaling. Next, the H.245 management session is established. While this session can take place on a separate channel from the H.225 setup, it can also be done using H.245 tunneling, which takes the H.245 messages and embeds them in the Q.931 messages in the previously established H.225 channel.

Next, the H.245 session opens dynamically assigned ports for the UDP-based RTP and RTCP video and audio data streams. The port numbers can range from 1024 to 65535. Because the port numbers are not known in advance, and because it would defeat the purpose of a firewall to open all these ports, a firewall must be able to *snoop* the H.323 data stream in order to open the additional ports needed for the call. This snooping is also known as stateful inspection.

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

Using the proxy in NAT or firewall environments allows administrators to target a single IP address to terminate all H.323 video calls. All incoming and outgoing video calls that access the public network will use the proxy. With the Cisco Pix Firewall, administrators can enable H.323 fixup and allow UDP port 1720 traffic to access the IP address of the proxy. If the proxy were not used, the administrator would have to configure UDP port 1720 to all videoconferencing devices. If you use a Cisco IP/VC 352X or 3540 gateway, port 1820 must be configured for the videoconferencing devices. Figure 8-12 illustrates the call flow in a network with NAT and a firewall.

In Figure 8-12, Public video #1 is registered with gatekeeper 206.165.201.1. (IP address dialing is not supported with the proxy.) All video devices in the private 10.1.1.X network are registered with gatekeeper 10.1.1.20. H.323 fixup has been configured on the PIX Firewall, and port 1720 has been opened to IP address 10.1.1.20 (the proxy). When a video call is placed from the public video device to any of the 10.1.1.X video devices, the call from the public video device terminates on the proxy.   The proxy then initiates a call to the 10.1.1.X device inside the firewall.

Note Polycom Viewstation version 7.2 does not work in this configuration.

Figure 8-12   Call Flow with NAT and a Firewall

9
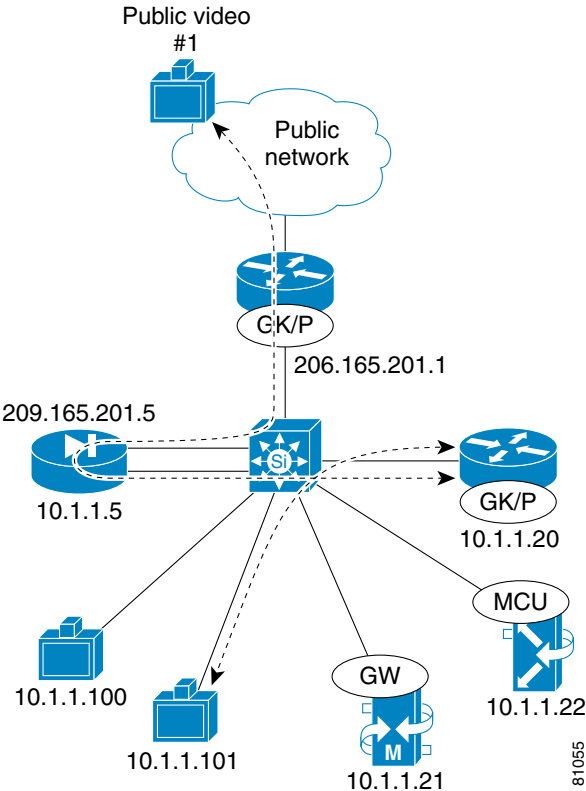
Multi-Zone WAN Case Study

This chapter provides an example of a typical WAN multi-zone model deployed in an enterprise environment.

In this case study, the enterprise is a health care provider with locations spread across the United States. Five locations currently use ISDN-based videoconferencing. The enterprise has a T1 to each site and would like to install new H.323 videoconferencing units and utilize their existing WAN bandwidth. Each site contains a minimum of three video units, and the enterprise has standardized on 384 kbps as their call data rate. The enterprise requires multipoint calls as well as the ability to call off-net to their clients.

Network Topology

Currently the enterprise in this example has five sites in the United States, consisting of Sacramento CA, Los Angeles CA, Dallas TX, Columbus OH, and Chicago IL. Each site connects back to Columbus, the headquarters, with a T1 link. The bandwidth utilization on all the connections is low. The enterprise has just upgraded their WAN routers at remote sites to Cisco 3640 routers to support voice, video, and data in the near future.   Currently, all videoconferencing units are directly connected to an IMUX with three BRI lines, allowing boded 384-kbps calls. The Columbus site contains an H.320 multipoint conference unit (MCU) with three PRI lines supporting multipoint calls among the sites. Figure 9-1 illustrates the current IP network, and Figure 9-2 illustrates the current videoconferencing network.

Figure 9-1   Current IP Network
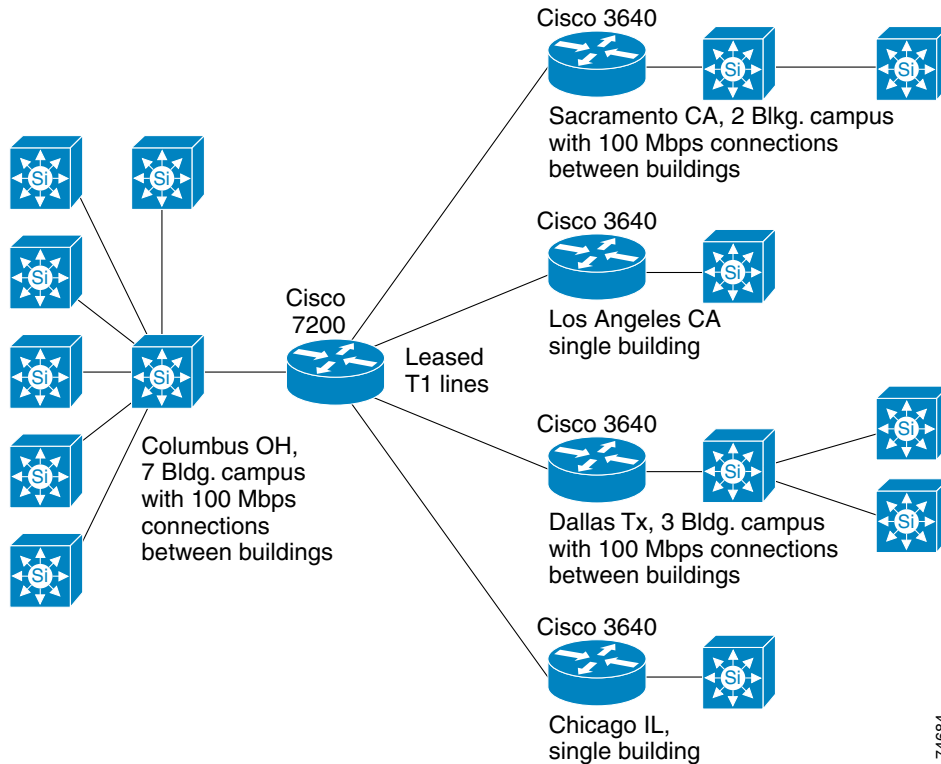
Figure 9-2   Current Videoconferencing Network
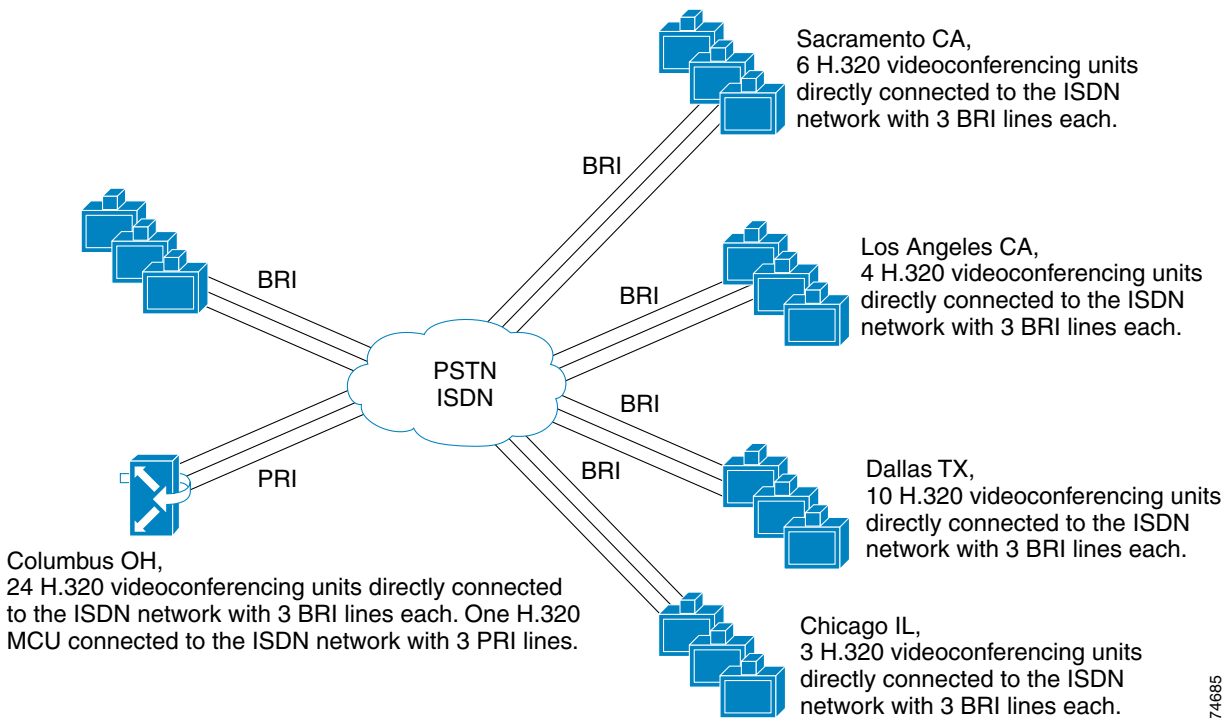
Network Design

The network outlined in the previous section is a classic WAN multi-zone model. There is sufficient WAN bandwidth, and each site contains three or more video terminals. In this network, a gatekeeper and proxy are located at each site. Directory gatekeeper services are configured, and Hot Standby Routing Protocol (HSRP) is used for gatekeeper redundancy at the Columbus site. Quality of service (QoS) and call admission control (CAC) need to be configured in the network to ensure video quality.

Quality of Service (QoS)

End-to-end QoS is a key factor in a successful deployment. The enterprise in this example has decided to use an H.323 video terminal that supports marking of IP Precedence. The Columbus, Sacramento, and Dallas sites have upgraded to Catalyst 6500 switches. In these three sites, LAN QoS will be configured; the remaining three sites will support LAN QoS when the switches at those sites are also upgraded. All video units will be connected to 10/100 Ethernet ports.

All video terminals are configured to mark IP Precedence 4. In Columbus, Sacramento, and Dallas, trust boundaries are set on the Catalyst 6500 switches. Video gateways and MCUs are also installed in Columbus, Sacramento, and Dallas. At this time, the gateways and MCUs do not support IP Precedence, so IP Precedence is marked and a trust boundary is set on the Catalyst 6500 ports to which the gateways and MCUs are connected. Gateways are also installed in Los Angeles and Chicago.

Priority queues are configured on all WAN routers and are provisioned for 920 kbps. This guarantees that bandwidth is available for two 384-kbps calls. An access list entry is also added on the WAN router to set the entrance criterion for the priority queue. Only video traffic received from the proxy is admitted to the priority queue.

The gatekeeper at each site is configured to use the local proxy for all inter-zone calls. The proxy rewrites IP Precedence 4 and provides a single access point to the configured priority queue.

For more information regarding network QoS, refer to the *Cisco AVVID Network Infrastructure Enterprise Quality of Service Design* guide.

Call Admission Control

Call admission control (CAC) must be implemented for inter-zone calls, and it is also a good idea to configure CAC for intra-zone calls. Enabling CAC for inter-zone calls guarantees that the bandwidth limits provisioned on the priority queues are not exceeded. If the provisioned bandwidth for the priority queue on the WAN route is exceeded, all video calls in the queue will be affected. The gatekeeper at each site contains the following three bandwidth statements for CAC.

```
bandwidth total {default | zone <zone-name>} <bandwidth-size>
bandwidth remote 1536
bandwidth session default 768
```

It is important to note that the bandwidth is calculated in half-duplex mode, so the call data rate must be doubled. A 384-kbps call is represented as 768 kbps in the bandwidth statements. The **bandwidth total** command allows administrators to limit the bandwidth within a single local zone, or for all local zones by adding the **default** statement. The remote bandwidth (available bandwidth to and from any remote zone) is limited to 1536 kbps, or two 384-kbps calls. The bandwidth per session is limited to 768 kbps, or one 384-kbps call.

Figure 9-3 illustrates QoS and CAC points for Columbus, Figure 9-4 illustrates QoS and CAC points for Dallas and Sacramento, and Figure 9-5 illustrates QoS and CAC points for Los Angeles and Chicago.

Figure 9-3   QoS and CAC for Columbus

2 Cisco 3620 routers providing gatekeeper and
proxy functionality and configured with HSRP.
CAC configured on gatekeepers and IP
Precedence set to 4 on the proxy.

CAC set for intra-zone
and inter-zone calls.

Video
terminal

Trust
boundaries
configured

Priority Queue (PQ)
configured and
provisioned for 920 kbps
on each T1.

IP precedence
set to 4

MCU or
gateway

IP precedence set
to 4 and trust
boundary configured

ACL added to set PQ entrance
Criterion. Only packets from the
proxy will have access to the PQ.

74686

Figure 9-4   QoS and CAC for Dallas and Sacramento

Gatekeeper and proxy functionality configured
on WAN router. CAC configured on gatekeeper
and IP precedence set to 4 on the proxy.

Video
terminal

Trust
boundaries
configured

Priority Queue (PQ) configured
and provisioned for 920k

IP precedence
set to 4

MCU or
gateway

IP precedence set
to 4 and trust
boundary configured

WAN
router

ACL added to set PQ entrance
Criterion. Only packets from the
proxy will have access to the PQ.

74687

Figure 9-5   QoS and CAC for Los Angeles and Chicago
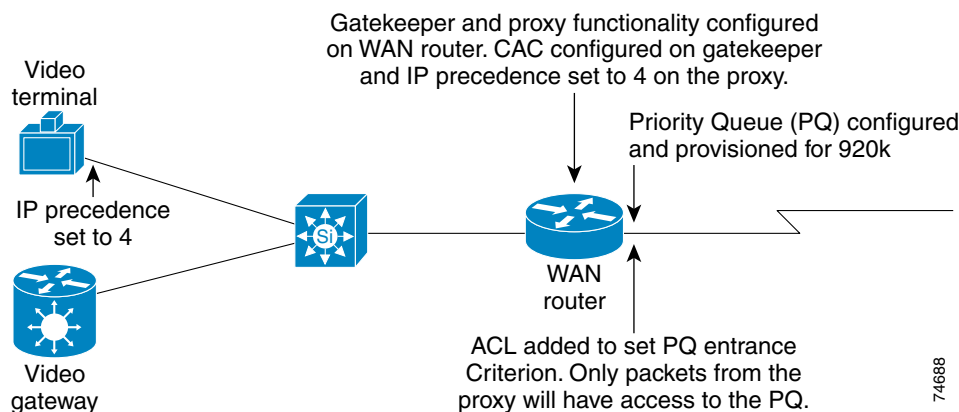
Dial Plan

When deciding on a dial plan, it is always a good idea to start with the incoming PSTN call routing. In our example, we have created five zones that all contain video gateways. DID is used to route incoming calls to video terminals.   IVR is used to route calls from the video gateways in Columbus, Dallas, and Sacramento, to their local MCUs. If, for some reason, one or more of the zones in our example did not contain a gateway, IVR for routing all incoming PSTN calls would have been a better choice.

Zone Prefixes

The zone prefix for each zone is based on the local area code. Area codes are unique, and users are familiar with the numbering structure. In our configuration there is a single zone in each site, so the zone prefixes are based on area codes. If more than one zone were required in a single area code, longer zone prefixes could be used. (Refer to Zone Prefix Design, page 6-6.) The zone prefixes in this network are.

- Columbus = 614
- Sacramento = 916
- Dallas = 972
- Chicago = 847
- Los Angeles = 213

Service Prefixes

Service prefixes must be configured for all MCUs and video gateways. As described in the chapter on Dial Plan Architecture, it is a good idea to reserve a block of numbers for video gateways and a block of numbers for MCUs. In this case, the enterprise has chosen to standardize on 384-kbps calls; this makes service prefixes for gateways very simple. The obvious choice would be to use 9 for all PSTN calls, but that would cause routing problems in the Sacramento and Dallas zones. The Sacramento zone prefix is 916, and overlapping gateway service prefixes and zone prefixes will cause routing problems. There are

two options; reserve another block of numbers other than 9*, or use a service prefix such as 9#. In this case, we have chosen 9# for PSTN access in all zones. Any time a user tries to access the WAN, the dial string will start with 9#.

For MCU service prefixes, 8* is reserved, and the zone prefix is appended to associate it with the zone in which the MCU resides. The MCU service prefix in the Sacramento zone is 9168*, and in Los Angeles it is 2128*. Table 9-1 lists the service prefixes chosen for different types of calls on the MCU. (These service prefixes are used in every zone, and the zone prefix is appended.)

Table 9-1   MCU Service Prefixes

Service Prefix	Data Rate	Number of Parties	Video Format	Continuous Presence
83	384 kbps	3	H.261	No
84	384 kbps	4	H.261	Yes
85	384 kbps	5	H.261	No
89	384 kbps	9	H.261	No

E.164 Addresses and H.323-IDs

The carrier provides E.164 addresses for this enterprise. Because DID has been chosen for incoming PSTN routing method, the enterprise will order blocks of DID numbers with each PRI line. Each video terminal is assigned a valid DID number for its E.164 address. In Columbus, there are 24 video terminals, and thirty DID numbers are ordered with the PRI line for Columbus. (The extra six numbers are for expansion.) In Los Angeles and Chicago, DID numbers off the BRI lines will be used as E.164 addresses. (See Video Infrastructure, page 9-7, for the video components at each site.)

H.323-IDs are based on the name of the conference room where the video system resides. Since the video terminals may be moved from room to room, H.323-IDs will not be used for dialing. The enterprise is using a global address book that will display all of the IP video terminals on the network. Users can choose to dial from the address book or manually enter the E.164 address of the unit being called. Figure 9-6 illustrates the Dial plan in Columbus.

Figure 9-6   Dial Plan for Columbus

Zone Prefix: 614

Video Terminal
E.164 Address
6143251010
H323-ID
Indians

MCU Service Prefixes:
61483
61484
61485
61486

Video Gateway
Service Prefix: 9#
DID Range: 6143251000-6143251030

74689
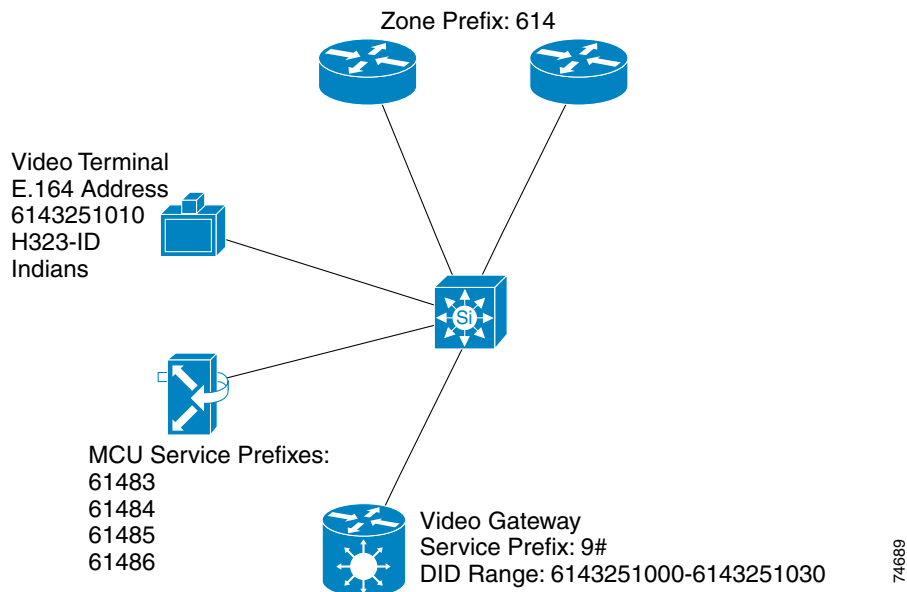
Video Infrastructure

When deciding on location and number of video components, it is important to understand the enterprise's needs. This enterprise made it clear that less than ten percent of all video calls were off-net calls. The number of video calls placed daily ranges from 10 to 15, and most calls are multipoint. For this reason, the enterprise decided to go with video gateways at each site and MCUs in Columbus, Dallas, and Sacramento. The following sections list the video components for each site.

Columbus

- IP Video Terminals, 24

  The current 24 H.320 video systems are over three years old and will be replaced with new H.323 group systems that support IP Precedence.

- MCUs, 4

  The four MCUs will be configured in a stack, allowing one set of service prefixes to be shared by all four MCUs.

- Video Gateway PRI, 1

  A single PRI gateway will be installed with 30 DID numbers that will be assigned to the IP video terminals.   All 10 digits will be passed to the gateway by the carrier, allowing each video terminal to register with a 10-digit fully qualified E.164 address. IVR will be enabled and used for PSTN access to MCU conferences.   One DID number will have to be reserved for IVR calls.

Sacramento

- IP Video Terminals, 6

  The existing six H.320 video systems are over three years old and will be replaced with new H.323 group systems that support IP Precedence.

- MCU, 1

  A single MCU will be located on the Sacramento campus for local, on-site multipoint calls. The Sacramento campus is in the process of adding another building and possibly adding two or three additional IP video terminals. The MCU will also allow multiple video terminals to participate in an off-campus multipoint call while consuming the bandwidth of only a single call. This will be done by cascading a Sacramento MCU conference with a Columbus MCU conference.

- Video Gateway, 1

  A single PRI gateway will be installed with 10 DID numbers that will be assigned to the IP video terminals. All 10 digits will be passed to the gateway by the carrier, allowing each video terminal to register with a 10-digit fully qualified E.164 address. IVR will also be enabled and used for PSTN access to MCU conferences. One DID number will have to be reserved for IVR calls.

Dallas

- IP Video Terminals, 10

  The existing 10 H.320 video systems are over three years old and will be replaced with new H.323 group systems that support IP Precedence.

- MCU, 1

  A single MCU will be located on the Dallas campus for local, on-site multipoint calls. The MCU will also allow multiple video terminals to participate in an off-campus multipoint call while consuming the bandwidth of only a single call. This will be done by cascading a Dallas MCU conference with a Columbus MCU conference.

- Video Gateway, 1

  A single PRI gateway will be installed with 15 DID numbers that will be assigned to the IP video terminals. All 10 digits will be passed to the gateway by the carrier, allowing each video terminal to register with a 10-digit fully qualified E.164 address. IVR will also be enabled and used for PSTN access to MCU conferences. One DID number will have to be reserved for IVR calls.

Los Angeles

- IP Video Terminals, 4

  The existing four H.320 video systems are over three years old and will be replaced with new H.323 group systems that support IP Precedence.

- MCU, 0

  Los Angles will not have a local MCU.

- Video Gateway, 1

  A single BRI gateway will be installed with four BRI lines. Each video terminal will receive a DID number from one of the BRI lines. IVR will not be enabled on the gateway.

Chicago

- IP Video Terminals, 3

  The existing three H.320 video systems are over three years old and will be replaced with new H.323 group systems that support IP Precedence.

- MCU, 0

   Chicago will not have a local MCU.

- Video Gateway, 1

   A single BRI gateway will be installed with three BRI lines. Each video terminal will receive a DID number from one of the BRI lines. IVR will not be enabled on the gateway.

Figure 9-7 illustrates the video components and dial plan for the new IP video network.

Figure 9-7   Dial Plan for Example Video Network
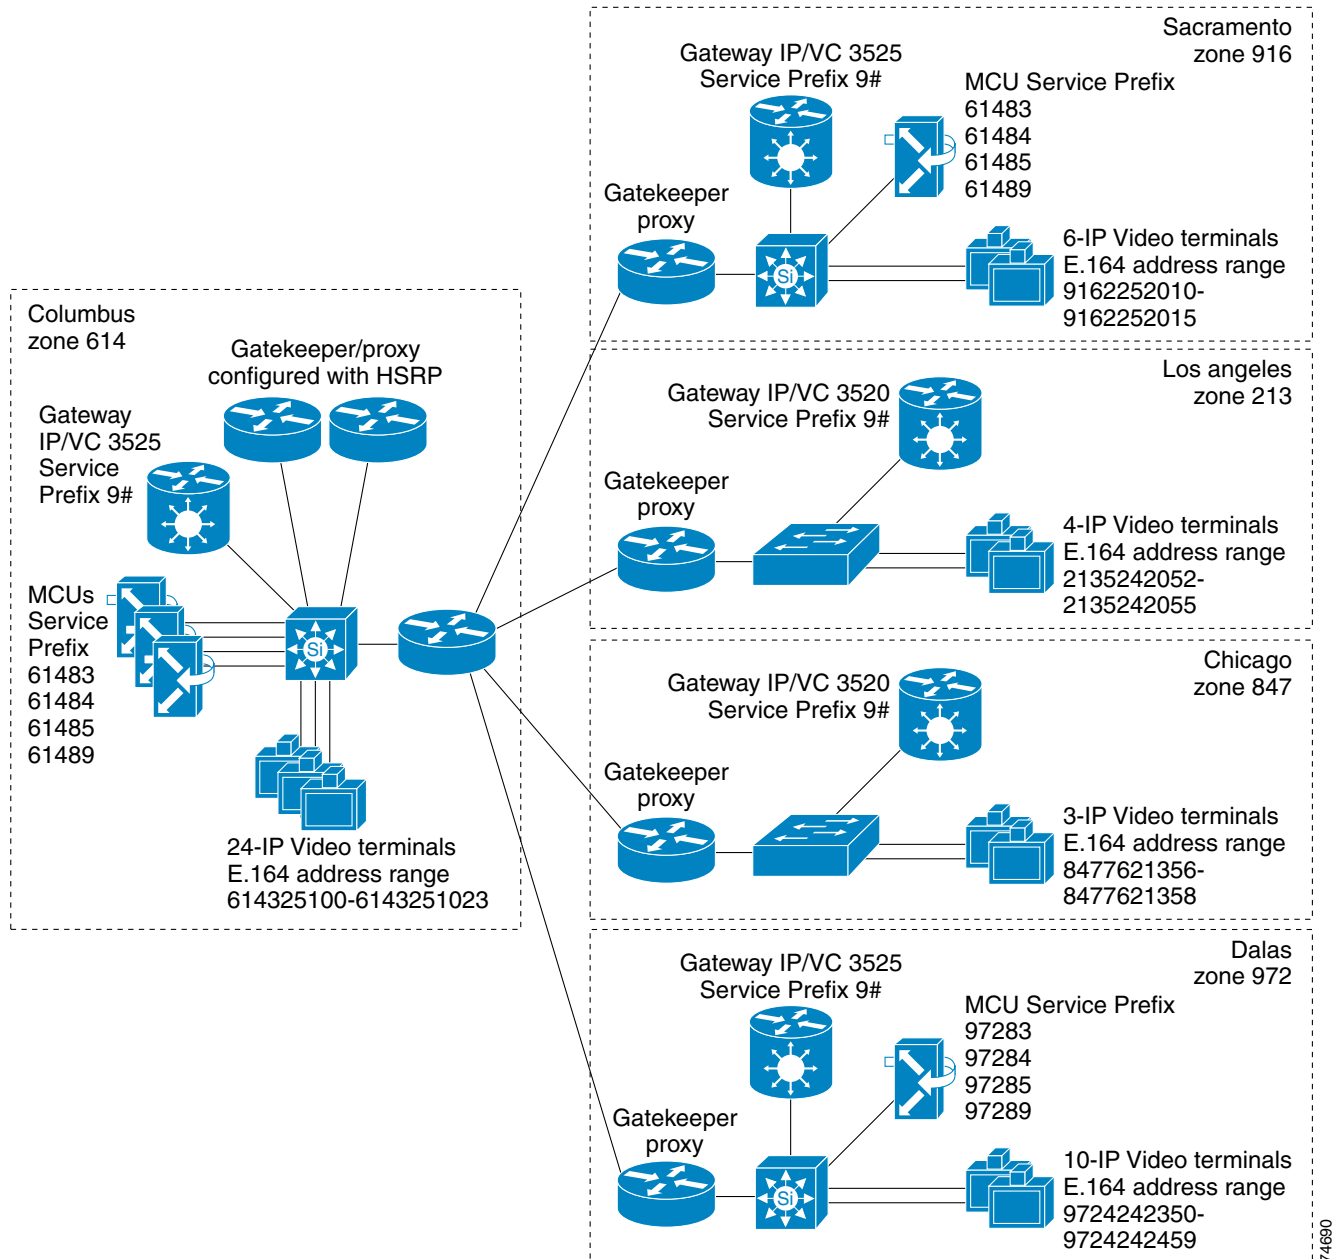
Resource Reservation Protocol (RSVP)

Enterprises are starting to deploy large-scale IP videoconferencing networks using Cisco AVVID solutions based on the Multimedia Conference Manager (MCM) and Cisco IP/VC products. One of the biggest issues for these deployments is call admission control (CAC), especially the limitations of the gatekeeper bandwidth controls. Currently, bandwidth management is limited to hub-and-spoke configurations, which do not allow video networks to scale adequately. However, by implementing call admission control with Resource Reservation Protocol (RSVP) and managing bandwidth on a hop-by-hop basis, you can scale IP videoconferencing networks to meet the needs of most enterprises.

There are two options available for implementing RSVP:

- Use RSVP for call admission control and queuing.
- Use of RSVP for call admission control, and use Differentiated Services (DiffServ) and Cisco modular QoS to service packet flows.

By decoupling the RSVP setup request from the servicing of media flows, network administrators can scale call admission control without maintaining RSVP state for every video call across the entire network.

Note   Currently, RSVP synchronization is not supported. This means that, even if an RSVP request fails, the video call will go through on a best-effort basis. With a future release of the Cisco IOS Proxy, RSVP synchronization will be supported.

802.1P **802.1Q**	802.1P and 802.1Q are the standards proposed by the inter-working task groups of the 802 standards committee. 802.1Q is the IEEE Standard for Local and Metropolitan Area Networks: Virtual Bridge Local Area Network (VLAN). 802.1P is the IEEE Standard for Local and Metropolitan Area Networks – Supplement to Media Access Control (MAC) Bridges: Traffic Expending and Dynamic Multicasting Filtering.
ARQ	Admission Request
AVVID	Architecture for Voice, Video, and Integrated Data
BRI	Basic Rate Interface
CAC	Call Admission Control
Cascade	The process of connecting two or more MCUs to create a larger conference
CODEC	Coder-Decoder, for digitizing voice and video. Compression algorithms can also be used during the digitizing process.
CoS	Class of Service
cRTP	Compressed Real-time Transport Protocol
DID	Direct Inward Dialing
DSCP	Differentiated Services Code Point is an Internet Engineering Task Force (IETF) standard that uses six bits in the TOS (Type of Service) field of the Ipv4 header to specify class of service for each packet.
DTMF	Dual Tone Multi-Frequency
E.164	Address format used for H.323 devices
Gatekeeper	Used for H.323 registration, call routing, and admission control
G.711	G.711 pulse code modulation (PCM) encoding provides 64 kbps analog-to-digital conversion using mu-law or a-law
H.261	Video codec
H.263	Video codec
H.323	Standard for audio, video, and data communications across IP-based networks
H.323-ID	Alphanumeric identifier assigned to an H.323 video terminal
Hopoff	Statement added to a Cisco gatekeeper for static inter-zone routing

HSRP	Hot Standby Routing Protocol
IMUX	Inverse Multiplexer
IP	Internet Protocol
IP Precedence	IP Precedence uses the three precedence bits in the TOS (Type of Service) field of the Ipv4 header to specify class of service for each packet.
ISDN	Integrated Services Digital Network
IVR	Interactive Voice Response
LAN	Local Area Network
LLQ	Low Latency Queuing is a QoS mechanism that ensures the timely queuing of critical, delay sensitive traffic.
LRQ	Location Request
MC	Multipoint Controller
MCM	Multimedia Conference Manager
MCU	Multipoint Conference Unit, used for videoconferences containing more than two endpoints
MP	Multipoint Processor
MSN	Multiple Subscriber Number
PRI	Primary Rate Interface
Proxy	H.323-to-H.323 gateway used for assigning QoS and security access
PSTN	Public Switched Telephone Network
QoS	Quality of Service
RAS	Registration, Admission, and Status protocol
RRQ	Registration Request
RSVP	Resource Reservation Protocol
RTCP	Real-time Control Protocol
RTP	Real-time Transport Protocol
Service Prefix	A digit string used to identify a service on an MCU or gateway
Stacking	Grouping MCUs to obtain a larger number of multipoint conferences
ToS	Type of Service
WAN	Wide Area Network

WRED	Weighted Random Early Detection
WRR	Weighted Round Robin
Zone	A logical group of H.323 infrastructure components managed by a single gatekeeper
Zone Prefix	A digit string used to identify a group of H.323 devices

Cisco IP Videoconferencing Solution Reference Network Design Guide

W

WAN

capacity planning   **5-2**

case study   **9-1**

infrastructure   **4-1**

multi-zone case study   **9-1**

multi-zone deployment model   **2-7**

multi-zone traffic classification   **4-7**

QoS   **5-1**

single-zone deployment model   **2-5**

single-zone traffic classification   **4-3**

traffic classification   **5-3**

traffic prioritization   **5-3**

weighted fair queuing (WFQ)   **5-3**

WFQ   **5-3**

wildcard

in service prefix   **6-2**

in zone prefix   **6-6, 7-11**

World Wide Web   **ix**

Z

zone prefix

defined   **6-1**

design   **6-6**

directory gatekeeper   **7-10**

dot (.) method   **6-6**

format   **6-6, 7-10**

wildcard   **6-6, 7-11**