



CHAPTER 3

Campus Infrastructure

Last revised on: October 30, 2009

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](#)

What's New in This Chapter

[Table 3-1](#) lists the topics that are new in this chapter or that have changed significantly from previous releases of this document.

Table 3-1 *New or Changed Information Since the Previous Release of This Document*

New or Revised Topic	Described in:
Traffic classification	Traffic Classification Types, page 3-4

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 control units (MCUs) connected to a dedicated 10/100/1000 switched Ethernet port. Cisco recommends using a 100/1000-Mbps full duplex connection to the Cisco gatekeeper to ensure adequate bandwidth on all router platforms. Some endpoints, however, do not support 100/1000-Mbps full duplex. For example, older Polycom ViewStations and the Cisco Unified Videoconferencing 3530 both support 10-Mbps half duplex only.



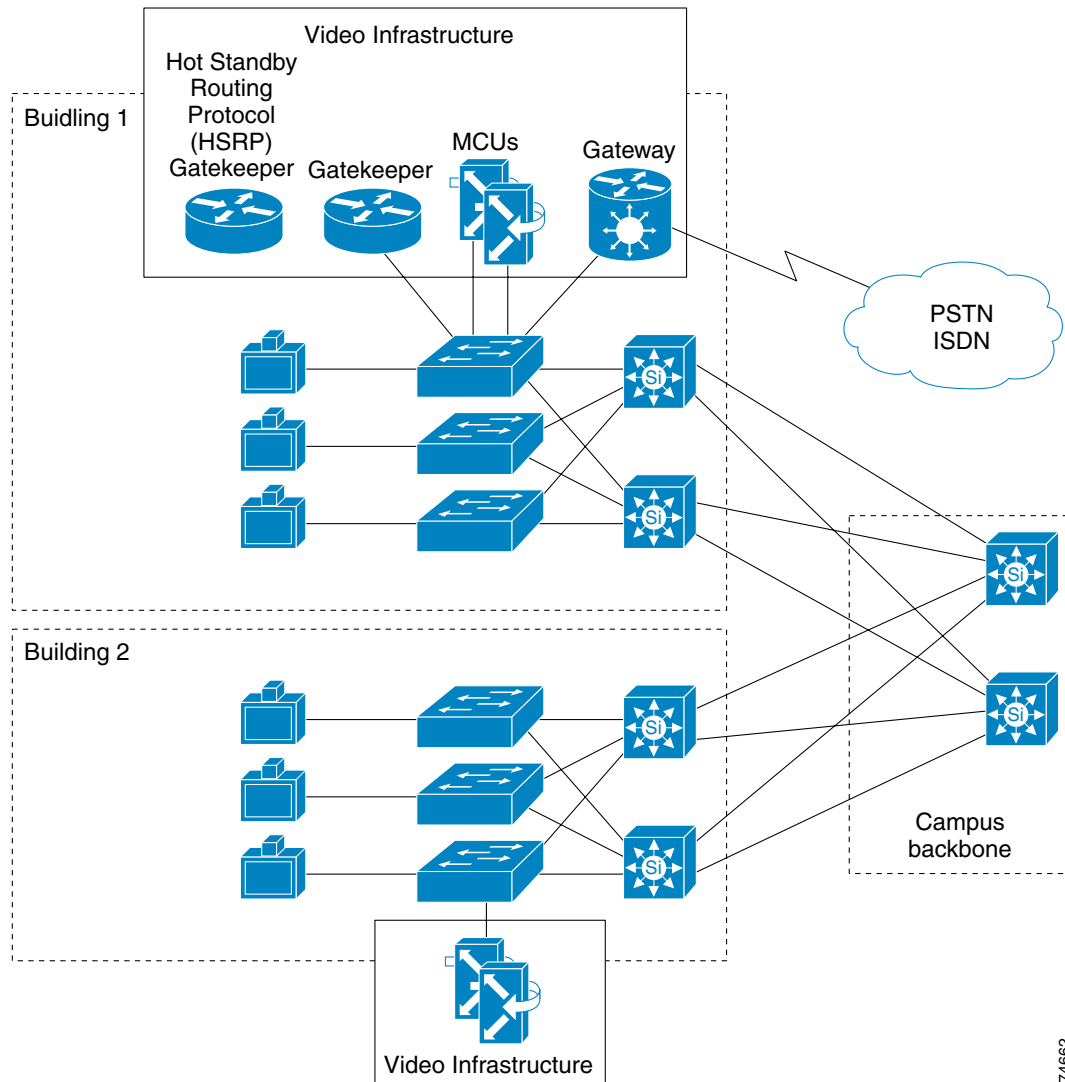
Note

Cisco recommends that you set all switch ports attached to H.323 video devices to 10/100/1000 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



74662

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:

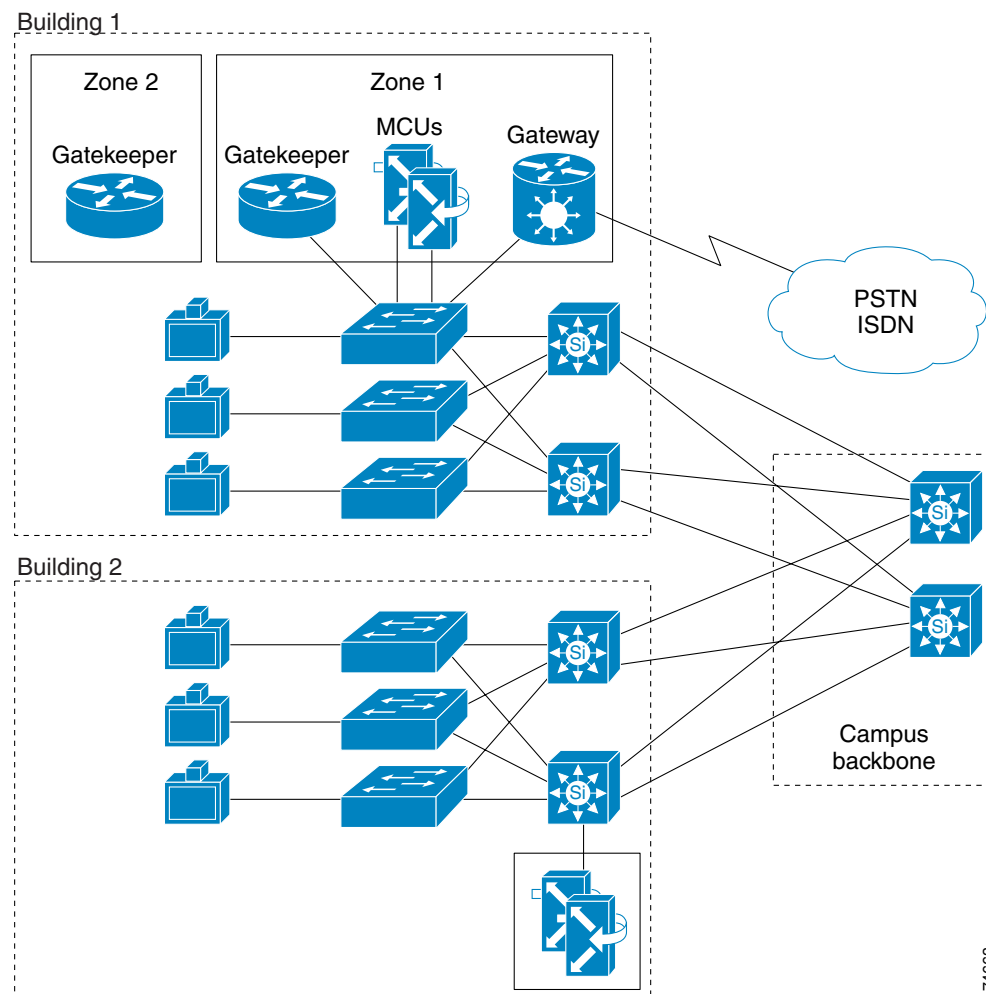
- Campus environment

- 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



74663

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 gatekeeper. If you configure multiple local zones on a single gatekeeper, 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, Tandberg, and Sony), 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 into various data classes with data queues 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 the Cisco Unified Border Element using a Cisco IOS feature. For H.323 signaling and RTP media, you can use access control lists to classify videoconferencing traffic by transport type and port ranges. [Table 3-2](#) lists the recommended traffic classifications for various applications.

Table 3-2 Traffic Classification Guidelines for Various Types of Network Traffic

Application	Layer-3 Classification			Layer-2 Classification
	IP Precedence (IPP)	Per-Hop Behavior (PHB)	Differentiated Services Code Point (DSCP)	Class of Service (CoS)
Routing	6	CS6	48	6
Voice Real-Time Transport Protocol (RTP)	5	EF	46	5
Videoconferencing	4	AF41	34	4
Streaming video	4	CS4	32	4
Call signaling ¹	3	CS3 (currently) AF31 (previously)	24 (currently) 26 (previously)	3
Transactional data	2	AF21	18	2
Network management	2	CS2	16	2
Scavenger	1	CS1	8	1
Best effort	0	0	0	0

1. The recommended DSCP/PHB marking for call control signaling traffic has been changed from 26/AF31 to 24/CS3. A marking migration is planned within Cisco to reflect this change, however many products still mark signaling traffic as 26/AF31. Therefore, in the interim, Cisco recommends that both AF31 and CS3 be reserved for call signaling.

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 *Enterprise QoS Solution Reference Network Design Guide*, available at <http://www.cisco.com/go/designzone>.

QoS Features Summary

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

Table 3-3 Supported QoS Features by Switch Platform

Platform	Auto QoS	Reclassify CoS/DSCP	Congestion Avoidance	Priority Queues	Multiple Queues	Traffic Management	Policing
Catalyst 2960	Yes	Yes	Yes (WTD ¹)	Yes	4 egress queues/port	SRR ²	Yes
Catalyst 3560	Yes	Yes	Yes (WTD ¹)	Yes	4 egress queues/port	SRR ²	Yes
Catalyst 3760	Yes	Yes	Yes (WTD ¹)	Yes	4 egress queues/port	SRR ²	Yes (64 policy rates or individual port)
Catalyst 4006 or 450x with Supervisor Engine IV	Yes	Yes	Yes (DBL ³ and QoS sharing on non-blocking Gb ports)	Yes	4 egress queues/port	SRR ²	Yes (64 policy rates or individual port)
Catalyst 4006 or 45xx with Supervisor Engine V	Yes	Yes	Yes (DBL ³ and QoS sharing)	Yes	4 egress queues/port	SRR ²	Yes (64 policy rates or individual port)
Catalyst 6000 with Policy Feature Card (PFC3)	Yes	Yes	Yes (scheduling and QoS sharing)	Yes	4 egress queues/port	SRR ² or DWRR ⁴	Yes (64 policy rates or individual port)

1. Weighted Tail Drop (WTD)
2. Shaped Round Robin (SRR)
3. Dynamic Buffer Limiting (DBL)
4. Deficit Weighted Round Robin (DWRR)

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.