IP Bandwidth Guide

Timothy M. O'Neil
Director of Technical Marketing
Polycom Video Communications
January 28, 2003

POLYCOM®

Table of Contents

List of Figures

IP Bandwidth Guide

Introduction

The purpose of this paper is to explore IP (H.323) bandwidth issues, as well as dispel some common misconceptions about the call bandwidth requirements of different call scenarios over TCP/IP networks.

Bandwidth for H.323 (IP) calls is based on the legacy call-quality math of switched circuit networking (SCN). This math is based specifically on the division of bandwidth into 64-Kbps increments or one DS0.

For reference, a standard ISDN line is composed of two 64-Kbps DS0s for a total transmission bandwidth of (2 x 64 Kbps) = 128 Kbps. There is also a signaling channel that is not part of the payload capability and beyond the scope of this document. Using historic SCN call quality measurements gives us relative parity in comparing the quality of H.320 (ISDN) calls to H.323 (IP) calls. However, call quality is not an accurate measurement of call bandwidth requirements as will be discussed in this paper.

The most common business quality video communications call quality used is 384 Kbps. Over ISDN, this requires six DS0s (or 6 x 64 Kbps = 384 Kbps). Therefore, the same number would be used (384 Kbps call quality) for an equal call quality of 384 Kbps over IP. However, additional considerations must be taken into consideration depending on the following factors:

- Is it a half or full duplex transmission?
- Which WAN technology is implemented?
- Is it a point-to-point call?
- If it is a multipoint call which particular multipoint control unit (MCU) is being used?

General H.323 Bandwidth Recommendations

H.323 traffic will use a slightly larger amount of bandwidth than the selected call quality or H.320 equivalent. Polycom, therefore, recommends that you allow for a 20% overhead for the H.323 signaling traffic on top of the media (audio, video, and T.120 data). This H.323 traffic overhead relates to the signaling that is required by the TCP/IP and H.323 protocols. As mentioned above, ISDN networks do not record this signaling in the payload calculations, as it is out of band. This is not the case in TCP/IP networks; all signaling must also be accounted for when making capacity decisions for the provisioning of LAN and WAN segments. For example, a 384-Kbps video call would consume approximately 384 Kbps + 20% = 460 Kbps of bandwidth on a TCP/IP network.

Even though H.323 videoconferencing is a bi-directional application, full-duplex network segments accommodate the transmission of packets in both directions simultaneously. For example, a full-duplex 100 Mbps Ethernet (100Base-T) segment actually has 100 Mbps of bandwidth in *each* direction. Therefore, bi-directional H.323 traffic does not require that full duplex network segments be provisioned for two times the bandwidth.

Please note that even if H.323 traffic starts out on a half-duplex network segment (thus requiring 2 x the bandwidth), it will take advantage of full-duplex segments as soon as it reaches them and, therefore, only consumes 1 x the bandwidth on those full-duplex segments. Whether or not you need to multiply the bandwidth by 2 is completely dependent upon the network, once it leaves the videoconferencing equipment. Understanding this is crucial because WAN segments (T1, Frame Relay, ATM) are typically full duplex.

Bandwidth Requirements for H.323 Traffic over Different LAN/WAN Technologies

Following are some simple equations to assist in determining the bandwidth required for H.323 traffic across various network segments:

- Full-duplex Ethernet = (Call Speed + 20%) x 1
- Half-duplex Ethernet = (Call Speed + 20%) x 2
- Wide Area Network = (Call Speed + 20%) x 1
- ATM (Using LANE) = (Call Speed + 35%) x 1

Point-to-Point Calls

The following example assumes a call quality of 384 Kbps. It should be noted that Polycom terminals can conduct videoconferences at both lower and much higher call quality speeds.

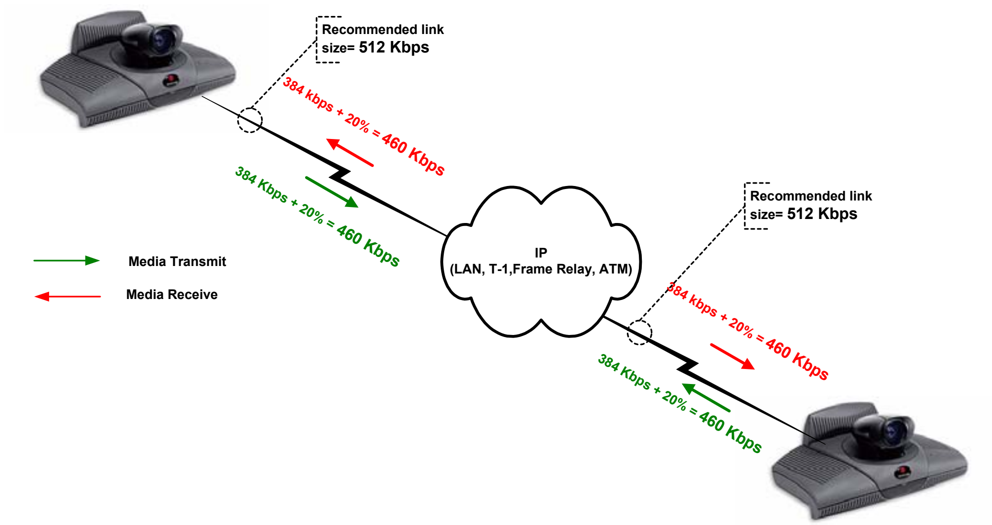

Figure 1-1. Point-to-Point Calls

Bandwidth Requirements Table

The Bandwidth Requirement Table is provided for reference.

Call Quality or Dialing Speed	Bandwidth Required over ISDN (H.320)	Bandwidth Required over IP (H.323)
128 Kbps	1 Basic Rate ISDN (BRI) line	153 Kbps
256 Kbps	2 BRI lines	307 Kbps
384 Kbps	3 BRI lines	460 Kbps
512 Kbps	4 BRI lines	614 Kbps
768 Kbps	Fractional T1* or full Primary Rate ISDN (PRI) line	922 Kbps
1.5 Mbps	1 PRI line	1.843 Mbps
2.0 Mbps	Multiple* PRI lines or E1 line (Europe)	2.4 Mbps

* Requires a third-party inverse multiplexer. Inverse multiplexers provide inverse multiplexing to transmit a single high-speed data channel over one or many T1 (PRI) or E1 links.

Multipoint Calls

ViewStation FX/VS4000 Multipoint Calls

This example assumes a speed of 384 Kbps.

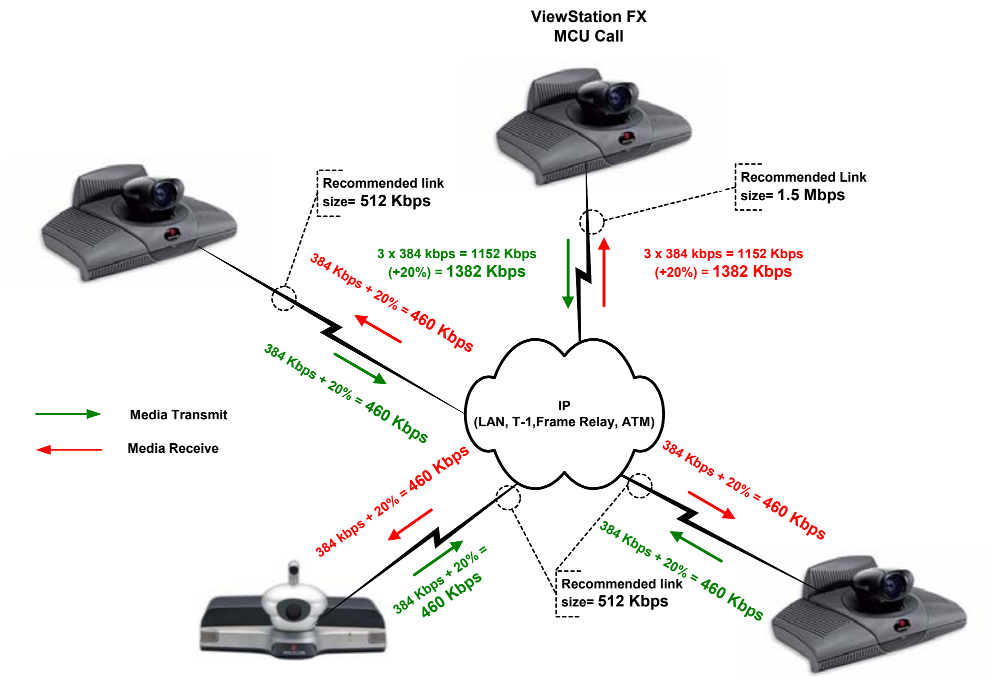

Figure 1-2. ViewStation FX/VS4000 Multipoint Calls

Each site connects to the ViewStation FX or VS4000 host site with a symmetrical speed of 384 Kbps. Whether it is a voice-activated or continuous-presence conference, only 384 Kbps is sent back to each of the remote sites.

The ViewStation FX or VS4000 host location will need to have the required bandwidth to accommodate the sum of all the remote participants.

For example, the ViewStation FX/VS4000 host location would need to have approximately 1382 Kbps of bandwidth to handle this multipoint call at 384 Kbps.

iPower 9000/900/600 Series Voice-Activated Multipoint Calls

This example assumes that the three remote locations are dialing at a speed of 384 Kbps (320 Kbps video + 64 Kbps audio = 384 Kbps).

The iPower MCU uses techniques to optimize the use of IP bandwidth. These techniques leverage flow control mechanisms. Flow control allows for the economical use of bandwidth. Flow control essentially tells endpoints to stop sending media when it is not required for retransmission.

The following example shows that when in voice-activated mode (a.k.a video-switched or full-screen mode), only two video streams (the current broadcaster and previous broadcaster at 2 x 320 Kbps = 640 Kbps) are accepted into the MCU and the three audio streams from remote participants (3 x 64 Kbps = 192 Kbps) are accepted into the MCU. The other sites on the conference are flow-controlled on the video stream—in other words the remote endpoints stop sending the video stream. These two parameters of audio added to video equal 640 Kbps + 192 Kbps = 832 Kbps, which is the maximum bandwidth accepted by the MCU. Add 20% to this number (832 Kbps + 20% = 998 Kbps) to obtain the actual IP bandwidth required to support this call.

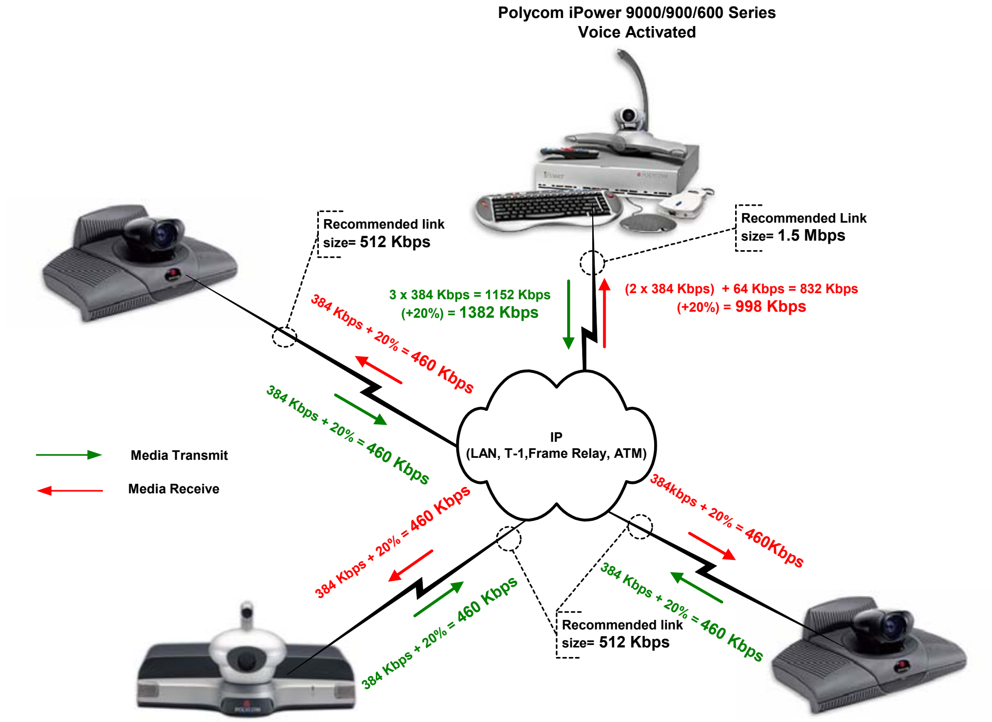

Figure 1-3. iPower 9000/900/600 Series MCU—Voice-Activated Multipoint Calls

Each site connects to the iPower host site with a symmetrical speed of 384 Kbps. Since it is a voice-activated conference, only 384 Kbps is sent back to the sites because each participant views only one site at a time.

The iPower host location will need to have the required bandwidth to accommodate the sum of all the remote participants. Specifically, the iPower host location would need to have approximately 1382 Kbps of bandwidth to handle this four-location (three remote + one local) multipoint call at 384 Kbps.

iPower 9000/900/600 Series Continuous Presence Multi-point Calls

This example assumes a speed of 384 Kbps. Inbound media to the MCU will be 518 Kbps; all remote sites will transmit media of 172.6 Kbps per site.

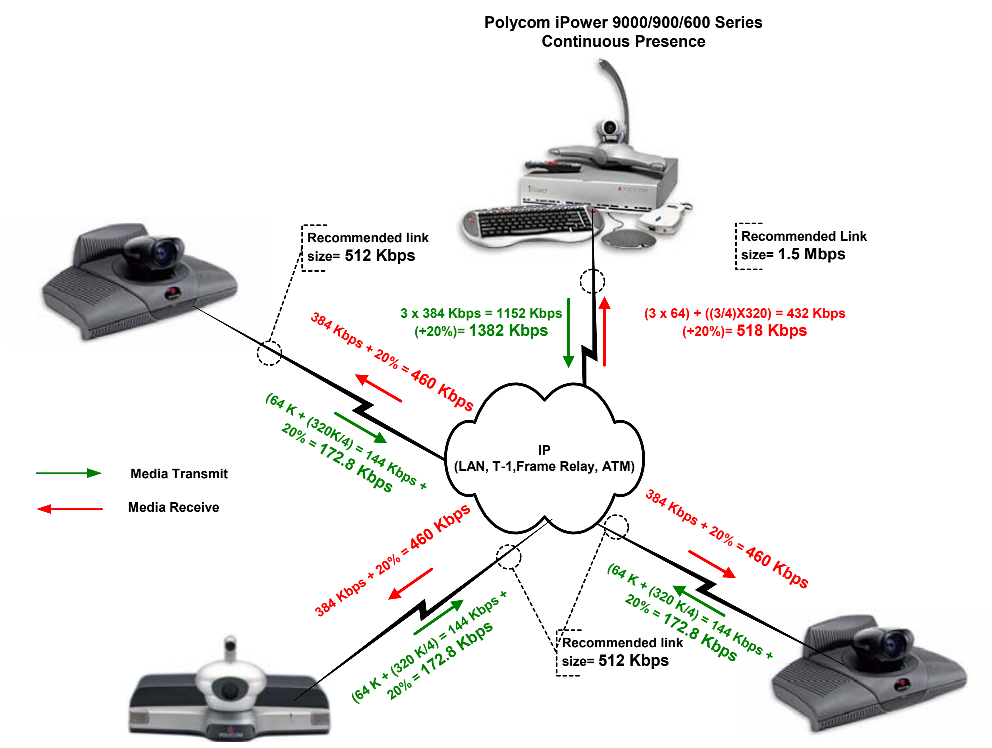

Figure 1-4. iPower 9000/900/600 Series MCU—Continuous Presence Multipoint Calls

Continuous Presence mode on the iPower series MCU is different from the voice-activated bandwidth usage described previously. Each site has an asymmetrical connection to the iPower host site. For example, for a 384 Kbps conference, the three sites are connected at 144 Kbps (64 Kbps + 320 /4= 144 Kbps, where the audio rate = 64 Kbps) to the iPower host site, with a QCIF resolution. Each remote site receives 384 Kbps and a CIF image back from the iPower host site. The other sites on the conference, the sites that are not part of the present continuous presence mix, are flow-controlled on the video stream, meaning that the endpoint stops sending the video stream. Note that these numbers do not take into account the required 20% overhead.

The formula for calculating outbound MCU bandwidth is:

(number of remote sites) x (audio rate) + ((the lower number or MIN of (# remote sites, or the constant 4) / 4) x (video bit rate of remote sites))

Note: MIN is used when selecting between variables. The MIN of (3, 4) would be 3. MIN is the lesser number.

Mathematically, this would look like the following:

(3X64K) + ((MIN (3, 4)/4) X 320K) = 432 Kbps

The iPower location that hosts the multipoint call needs to have the required bandwidth to accommodate the sum of all the remote participants. For example, the iPower host location would need to have approximately 1152 Kbps of bandwidth (plus 20%) to handle the three sites at 384 Kbps that are connected to it.

Polycom MGC Multipoint Call Types

The Polycom MGC platform supports three distinct variations of multipoint calls. These modes include video-switched mode (additional media processing hardware is not required), hardware continuous presence with transcoding (media processing hardware is required), and software-based continuous presence (additional media processing hardware is not required).

Video-Switched Mode Call on Polycom MGC

In video-switched mode (also known as voice-activated switching), the MCG does not process the video stream; it merely switches the video stream. Switching refers to selecting a video site to become the broadcast site. The active speaker is the broadcast video source displayed on all remote sites. When a new loudest speaker begins talking, the MCG instantly switches to broadcasting this new speaker to all remote sites.

Bandwidth requirements for video-switched mode are symmetric for inbound (media receive) and outbound (media transmit) streams. Bandwidth is negotiated at conference start through an H.245 capabilities exchange. The negotiated common parameters now force all endpoints to either comply with the common capability, or in the case where the endpoint cannot meet the common conference capabilities, the connection will be audio only. Audio algorithm and data bandwidth (T120) define the video portion of the stream.

It should be noted that in this scenario, the media processing capabilities of the MGC are not being used. The MGC has the unique capability of transcoding video calls on five different parameters (call speed, video protocol, audio protocol, video frame rate, and T.120 data speed) to ensure all participants can conference with their unique capabilities. For more information on the transcoding capabilities of the MGC, please refer to the product documentation posted on www.polycom.com.

The following example assumes a speed of 384 Kbps, although the Polycom MGC can support both higher and lower call speeds.

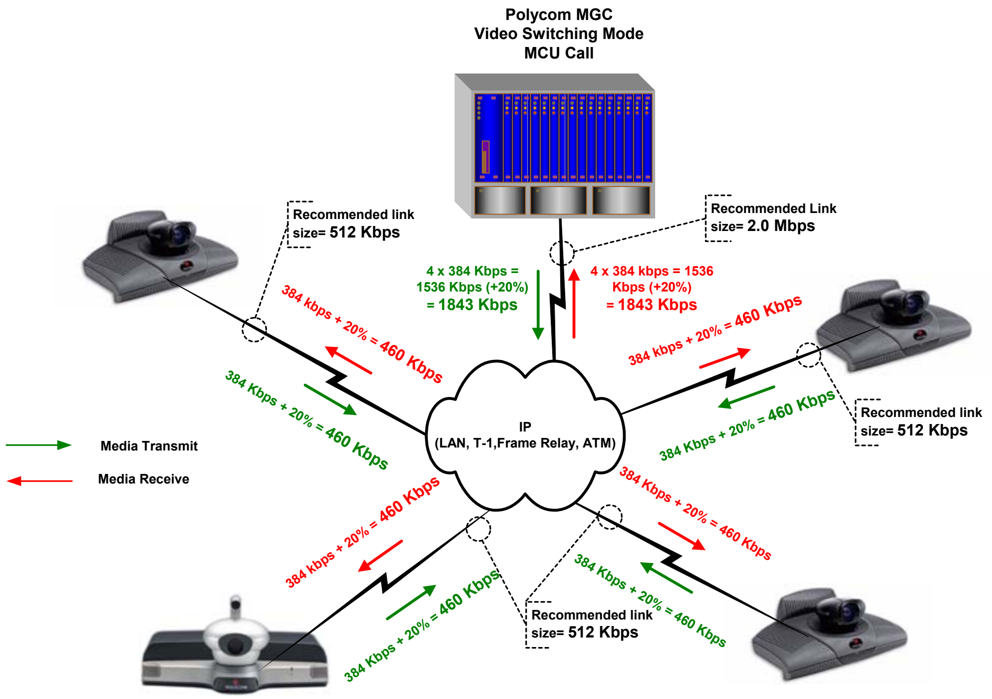

Figure 1-5. Video-Switched Mode Call on Polycom MGC

In this example, each site connects to the MCU with a symmetrical speed of 384 Kbps. Since it is a voice-activated conference, only 384 Kbps is sent back to the sites because each participant can only view one site at a time.

The MCU location that hosts the multipoint call needs to have the required bandwidth to accommodate the sum of all the remote participants.

For example, the MCU location would need to have approximately 1.8 Mbps of bandwidth to handle the four remote sites at 384 Kbps connecting to it.

Hardware-Based Continuous Presence MCU Call on Polycom MGC

The second multipoint conferencing mode is hardware-based and can also include transcoding. In our example, we will be leveraging two distinct and unique features to the Polycom MGC platform. The first distinct and unique feature is hardware-based continuous presence which allows for multiple video layouts. This is often referred to as “Hollywood Squares.” This feature is described in great detail in the documentation posted on www.polycom.com.

The second distinct and unique feature about the MGC discussed is transcoding. Transcoding allows for endpoints with different connection and protocol capabilities to conference together.

This allows for the highest level of call connectivity in the videoconferencing industry. Not every location has the same bandwidth, video protocol, and audio protocols capabilities. The only limitation is that the maximum bit rate can be equal to or lower than the conference setting. Note that the specific link between the MCU and endpoints is symmetric.

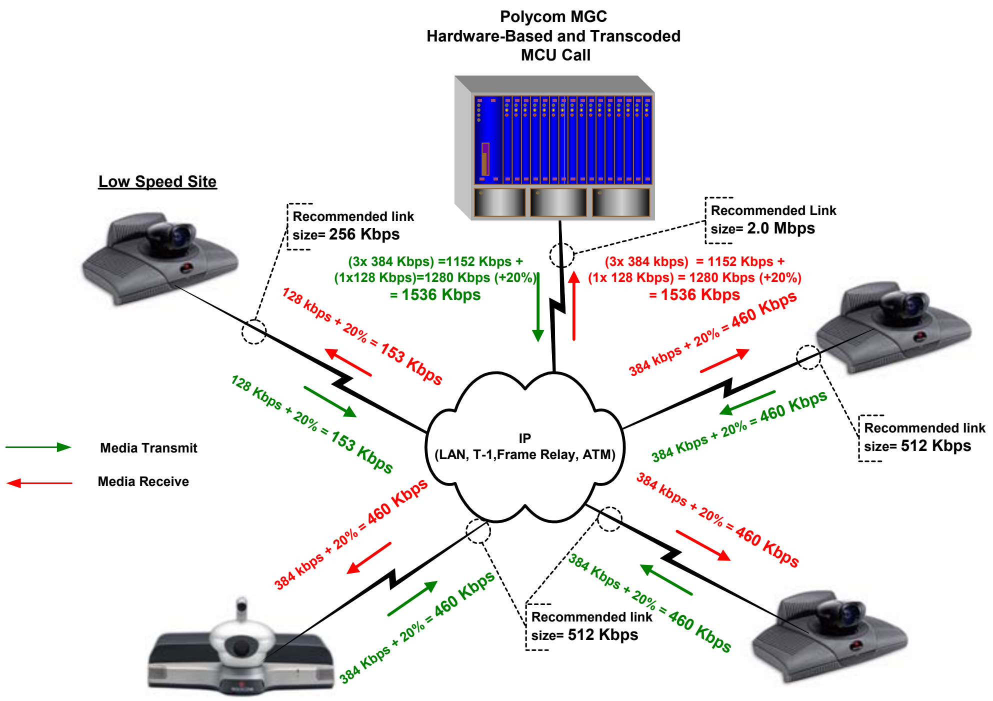

Figure 1-6. Hardware-Based Continuous Presence MCU Call on Polycom MGC

Hardware continuous presence works in a very similar manner as the voice-activated conference. Each site connects to the MCU with a symmetrical speed. The continuous presence hardware in the Polycom MGC decodes all the sites and builds the continuous presence screen, then sends the composite stream of the built continuous presence screen to all the participating sites. In the scenario above, the slow speed site is transmitting a symmetrical 153 Kbps, and the three other locations are transmitting a symmetric 460 Kbps.

Each site receives the same media rate that it transmits. This media rate is inclusive of both the audio video media. If all sites were transmitting a 384-Kbps call speed, the math would be exactly the same as that used for video-switched mode.

The MCU location that hosts the multipoint call needs to have the required bandwidth to accommodate the sum of all the remote participants. For example, the MCU location would need to have approximately 1.54 Mbps of bandwidth

Software-Based Continuous Presence MCU Call on Polycom MGC

The third multipoint conferencing mode supported on the MGC platform is called software-based continuous presence. In this mode, which is only supported in IP-based calls, hardware media processing resources are not required to support the call. This is a cost saving compared to hardware-based continuous presence. It is less expensive to purchase an MGC platform that supports only software-based continuous presence. As with anything that costs less, the feature support is also less. Software-based continuous presence does not allow for endpoints with dissimilar capabilities. It operates just like video-switched mode.

Software-based continuous presence does not support the multiple video formats supported in Hardware-based continuous presence with transcoding. If a site cannot meet the common conference parameters negotiated in H.245 during conference setup, it will be relegated to an audio-only participant. This works the same as the video-switched mode mentioned previously.

The one important detail to understand, when not using hardware, is the change from a symmetrical to an asymmetrical broadcast mode. Asymmetrical broadcast mode requires that the receiving speed on remote sites be four times their transmitting speed. This mode relies on the media processing capabilities in the remote terminal to handle the processing instead of the centralized MCU. It is calculated as follows:

(Return video bandwidth) = 4 x (Transmit bandwidth) + 20%

This rule can be applied to any call-quality speed. For example, if the remote terminals were all dialing 128 Kbps and there were six of them, the maximum return would be (4 x 128 Kbps) + 20% = 614.4 Kbps.

The other important difference is in the video transmit versus receive. All remote sites must negotiate to transmit in QCIF format. Then, the MGC will transmit in FCIF format. This allows the MCG to simply switch the incoming QCIF streams into one outgoing FCIF stream. It should be noted that the maximum video display format is four sites (quad screen).

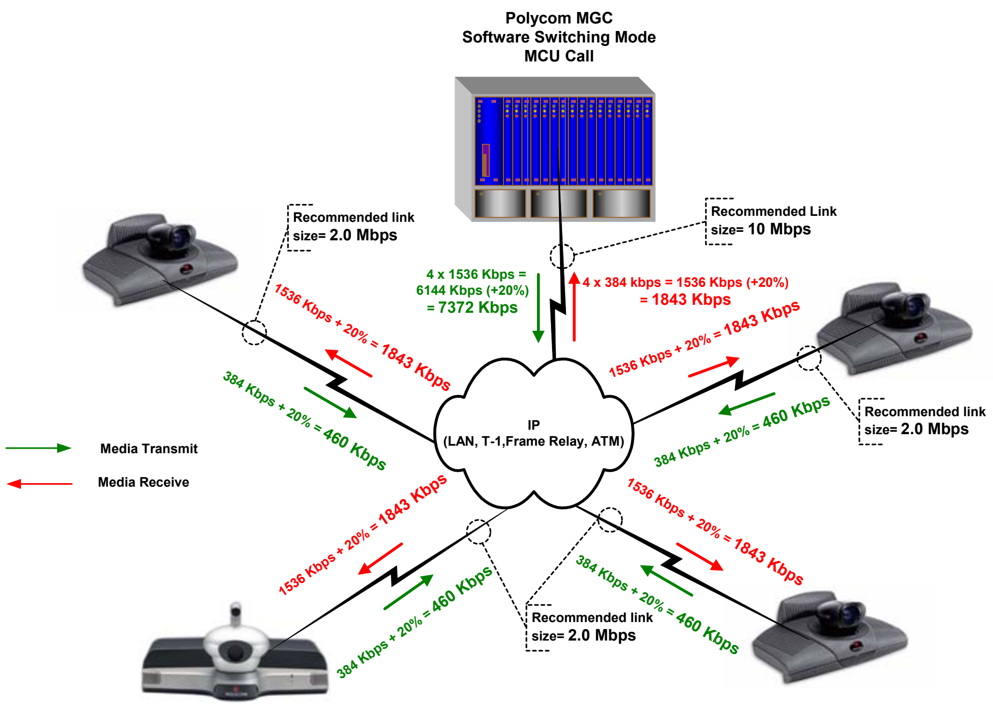

Figure 1-7. Software-Based Continuous Presence MCU call on Polycom MGC

Software-based continuous presence works differently than either of the other two methods mentioned previously. Each site has an asymmetrical connection to the MCU: if we consider our diagram above, the bandwidth coming from the endpoint to the MCU is 1/4 of the defined multipoint conference speed. The speed defined for the conference is the speed that is sent back to each site.

For example, the conference speed was defined as 384 Kbps. The MCU receives each site's connection at 96 Kbps rate and then sends the four 96 Kbps streams that are to be displayed in the continuous presence layout to each site. The MCU then builds the screen at each endpoint, which is why the connection is asymmetrical. Note that these numbers do not include the required 20% overhead.

The MCU location that hosts the multipoint call needs to have the required bandwidth to accommodate the sum of all the remote participants.

For example, the MCU location would need to have approximately 1.5 Mbps of bandwidth (plus 20%) to handle the 4 remote sites at 384 Kbps connecting to it.

H.323 Bandwidth Recommendation Summary

We hope that the math illustrated with various diagrams has helped you gain a better understanding of the actual bandwidth requirements for business-quality video communication over TCP/IP networks.

Bandwidth Provisioning. Remember to provision your WAN link with the adequate amount of bandwidth. As an example, provisioning a WAN link for 384-Kbps data service and expecting it to be able to support the actual 460 Kbps that is required will not work. 512 Kbps would be the minimum standard link size to accommodate this call. This 512-Kbps size is recommended only if video communications is the only application that will be traversing this link. If the link is shared with other applications, please review the white papers on www.polycom.com for more recommendations on best practices.

Half- or Full-Duplex Misconception. The actual bandwidth available to a half- or full-duplex TPC/IP network interface card on the switch port or video communications device will be barely impacted by even the highest call-quality speed. As an example, if a ViewStation were to make a 384-Kbps call over a half-duplex, 10-Mbps connection, it would use less than 10% of the available link capacity. If a ViewStation were to make a full-duplex call over a 10-Mbps switch port, it would use less than 5% of the link capacity. Therefore, whether or not a video communications terminal or Ethernet switch is half or full duplex doesn't really matter in the larger scheme of deployment issues. Ultimately, this comes down to plus or minus 5% of 10,000,000 available bits....

Bandwidth Consumption Determined by the MCU Model. The actual bandwidth used by an MCU depends on the particular model and the way it is handling media. Three modes are available to the MCU to handle video:

1. Video-switched mode
2. Hardware-based continuous presence with or without transcoding

3. Software-based continuous presence

Cost-Saving Suggestions. Because software-based continuous presence is asymmetric, it can cost more in WAN link capacity expenses than using a media processing MCU. As an example, if an organization needed to support a four-site, 384-Kbps call once a week, the actual bandwidth cost for a network would be drastically different depending on the chosen MCU.

Rather than having to set up a full PRI link (1.5 Mbps) at each location, as would be required of a software-based MCU, you could set up 512 Kbps circuits at three locations and set up one site with 1.5 Mbps to handle the media processed MCU. This could end up costing a lot less over a three-year period, depending on the actual call frequency. This would equate to 66% less in WAN link costs for three sites.

Contact Information

Corporate Headquarters
Polycom, Inc.
4750 Willow Rd.
Pleasanton, CA 94588-2708
USA

Phone: +1.925.924.6000
Fax: +1.925.924.6100

European Headquarters
Polycom Ltd.
270 Bath Road
Slough
Berkshire SL1 4DX
United Kingdom

Phone: +44 (0) 1753 723000
Fax: +44 (0) 1753 723010

Asia Pacific Headquarters
Polycom Solutions Pte. Ltd.
16 Raffles Quay
#40-02A, Hong Leong Building
Singapore 048581

Phone: +65.323.3882
Fax: +65.323.3022

For additional information, please visit the Polycom web site at: www.polycom.com

POLYCOM®