PictureTel

H.323 Videoconferencing Network Bandwidth Analysis

by

John Bartlett

NSD Engineering

PictureTel Corporation

August 25, 1997

Introduction

This document evaluates the network bandwidth required for various configurations of H323 videoconferencing clients and MCUs. The merits of each approach are discussed. The bandwidth impact on local links, the corporate backbone, and the WAN access link is examined.

Link Bandwidth and Point to Point Videoconferencing Sessions

Client Bandwidth Requirements

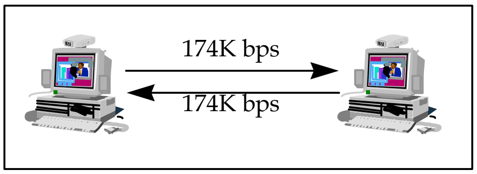

Figure 1 - Simple Videoconferencing Session

Each client in a videoconference sends audio and video information into the network, and receives information from the network. A simple point-to-point videoconference is shown in Figure 1. LiveLAN 3.0 uses approximately 174Kbps of bandwidth in each direction when using its 174K bandwidth option. The two directions are shown here as two different traffic streams. Each client sends audio and video data to the other client. Bandwidth utilization for this configuration is 174Kbps on a full duplex link (for example, a T1 or a full duplex Ethernet), and 348Kbps for a half duplex link (for example, a shared Ethernet.)

Table 1 shows the bandwidth requirements for a number of H323 clients. The two columns indicate one-way bandwidth, and bidirectional bandwidth. One-way bandwidth is what is seen on each side of a full duplex link. The bidirectional bandwidth is what is seen on a shared LAN segment, such as a normal Ethernet.

	One Way	Bi-directional	Audio	Video
NetMeeting 2.0	64	128	G.723 3.4 KHz - 6.4Kbps	H.263 QCIF 176x144
LiveLAN 3.0 (@174Kbs)	174	348	G.722 7 kHz - 64Kbps	H.261 FCIF 352x288
LiveLAN 3.0 (@384Kbs)	384	768	G.722 7 KHz - 64Kbps	H.261 FCIF 352x288
Intel Internet Video Phone	TBD	TBD	TBD	TBD

Table 1 - Videoconferencing Client Bandwidth Requirements (Kbs)

These bandwidth numbers are the worst case bandwidths for both utilities, with constant motion across the entire video field. In a more typical application, much of the background field is constant, and the videoconferencing client will reduce its bandwidth consumption by eliminating unnecessary updates. Note that different utilities are running different video and audio protocols. Each utility is optimized for a different environment. The quality of their audio and video reproduction is heavily dependent on the bandwidth they consume.

Given the bandwidths in Table 1, Table 2 indicates the percentage of these link types consumed by a single LiveLAN or NetMeeting session. Bandwidth consumption drops for the full duplex links because these videoconferencing clients are by nature full duplex applications, creating symmetric traffic. The T1 and T3 links are full duplex.

	Enet-10 Shared	Enet-10 FD	Enet-100 Shared	Enet-100 FD	T1 FD	T3 FD
NetMeeting	1.28%	0.64%	0.13%	0.06%	4.12%	0.14%
LiveLAN 174K	3.48%	1.74%	0.35%	0.17%	11.27%	0.39%
LiveLAN 384K	7.68%	3.84%	0.77%	0.38%	24.87%	0.86%
Intel Internet Video Phone	TBD	TBD	TBD	TBD	TBD	TBD

Table 2 - Percent of Links Used by Videoconferencing Client

Available Link Bandwidth

Before calculating the number of videoconferencing sessions we can support on each of these links, we need to degrade the available bandwidth of each link to account for its efficiency.

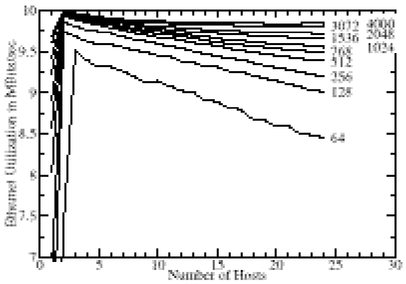

Figure 2 - Ethernet Efficiency

Figure 2 is a chart showing a fully loaded 10 Mbps Ethernet under different configurations of packet size and number of hosts on the segment. This chart (from WRL Research Report 88/4, Digital Equipment Corporation) shows that as the number of hosts increases, the overall utilization of the Ethernet drops. This drop is due to the collisions that occur on the Ethernet as multiple stations attempt to access the network simultaneously.

Figure 2 has information for a 12 µs network, corresponding to an Ethernet cable length of over 3000 feet. The majority of today's Ethernet installations do not use the old Ethernet cable, but instead use 10BaseT wiring and hubs. If end stations are 300 feet from the hub, it is equivalent to a 1.2µs network. Ethernet performance improves as the network gets shorter, because collisions are detected more quickly.

Note also that the majority of traffic created by videoconferencing is in packets of 600 bytes or larger. The Ethernet is better behaved for longer packet sizes.

To be conservative in this document, we will limit the use of the Ethernet to 70% for half duplex operation. However, when an Ethernet is operating in full duplex mode, no collisions can occur. The following calculations allow a 90% utilization for a full duplex link. So the available bandwidth in these links is shown in Table 3.

	Enet 10 Shared	Enet 10 FD	Enet 100 Shared	Enet 100 FD	T1	T3
Full Performance Link Speed	10	10	100	100	1.544	44.736
Allowed Percentage	70%	90%	70%	90%	90%	90%
Available Link Bandwidth	7	9	70	90	1.389	40.262

Table 3 - Available Link Bandwidth (Mbps)

Number of Videoconferencing Sessions per Link

Using the bandwidths listed for videoconferencing clients in Table 1 and the link bandwidths listed in Table 3, the following chart calculates the maximum number of simultaneous full duplex videoconferencing sessions that can take place using that particular link. Note that the full duplex links have the advantage of twice as much bandwidth and higher efficiency, making them capable of more than twice as many videoconferencing sessions.

Table 4 shows that a standard 10Mbps Ethernet can support 20 LiveLAN 174K video-conferencing sessions involving 40 LiveLAN clients.

	Enet 10 Shared	Enet 10 FD	Enet 100 Shared	Enet 100 FD	T1 FD	T3 FD
NetMeeting	54	140	546	1406	21	629
LiveLAN 174K	20	51	201	517	7	230
LiveLAN 384K	9	23	91	234	3	104
Intel Internet Video Phone	TBD	TBD	TBD	TBD	TBD	TBD

Table 4 - Maximum Videoconferencing Sessions per Link

Bandwidth Effect of Multipoint Videoconferences

MCU Connections

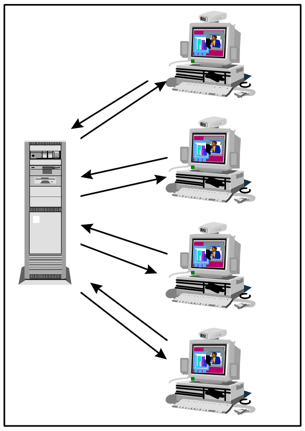

Figure 3 - Videoconferencing Streams

The current implementation of NetConference MCU receives a unicast stream from each videoconferencing client, and delivers back to that client a separate unicast stream. Thus the MCU becomes one end of a full duplex session between each client involved in the call and the MCU (Figure 3).

To determine how many clients can be using an MCU on a single LAN segment, we can refer directly to Table 4. Each session in Table 4 now represents one client connected to the MCU. For an MCU connected via a 10Mbps Ethernet, this would indicate that 20 LiveLAN 174K clients or 54 NetMeeting clients can be connected simultaneously.

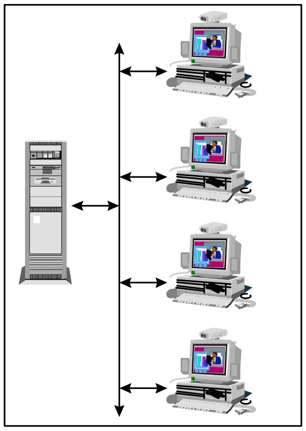

Figure 4 - Physical connectivity

This configuration drawing assumes that all of the clients and the MCU are on the same network segment, (Figure 4) and thus are limited by its bandwidth. With this physical configuration, all traffic streams are passing through the same 10 Mbps Ethernet segment.

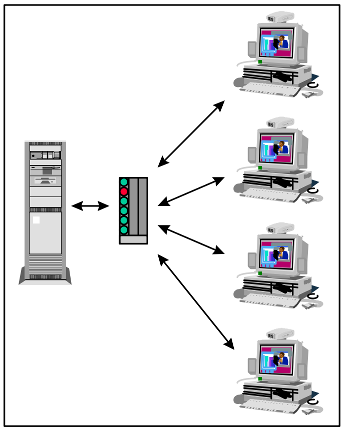

Figure 5 - Workgroup with Ethernet Switch

Adding an Ethernet switch to this configuration changes the picture (see Figure 5). When clients are connected through a switch, only their own traffic is on their individual segments. In Figure 5 each segment on the right of the switch only sees 348Kbps for a LiveLAN 3.0 session.

However, the link between the switch and the NetConference MCU still supports all of the videoconferencing streams simultaneously. For a shared 10Mbps link, this would limit the size of the videoconference to 20 LiveLAN 174K members, or 9 LiveLAN 384K members (see Table 4.) This link can now be upgraded to either a full-duplex 10Mbps link or to a 100Mbps link. With these higher bandwidths the MCU can now support many more clients in a single videoconference.

A more typical configuration in a medium to large office setting is shown in Figure 6. The intermediate boxes are now either switches or routers. In either case they act to manage the bandwidth, only allowing through traffic destined for the client or clients on that particular segment. The backbone can then be sized appropriately to carry the combined traffic requirements of the clients.

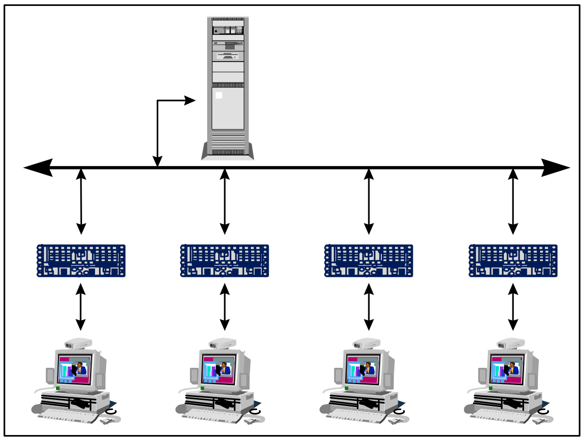

Figure 6 - Building LAN

In this case the NetConference MCU is shown attached directly to the backbone with a high speed interface. The MCU can now support a large number of conferences with clients spread throughout the organization.

The MCU need not be directly attached to the backbone. It could also be behind a switch or a router, as long as its connection is of sufficiently high a bandwidth to support the desired number of videoconferencing sessions.

The next logical extension is to include some videoconferencing clients that are at remote sites. In Figure 7, four clients are attached to a Virtual Private Data Network (VPDN) provided by an Internet Service Provider (ISP). The connection between the VPDN and the main office is a T1 link.

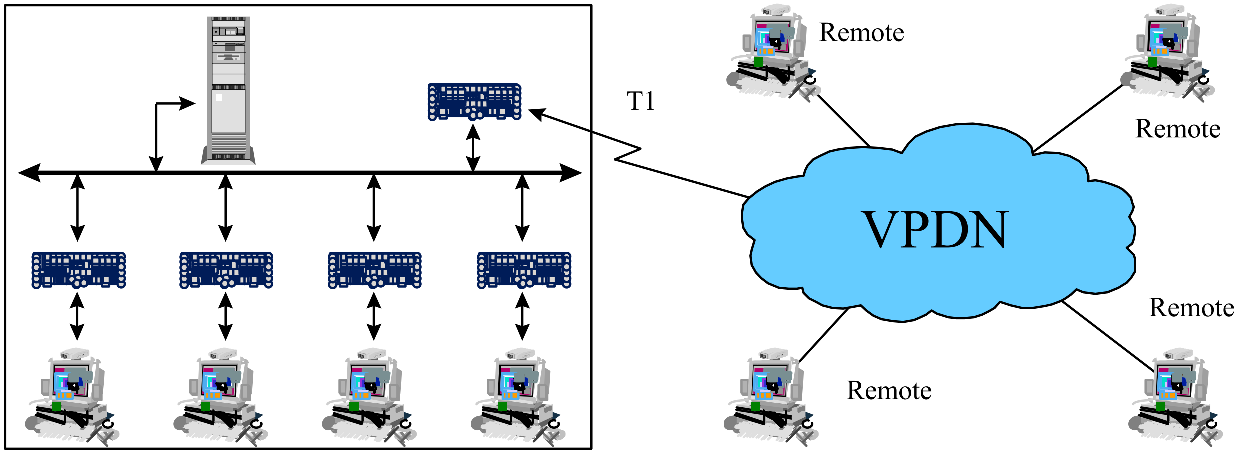

Figure 7 - LAN Attached to Virtual Private Data Network (VPDN)

In this configuration, the primary bandwidth concern is the T1 link connected to the VPDN. Each remote client establishes a connection with the MCU. Four sessions from the remote clients means 690Kbps for four LiveLAN 174K sessions. The T1 is a full duplex 1.544 Mbps link. Use of 690 Kbps in each direction is consuming just under half of the available bandwidth. This is about the limit of safe utilization of this link if the link is being used for simultaneous data traffic. Four LiveLAN 384K remote clients would overload this link. Three LiveLAN 384K clients would consume 74% of the total available bandwidth.

An alternative configuration that supports more remote clients is shown in Figure 8. In this configuration, the MCU is connected directly into the wide area network via a high speed link. If the wide area service is being provided by an ISP, the MCU can be connected within the ISP. The MCU may be a service provided by the ISP for its customers. In this particular configuration we have not reduced the T1 bandwidth requirement; there are still four clients on each side. However, if the number of remote clients is increased in this configuration, the T1 bandwidth will not change.

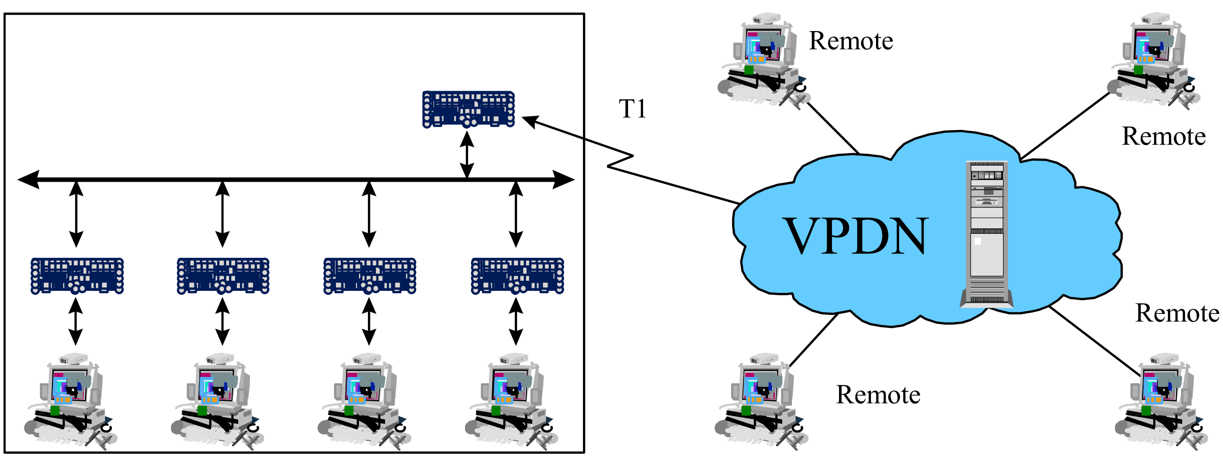

Figure 8 - MCU Service Within VPDN

When to Upgrade the Network

Limits of Network Segments

Table 4 shows that 20 LiveLAN 174K clients and an MCU can be supported on a single 10BaseT shared segment. This combination will consume about 70% of the available bandwidth on that segment. Although this configuration will work, it leaves little room for other uses of the network, and makes no concessions to the trade off between latency and throughput. For practical purposes we recommend that a segment supporting clients should never be configured to use more than 50% of the available bandwidth for videoconferencing.

A dedicated segment, such as the one shown in Figure 5 between the Ethernet switch and the MCU, can be more heavily loaded than a shared link. The predominance of traffic on this link is videoconferencing traffic. Users are not relying on this link for other services such as DNS, Lotus Notes servers, NFS, etc.

Full duplex Ethernet

Full duplex Ethernet can be used when a connection is strictly point-to-point, such as the link in Figure 5 between the switch and the MCU. Full duplex Ethernet can send traffic in both directions simultaneously with no interference. If traffic is fairly symmetric, like videoconferencing, then the effective bandwidth is doubled. Many Ethernet cards and Ethernet switches today support full duplex links, for 10Mbps and for 100Mbps data rates.

Using Switches to Partition Traffic

As shown in Figures 5 and 6, inserting Ethernet switches into the network can help reduce the traffic on client segments. An Ethernet switch, like a bridge, makes an intelligent decision about how to forward traffic, only forwarding it onto the link that has the destination address for those data packets.

An Ethernet switch is also valuable for connecting LANs of different speeds. If client segments are 10Mbps Ethernets and the backbone is a 100Mbps Ethernet, the switch can be used to move traffic between the backbone and the individual client segments.

Routers can also be used to do this kind of bandwidth segmentation. If a network is already heavily routed, using the routers may provide sufficient bandwidth isolation to solve the problem. Switches can often provide bandwidth segmentation at a lower cost and with higher performance for new installations.

Upgrading to 100Mbps Ethernet

Upgrading a backbone to 100Mbps Ethernet, FDDI or ATM becomes necessary when the aggregate traffic demand on the backbone exceeds the current configuration's ability to handle it. A small network configuration can move quite easily to 100Mbps Ethernet from 10Mbps Ethernet; adapter cards and switches are readily available. The cost of a 100Mbps network is higher, but not substantially higher than its 10Mbps equivalent.

The decision to upgrade to FDDI or ATM should be carefully considered, as many other factors weigh into the decision. These may include WAN ATM support, voice/data integration, legacy devices that need to be connected, etc.

Network Measurement Tools

Following is a discussion of a few tools that can be used to determine the current state of your network. They are listed from simple (and inexpensive) to complex (and expensive).

PictureTel NetTach

NetTach is a software utility from PictureTel, that runs on a Windows95 or NT 4.0 platforms. NetTach can be used to characterize a network path for its ability to support videoconferencing.

NetTach is installed on two PCs, at either end of the network path to be tested. NetTach sends data packets back and forth between the two clients, and analyzes the results. NetTach will show the available bandwidth, packet loss, packet jitter, and out-of-order packet counts. NetTach can be run once, or can be left running to log the state of a network path over an extended period of time.

NetTach is available from PictureTel on the NetConference H.323 MCU CDROM, and for internally from within PictureTel on the NSD Engineering web server at http://www.nsdeng.pictel.com. If you do not have access to either of these, please contact your PictureTel representative to obtain a copy.

LiveLAN Built in Diagnostics

LiveLAN 3.0 has built-in diagnostics to help determine the status of a network path. During operation, a diagnostics window can be opened (Tools/Diagnostics/Network). This page shows a set of running statistics for the current videoconferencing session, including packet loss, packet jitter, out-of-order packets, and round trip delay time.

NetScout Systems Probe

The NetScout product from Frontier is a remote probe that can be attached anywhere in a LAN or WAN environment. These probes are all managed by a central management station through in-band communications. The NetScout is capable of logging and graphing the usual set of network diagnostics, including utilization, packet sizes, talkers, etc.

Working in conjunction with PictureTel, Frontier has added features to the NetScout that allow it to monitor a videoconferencing stream, and record packet loss, packet jitter, and out-of-order events. This information can be logged over an extended period of time, and can be displayed on a moving graph on the management station.

References

Measured Capacity of an Ethernet: Myths and Reality. Digital Equipment Corporation, Western Research Laboratory, WRL Research Report 88/4. http://www.research.digital.com/wrl/publications/abstracts/88.4.html