Planning for Video on the Enterprise Net

Net Forecasts – Peter J. Sevcik
BCR Volume 33, Number 7
July 2003

I know you have heard about video being the next "killer app" for a long time. But before you dismiss it again, consider the following:

Polycom and Tandberg -- the top vendors of traditional ISDN-based conference room equipment -- are leading the transition to cheaper, better IP-based systems. This next-gen conference room equipment can be purchased for one-tenth the price of the older systems and no longer requires special ISDN connectivity. If there are videoconference rooms in your enterprise, someone is thinking about upgrading them to IP video.

Desktop video is rapidly evolving from an *ad hoc*, peer-to-peer connection model to true conferencing. First Virtual Communications, which started selling conference bridges, has seen sales grow of its Click to Meet, enterprise server software for managing desktop video meetings through a browser. Not to be left out of the picture is Microsoft, which recently purchased Place Ware, an operator of web-based, collaborative meetings.

Planning is Essential

Video will invade your enterprise network, so you face a basic choice: Wait until video runs you -- and your network -- over, or get ahead of the wave and guide its impact. Since video traffic can swamp a network, even if there are only a handful of users, planning and managing for its arrival is the only option. Planning for an IP video architecture will involve several steps, some of which are highlighted below.

<u>Connectivity and Scheduling</u> The naming and numbering plan should make calling between endpoints simple and, to the extent possible, consistent with existing naming and numbering systems (e.g. email addresses or telephone extensions). Currently deployed ISDN-based conferencing units can -- and should -- continue to operate, and interoperate with IP-based conferencing.

A scheduling system needs to be chosen that makes it easy for users to reserve conference rooms, bridge ports and bandwidth. The scheduling system must integrate with tools such as Microsoft Outlook or Lotus Notes.

<u>Traffic and Network Assessment</u> Create a traffic profile for each campus -- the expected number of systems, times of day they'll be used, bandwidth per session and connection patterns. This data will help you determine if there's sufficient capacity in the network and set QOS goals.

The network review should include link bandwidths, router devices, router OS levels and QOS capabilities in routers and switches. Check end-points to determine if they're on full duplex, dedicated switched connections, or on older hubs. Measure utilization levels on key links during busy hours to determine the impact of IP video on existing applications.

A key output of the network assessment should be a service level agreement (SLA) between the videoconferencing and network support teams. This SLA defines limits to the demands of IP video, and insures that there is proper network service to support the video sessions. Unless both these teams comply with the SLA, mission-critical applications will be impacted.

<u>Management Plan</u> You'll need a management system that tracks end-point status, capabilities and versions, and provides statistics on utilization, call set-up success rates, conference interruptions and the quality of connections. Internal processes must be in place for ongoing use of this information to monitor and manage the quality of the video experience.

Mixing Video with Other Applications

Most routers support QOS techniques that help data and real-time traffic coexist, but engineering a proper solution is still complex. The most common solution is to use Low Latency Queuing, a feature that brings strict priority queuing to Class-Based Weighted Fair Queuing (CBWFQ). But there still needs to be a separate way to tag the video packets when they leave the campus, and there's less consensus about the best approach; some use the Class of Service (COS) tag in the IP header, others recommend the DiffServ protocol.

None of the QOS mechanisms work, however, unless the video flows can be identified. The H.323 videoconferencing protocols assign dynamic UDP ports for the audio and video flows. Gatekeepers, MCUs, proxies or other application aggregation servers have well-known IP addresses that can be used to identify the video traffic (all traffic going into and out of them is video).

Once video traffic is identified, tagged and assigned to a QOS mechanism, a crucial design decision still remains: How much video? The amount of bandwidth each video session requires varies based on the picture size, quality and frame rate. However, most enterprise video sessions operate at 384 kbps, which translates to 400 kbps with IP overhead (desktop and Web-based systems often require less).

For many networks, this is a significant continuous load. A traffic plan must be developed that shows how many sessions can be supported at each major section of the network (e.g., feeder LAN, backbone, WAN access). IP video operates over UDP and will keep pumping packets into the network at its assigned rate, contending for bandwidth with the data applications. Data apps operate over TCP, which will throttle performance as it sees the impact of competing traffic. So, the data applications have a smaller pipe during the busy hour. Given that video gets a higher precedence, how much video can the network take and how much video contention can the data applications live with?

There is little hard data on this issue. Raymond Kneipp of The Burton Group recommends that real-time traffic should never exceed 25 percent of the bandwidth, and at 40 percent real-time you will not get any meaningful TCP throughput.

I have a different rule of thumb: One VOIP call displaces the load of 4 to 6 "typical" data application users, while IP video will need to kick off 25 to 35 users (these displacement values are higher with many older client-server applications). So your current network either needs that much spare capacity, or you will have fewer satisfied data application users.

Any way you look at it, the portion of the network that is dedicated to video must be carefully managed. For example, the total number of sessions that can be supported on a 100-Mbps campus network will severely degrade application performance on even a 45-Mbps WAN access line, and T1 lines to field offices will be totally overrun. Video calls can be stopped using an age-old mechanism: busy signals. Since there will be a different maximum number of sessions by location and time of day, use gatekeepers or MCUs to schedule conferences and manage the number of sessions.

The Multi-High-Priority Problem

In addition to VOIP and IP video, some enterprises also have real-time, machine-to-machine communications for instrumentation, machine control, telemetry, surveillance, etc. All these applications want -- and need -- high-priority handling, but they are migrating onto a common IP network that was originally designed for data applications.

The major ISPs have been wrestling with this problem for some time; for example, both Sprint and Genuity (now part of Level 3) offer prioritized services on an enterprise access line – Premium, Critical, Business and Standard or best effort (these are the Sprint labels but they apply to Genuity as well).

Sprint provides the following guidelines: Premium is reserved for VOIP or IP video traffic. If customers need both VOIP and IP video, the video traffic will be placed into second class – Critical. Sprint also recommends that no more than 35 percent of a customer's bandwidth be assigned to the Premium class when using VOIP, and no more than 17 percent when using video. There is a separate recommended limit of 25 percent in the Critical class.

Genuity does not make its bandwidth recommendations public, but it also restricts the top priority to VOIP, and assigns video and other real-time traffic to the second priority level.

The problem of the need for limits on real-time traffic has been attacked by Cetacean Networks -- it claims to eliminate the problem by sequencing packets in the network. I saw a demonstration of its technology at Mitre, which has integrated it into their Assured Delivery Communications Network (ACDN) project, which supports multiple, mission-critical, military applications.

The demonstration included data, VOIP and an IP videoconference, all converging on the same network and each needing "top" priority. Each of the applications performed properly when sequenced, but conflicted and fell apart without it. With sequencing, the video session retained perfect fidelity and frame rate despite the highly loaded network.

Conclusion

Video is going to hit most enterprises sooner rather than later, and you're going to need a good strategy and architecture to handle it. However you've tackled some of these issues for VOIP won't be sufficient for video, and many enterprises will have to contend with both voice and video in addition to other real-time traffic. The problem of managing multiple classes of real-time traffic on one network needs careful engineering and probably new approaches to succeed as a scalable solution.

<u>Companies Mentioned</u>
Burton Group (www.tbg.com)
Cetacean Networks (www.cetacean.com)
First Virtual Communications (www.fvc.com)
Genuity (www.genuity.com)
IBM Lotus (www.lotus.com)
Microsoft (www.microsoft.com)
Mitre (www.mitre.org)
Place Ware (www.placeware.com)
Polycom (www.polycom.com)
Sprint (www.sprint.com)
Tandberg (www.tandbergvision.com)

Peter Sevcik is president of NetForecast and is a leading authority on Internet traffic, performance and technology. Peter has contributed to the design of more than 100 networks, including the Internet, and holds the patent on application response-time prediction. He can be reached at peter@netforecast.com.

NetForecast provides custom consulting services on the value of new network technology. Our seasoned consultants draw on decades of experience to help clients worldwide choose new technologies, improve performance, and align infrastructure to business. We have helped leading enterprises, service providers, and vendors navigate the changing competitive landscape of the Internet economy. Please call us to discuss how we can help your information network succeed.

NetForecast
www.netforecast.com