MULTIMEDIA

Moving ISDN-Based Videoconferencing Onto The IP Network

John Bartlett

Videoconferencing over IP has had potential for many years. Finally, the installed base of ISDN users is considering a move.

Driven by cost, efficiency and quality concerns, the installed base of ISDN (H.320) videoconferencing users is moving to IP (H.323). Often this comes as part of a broader consolidation of circuit-switched traffic onto IP packet backbones (see "CalVIP Learns New QOS Rules"). This article briefly explains why this move is afoot, and discusses factors that contribute to a successful transition.

First of all, corporate decision-making in the post bubble telecom world is driven primarily by cost. ISDN-based videoconferencing is a visible and useful service in most large corporations, but it can be expensive. ISDN BRI and PRI lines cost roughly $30 per month, whether they are used or not, and usage charges run about 20 cents per minute. Each ISDN videoconferencing room or office needs three BRIs (for a 384-kbps call), and most calls are an hour to an hour and a half long. This means 60 cents × 75 minutes = $45 for the call, and a flat $90 per month to maintain the room. For an active videoconferencing environment, the per-call costs add up quickly, especially when users are calling overseas.

Many organizations look at the unused capacity in their IP packet data networks as a potentially "free" substitute for the dedicated ISDN lines. Some also think their IP networking staff could take over for the people who have been running their ISDN videoconferences. And features of the newer H.323 systems, like centralized management and simplified call setup for users, also help cut costs and improve efficiency.

These features improve on the older, ISDN-based systems that required trained staff at each videoconferencing site to maintain the equipment, establish calls and monitor call quality on finicky ISDN lines. With H.323, the requirements for IT support staff can be reduced, and smaller centralized staffs can support larger deployments.

Call Quality: Better With IP

Another major driver for conversion to IP is call quality. ISDN lines have a habit of developing bit loss, and of then dropping out, requiring either manual or automatic redialing. Meanwhile the conference is either on hold or operating at a much poorer level of video quality, because it is trying to function on the remaining bandwidth.

Even in very well-managed ISDN videoconferencing installations, where users actively track and pursue the quality of their calls and ISDN lines, call success rates run only to 90 and 95 percent. For organizations that do not have a rigorous quality program, success rates can fall to as low as 60 percent. (Successful calls start on time, the quality is sufficient for the meeting to proceed, and the call is not dropped or interrupted.)

In the past, IP networks have not been considered super reliable, because they do not yet hold up to the 5-nines standard of voice telephone systems and the PSTN. But compared to the call success rates of videoconferencing on ISDN, IP videoconferencing's simplified call model, corporate network control and centralized management can offer a substantial quality improvement.

Another factor to consider is that ISDN is no longer a primary focus for the telecom providers, so its quality is not likely to improve dramatically, whereas IP is a big focus both for the carriers and for the enterprise. So future trends also lead us to the conclusion that the more reliable transport is and will be IP.

IP-based conferencing also promises increased functionality—for example, higher video rates, better data handling and integration with other enterprise databases for scheduling and user identification. These features are of interest to the enterprise, but it is cost, efficiency and quality that drive conversions from ISDN to IP today.

John Bartlett is a vice president at NetForecast, where he focuses on real-time traffic, Internet performance and quality of service (QOS) techniques. He can be reached at john@netforecast.com.

Before making this conversion, however, four areas must be addressed:

- Business case
- Videoconferencing architecture
- IP network evaluation
- Transition planning

Attending carefully to all four will ensure a smooth transition, and that no conferencing or networking functions will be lost along the way. Let's take them one at a time.

The Business Case

Most mid- to large-sized companies that are migrating to IP for other reasons, and firms that

ISDN quality is unikely to improve, since it is no longer a primary focus for service providers

CalVIP Learns New QOS Rules

David Stein

California educators, many of them accustomed to ISDN-based videoconferencing over an ATM backbone, are getting ready for the conversion to a production IP video network. Many are replacing existing H.320 codecs with new H.323 endpoints.

The change comes as the California Video Over IP (CalVIP) Consortium of educational institutions moves to put its videoconferencing traffic on the new Corporation for Education Network Initiatives in California (CENIC) IP backbone (see www.cenic.org). Migration to the CENIC backbone is in progress, and this transition promises higher performance at a lower cost.

The CENIC backbone and IP video infrastructure will support academic and administrative videoconferencing and some streaming applications. Besides saving money and improving performance, the move to IP has four specific goals:

- Implement a global dial plan.
- Provide equal or better video and audio quality.
- Provide the ability to conference with legacy H.320 (ISDN) callers.
- Provide the ability to conference internationally via Internet 2 and the Video Development Initiative Network Consortium (ViDeNet).

The biggest challenge in meeting these goals has been figuring out how to provide end users with audio and video of quality comparable to the existing network, which is known as the California Community Colleges and California State University (4CNET) ATM network. Most CalVIP constituents have made extensive use of H.320 videoconferencing (ISDN fixed and dial-up connections to ATM) for routine remote classes and administrative functions.

To provide commensurate quality of service (QOS) in the new H.323 and CalVIP environment, successful implementations were evaluated and four key factors were determined:

- **Scheduling System—**All video calls requiring QOS will be scheduled in advance. If adequate bandwidth is available for the day/time of the call, the scheduling system will reserve the bandwidth. The scheduler also uses a "Least Cost Routing" approach that minimizes Multi-Channel Unit (MCU) cascading. Best-effort calls can be made without the scheduler.
- **Local Gatekeeper/Proxy—**The WAN will only grant QOS to calls originating from the known Gatekeeper/Proxy IP address. This prevents unauthorized users from setting their own priorities and minimizes firewall and Network Address Translation (NAT) issues.
- **LAN QOS support and configuration—** The H.323 codecs are connected to switches supporting QOS tagging and hardware priority queues. In addition, the LAN switches are appropriately configured to support QOS.
- **WAN QOS support and configuration—** Using router queuing algorithms, the tail circuits of the WAN are provisioned with an adequate amount of "QOS-able" bandwidth consistent with the scheduler and local gatekeeper information.

Although the combination of the four elements described above ensures production quality videoconferencing for all designated calls, the CENIC NOC is also equipped with tools for QOS troubleshooting. These include help desk trouble ticketing, a traditional SNMP network management console and proprietary tools that allow for management of devices that don't support SNMP (such as lighting controls). The NOC can also monitor the WAN circuits and in-progress videoconferences, adding and dropping users and adjusting system parameters as needed.

Actual cost savings associated with converting to the CENIC backbone are confidential, but the amount is expected to be significant. At press time, deployment was progressing, pending the installation of remaining CENIC tail circuits□

David Stein, a principal with PlanNet Consulting, provides voice, data and video infrastructure design and project management implementation services for a variety of private and public sector clients. He can be reached at 714/982-5814 or dstein@plannetconsulting.com

Include softer costs in your business case, but discount them by some percentage

want to increase the use of videoconferencing, will find the ISDN-to-IP conversion makes dollar sense. Making the business case to justify the conversion will require information from the architectural design work and the IP network evaluation, both of which are discussed below, but here are a few general guidelines:

Build your business case with hard costs. These include the current monthly and usage costs for ISDN lines, the costs for the new H.323 equipment and perhaps additional WAN bandwidth. Hard costs also include maintenance and support. If you hope for increased utilization of videoconferencing, include costs saved by reduced travel time. A videoconference takes much less time and expense than travel to another site.

Softer costs such as increased productivity, reduced product development cycles and increased revenue are harder to justify, because their relationship to conferencing usage is harder to prove. Include these in your business case, but then discount them by some percentage to account for the risk. This ensures that management will consider these costs, even if they don't allow them in the final calculation.

Lastly, consider performing a best practices analysis of your videoconferencing usage. This exercise will provide direct cost numbers for the service you are currently providing, which you can compare against benchmark numbers from a peer group of companies. These results immediately show targets for improvement, and these areas can be addressed during the transition process to ensure better service and higher productivity after the transition.

Designing A Videoconferencing Architecture

Ten years ago, videoconferencing was about buying $50,000 endpoints and connecting them through a point-to-point ISDN network. Today, IP-based endpoints are rapidly approaching free, while the cost and complexity have migrated into the infrastructure. Spending on the new infrastructure components will pay back, however, because the incremental cost of the IP-call is low. Five components must be considered, as follows:

1. The Gatekeeper is the central registration point, phone book and access manager for the videoconferencing system. It translates phone numbers or names into IP addresses, verifies permission of endpoints and manages bandwidth utilization. Gatekeeper design issues include determining how many gatekeepers to use and the geographic areas (zones) they will cover, how they interoperate, and what happens in failure situations.

Because gatekeepers are involved in each call, redundancy is necessary to ensure a reliable system. Gatekeeper design also should consider how to cross firewalls, and how to manage user access. A dialing plan that the gatekeeper will manage must be developed, including the names of conference endpoints and their associated IP addresses and/or E.164 phone numbers. Lastly, the gatekeeper maintains knowledge of network connectivity and bandwidth allocations so that it can control the bandwidth used by videoconferencing.

Cisco offers a gatekeeper function built-in to its router software. Higher-functionality gatekeepers include the Path Navigator from Polycom and the Enhanced Communications Server (ECS) from Radvision.

2. Multipoint conference capability brings together more than two endpoints in a common meeting. This can be handled by an external service provider or with one or more multipoint conferencing units (MCUs) in the enterprise IP network. Consider the bandwidth demands of the MCU when determining its location and implementation. Because the MCU must sustain the full bandwidth of all calls concurrently in conferences, its network demand can be very high. Mapping out where conference users are located will help determine if a central MCU is appropriate, or if smaller units should be deployed in a more distributed fashion.

Radvision offers the ViaIP MCU, Polycom has the MGC line from the former Accord Systems, and Tandberg also offers an MCU. Small conferences (e.g., four users) are supported by some endpoints from Tandberg and Polycom as well.

3. Gateways provide the interface between the IP-based and the ISDN-based conferencing worlds. This device has both LAN and BRI interfaces, and speaks the protocols of both networks. The gateway should provide full feature transparency, at least for all the ISDN features in use, and as such it is critical to a smooth transition to IP.

Gateways are also important because ISDN-based conferencing will not be eliminated from some sites. For example, in small, stable branch locations that make only one or two videoconferencing calls per week, it might make more sense to stick with ISDN—especially if the branches are served by T1 lines that are nearly full.

Gateways also allow users to conference with vendors, partners or customers who are not connected to the enterprise via IP. In this case, the gateway connects the internal IP network to the PSTN, and calls can be completed through an ISDN connection.

Finally, gateways can also provide interesting options for managing the videoconferencing environment, such as toll bypass, peak demand management and back-up paths for network failure or congestion situations. Gateway providers include Radvision, Polycom and Tandberg, as well as a gateway service offered by AT&T.

4. Management Systems comprise administrative tools that let just a few people keep systems with hundreds of endpoints operating smoothly. The management system should support software revision control, system availability monitoring, remote dialing and gathering of statistics.

Statistics from this system can be used to find calling patterns, track call failures (could not connect, interrupted call, call started late, etc.), and monitor historical bandwidth utilization. This information then feeds the quality process and the future planning efforts to ensure the conferencing system continues to provide excellent service.

Management systems include the GMS system from Polycom, the Tandberg Management Suite and the iView Network Manager from Radvision.

5. Scheduling is another critical software tool, responsible for coordinating the use of equipment and bandwidth. Preferably, scheduling will integrate with other enterprise scheduling tools such as Lotus Notes or Microsoft Outlook. The scheduling software should manage the physical rooms where conferencing endpoints are located, the virtual rooms on the MCU where conferences take place, and the bandwidth that will be needed on specific links to hold a scheduled conference.

Scheduling solutions include Network Aware Scheduler from Polycom, Tandberg Scheduler, Tandberg Management Suite and iView Conference Scheduler from Radvision.

A critical exercise in developing the videoconferencing architecture is determining the bandwidth required on each wide area link to support the expected call volume. First plot the location and calling patterns of endpoints, then build a spreadsheet to see how many calls will be occurring simultaneously, and at what time of day. Add 20 percent to the ISDN call bandwidth (a 384-kbps call becomes 460 kbps) to get the IP bandwidth required for each link.

To help manage the interface between the videoconferencing and networking support groups, some companies have implemented a service level agreement (SLA). This document, which is agreed to by both teams, spells out the network requirements for video, including bandwidth, packet loss, jitter and latency.

The SLA serves both parties well. The videoconferencing group can test the network to ensure their requirements are being met, and to isolate problems to either the equipment or the network. The networking group can use the SLA to forecast the demand on their network, to understand the videoconferencing requirements, and to allocate a percentage of their costs to videoconferencing, either as a bill-back or just to help justify additional IT costs to management.

Getting The IP Network Ready For Videoconferencing

Real-time traffic is different. Voice and videoconferencing are real-time, interactive applications, which create the most demanding type of traffic. If voice over IP (VOIP) has already been implemented, the network folks have learned this lesson. If not, it is time to learn it now.

Unlike transaction-based data applications, which can wait a few milliseconds while TCP's retry mechanism asks for lost packets, real-time traffic must arrive promptly and in order, like the frames of a movie. Packets arriving late don't fit properly into the playback and might as well not have arrived at all.

Two questions usually arise when adding videoconferencing to the existing IP network: Will the existing network support the bandwidth, packet loss and jitter requirements of videoconferencing? And will the introduction of videoconferencing adversely affect existing applications?

To answer these, the first step is to understand the existing network utilization. The WAN links are usually the critical resource. The bandwidth-demand study that you did during the architecture phase should have described the expected videoconferencing demand on each WAN link, for each period of the day. Videoconferencing will reduce available link bandwidth by the amount of bandwidth it consumes.

If the busy-hour for videoconferencing and the busy-hour for data applications do not overlap, then problems are avoided. But if they do overlap, the new constraint may adversely affect the applications. If QOS is implemented, the bandwidth for the videoconferences will be protected, but the data applications could suffer reduced throughput; if QOS is not implemented, bursts of data traffic could impair the videoconferences.

To predict the impact of the new traffic, either the scenario should be modeled, or synthetic video traffic should be introduced to the network in a controlled manner, and the actual responsiveness of the data applications tested. IP network capacity and QOS upgrades might be needed, which must be added to the costs associated with videoconferencing deployment.

Like any priority mechanism, QOS works best when the volume of traffic getting priority is relatively small. In addition, care must be taken to limit the amount of video and audio traffic on any given link to 30 or 40 percent of the available bandwidth. If a higher percentage of high-proirity traffic is allowed, then it will begin to interfere with itself.

To understand this concept, think about the ski patrol's priority privileges at the ski area. They always go to the front of the line, and get the next chair lift up the mountain. Think of how disruptive it would be if half or more of all the skiers waiting for the lift were ski patrol.

Many enterprise campus networks already have enough bandwidth to support video and audio streams without implementing QOS, but this is a game of chance. According to Murphy's Law, congestion is certain to arise on the network at the exact same time that an important videoconference is occurring.

Test The Network For Jitter And Packet Loss

Before deployment, you must test the network for jitter and packet loss to ensure it is clean. Many

Like any priority mechanism, QOS works best when a small percentage of the traffic gets prioritized

Video-conferencing failures are very visible, so train end users and administrators to succeed

small issues can cause low-level data loss in the network that goes unnoticed by transaction-based applications. Test end to end, using synthetic videoconference traffic that flows down the same paths the real videoconferences will use.

Tools from companies like NetIQ or Brix can be used to simulate videoconferencing flows and deliver test results on the received traffic. When problems are found, isolation can search out the router, the old bent Cat3 wire or the segment of shared 10-Mbps Ethernet that needs an upgrade.

After deployment, keep testing the network on an ongoing basis. Whether using home-grown tools or those from outside vendors, it is important to monitor changes in network behavior, and to be able to correlate those changes to trouble tickets from videoconferencing users. Plan now for a methodology that will make maintenance simple in the future.

Writing A Transition Plan

A transition plan should be written to ensure each step of the conversion is well thought-out. The conversion process has to keep the old system going, since videoconferencing is in use and a disruption of service may affect the business. Several approaches will ensure a smooth process.

One approach is to upgrade the IP network fully and build the infrastructure for the whole videoconferencing deployment. After thoroughly testing the infrastructure, management and scheduling processes with a few friendly endpoints, additional endpoints can be rolled out.

A second approach is to convert only one building or campus at a time. Deploy the components of the IP infrastructure necessary to support the first building/campus, and install IP-endpoints. Connect to the rest of the company via gateways, dialing through the ISDN network as before. With careful design of the dialing plan, this can be transparent to users.

Once this campus is stable, move to the next campus, bring them up on IP, and establish IP connectivity to the first campus. The gateways can be moved from location to location to support the connections to ISDN as necessary to bring each new campus on line. This incremental approach works well for a company with limited funding for the conversion, because it allows a more gradual upgrade.

In either approach, be sure to provide comprehensive training of both administrative personnel and end users. Videoconferencing failures are very visible, and usually come in front of a group of peers. Good training will help both end users and administrators succeed.

Conclusion

IP-based conferencing will be the standard in just a few years, so the conversion is inevitable. In the same way that VOIP is gaining momentum, video will soon be expected to be on the IP network. Collaboration technologies will come together to provide a much richer and simpler environment for doing business across geographic boundaries, and this will be driven by global IP connectivity.

The enterprise decision is not if, but when to convert. Understanding the parameters of the conversion process and working out a careful plan will ensure users a smooth transition into the new environment□

Companies Mentioned In This Article
AT&T (www.business.att.com)
Brix Networks (www.brixnet.com)
Cisco Systems (www.cisco.com)
NetIQ (www.netiq.com)
Polycom (www.polycom.com)
Radvision Corp. (www.radvision.com)
Tandberg (www.tandbergusa.com)