Multimedia Conferencing Over IP

In the IP environment, many obstacles stand in the way of establishing a video connection.

By Andrew Ward

Although videoconferencing has very successfully taken off over ISDN lines - particularly in certain vertical markets such as finance - the time, trouble and investment required to install and use it have been amongst the factors that have restricted its widespread adoption. In particular, the need to install ISDN lines for H.320 videoconferencing - often three lines were used, to provide six channels for 384 Kbits / sec connections - was a major inhibitor, and severely curbed flexibility.

By contrast, the Internet provides almost universal IP connectivity, and should make it easy to connect virtually any two conferencing endpoints. It also removes the fixed bandwidth limitation of ISDN lines. However, this and other advantages of IP are only gradually having an impact, and according to "Conferencing Markets & Strategies 2001 - October 2001" by Wainhouse Research, it won't be until the end of 2003 that IP calls will represent 50% of all videoconferencing calls made.

Slow Growth

This slow growth can be partly attributed to the technical issues surrounding multimedia conferencing over IP. Anyone attempting to establish a video connection across the Internet - even between just two endpoints, never mind a large-scale conference - may have encountered one or several of the numerous difficulties that can attend such an endeavour. Although the H.323 standard for multimedia calls over IP networks is now generally accepted, that doesn't solve all the problems. Fortunately, academics and equipment vendors have been working to overcome the remaining issues, and with judicious choice of hardware and software systems, they can all be avoided.

Network Compatibility

The first and most obvious point is that for some organisations, it will be economically undesirable to move to IP overnight, owing to an existing investment in ISDN endpoint equipment. Also, some locations may not yet be served by high-bandwidth IP connections and are therefore restricted to ISDN. And if a videoconference is held that includes external participants, then clearly the organisation has no control over what technologies they may want or need to use, and so once again ISDN may enter the picture.

This problem is solved by the availability of videoconferencing servers, also known as multipoint conference (or control) units, that provide support for both network types. Sometimes this is accomplished via plug-in cards, and in some cases existing ISDN videoconferencing servers can be upgraded to include IP interfaces simply by the addition of new connection cards. Over time, as endpoints switch technologies - most endpoints today have dual network capabilities, in fact - ISDN cards can be removed to make way for further IP adapters. An alternative solution is to use an ISDN to H.323 conversion gateway that can convert multimedia data between the two networks and standards.

IP Addressing

IP addressing probably provides the biggest challenge to anyone attempting to link two or more users. To make a connection between two endpoints means knowing both IP addresses. However, in many cases an IP address will be inaccessible -

perhaps because NAT, PAT (Port Address Translation) or some other address translation scheme will be in use. These schemes provide some protection for internal IP addresses, which are then usually completely inaccessible by external users.

Thus NAT schemes are nearly always incompatible with H.323 dialling schemes. In some cases, vendors of NAT-type devices make them to a certain extent H.323-aware. This means that outgoing calls can be established from a single endpoint on the internal network, but that's as far as it goes. Another solution to this problem is to use the Accord VGC-20 gateway, described below under Firewalls, to bridge the NAT device. Cisco also has a solution - the Cisco Multimedia Conference manager, using both the Cisco Proxy and Cisco gateway.

Secrecy

Another problem also related to IP addressing is that one or more of the parties involved in the call may not know their IP address or may not be prepared to divulge it. Kenneth Tanner, Information Technology Analyst at the Louisiana State University Health Sciences Centre (LSUHSC), gives an example of this situation.

"I was asked to bring in people from various locations via videoconferencing to discuss the establishment of the Production network,, he said. "The other H.323 systems were either on our own network and thus used the network dial scheme, or the site gave me the IP address of the endpoint and I was able to dial them directly. Except for one H.323 site. The person we needed to participate in the conference was not agreeable to giving me the IP address of his endpoint."

Gatekeepers are devices responsible for establishing H.323 calls, and perform the translation from dial schemes to IP addresses. But without the IP address of the endpoint, the only other way that Tanner could connect to the site was to register the LSUHSC resources with the gatekeeper of third-party network, Vide. Then, he would have been able to access the site using their dial scheme, instead of directly via the IP address. But to do that would also have meant that the LSUHSC dial scheme would have to change to match the Vide dial scheme - clearly, a major and impractical task.

Tanner was able to solve the problem using a feature of the Accord videoconferencing server. "One setting on each of our three H.323 interface cards is the IP address of the gatekeeper you want the card to register with. In our case, all of the cards had always been assigned to a gatekeeper within our network. However, there is no reason why you can't assign a different gatekeeper per individual cards - so what I did was go in and reconfigure one of the three H.323 cards to register with one of the Vide gatekeepers."

Negotiation

Another way the problem can be solved in some circumstances is by negotiation between the gatekeepers. The LSUHSC gatekeeper could have communicated with the Vide network gatekeeper to determine how to address the user's endpoint. However, the complexities involved with this solution mean that it's unlikely to be practical for an *ad hoc* arrangement; and because the protocols used involve broadcast messages from the gatekeepers, security concerns may prohibit it altogether.

Firewalls

"H.323 presents difficulties for firewalls because it can use several out of a large number of different ports."

H.323 presents difficulties for firewalls because it can use several out of a large number of different ports. In fact, H.323 is the protocol used to establish the call. The actual multimedia data is transmitted using RTP, which uses UDP and has no fixed port assigned to it. Firewalls must either have huge holes blown through them to allow any of the possible ports - something that is clearly undesirable - or must have the intelligence to understand the H.323 protocol and therefore work out, dynamically, which ports need to be open. However, a protocol-aware firewall uses processing power to perform this task, and therefore has scalability issues.

It is not unusual for firewalls to be the source of performance bottlenecks anyway - Mercury Interactive's ActiveTest service found that one major insurance company was only obtaining 50% of its Internet connection's bandwidth owing to firewall performance issues - and so it's clearly undesirable to exacerbate the situation.

There are several other possible solutions, such as setting up a VPN, but they all have limitations. In particular, the VPN solution is today really only practical for communication within a single organisation, and not to third-party networks.

One vendor at least has found an intriguing way around this problem. In the Accord gateway products from Polycom, all incoming H.323 audio, video and data packets are not just inspected but are actually decoded back to their original content. For retransmission they are then encoded back to an H.323 signal and new IP packets are constructed. No IP packet can ever cross the gateway, and nothing other than H.323 traffic can ever be decoded, encoded and then retransmitted. Thus the gateway is configured to straddle the firewall so that all H.323 traffic bypasses the firewall and passes through the gateway instead. Not many network managers will be overjoyed at the thought of breaching the firewall in this way, but the Accord products have been tested for security by the NSS Group and awarded the NSS Approved status.

Codecs

Once communication has been established at an IP level, the next challenge is for the audio and video devices to talk to each other successfully. One of the benefits of the H.323 standard is that there are many different possible codecs in use - the software or hardware that encodes and decodes the video and audio signals. For example, for video there is H.261 or H.263, CIF or QCIF, 7.5 frames per second up to 30 frames per second; for audio there is G.711, G.728, G.722, G.723 or G.729. This wide choice offers different trade-offs between various parameters such as bandwidth, video resolution and frame rate, and audio quality.

Codecs may either be hardware or software - or can sometimes be a combination of the two.

Unfortunately, the wide choice of codecs presents further problems. Firstly, it is possible to have two endpoints that have no codecs in common, or maybe only a video codec in common but no audio. For a two-party call, there would be no way they can communicate at all. And when trying to set up a videoconference with several parties, it may mean that one party can't participate. Sometimes it can be possible to install a different codec, but this takes time, and in a conference is clearly disruptive. Still disruptive, but slightly less trouble, is where one device needs reconfiguration to select the right codec.

Furthermore, in the IP environment, even though all users may have the same codec, there are often occasions when not everyone would want to use it. If one user is connected via a modem link and restricted to using a very low bandwidth, poor quality and low frame rate video signal, it is unlikely to be acceptable to reduce all other conference participants to the same level.

At the other extreme, although some endpoints may be happy communicating at 1 Mbit / sec, this may represent an unwarranted use of precious bandwidth for others. Ideally, each user would be able to select the appropriate codec to provide the optimum compromise between quality and bandwidth, taking into account their endpoint equipment and network connection.

"Once communication has been established at an IP level, the next challenge is for the audio and video devices to talk to each other successfully."

Transcoding

A technique known as transcoding provides the solution. It involves decoding every video and audio stream that comes into the multipoint control unit, and then recoding them using the codec that's appropriate for each participant. Work on transcoding has been going on for some years at various academic and research institutions. Further information on the University College London transcoding gateway is at **www-mice.cs.ucl.ac.uk/multimedia/projects/utg**. The Planète group, located at the Sophia Antipolis and Rhône-Alpes research units of the Institut National de Recherche en Informatique et en Automatique, has developed its own tool, Rendez-Vous, that performs transcoding. Information is at **www.gaia-interactive.com/rv**.

The MASH research group at the University of California, Berkeley also provides transcoding proxies that were originally developed for DARPA's global mobile information systems program.

However, the commercial world is catching up, and many vendors do now include audio transcoding in their control units. The Cisco IP / VC 3520 and 3525 Videoconferencing Gateways support audio transcoding between audio codec standards G.711 and G.728, and between G.723 and G.711. RADVision's L2W-323 Multimedia Gateway has optional hardware modules that perform the same transcoding functions. Accord products from Polycom transcode all audio and video parameters as standard.

One of the challenges with transcoding is performance. All coding and decoding operations take time, simply because of the nature of the algorithms in use. They must look at a certain minimum length of analogue signal before any encoding operation can happen, in addition to any processor delay that may take place. Although the first delay cannot be avoided, the second can be reduced by using as powerful a processor as possible. Attention also needs to be paid to any possibility of loss of synchronisation between audio and video signals.

Quality

Unlike an ISDN connection, an IP multimedia data stream can be subject to many different types of disruption. In particular, audio and video packets may be lost altogether. Endpoints have a variety of different ways of dealing with this problem, but the multipoint conference unit can also help by inserting either white noise or blank video frames, as appropriate, to maintain the integrity of the audio and video streams. If the endpoints in use don't handle dropped packets well, then looking for this facility in the conference unit is a good idea.

A more serious problem is out-of-synch audio and video - something that can be very irritating during a conference - and can be caused by audio and video packets taking different paths across a network. Again, there is scope for the conference unit to improve matters, by resynchronising audio and video time stamps as the packet leaves the unit, although disruption can still take place en route to the receiving endpoint. This technique also helps reduce jitter, but once again further jitter can be introduced to the retransmitted signal, so it's not a complete solution.

Although techniques for setting priority, class of service or quality of service are only slowly being adopted, there is potential for the multipoint conference unit to play a role. Vendors are working on adding functionality to their products that will support standards such as Cisco's IP Precedence, DiffServ and RSVP. Servers could add the IP Precedence bit for endpoints that do not support it and respond to DiffServ commands, for example.

Conference Management

Being aware of the technical issues, and choosing conference equipment with these in mind, will help to ensure that the technology is in place to establish a successful IP conferencing connection. Paying attention to various conference management techniques will also help to ensure that conferences are set up with as little pain as possible.

Firstly, it's not unusual during any conference - IP or ISDN - for a participant to have to leave early, a new participant to join, or something to fail - either endpoint equipment or a communications link. A further complication with IP networks is that it may be desirable to reconfigure the settings in use by one of the conference members after the conference has started. In these circumstances, it's desirable to be able to add and remove people from a conference on-the-fly, to avoid having to disrupt everybody by stopping and restarting with a different configuration. Of course, one huge advantage of IP conferencing is that management can potentially be carried out from anywhere, and not just a device - such as a monitor, terminal or PC - locally attached to the multipoint conference unit.

PCNA

Useful Resources

Conferencing Markets & Strategies 2001
www.wainhouse.com/cmands.html

Traversal of IP Voice and Video Data through Firewalls and NATs
www.forgent.com/pdf/IP_Data_through_Firewalls_and_NATs.pdf

Interactive Video on Enterprise Networks, A Perey Research and Consulting paper
www.nwfusion.com/media/InteractiveVideo001114.pdf

Deploying H.323 Applications in Cisco Networks
www.cisco.com/warp/public/cc/pd/iosw/ioft/mmcm/tech/h323_wp.htm

RADVision
www.radvision.com

Accord / Polycom
www.polycom.com

Cisco
www.cisco.com

Ezenia!
www.ezenia.com

New Reviews from Tech Support Alert

Anti-Trojan Software Reviews

A detailed review of six of the best anti trojan software programs. Two products were impressive with a clear gap between these and other contenders in their ability to detect and remove dangerous modern trojans.

Inkjet Printer Cartridge Suppliers

Everyone gets inundated by hundreds of ads for inkjet printer cartridges, all claiming to be the cheapest or best. But which vendor do you believe? Our editors decided to put them to the test by anonymously buying printer cartridges and testing them in our office inkjet printers. Many suppliers disappointed but we came up with several web sites that offer good quality cheap inkjet cartridges with impressive customer service.

Windows Backup Software

In this review we looked at 18 different backup software products for home or SOHO use. In the end we could only recommend six though only two were good enough to get our “Editor’s Choice” award

The 46 Best Freeware Programs

There are many free utilities that perform as well or better than expensive commercial products. Our Editor Ian Richards picks out his selection of the very best freeware programs and he comes up with some real gems.