Video Teleconferencing and Firewalls

Firewall traversal is the toughest obstacle inhibiting video teleconferencing (VTC) over the public Internet. This paper describes the specific firewall traversal challenges and discusses methods developed by Infrasupport using open source tools to meet them.

A few brief definitions are important. Several excellent textbooks are available for deeper discussions of all these concepts.

First is an IP (Internet Protocol) Address. An IP Address represents a computer, or host, in an IP network such as the public Internet. Some IP Addresses are designated as private, meaning they will never route over the public Internet. A public IP Address is roughly analogous to the physical street address of a building. A private IP Address is roughly analogous to a room inside that building. Private IP Addresses are never directly visible on the public Internet. Most business and home networks advertise a small number of public IP Addresses and use private IP Addresses for internal hosts, with a NAT firewall/gateway regulating traffic between the internal network and public Internet.

NAT, or Network Address Translation, is a technique to enable communication between hosts with private IP Addresses – not visible on the public Internet – and other hosts across the public Internet. A NAT firewall/gateway handles all the NAT chores.

A firewall is a system that makes decisions, based on a set of rules, whether or not to forward packets. A packet is just a set of information, similar to a sentence in a dialogue between people. Firewalls live at the boundary between private home/business networks and the public Internet. Most firewalls are also NAT gateways. They disguise internal hosts that wish to communicate with the Internet and handle the details routing data to/from private hosts and the public Internet.

By now, NAT may be the single most widely deployed technology on the Internet. Unfortunately, the people who developed NAT and the people who developed the VTC protocol standards worked independently of each other. The standards were never meant to work together and do not interoperate well. NAT wreaks havoc with VTC communications and is major impediment to wide scale adoption.

TCP and UDP are the most common IP session protocols. TCP is stateful, meaning the sender and listener incur some overhead setting up, maintaining, and finishing a conversation. UDP is stateless, meaning a sender sends streams of packets to a listener without setting up a conversation and without waiting for any acknowledgement. The tradeoff: TCP is reliable, UDP is fast. VTC systems use TCP to set up and control calls, and UDP for call content.

Both TCP and UDP use a concept called a port. The word “port” might be one of the most overused words in the English language. In this context, it means a communications channel between a sending program on one computer and a listening program on another computer. Both sending and listening programs have port numbers. While IP Addresses define computers, or hosts, port numbers identify the sending and receiving programs inside those hosts.

Many application protocols are simple and use a single port. These are easy to model with firewalls. VTC protocols are not so simple. H.323 is the most common VTC standard, with another standard, SIP, coming on strong. Both use a call setup mechanism with a well known TCP or UDP port, followed by random ranges of UDP and/or TCP ports to carry call data. H.323 uses TCP port 1720 for call setup while SIP uses UDP port 5060.

Modeling the call setup is easy. Modeling the random ports carrying the call data stream is much more difficult because firewalls cannot know in advance which ports to accommodate. Classic firewalls have no way to know that a stream of UDP packets using a range of random ports is really part of the same conversation defined at call setup. Without special handling, firewalls will block this traffic and VTC communications will be garbled or nonexistent.

Many conventional firewall systems work around this challenge by simply opening all TCP and/or UDP ports to and from VTC equipment. Although this tactic does route the data streams correctly, it also passes unwanted traffic, which completely defeats the purpose of a firewall and opens the private network to significant security risks. It is an unacceptable workaround.

Some vendors also try to work around the problem by offering an option to use fixed instead of random ports with their equipment. With fixed ports, firewall modeling is easy. This works, as long as all VTC communications use equipment from the same vendor using the same fixed ports. It is analogous to a telephone from one company unable to interoperate with telephones from other companies.

NAT represents an even greater challenge. In a glaring violation of good network layering practice, the H.323 VTC standard imbeds IP Addresses in the data stream. Without some means of special handling, NAT and H.323 are fundamentally incompatible because NAT mangles IP Address fields in the packet headers while H.323 imbeds IP Addresses in the packet data. When these do not match, VTC communications break down.

Vendors have tried with varying levels of success to work around this problem. Most modern equipment built in the past few years includes special processing providing the ability to live behind a NAT gateway and the most recent firmware upgrades seem to be reasonably debugged. Older equipment, especially shipped before 2005 or so, is generally buggy.

After solving the NAT and random port challenges, the next issue is QOS, or Quality of Service. VTC and VOIP (Voice Over IP) communications are for real-time meetings among people. Unlike email and web downloads, milliseconds count with VTC communications. In order to handle VTCs properly, firewalls need to take appropriate steps to ensure the best quality among VTC and VOIP participants.

After meeting the technical challenges, the next and possibly greatest challenge is political. After years of building reasonably secure network boundaries, many IT departments are unwilling or unable to disrupt those boundaries to support VTC communications. This means VTCs need to fit into the existing infrastructure, even though that infrastructure was never designed with VTCs or VOIP in mind, and even though they may be fundamentally incompatible.

With no single commercial product available to address these challenges, Infrasupport turned to the open source community and found tools to meet each of these challenges. Standing on the shoulders of giants, Infrasupport builds firewalls based on the Linux operating system and publicly available open source tools. Infrasupport packages these tools along with customized scripts to set up rulesets tailored for individual customers. Infrasupport firewalls provide all the network protection and key features of much more expensive and proprietary equipment, but with enhanced flexibility at a fraction of the cost.

Linux provides a rich set of kernel modules for ruleset scripts to model VTC and VOIP conversations. Infrasupport uses these tools to elegantly solve the random port problem by allowing firewalls to model entire conversations, not just call setup. Organizations can use Infrasupport firewalls to deploy VTC equipment with no need to configure fixed ports. This eliminates the interoperability challenge with equipment from one vendor communicating with equipment from another vendor.

Infrasupport also found a way to eliminate the NAT challenge. While traditional firewalls are strictly routers, operating only at layer 3 of the OSI networking model, Infrasupport firewalls can also operate as bridges, or at both layers 2 and 3 of the OSI model. This means that organizations can set up VTC equipment using public IP Addresses, instead of private addresses, with the equipment safely and now transparently behind Infrasupport's industry leading firewall. Bridging solves the NAT challenge by eliminating the need for NAT. Bridging is not appropriate for all situations and that is why Infrasupport firewalls provide both bridging and routing.

The next challenge is QOS. The Infrasupport firewall uses standard Linux kernel bandwidth management tools to prioritize VTC communications over other traffic at the network boundary, which means web downloads and email might wait a few extra milliseconds while VTC traffic jumps to the head of the queue. Although nobody can control packet timing across the public Internet, Infrasupport makes an implicit assumption that the public Internet has plenty of capacity and the bottleneck is the boundary between organization private networks and the public Internet. Controlling that boundary provides reasonable performance, provided the bandwidth coming into the site is adequate. QOS is only as good as the bandwidth feeding any particular site.

The toughest challenge of all is political. Although reliable VTC communications depend on an Infrasupport firewall at the boundary between the private network and public Internet, many organizations are unwilling to replace legacy equipment with an Infrasupport firewall. How can an Infrasupport firewall meet these conflicting goals?

The answer: Keep the legacy equipment in place and add an Infrasupport firewall configured as a bridge/router in front of it. The Infrasupport firewall will direct VTC traffic appropriately and bridge all other traffic unmodified to the legacy firewall equipment already in place. Setup in this manner, the Infrasupport firewall can provide QOS for VTC and/or VOIP traffic without disrupting traditional traffic and without breaching any security policies. Organizations can add new VTC capability while keeping familiar legacy equipment in place.

The Infrasupport firewall also provides all the traditional stateful packet filtering expected from any modern firewall, and Infrasupport firewalls are protecting dozens of networks today, ranging from very small to very large.

VPNs (Virtual Private Networks) are also important, and the Infrasupport firewall supports several VPN approaches, including IPSEC and PPTP. Dozens of organizations are using Infrasupport VPNs today.

All Internet connections eventually experience problems, and the Infrasupport firewall system includes a wealth of Linux troubleshooting tools to help quickly diagnose and fix these problems. Infrasupport frequently uses these tools to help Internet Service Providers find problems in their networks and to debug VTC, VOIP, VPN, and other conversations.

Infrasupport also offers HA (Highly Available) firewall systems, configured in an active/standby pair with automated failover. The systems in a failover set monitor each other, along with an array of Internet connections and optional key customer equipment. Whenever any of these change state, the system notifies a list of network administrators via email and/or cell phone text message.

Once installed, Infrasupport firewall systems are completely automated. They require no manual intervention during normal network operations.

Infrasupport spent significant engineering effort pioneering, testing, and improving its existing firewall technology to meet the new challenges posed by video teleconferencing. Infrasupport has already overcome the technical hurdles that most others have yet to learn even exist.

InfraSupport
CORPORATION