Traversing Firewalls and NATs With Voice and Video Over IP

An Examination of the Firewall/NAT Problem, Traversal Methods, and Their Pros and Cons

Traversing Firewalls and NATs With Voice and Video Over IP

An Examination of the Firewall/NAT Problem, Traversal Methods, and Their Pros and Cons

April 2002

Executive Summary

Strategic advantages exist for organizations and enterprises when their voice, video, and data communications run over a single converged IP network infrastructure. Unfortunately, the ability to capitalize on IP communications systems has been severely delayed because nearly all corporate networks have firewall and network address translation devices (NATs) that effectively block IP voice and video calls. Firewalls block IP voice and video traffic by placing a barrier to any unsolicited, incoming communications. NATs block IP communications traffic because the IP voice and video devices behind the NAT have private IP addresses that are not routable outside their local domain or on the public Internet.

Several solutions exist for overcoming the NAT and firewall problem for IP communications including bypassing the firewall and NAT, upgrading the network infrastructure devices using an application level gateway (ALG), and navigating across the firewall and NAT using a semi-tunneling traversal method.

Bypassing the firewall and NAT is clearly not an option for most organizations. Removing firewall protection or employing a device such as a proxy or MCU at strategic locations in the network to bridge around the firewall/NAT may compromise network security. These "device" solutions may also be costly, and they require physical, political, and/or intellectual access to the enterprise-critical network firewalls and NATs. In addition, one of these bypass devices will needed at every location along the communications path where a firewall or NAT presently exists.

Upgrading the firewall/NAT with an ALG is another possibility, albeit intrusive and potentially expensive. ALGs are essentially vendor specific software upgrades to the firewall devices that examine each data packet attempting to cross the firewall to see if it is of a known protocol type, such as H.323 or SIP. If packets contain the known protocol type, the firewall allows the packets to pass. However, like the proxy or MCU bypass solutions, ALGs require political and intellectual access to the firewall, and every firewall/NAT in the call path must be upgraded with the ALG software. Furthermore, as new protocols are developed, a new vendor specific firewall ALG software upgrade will be required.

Ridgeway Systems' IP Freedom™ transparent traversal method is the only available method for enabling IP voice and video communications that neither bypasses the firewall and NAT nor requires firewall/NAT software upgrades. It is also the only method that is global in scope, functioning irrespective of the number of number of firewalls or NATs in the communication path. The Traversal method requires no firewall or NAT device configuration modifications in the overwhelming majority of cases. This is accomplished using Ridgeway's client software on the inside of the firewall that establishes outbound communication connections through the firewall with a Ridgeway server on the outside of the firewall. All IP voice and video connections pass through the server, making it possible to seamlessly traverse any number of firewalls and NATs in the actual call path. Inbound calls are received through the same client software and are routed to the appropriate IP voice or video device for which the call was intended. The whole process is transparent to the IP voice or video device. One shortcoming of the Traversal method is that minor call latency is added since all connections pass through the secure communications server.

Each firewall and NAT traversal scheme has its own pros and cons; organizations seeking a global, non-intrusive, non-upgrade method may consider Ridgeway's IP Freedom transparent traversal solution an excellent alternative.

Traversing Firewalls and NATs with Voice and Video Over IP

Using existing computer network infrastructure for voice, video, and collaborative data communications promises compelling strategic advantages for organizations of all sizes. Collectively known as rich media communications or Internet Protocol (IP) communications, these *converged networking* technologies offer unprecedented opportunities to communicate, coordinate, and collaborate with customers, suppliers, business partners, colleagues, and associates around the globe.

Unfortunately, the protocols used for communicating rich media over IP networks conflict with most network security mechanisms like firewalls and network address translation devices (NATs), resulting in slower or delayed deployment of IP voice and video applications. In this paper, we briefly review how firewalls and NATs work to protect the network and why real-time rich media communications protocols are a security challenge. We then review the advantages and disadvantages of different methods for allowing IP voice and video to negotiate the firewall and NAT device. We also provide recommendations for those methods that seem best suited for allowing enterprises and institutions to securely implement IP rich media communications.

How Network Firewalls and NATs Work

On an IP network, every device is assigned a unique IP address. Computers, IP telephones, and videoconferencing devices (often called terminals or endpoints) also have approximately 65,000 network data ports, which are used to establish communication channels for transmitting data between devices on the network.

Messages between network devices consist of data packets containing the following elements:

1. The IP address of the device originating the message and the port number where the message originated from,
2. The IP address of the device to receive the message and the port number on that device where the message is to go, and
3. The data to be transmitted.

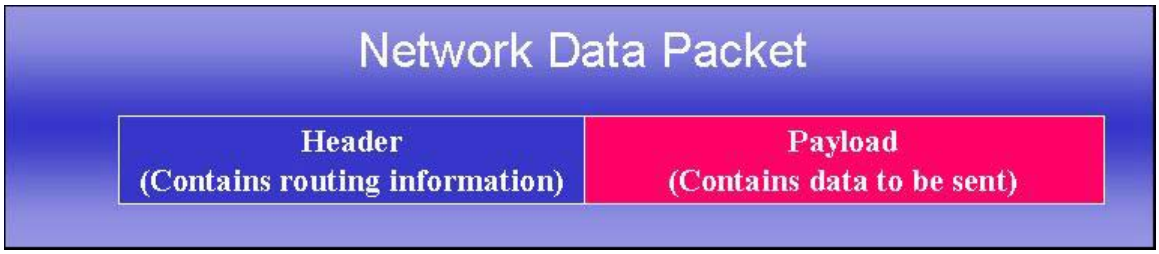

Figure 1. Network data packets contain a header with routing information and the payload containing the data to be sent.

Firewalls

Organizations that make the Internet available to their workers typically install a firewall to prevent intruders from getting into the organization's private data resources. A firewall is a device located on a private network that protects the resources of the network from outside malicious intent.

© 2002 Wainhouse Research, LLC & Ridgeway Systems and Software, Inc.

Firewalls examine the IP address and destination port of each data packet received from the outside world. Firewalls are often configured so that if a computer on the *inside* of the firewall requests data from a computer on the *outside* of the firewall, the firewall will let the data from the computer outside the firewall pass, but only if it sends the data packets to the same IP address and port number of the computer on the inside of the firewall that originated the request. If the firewall receives a packet destined for a computer on the inside of the firewall, and it determines that the destination computer did not first initiate a request for data on that port number, the firewall will typically discard the incoming data packet.

Firewalls are almost always configured to block all unsolicited incoming network traffic. One exception is providing a web server inside the firewall for access by the outside world. In this case, the organization will configure the firewall to allow packets destined for port 80 and the web server's IP address to pass through. This enables those outside the organization to send unsolicited packets to the organization's web server requesting some type of data the organization has hosted on that server.

Network Address Translation (NAT)

Network address translation is an Internet standard that enables a local-area network (LAN) to use one set of IP addresses for internal traffic and a second address or set of addresses when connecting to services on an external network, such as the Internet. NAT devices are located where the LAN meets the Internet and are designed make all the necessary IP packet address translations. NAT serves two main purposes:

- Many organizations use NAT as a network security device because it hides internal IP addresses – if hackers do not know the IP address of a particular machine, it is much more difficult to break into that machine.
- NAT enables a company to use more internal IP addresses. Since these addresses are used internally only, there is no possibility of conflict with the IP addresses used by other companies and organizations[1].

Firewalls and NATs Obstruct IP Voice and Video Communications

IP-based voice and video protocols, like Session Initiation Protocol (SIP) and H.323, require voice or video endpoints to establish data communication channels with each other using IP addresses and data ports. Herein lies a dilemma: the endpoints must be "listening" for incoming calls in order to establish a data connection, but the firewall is usually configured so that unsolicited data packets are blocked.

[1] When the present IP address standard, IPv4, was adopted some years ago, it was assumed that 32-bit IP addresses supporting 4,294,967,296 addresses would provide far more IP addresses than the world would ever need. However, as computer use has exploded, the demand for IP addresses has grown beyond the present standard's capacity, and routable IP addresses are becoming scarce. According to calculations by the IETF, IPv4 addresses will be exhausted by about 2008. The IP version 6 (IPv6) standard will solve the problem of not enough IP addresses. IPv6 supports 3.4 x 10^38 addresses because its length is 128 bits. However, even with IPv6, NAT traversal issues will not disappear. Using NAT for security will likely not change even when IPv6 becomes commonplace. Another benefit of using private IP addresses is that IT managers do not need to buy or manage a pool of IP addresses every time they have a new device they want to enable. Private IP addresses are like PBX extensions where the manager can add, delete, or remove addresses as needed without having to contact an external entity or paying to reserve a block of numbers that may never be used.

© 2002 Wainhouse Research, LLC & Ridgeway Systems and Software, Inc.

Even if the network administrator opened up one firewall port to receive call initiation packets, such as "well-known port" 1720, the IP voice and video communications protocols require additional open ports to receive call control messages and to establish the voice and video data channels. These additional port numbers are determined dynamically, not in advance, which implies that the network administrator would have to open up all the firewall ports to allow voice and video communications, effectively disabling the firewall. Few organizations would allow a wide-open firewall for network security reasons.

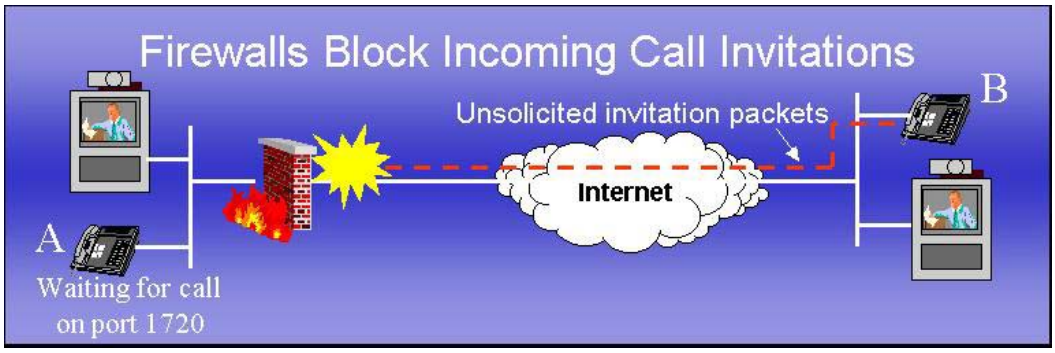

Figure 2. Most firewalls block call setup packets system B must send to system A when initiating a call.

NATs also cause challenges for IP voice and video communications. A NAT allows an organization to assign private IP addresses to devices on the local area network (LAN). Unfortunately, routing devices that control the flow of information across the Internet can send data only to devices with routable or public IP addresses.

An endpoint behind a NAT can initiate an IP call with any other endpoint on the same LAN because the IP addresses on the inside of the LAN are internally routable. However, because their IP addresses are private and are not routable beyond the LAN, endpoints behind the NAT cannot receive calls from endpoints outside the LAN.

Figure 3. Endpoints behind NAT cannot receive incoming Internet call invitations. If IP telephone B tries to call IP telephone A at its private IP address, the routers on the Internet are not able to route the packets, and they are discarded, causing the call to fail.

Even if the endpoint behind the NAT initiates the call to an endpoint outside the NAT, there is still a problem. When an IP call is initiated, the IP address of the endpoint initiating the call is embedded within the data packet payload. The endpoint being called receives the call initiation packets, opens them, and begins transmitting audio and video data back to the initiating endpoint

© 2002 Wainhouse Research, LLC & Ridgeway Systems and Software, Inc.

at its IP address obtained from the packet payload. If this IP address is private, Internet routers will discard audio and video data packets sent from the external endpoint to the internal endpoint because they are being sent to an un-routable IP address. The call will appear to have connected, but the endpoint behind the NAT never receives the external endpoint's audio and video.

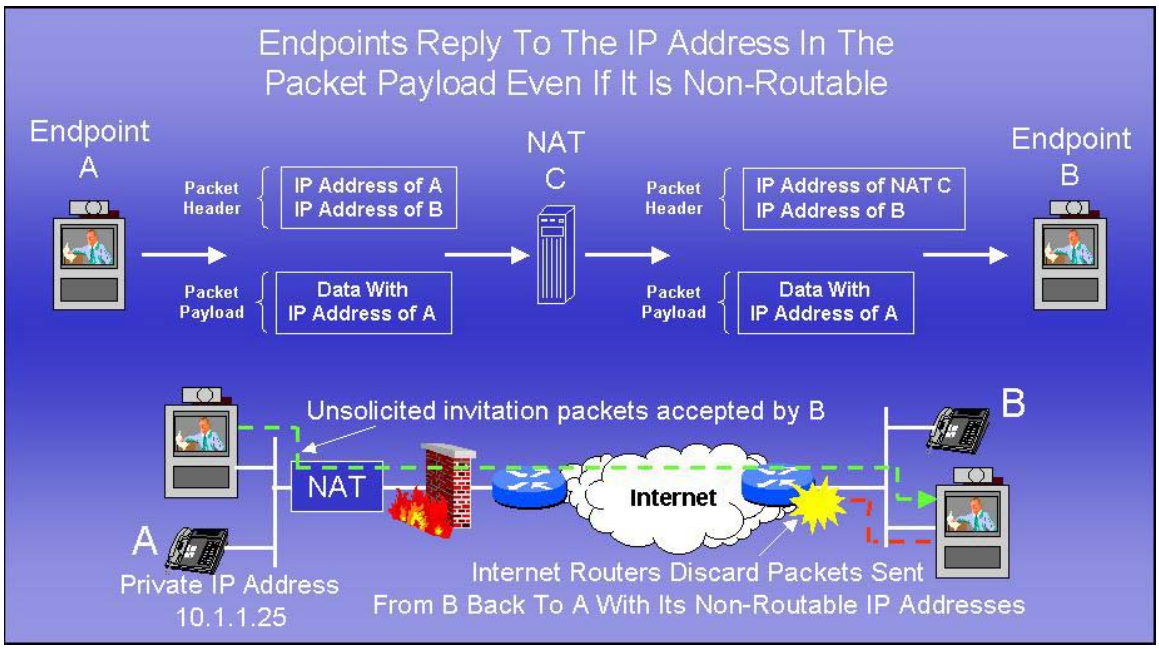

Figure 4. A NAT modifies the message data packet header information. When endpoint A creates a packet intended for endpoint B, the IP addresses of both A and B are placed in the packet header. The IP address of A is also placed in the data packet payload. NAT C substitutes its own IP address, C, in the message header. B receives the message and uses the IP address of A from the packet payload as the destination for return packets. However, since A has a non-routable IP address, the return packet from B cannot be routed back to A.

Methods for Firewall and NAT Traversal

Most organizations wishing to reap the benefits of IP rich media communications will ultimately face the firewall/NAT challenge. In practice, most organizations implement firewalls and NATs concurrently; consequently, traversing one does not necessarily negotiate both. A number of ways exist for traversing these devices as described below.

Open Firewall/No NAT

One obvious way for organizations to overcome the Firewall/NAT problem is to avoid using them. For most organizations, the security risks of this solution are too great; furthermore, obtaining enough routable IP addresses for the entire organization may prove difficult and expensive. There are, however, a number of organizations, particularly among educational institutions, that have little firewall protection, and they do not use NAT.

© 2002 Wainhouse Research, LLC & Ridgeway Systems and Software, Inc.

Gateway to PSTN

Rather than have any concern with employing IP communications outside the LAN, organizations can use a gateway to convert from IP voice and video on the LAN to PSTN voice and video over the public circuit-switched network. Use of a gateway eliminates the concern for network firewall traversal because no data packets cross the firewall. It also overcomes the NAT issue because all calls made to endpoints on the LAN are routable, and calls coming into the LAN through the gateway are routable. Today most IP telephones use a gateway to communicate with non-IP telephones both within and without an organization.

Gateway approaches are local solutions, however, that requires all locations participating in the call to have a corresponding gateway behind the last layer of NAT and firewall they have deployed. Using PSTN gateways also removes the converged network cost savings and mobile use benefits an all-IP solution provides.

Full Proxy

SIP or H.323 proxies can be used to negotiate the NAT or both the NAT and the firewall depending upon how they are configured. Proxies act like a gateway, but instead of converting from one IP communications protocol to another, the same protocol is used on both sides of the proxy. A proxy has knowledge of both the public and private IP networks and makes the IP call effectively look like two separate calls: one from the endpoint in the private network to the proxy and a second call from the proxy to the endpoint in the public network. Internally, the proxy puts these two calls together thus resolving the NAT issue.

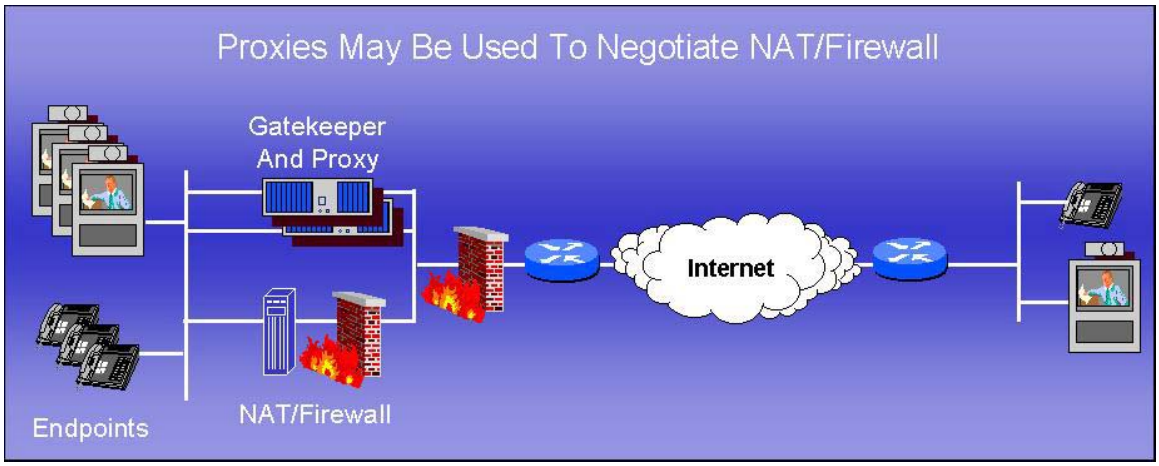

Figure 5. Proxies can be used to negotiate only the NAT or both the NAT and the firewall, as in this figure, depending upon how they are configured.

Proxies may have several different configurations in the network including being built into a gatekeeper or firewall, but they always require a gatekeeper for H.323 or a SIP registrar in order to resolve where to properly route the voice and video data packets.

In some cases NAT is deployed at multiple locations along the network path: multiple points within the enterprise, and even within the external network at the ISP. For a proxy to work, it needs to be deployed at every NAT.

© 2002 Wainhouse Research, LLC & Ridgeway Systems and Software, Inc.

Application Level Gateways

Application level gateways (ALG) are firewalls that are programmed to understand specific IP protocols, like H.323 and SIP. Rather than simply looking at packet header information to determine if packets can or cannot pass, ALGs go deeper by parsing the data in the packet payload. H.323 and SIP both put critical control information in the payload, such as which data ports the voice or video endpoint is expecting to use to receive the voice and video data from the other endpoint in the call. By understanding which ports need opening, the firewall dynamically opens only those ports needed by the application, leaving all others securely closed. This technique of opening small numbers of ports in the firewall dynamically is called "pinholing."

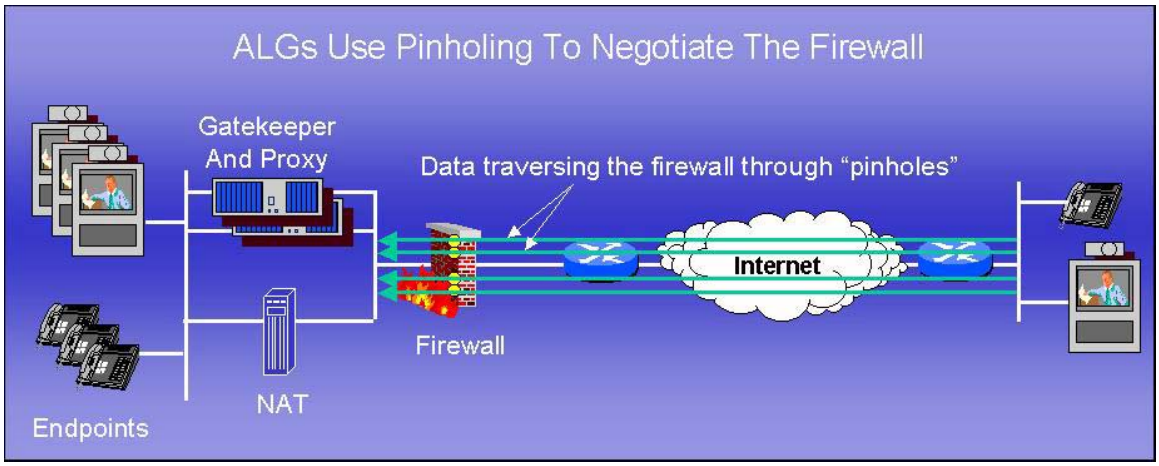

Figure 6. Application layer gateways parse the packet payload and open "pinholes" in the firewall based upon the IP protocol being used to allow voice and video data to pass through.

ALGs require a proxy if a NAT is being used to hide internal addresses. Some firewall manufacturers build the Proxy into the ALG, but it must be there to negotiate the NAT. As firewalls are mission-critical components for most enterprise networks, adding an ALG may prove difficult in some organizations because it requires both physical and political access to the firewall.

Major firewall vendors like Cisco, Checkpoint, Raptor, and Gauntlet, have developed H.323 ALG upgrade capability for their firewall products, but ALGs are not available for all firewalls. ALGs are not consistently implemented from vendor to vendor; for example, at least one vendor's ALG does not allow T.120 data sharing. ALGs can affect network performance, placing a heaver load on the firewall due to the parsing of the packet payloads. Moreover, if there are multiple levels of firewall/NAT combinations, each firewall/NAT in the call path must be upgraded to support ALG functionality.

MIDCOM Devices

Middlebox Communication (MIDCOM) is a scheme very similar to the ALG method; however, in the MIDCOM approach to firewall and NAT traversal, protocol intelligence is not built into the firewall. Instead, the SIP or H.323 protocol knowledge is built into a different device, called a "trusted system," that tells the firewall which ports to open for a given voice or video call. The advantage of this technique, in principle, is that firewalls do not have to be continually upgraded as protocols change or as they come in and out of fashion. Its disadvantages are similar to those

© 2002 Wainhouse Research, LLC & Ridgeway Systems and Software, Inc.

of ALG. Additionally, the firewall would require an initial "forklift" upgrade to implement the MIDCOM strategy. This method is still under development in an Internet Engineering Task Force working group[2].

DMZ MCU

Some organizations overcome NAT and firewall traversal issues by placing a multipoint control unit (MCU) in what is known as the demilitarized zone, or DMZ. The DMZ usually sits between the Internet and an internal network's firewall. Organizations that want to host their own Internet services, such as web servers, ftp servers, email servers, and domain name servers, without sacrificing unauthorized access to their private networks, place these servers in the DMZ.

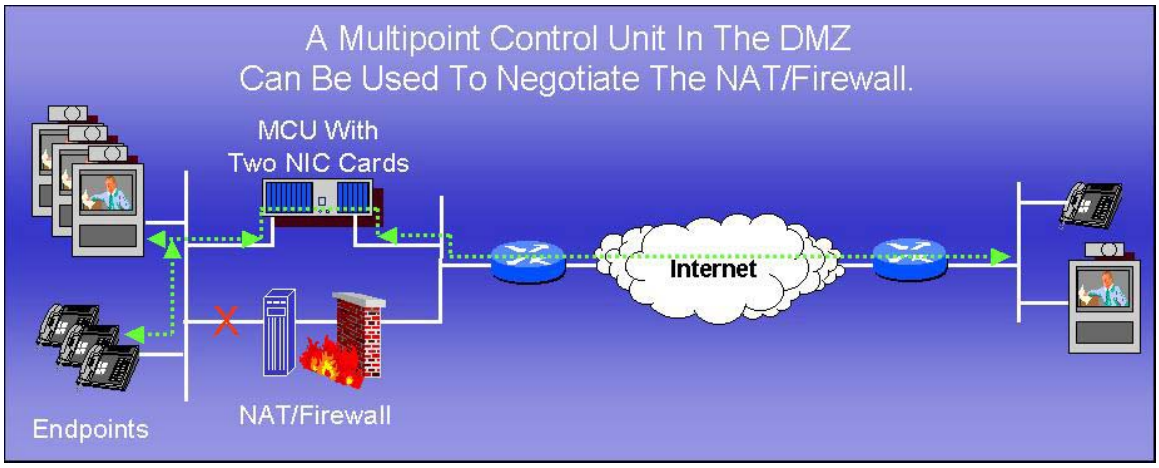

Figure 7. Multipoint Control Units may be used to bypass the firewall and NAT.

An MCU can be configured in the DMZ with two network interface cards such that one card provides access to the private network, and the other card gives access to the public Internet. One of the big disadvantages to this solution is that it uses ports on the MCU unnecessarily if the call is only point-to-point. As with a proxy or an ALG, if there are multiple NATs in the network, an MCU would need to be placed across each individual NAT. This solution does not scale very well either.

Semi-Tunnels/Transparent Traversal

The semi-tunnels/transparent traversal method for allowing IP voice and video to negotiate firewalls and NATs is based around the idea that organizations do not want to upgrade or modify their current NAT and firewall configurations, nor do they want to bypass them. IP Freedom, a semi-tunnel/transparent traversal method patented by Ridgeway Systems & Software[3], borrows concepts from the proxy method discussed above to "funnel" data traffic through two "well known ports[4]" in the firewall. IP Freedom is marketed and sold in two configurations:

[2] More information about MIDCOM can be found at http://www.ietf.org/html.charters/midcom-charter.html.
[3] For more information on Ridgeway Systems, go to www.ridgewaysystems.com.
[4] These ports, numbers 2776 and 2777 have been assigned to Ridgeway Systems by the Internet Assigned Numbers Authority (www.iana.org).

© 2002 Wainhouse Research, LLC & Ridgeway Systems and Software, Inc.

VoiceFreedom for voice-over-IP systems, and VideoFreedom, designed for systems using both video and voice over IP.

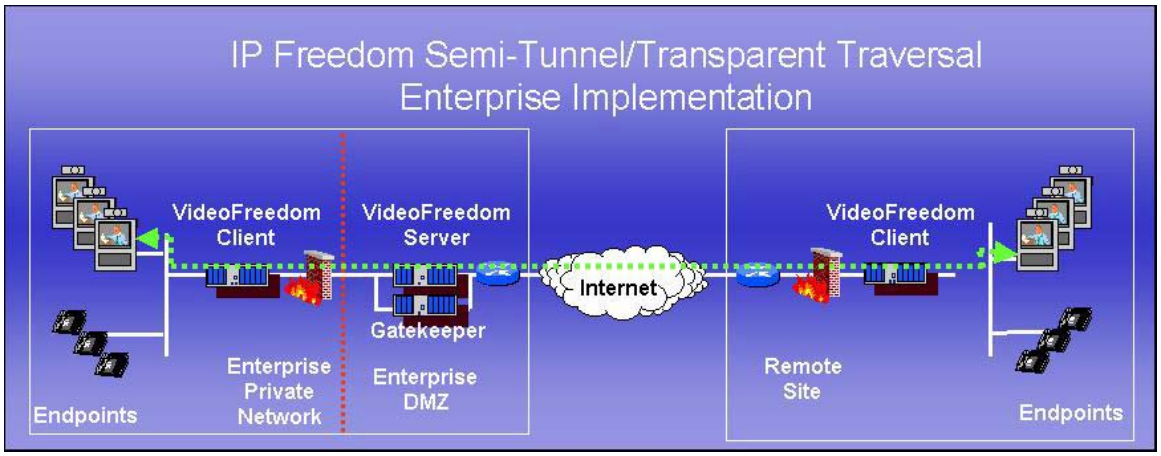

Figure 8. Ridgeway System's IP Freedom traversal method for enterprises places a single server in the DMZ and a client within each private local area network.

In this solution, a Ridgeway VideoFreedom Server (server) in the DMZ plays the role of a full proxy. A Ridgeway VideoFreedom Group Client (client) inside the firewall, in the private address space also acts as a proxy in that it substitutes its own address and port numbers within packets sent to and received from the endpoints.

When the VideoFreedom client starts,

1. It creates a single connection to port 2776 on the server for control and status information.
2. It listens for H.323 gatekeeper registrations and inquiries or SIP proxy/registrar registrations and inquires.

As an IP endpoint boots up,

1. Endpoints send registration information, typically consisting of a telephone number and/or an email address, to the gatekeeper or registrar through the client/server connection.
2. The server allocates each registering endpoint a unique port on the server IP address and registers that endpoint with the gatekeeper.

When an endpoint makes a call to another endpoint outside the firewall, all data packets are routed through the client to the server and from the server through the client back to the endpoint. As the call is setup, the client insures that all needed voice and video connections through the firewall are opened in an outbound direction. Voice and video data then flow in both directions through the firewall via these open ports.

IP address information is well hidden using this method. Because all packets are routed through the server, each endpoint behaves as if it were communicating directly with the server, not with another endpoint. This assures that the endpoint IP addresses are not available outside the network to packet sniffing devices.

Because most firewalls trust opening connections originating from devices inside the firewall, Ridgeway's IP Freedom traversal method can be used with no modification to the firewall

© 2002 Wainhouse Research, LLC & Ridgeway Systems and Software, Inc.

configuration in the majority of cases. In those instances where firewalls are programmed to restrict opening outbound ports, the administrator can create simple rules allow outbound connections from the client to the two well known ports, 2776 and 2777, on the server.

This method is not restricted to enterprises. Service providers can provide the IP Freedom NAT and firewall traversal method to small businesses and telecommuters by placing the IP Freedom server on the ISP's network backbone. Ridgeway has also developed a version of the software client that runs on an individual workstation if a small organization has endpoints that are all desktop computer based.

In either the enterprise or service provider deployment scenario, reaching a new endpoint can be as simple as downloading a software client, loading it on a PC, and placing an ad-hoc call. This means IT managers and end users don't care about how many NATs or Firewalls are in the path between the parties in the call or whether IT administrators have the physical, political, or even intellectual ability to access the infrastructure components.

As with the proxy method, the biggest disadvantage of this approach is that all of the traffic crossing the firewall goes through the IP Freedom server, causing a potential bottleneck. This is necessary because the server is the only device that the firewall trusts. Additional processing through the client and server typically adds less than 5 ms of latency, however.

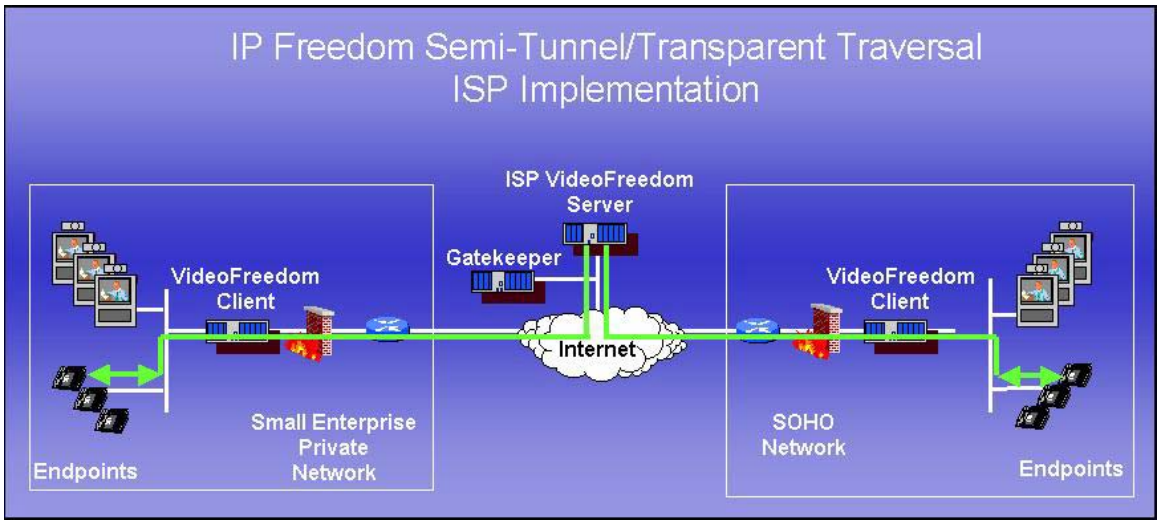

Figure 9. Ridgeway System's IP Freedom firewall/NAT traversal method for ISPs places the Ridgeway VideoFreedom or VoiceFreedom server in the ISP's network operations center. Packets from each small enterprise LAN or SOHO endpoint traverse the ISP's server.

© 2002 Wainhouse Research, LLC & Ridgeway Systems and Software, Inc.

Summary of Method Advantages and Disadvantages

Method	Advantages	Disadvantages
No Firewall No NAT	• Simple to use IP voice and video applications	• No network security whatsoever
PSTN Gateway	• No firewall/NAT issues to worry about • Relatively straightforward to use	• Removes advantages (reach and ad-hoc character) of using IP • Can incur costly tolls
Full Proxy	• No firewall upgrade • Method understood by IT managers	• Requires access to every NAT/Firewall on the call path for successful implementation • IP address published • Permanent inbound connection to proxy server required • Additional IP addresses required • Requires one proxy per protocol • Additional media hop required causing some latency
Application Level Gateway	• No additional equipment required • Method understood by IT managers	• Requires access to every NAT/Firewall on the call path for successful implementation • Firewall/NAT upgrade likely required • Additional processing on firewall • IP address published • Permanent inbound connection required
MIDCOM	• Once firewalls upgraded, will not need further upgrades if protocols change	• Requires access to every NAT/Firewall on the call path for successful implementation • Firewall/NAT upgrade required. • Some additional processing on firewall • IP address published • Permanent inbound connection required • Complex and unproven • Does not scale down
DMZ MCU	• Straightforward to implement • Does not require modifications to NAT/firewall to implement	• Deploy at every NAT/Firewall on the call path for successful implementation • Bypasses NAT and firewall • MCU IP address published • Does not scale down • Consumes expensive MCU ports
Semi-Tunnel/ Transparent Traversal	• Straightforward to implement (particularly client on workstation) • Does not require access to NAT/firewall to implement • Does not bypass NAT/firewall • IP addresses not published • Works for all protocols • Works for all deployments	• All traffic routed through single server • Additional hops could increase latency

© 2002 Wainhouse Research, LLC & Ridgeway Systems and Software, Inc.

Conclusion

NAT and firewall traversal is an issue all organizations wanting the benefits of IP voice and video will ultimately encounter. A number of methods exist for allowing voice and video data to traverse firewalls and NATs. Some organizations can avoid firewall and NAT traversal temporarily by using a gateway to PSTN, particularly for IP telephones. However, those organizations wanting to avoid the PSTN tolls and quality issues will need an immediate secure NAT and firewall traversal method.

IT professionals will likely favor the semi-tunnel/transparent or the full proxy methods because they do not require firewall upgrades nor do they compromise network security. Furthermore, these solutions are fairly straightforward to implement, although setting up and managing the full proxy method will likely add some overhead to a network administrator's already full schedule.

In the traversal method employed in the Ridgeway IP Freedom solution, IT professionals need not worry about how many NATs or firewalls are deployed along the path in their IP communications system, nor will they need to be concerned about what brand of firewall or NAT is used. This gives enterprises the ability to connect to other organizations, the telecommuter, or the road warrior using a single solution. Organizations that do not want to create their own firewall/NAT traversal solution, should consider Ridgeway's IP Freedom traversal method.

###

© 2002 Wainhouse Research, LLC & Ridgeway Systems and Software, Inc.