

NAT Traversal for Multimedia over IP



Simply Better Connected

Solving the Firewall and NAT Traversal Issues for MoIP

Executive Summary

Service Providers are looking at IP-based voice and multimedia services to solve some of their business challenges. For local access companies, including incumbents, LECs etc. this is about retaining business customers. For ISPs, it is about increasing revenue and profit by offering voice and multimedia services in addition to existing data/Internet services. For long-distance carriers, it is about reducing costs – particularly for interconnection with other providers. Finally, wholesale carriers want to offer enhanced services to their Service Provider customers and reduce their interconnection costs.

A number of technical problems have to be overcome before Service Providers and carriers can benefit from these new services. Today, it is important to be able to provide secure connection to subscribers behind NAT (Network Address Translation) devices and Firewalls. Overcoming this challenge will lead to continued deployment of profitable voice and multimedia services over IP to any subscriber with a broadband connection.

This White Paper describes the various proposals for solving the NAT and Firewall traversal problems and shows how the session border controller provides the preferred solution for Service Providers.

The NAT and Firewall Traversal Problems Defined

Firewalls and NATs are located at the edge of virtually all business networks. Often software-based Firewalls and NATs are bundled in residential DSL packages as well, so this problem affects both business users and residential users.

The problem actually consists of two components. While today's Firewalls are able to dynamically open and close multiple ports as required by VoIP signalling protocols, such as SIP, they remain ineffective at securely supporting unsolicited incoming media flows. NATs prevent two-way voice and multimedia communication, because the private IP addresses and ports inserted by client devices (IP phones, video conferencing stations etc.) in the packet payload are not routable in public networks. Thus, incoming calls that are essential in any service that is intended to replace the PSTN are not possible with existing NAT/Firewalls.

The 'Firewall Problem'

The role of the Firewall is to protect the network from being accessed by unauthorised sources. It does this by blocking traffic based on three parameters: the source, the destination and the traffic type. Firewalls also make decisions based on the direction of traffic flow. Typically, incoming traffic (from the un-trusted, public domain) is only allowed if that session has been initiated from a device on the trusted, private domain

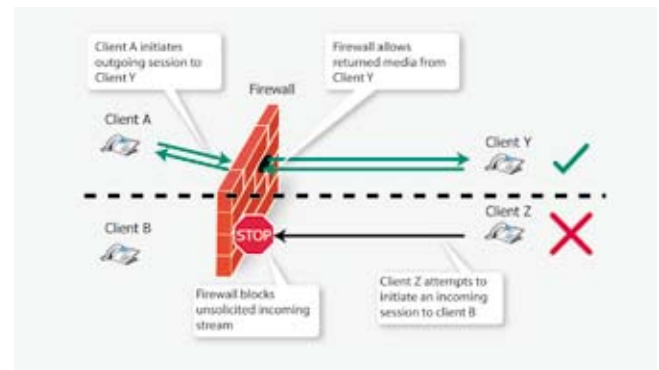


Figure 1 The Firewall Problem

SIP-based communication, much like traditional telephony, is based on unsolicited incoming calls, from a wide range of unknown (and therefore un-trusted) sources – as it must be to support true public services. However, this is at odds with the Firewall filtering policies described above. Most Communication Managers are rightly reluctant to change these policies to allow unrestricted two-way communication because of the serious security risks created.

Any approach to solving this problem must allow secure two-way communication – *without* major changes to Firewall filtering rules. In addition, it should not reduce the current level of security provided by Firewall.

The 'NAT Problem'

NATs translate IP addresses and port numbers in private address ranges into public addresses when traffic flows from a private to public network. This allows a limited number of public IP addresses to serve the needs of even the largest corporation.




Figure 2 The NAT Problem

Each device in the private network has its own private IP address. Traffic (a media stream, for example) sent to a device on the public network will be dynamically assigned a specific port number at the public address by the NAT. The NAT maintains a 'table' that links private addresses and port numbers to the public port numbers and IP addresses. It is important to note that these 'bindings' can only be initiated by outgoing traffic.

The NAT acts in a similar way to a PABX. Initiating an outbound connection is easy. Users of the PABX can dial out using one of the few public telephone lines (equivalent to public IP Addresses) that are available. The line that is used (the port number) is automatically selected and invisible to the user. Receiving an incoming call is more difficult however, because the internal extension numbers are non-routable from the public network. Users dialling in must be routed to an attendant to be connected to the correct extension. Clearly in the case of a NAT, there is no equivalent of an attendant so unsolicited incoming calls cannot be supported.

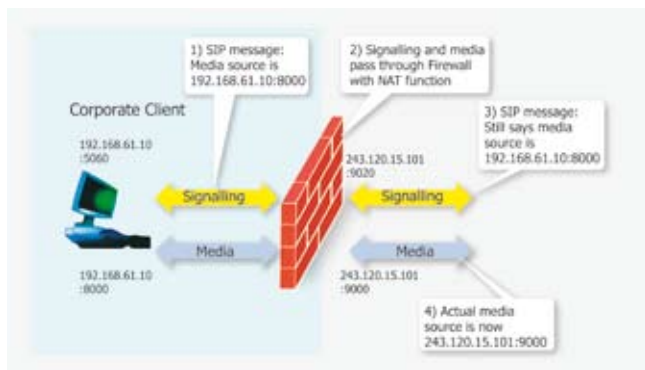


Figure 3 NAT breaks end-to-end media flow

To complicate matters further, the end-to-end SIP messaging between clients contains details of the private IP addresses and ports that the clients (User Agents) want to use for the media flows. When the clients attempt to use these private addresses to send/receive media, the connection fails because they are un-routable. Note that this issue also applies to other signalling protocols such as H.323 and MGCP.

Any approach to solving this problem must allow secure two-way communication – including unsolicited incoming calls – and minimise dependence on upgrading NATs or even using any specific vendor's NAT device.

Methods of Solving the 'NAT Problem'

In this White Paper, we will describe a range of proposed solutions to the 'NAT problem' and consider the implications of each with regard to security (i.e. how it solves the 'Firewall problem').

The current proposals for solving the 'NAT problem' are:

- ♦ Universal Plug and Play (UPnP)
- ♦ Simple Traversal of UDP Through Network Address Translation devices (STUN)
- ♦ Application Layer Gateway
- ♦ Manual Configuration
- ♦ Tunnel Techniques
- ♦ Automatic Channel Mapping™ (ACM)

Universal Plug and Play

UPnP is a technology that is predominantly targeted at home-office users and domestic residential installations etc. One of the driving forces behind UPnP is Microsoft Corporation.

The UPnP architecture is designed to address a number of general issues – not just VoIP – and is designed to allow the ready configuration of small networks by typically un-skilled people. UPnP allows client applications to discover and configure network components, including NATs and Firewalls, which are equipped with UPnP software.

The fundamental need in VoIP applications is to discover and use the external IP address and port that the NAT selects for signalling and media flows. Once this information is known, the VoIP client (a SoftPhone or a standalone SIP phone) can put this information into the VoIP signalling to establish the call. This ensures that the call is established using public, routable addresses and ports, and ensures end-to-end connectivity.

To achieve this, the NAT and VoIP clients must be UPnP enabled. While many small end NAT vendors are committed to supporting UPnP for VoIP purposes, there are few UPnP VoIP clients available from many manufacturers yet. However, it is only a matter of time before these devices are available and many small companies (and residential subscribers) will find them useful.

The main disadvantage of this approach is related to security. It does not satisfactorily solve the 'Firewall problem'.

UPnP relies on the NAT opening pinholes to the outside world under the dynamic control of the UPnP client – maybe a SoftPhone on a PC. This capability is most likely contrary to most security/Firewall policies and therefore may not be accepted by Communications Managers of large corporate customers.

In addition, there are presently only a small number of NAT and Firewall vendors committed to supporting UPnP.

In summary, this method is likely to be limited to small installations.

Simple Traversal of UDP Through Network Address Translators (STUN)

The STUN protocol enables a SIP client to discover whether it is behind a NAT, and to determine the type of NAT. STUN has received a lot of attention as a technique in the IETF, but suffers from a flaw, which means that it will only work with some NATs. In fact, STUN does not work with the type most commonly found in corporate networks – the symmetric NAT.

The IETF Midcom Working Group has undertaken an investigation of residential NAT devices. This indicates that some NATs will not work with STUN.

Also, STUN does not address the need to support TCP based SIP devices. As SIP User Agents and Call Agents become more complex, the use of TCP will increase. In fact, support of TCP is mandatory in the SIP RFC3261.

The STUN proposal defines a special server (STUN server) in the public address space to inform the STUN-enabled SIP client in the corporate (private) address space of the Public NAT IP address and port being used for that particular session. Having to use STUN-enabled clients, or upgrade existing clients to support STUN, makes this method unpopular. In fact, very few vendors have announced STUN support for their clients.




Figure 4 STUN

STUN identifies the public side NAT details by inspecting exploratory STUN messages that arrive at the STUN server. The STUN-enabled client sends an exploratory message to the external STUN server to determine the receive ports to use. The STUN server examines the incoming message and informs the client which public IP address and ports were used by the NAT. These are then used in the call establishment messages. Note that the STUN server does not sit in the signalling or media data flows.

As mentioned previously, however, there is a problem with this approach. Most NATs in use today are 'symmetric NATs'. This means that they create a mapping based on source IP address and port number as well as the destination IP address and port number.

The destination VoIP client address is different from that of the STUN server. This means that the NAT will create a new mapping using a different port for outgoing traffic, which in turn means that the information contained in the call establishment messages is incorrect and the call attempt may fail. The same problem occurs incoming traffic.

Therefore, STUN relies on the fact that once the outgoing port has been mapped for the STUN server traffic, any traffic appearing from any part of the network, with any source IP address, will be able to use the mapping in the reverse direction and so reach the receive port on the client.

NATs that do work in this way are susceptible to port scan attacks and generally create major security concerns. This method, therefore, fails to solve the 'Firewall problem' because it introduces additional security risks that are unacceptable to Communication Managers.

The IETF have proposed an additional mechanism – TURN – that is designed to solve the media traversal issue for symmetric NATs.

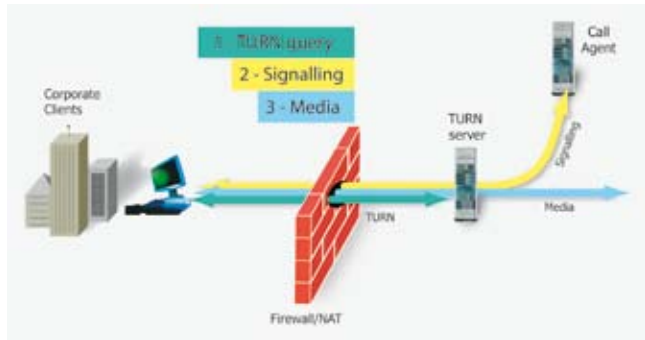


Figure 5 TURN

TURN relies on a server that is inserted in the media and signalling path. This TURN server is located either in the customers DMZ or in the Service Provider network. The TURN-enabled SIP client sends an exploratory packet to the TURN server, which responds with the public IP address and port used by the NAT to be used for this session. This information is used in the SIP call establishment messages and for subsequent media streams. The advantage of this approach is that there is no change in the destination address seen by the NAT and, thus, symmetric NAT can be used. TURN has recently been extended to address some serious security issues associated with TURN, which may have held back its acceptance.

Many Service Providers expect to be able to manipulate QoS information and provide enhanced security at the entry point to the network. TURN does not support this requirement, because details of the SIP session are not revealed to the TURN server through the TURN protocol, so its acceptance by the Service Provider community is not certain at this stage.

Both these methods add significant complexity to the CPE installation and TURN, like STUN, requires SIP clients (such as SoftPhones or IP telephones) to be upgraded. There is considerable reluctance by client vendors to undertake this work, making STUN and TURN non-ideal solutions. SNOM are clearly an exception to this reluctant attitude

Application Layer Gateway (ALG)

This technique relies on the installation of a new, enhanced Firewall/NAT – called an Application Layer Gateway – that ‘understands’ the signalling messages and their relationship with the resulting media flows.




Figure 6 Application Layer Gateway

The ALG processes the signalling and media streams so it can modify the signalling to reflect the public IP addresses and ports being used by the signalling and media traffic.

As suggested, this technique requires replacement of the existing NAT/Firewall with an ALG. Alternatively, some vendors provide software upgrades to their NAT/Firewalls to support ALG functionality.

ALGs require similar, if not more advanced, configuration and management skills to NATs and Firewalls, which means that upgrades or new installations will not be undertaken lightly. These issues mean that deployment of ALGs is likely to be slow and restricted to larger corporate networks with the associated support staff.

Manual Configuration

In this method, the client is manually configured with details of the public IP addresses and ports that the NAT will use for signalling and media. The NAT is also manually configured with static mappings (or ‘bindings’) for each client.

This method requires that the client must have a fixed IP address and fixed ports for receiving signalling and media.




Figure 7 Manual configuration

Due to the manual and knowledge-based configuration process, as well as the fixed configuration, this is only suitable for very small networks where there is a great deal of experience in configuring and managing the NAT/Firewall.

It is very likely that UPnP, when available, will supersede this manual method.

Tunnel Techniques

This method achieves Firewall/NAT traversal by tunnelling both media and signalling through the existing Firewall/NAT installations to a public address space server.

This method requires a new server within the private network and another in the public network. These devices create a tunnel between them that carries all the SIP traffic through a reconfigured Firewall. The external server modifies the signalling to reflect its outbound port details, thus allowing the VoIP system to both make outgoing calls and accept incoming calls. The tunnel through the existing infrastructure is not usually encrypted.




Figure 8 Tunneling

While this method provides only minimal changes to the existing security policy, it does create additional risks. In particular, the external server is a point of vulnerability, which, if breached, can provide an easy way of reaching the private network.

In addition, this method can create additional delay in the media path – which may reduce voice quality.

Automatic Channel Mapping™ (ACM)

The Newport Networks 1460 session border controller, equipped with the ACM Application Pack, is specifically designed to address the NAT and Firewall problems – without requiring any changes to the existing security rules or to the clients.

Solving the 'NAT Problem'

Solving the 'NAT problem' means that un-routable private addresses need to be replaced with public, routable ones so that the media flows can find their way through networks (both public and private) to reach the client devices. This is the primary function of the SignallingProxy™.

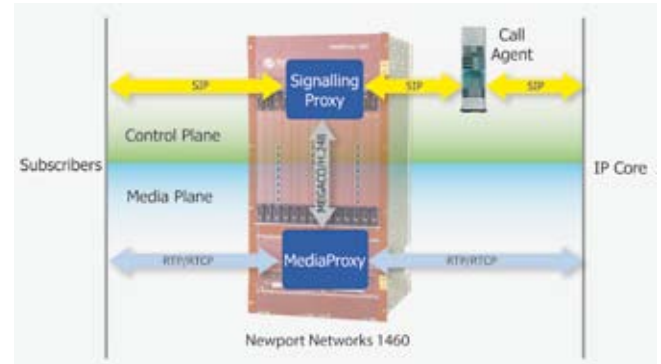


Figure 9 Newport Networks 1460 session border controller

The SignallingProxy™ acts as a high-performance B2BUA (Back to Back User Agent). It is configured as a transit point for SIP signalling messages between the client (User Agent) and the Call Agent (and vice versa). In this way, it acts as a proxy for both client and server – ensuring that all signalling messages pass through it. This provides complete visibility and control of call establishment. SIP messages from the client are directed to the SignallingProxy™ by making minor changes to the Service Provider DNS entries for the Call Agent.

The MediaProxy™ operates under the control of the SignallingProxy™ to provide a transit point for RTP and RTCP media streams between User Agents. All media is directed to the MediaProxy™ ensuring that the Service Provider has full visibility and control of the media stream to ensure service quality, and for charging purposes. Finally, the MediaProxy™ performs dynamic NAT to hide details of the network and other users from subscribers and other networks – helping to provide protection against Denial of Service attacks.

The SignallingProxy™ and MediaProxy™ exchange information using an internal Megaco/H.248 protocol. This approach makes the Newport Networks 1460 'SIP-ready' and compatible with all SIP User Agents and Call Agents. This ensures easy integration with existing SIP systems and fast time to market.



Figure 10 Interaction with External Call Agent

Call Agents that have been suitably enhanced can control the MediaProxy™ through the Megaco/H.248 interface, which, in turn, controls the media plane. This ensures support for a broad range of call control protocols including H.323, MGCP and of course SIP. Secondly, it means that other call control devices (such as Softswitches) can directly control the MediaProxy™. The result is that Service Providers can implement specific features that will differentiate their service from competitors.

Taking Control of the Signalling Path

SIP signalling messages destined for the SignallingProxy™ exit the private network using a public IP address and port allocated by the NAT. When the SignallingProxy™ receives the initial REGISTER message from the User Agent, a source address on the SignallingProxy™ is allocated for signalling messages to this client. A modified REGISTER message is then forwarded on to the Call Agent with the CONTACT and VIA fields indicating that the SignallingProxy™ is the source.

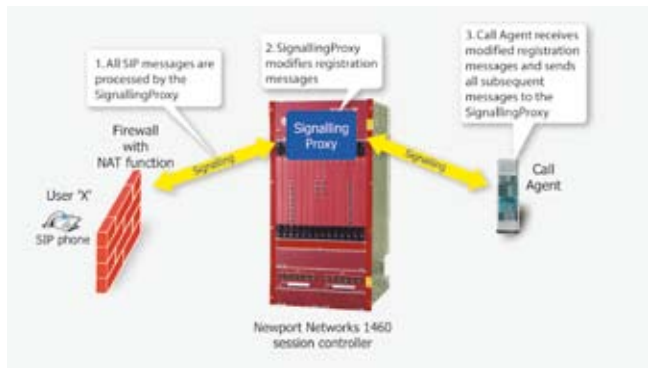


Figure 11 SignallingProxy™ controls the signalling path

Taking Control of the Media Path

The consequence of the NAT public port allocation method is that the ports that are allocated for media flow from each client will be unpredictable.

In the Newport Networks solution, based on another IETF draft, the SignallingProxy™ manipulates the signalling messages to ensure that the media streams are directed to specific, dynamically allocated ports on the MediaProxy™.

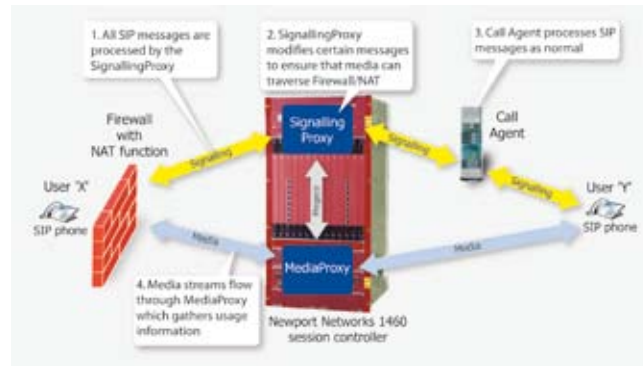


Figure 12 Ensuring end-to-end media flow

When the User Agent initiates a call, the SignallingProxy™ receives the INVITE message. It communicates with the MediaProxy™ to obtain NAT information for this call. It then modifies the source IP address and the SDP fields to define the SignallingProxy™ as the return path for signalling and the MediaProxy™ address as the return address for the media. The INVITE is then forwarded to the Call Agent. To the Call Agent it appears that the message has originated from a user with a port and IP address belonging to the MediaProxy™. The receiving User Agent will return an ACK via the SignallingProxy™, which will modify the message so that the originating User Agent directs the media to the port dynamically allocated to this call on the MediaProxy™. The IP address and port used by the NAT can now be easily determined by reading the IP address and port from the actual media stream. Thus, all signalling messages flow through the SignallingProxy™ and all media streams will flow through the MediaProxy™ allowing the Service Provider to connect, control and charge for the connection.

Solving the 'Firewall Problem'

Solving the 'Firewall problem' means allowing secure incoming, unsolicited media from unknown IP addresses and ports. This is in clear conflict with sensible security policies. In the Newport Networks solution, the MediaProxy™ acts as a transit point (or meeting point) for all media sessions. Media sessions are always initiated from inside the Firewall – sent to a specified IP address and port on the MediaProxy™ that has been dynamically allocated for that session. The MediaProxy™ learns the originating public address from this in order to return the incoming stream to the same address and port.

Thus, receiving an incoming call is achieved through always establishing outgoing paths first, complying with typical Firewall security policies.

NAT traversal with IPsec

The increasing need to provide security for SIP signalling has led to bodies such as 3GPP and TISPAN to evaluate and select suitable security protocols. 3GPP selected IPsec ESP, however this was not suitable for TISPAN's use in fixed line networks. IPsec encounters problems when traversing NAT devices which lead TISPAN to select UDP encapsulation of IPsec. This overcomes the NAT traversal problems whilst providing the required encryption and authentication and still complies with 3GPP's overall security architecture. A separate White Paper "IPsec in VoIP Networks" examines the flavours and IPsec and TISPAN's selection of UDP encapsulation.

Conclusion

Service Providers are looking at IP-based voice and multimedia services as major sources of new revenue. Unfortunately, a number of technical problems have, to date, prevented Service Providers and carriers alike from realising these benefits. Today, the most significant of these is to provide secure connection to subscribers behind NAT (Network Address Translation) devices and Firewalls.

The Newport Networks 1460 Automatic Channel Mapping™ (ACM) Application Pack solves these problems by enabling secure traversal of ALL corporate Firewall/NATs. This solution does not require additional customer premise equipment, nor does it require the replacement of existing Firewalls and NATs. This removes the necessity to visit customer premises to install new equipment, reduces the cost of connecting new subscribers and significantly simplifies the subscriber registration process.

Other White Papers

You may find the following White Papers of interest:

- SIP, Security and Session Controllers
- IPsec in VoIP Networks
- SIP, Security and Session Controllers
- Lawful Interception Overview

These and other papers can be found at:

<http://www.newport-networks.com/whitepapers>