NETWORK ADDRESS TRANSLATION IN HOME GATEWAYS

Jed Johnson

Art Harvey

Motorola, Inc.

Abstract

Network Address Translation (NAT) has become a common feature in Home Gateways because it reduces the number of IP addresses a service provider needs to manage, it bounds the domain of the network and it provides a modicum of security for the home owner. NAT however does introduce problems because it breaks the end-to-end addressing assumption built in to many applications. In addition, the applications that have the most difficulty are generally the applications that deliver advanced services like IP telephony and streaming media.

NAT also makes network management difficult because the devices behind NAT cannot be addressed directly. If a service provider wants to be able to diagnose a network problem through NAT, many standard tools and procedures will not work.

Because of these problems, NAT in many circles has been equated to "a bad thing" that must be eliminated. This paper takes the position that NAT cannot be eliminated from all, and some might say most, home networks so we should learn to deal with it. This paper specifically looks at how end-to-end management and advanced services can be delivered with NAT in place in the home network.

This paper will review how NAT causes problems and then go on to show how extensions or work-arounds to NAT can recover the end-to-end addressing assumption that applications require to work properly.

INTRODUCTION

NAT has become prevalent in home networks so it needs to be discussed

We are not NAT fanatics but believe it can serve a useful purpose in home gateways

This paper will provide some background on NAT including a brief overview of how NAT works. The paper will discuss some advantages that NAT has and a section on debunking myths about NAT is included. Finally the paper covers some system level descriptions of how NAT can be used to solve unique home networking problems.

NAT OVERVIEW

The most concise definition of NAT is, "Network Address Translation is a method by which IP addresses are mapped from one realm to another, in an attempt to provide transparent routing to hosts." (RFC 2663) This definition is typically implemented as a gateway device that connects a private address space, such as that of a homeowner or business, to the public Internet address space through an Internet Service Provider (ISP). The NAT gateway replaces a private network address with a public one in packets sent from a system in the private network to the public network, and performs the inverse replacement for packets flowing in the reverse direction. This mapping of addresses between addressing realms is called "transparent routing".

RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations,* presents an overview of the variants of NAT

and the standard terminology. It also describes the characteristics of NAT, typical usage, operational characteristics and limitations. We only briefly describe some of these here and then mainly in the context of home networking.

There are many reasons for using NAT. A common one for large organizations is to avoid the problem of changing the network address of every system in the private corporate network if there is a change of the set of addresses provided by their ISPs. For the home network, the prime motivation for NAT is to share the single public network address provided by the ISP among multiple systems in the home so that all the devices in the home have Internet access.

There is no single method or standard for NAT and the variations are many. We present conceptual overviews of three of the main variants relevant to home networking. We skip many of the details, but more complete descriptions and other variants can be found in the appropriate RFCs and IDs. The methods we describe are known as:

- NAT with dynamic address assignment, see RFC 3022, *Traditional IP Network Address Translator (Traditional NAT)*
- NAPT - Network Address Port Translation, see RFC 3022, *Traditional IP Network Address Translator (Traditional NAT)*
- RSAP-IP - Realm Specific Address and Port IP, see draft-ietf-nat-rsip-framework-05, *Realm Specific IP: Framework*; and draft-ietf-nat-rsip-protocol-07, *Realm Specific IP: Protocol Specification*

Consider the case of a home with two PCs, a Home Gateway, and a connection to the Internet as shown in the following figure. Further suppose the ISP providing connection to the Internet assigns a single IP address to the Home Gateway.

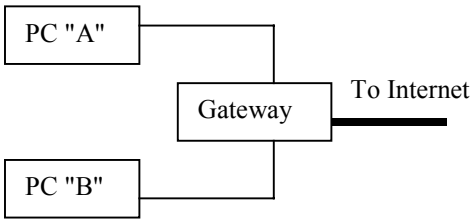

While data can be exchanged within the home using privately assigned addresses (RFC1918, *Address Allocation for Private Internets*), these addresses may be duplicates of those used in some other private realm and cannot be routed outside the home LAN. Suppose the Home Gateway supports NAT with dynamic address assignment. If PC "A" sends a packet to a remote system in the public network that has a globally assigned (unique) address, then the NAT gateway binds (typically stores in a mapping table) the private IP address of "A" to the IP address assigned by the ISP. Next, it replaces "A's" private IP address in the packet with the IP address assigned by the ISP. It then forwards the packet to the public network. If the gateway subsequently receives traffic from the remote system, it performs the inverse mapping (replaces the destination IP address with the address of "A") and forwards the packet on the private network. Based on some heuristic, (for example, receipt of a TCP FIN message indicating the connection is terminated and allowing some time for retransmission of lost packets) it unbinds (removes from the mapping table) the address of "A" from the IP address assigned by the ISP. "B" can now go through the same steps to communicate outside the private network.

Note that while "A" is using the single ISP assigned address, "B" cannot send or receive data from the public network and vice versa.

NAPT, Network Port Address Translation, enables multiple concurrent systems to communicate with remote systems across the public network. It does this by replacing not only the IP address of packets crossing the border between the private and public networks but also the TCP or UDP port address. The following example demonstrates how this typically works.

Suppose "A" has a locally assigned IP address of **IP-A**, and "B" has locally assigned IP address of **IP-B**. Denote the single, shared IP address assigned by the ISP as **IP-External**.

Consider the case where both hosts, "A" and "B", establish a connection with the same remote host, that has a global IP address of **IP-Rem,** and to the same application on a well-known TCP port denoted **TCP-ServerPort**. Also, assume both "A" and "B" choose the same source TCP port number **TCP-ClientPort**. [The algorithm operates in the same fashion when "A" and "B" choose different source port numbers and talk to different systems.] To initiate the connections, "A" and "B" send packets that include the following IP and TCP header information.

"A" sends
Destination IP address = **IP-Rem**,
Source IP address = **IP-A**,
Protocol = TCP,
Destination Port = **TCP-ServerPort**,
Source Port = **TCP-ClientPort,**
TCP message type = SYN

"B" sends
Destination IP address = **IP-Rem**,
Source IP address = **IP-B**,
Protocol = TCP,
Destination Port = **TCP-ServerPort**,
Source Port = **TCP-ClientPort**
TCP message type = SYN

The gateway NAPT receives these packets and notices the destination IP address in each packet is in the public address space and the protocol is TCP. The SYN indicates this is a new connection. These packets cause the NAT to create a mapping table entry that will last until it sees FIN messages that terminate the TCP connections or by some other heuristic. The NAT uses this table to change the outgoing packet headers above to:

Packet sent by "A" is translated to
Destination IP address = **IP-Rem**,
*Source IP address = **IP-External**, (the external address to be shared)*
Protocol = TCP,
Destination Port = **TCP-ServerPort**,
*Source Port = **TCP-A1** (a different source port)*

Packet sent by "B" is translated to
Destination IP address = **IP-Rem**,
*Source IP address = **IP-External**, (the external address to shared)*
Protocol = TCP,
Destination Port = **TCP-ServerPort**,
*Source Port = **TCP-B1** (a different source port)*

The modified packets are then forwarded to the public network. To the remote system, these appear to be two different connection requests from a single host with IP address **IP-External**. Packets received by the NAT device from the remote system undergo the analogous inverse translation, and are sent to the private network.

The nice thing about NAT is that it performs transparent routing. The PCs "A" and "B", and the remote system have no idea that the gateway is modifying the addresses. Unfortunately, this does not always work as desired. For example, the FTP application has messages ("PORT" and the "PASV" response) in which it includes the local system IP address in its data. Since a private address is of no use to a system outside that private network, a NAT ALG (Application Layer Gateway) must do the same type mapping of these IP addresses inside the application data as it does to the packet headers. This can get

tricky because these addresses are encoded in ASCII so that changing the address may also change the size of the packet. That means the ALG must also modify checksums, the TCP fields, and maintain state for TCP sequence numbers and acknowledgments. Throw IP fragmentation into the mix and it becomes apparent that maintaining transparency is not trivial. Furthermore, IPsec transport mode, both AH and ESP, include an integrity check over the entire payload including the TCP and UDP checksum. Modifying headers in protected packets will cause the receiving IPsec to discard the packet as having failed the integrity check. While these problems may seem insurmountable, they can all be addressed as described in subsequent sections.

Partially relaxing the transparency constraint eliminates many of these problems and is one of the motivations behind RSAP-IP (Realm Specific Address and Port IP). With RSAP-IP, address translation remains transparent to the application, but the network stack at the end systems are aware of the address mapping. Here is one possible implementation.

A system in the private network, say "A", queries an RSAP-IP server in the Home Gateway asking for an IP address and port number. The Home Gateway establishes the binding between the private IP address plus port of "A" and the external IP address plus port just as it does for NAPT. The Home Gateway responds to the query from "A" by returning this binding. Now when "A" sends a packet to a system outside the private network, "A" uses the external IP address and port numbers in the packet header. Also, if an application (such as FTP) asks for an address, the local stack can return an external address. Packets sent from "A" might be tunneled (encapsulated in another IP header) to the Home Gateway for decapsulation and transmission on the public network. They could also be tunneled from "A" directly to a remote system. They may even be sent just as they are since the Home Gateway knows the address binding and can know how to route the packets. Other alternatives are possible as well.

ADVANTAGES OF NAT IN A HOME GATEWAY

Firewall by nature

While many networking purists cringe at the thought, NAT in a broad sense can be regarded as a firewall and is sold as such in some Home Gateway products. Because most NAT implementations in Home Gateways only open up ports based on traffic that is initiated in the home, the only ports that are open are to support applications that reside in the home. Unsolicited traffic to any another port is dropped. The only traffic that gets through is for the ports associated with applications in the home and only while those applications are running.

There are variations of NAT that open up ports for UDP traffic and for servers in the home. These features obviously reduce the effectiveness of NAT as a firewall and need to be used with this fact in mind.

Natural demarcation point of ISP

The ISP providing Internet access to the home might not want to get involved with supporting the home network. If a homeowner has multiple PCs and wants each PC on the network then without NAT each PC needs a separate address and in all likelihood these PCs are located in different rooms in the house. This means that from an IP standpoint the ISP can be expected to managed connectivity through the home network to each PC.

The problem for the ISP is the lack of physical access to he home network and local of configuration control.

By using NAT in the Home Gateway the ISP can terminate management of IP connectivity in the Home Gateway and leave the

management of connectivity in the home to the homeowner.

Allows device provisioning in the home without direct involvement of ISP

Without NAT in the Home Gateway each device that is added to the network needs to be provisioned by the ISP. There are automated provisioned systems to reduce the manual effort for this process but resources like IP addresses are required in greater abundance.

With NAT in the Home Gateway IP management can be terminated in the Home Gateway as discussed in the previous section. This allows device provisioning up to the IP layer to occur in the home under the control of the Homeowner and transparent to the ISP.

It is possible, by the way, for the ISP to limit the number of devices that access the Internet from the home in this scenario. This means the ISP does not lose out on implementing a potential business model by not participating in network provisioning.

It also needs to be pointed out that service provisioning at the applicaityon layer is another subject and it can be independent of network provisioning.

Eases migration to IPng.

NAT can translate between different network layer protocols. Any migration from IPv4 to another address format will require this for existing systems to continue working. RFC 2766 *Network Address Translation - Protocol Translation (NAT-PT)*

Since proxies, firewall, other gateways will be at the boundary, NAT is easily included with little additional overhead since they are either terminating the connections, or already inspecting and modifying headers and application data. This also relates to myths that NAT is too slow.

DEBUNKING THE MYTHS

Myth 1: Devices behind NAT cannot be managed

It is true that devices behind NAT cannot be managed using traditional methods that depend on public IP addresses. However, there are mechanisms that get around this issue.

There are RFCs and Internet Drafts that describe how to managed devices with SNMP through NAT. It can be complicated in some cases but it can be achieved.

There are also mechanisms that use SNMP Proxy Agents, discussed later in this paper, that also provide the ability to manage devices behind NAT. The SNMP commands are terminated on the public side of NAT and SNMP is not required in the end devices. This can be an advantage for some devices.

Myth 2: Can't support multiple ISPs with NAT

It has been claimed that mechanisms like source routing that are used to support multiple ISPs don't work with NAT. While we have not investigated every possible means of supporting multiple ISPs we have shown that source routing can be used with NAT to support multiple ISPs. A section later in this paper goes into more detail.

Myth 3: NAT breaks end to end security

Security mechanisms may be applied to any layer of the communication architecture. All mechanisms above the network layer, Transport Layer Security (TLS) for example, do not interfere with the NAT address mapping in the network layer. However, TLS does preclude ALGs from modifying the IP address information that may be present in some application protocols. As described later, this is not a problem when using RSAP-IP.

IPsec secures packets at the network layer. IPsec security includes protecting the integrity

of the parts of the IP header not modified by routers and all data contained in the IP packet. Integrity protection guarantees that any modification of the data will be detected and the packet will be discarded by the receiver. As the main function of NAT is to modify the IP (and possibly TCP or UDP) addresses and header, IPsec and NAT directly conflict.

There are two common solutions to the problem. If the NAT gateway is at border between a trusted, private network, such as the home, and the untrusted, public network, then tunnel-mode IPsec alleviates all difficulties (RFC2709). Data is encapsulated and protected between the NAT gateway and the system across the public network. Packets sent from the home are first subjected to NAT, then IPsec protections are applied. Incoming packets for the home are received by the gateway, the IPsec protection is checked and removed, the address is translated, and lastly the packet is forwarded on the home network. If the home network is subjected to threats that demand safeguards, then the gateway can establish an IPsec tunnel to the system in the home so that all packets are protected there as well.

Where protection must extend from the system in the home all the way to the remote system with no intermediaries, then the RSAP-IP is an alternative to gateway tunneling (draft-ietf-nat-rsip-ipsec-04, *RSIP Support for End-to-end IPsec*). As described in the overview, the system in the private network acquires an external IP address (and TCP or UDP port number if needed) from the RSIP server. The system forms packets in the normal way using this external address. Any IPsec protections are applied by the system before transmitting the packet. Since the packet already contains the externally routeable IP address, the gateway no longer modifies the IP or TCP/UDP headers, and IPsec operates end-to-end.

Myth 4: NAT means you can't deploy advanced services

We have heard that service providers resist or reject using NAT in Home Gateways because advanced service cannot be deployed using NAT. Typically these services require multiple ports, some using UDP and typically these services need to support asynchronous traffic into the home.

It is true that basic NAT cannot support these applications but there are two well known approaches to supporting applications like these that have been able to support every application we have encountered.

One approach is the Application Level Gateway for ALG. This is an application aware algorithm that runs within NAT to provide assistance in address translation and port binding. This algorithm knows the content of the messages and provides the translation within the messages as required.

Another approach is to put part of the application into the Home Gateway. A section later in this paper shows an example using SNMP. Another popular example to support IP telephony using protocols like H.323 by putting part of the telephony application in the gateway.

We will grant that these approaches are different and by no means traditional but they preserve the advantages of NAT without breaking the application.

Myth 5: NAT is too slow.

There is no doubt that address translation takes time and reduces throughput, however, if Home Gateways become the residence of a firewall to protect the Home Network from the Internet then we claim the overhead to examine IP packets has already been put in place and the incremental time needed for address translation is negligible.

The assumption of course is that IP packet headers are being processed for another reason

independent of NAT. If this is not the case then NAT will reduce throughput especially when compared to a level 2 switch. NAT in IP routers can also be shown to reduce throughput but to a lesser degree.

SYSTEM LEVEL NAT EXAMPLES

Managing devices behind NAT

Adding capabilities around NAT can create a full set of management capabilities for a range of device types in the private address space behind NAT. One approach is to deploy an SNMP Proxy Agent that as an application has access to both the private and public address spaces on each side of NAT. The figure below shows how this could be designed.

In this approach an SNMP Proxy Agent opens up two network interfaces; one for connecting to the public address space and the other for connecting to the private address space.

The port on the public side is used to send and received SNMP messages to and from and a management system also in the public address space, presumably located at the MSO head end or network control center.

The port on the private side is used to send messages in an arbitrary format to the appropriate devices or objects in the private address space.

In this configuration NAT is not part of the data flow. NAT is in the system to act as a quasi-transparent address translator for end-to-end applications. In the SNMP Proxy Agent case, the TCP connections are terminated and appropriate addresses that do not need translating are used.

The SNMP Proxy Agent can use an approach that provides a separate Object Identifier (OID) for each managed object or device and thereby give the look to the management system that each object has its own SNMP agent. A private MIB is created for each Object class.

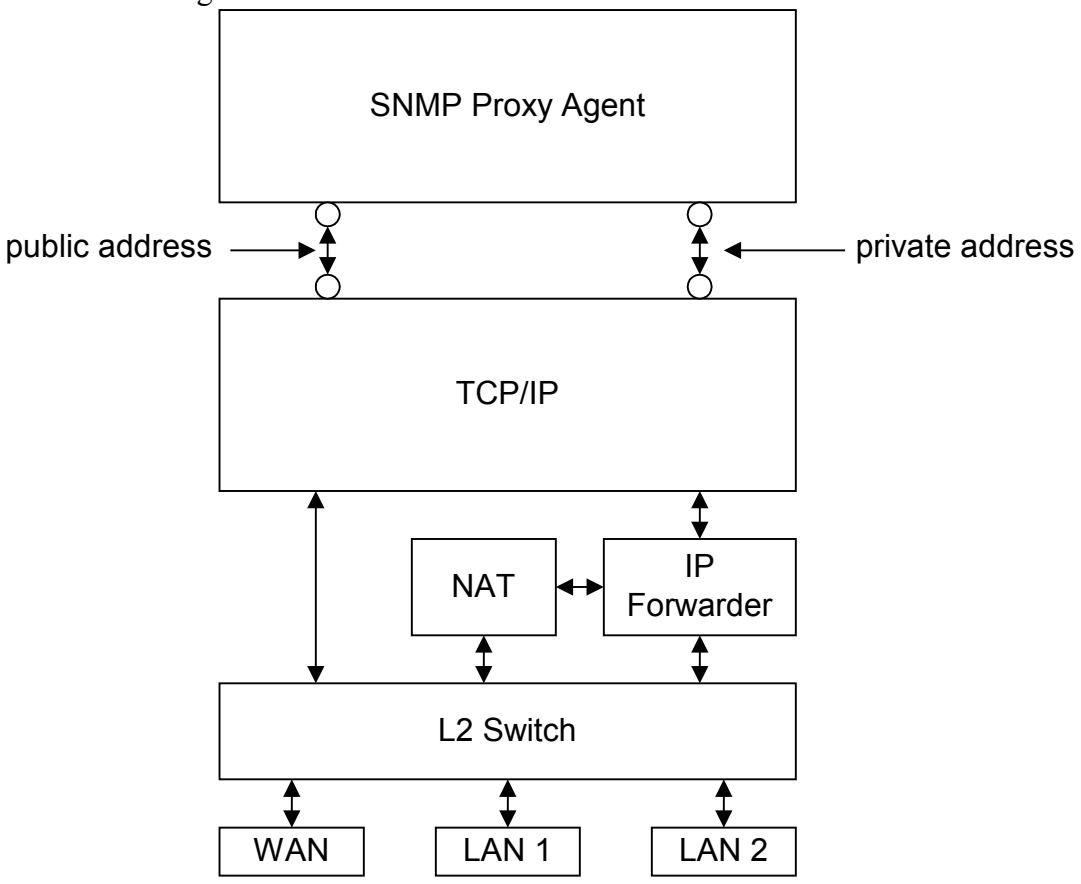

Multiple ISPs with NAT

Overview

In a Home Gateway with NAT, the Home Gateway translates and forwards all IP traffic to and from the CMTS and the customer premise equipment (CPE). The CPE is a home owner PC. The following list describes the provisioning process when multiple ISPs are used:

1) At startup, the PC transmits a DHCP request for a private IP address from the Home Gateway's DHCP server.

2) The PC binds its MAC address with that private IP address assigned to it.

3) Since no ISPs have been provisioned, the only web site the end user can get to is a web site to provision and ISP. The end user provisions an ISP using this web page.

4) After the ISP is provisioned, the Home Gateway is forced to have NAT get a new IP address. This can be done having the provisioning server send a command to the Home Gateway forcing it to release and renew the IP address.

5) When NAT attempts to get a new IP address, the MAC address associated with NAT is assigned an IP address for the provisioned ISP. This is because the DHCP server associated with the provisioned ISP has been configured with the MAC address.

6) The NAT in the Home Gateway now translates all IP traffic from that PC to the selected ISP IP address

7) At this point all the PCs in the house would use this ISP.

8) If the End User would like to provision another ISP the user will open a management window to the Home Gateway and request to provision another ISP

9) The NAT function would use the default IP address for the provisioning server. The PC in use by the end user would at this point only be able to go to the provisioning web page.

10) The end user would provision another ISP and NAT would be forced to renew the IP address for the 2nd ISP.

11) At this point NAT would have two IP addresses. All the PCs except the one use to provision the new ISP would be connected to the old ISP. The PC used to configure the new ISP would be connected to the new ISP. A management window to the Home Gateway would be available for the end user to move PCs between ISPs

12) Additional ISPs can be configured in the same way.

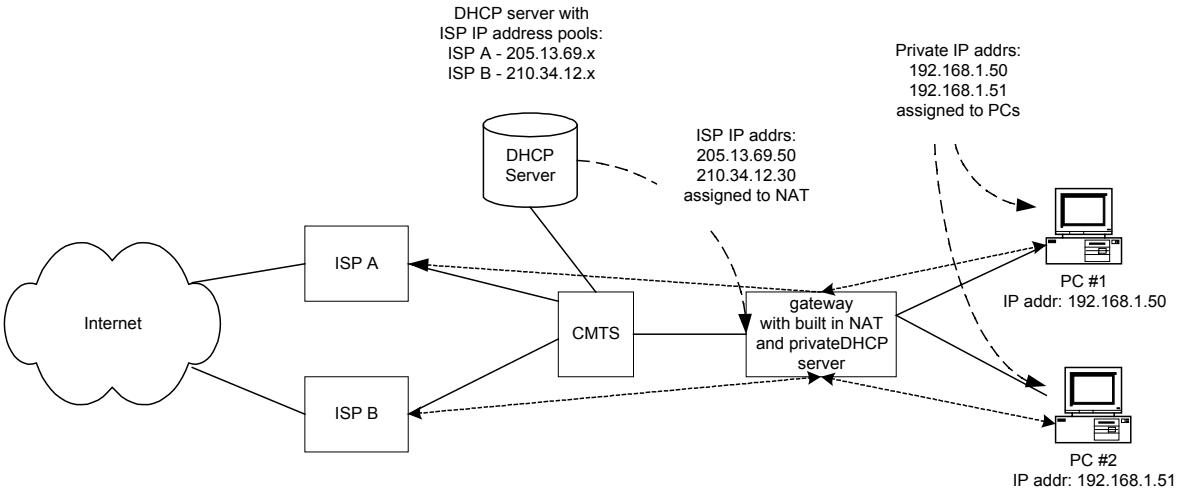

Discussion

IP addressing on the Home Network is handled using a local DHCP server. The address space is private and NAT is used in the Home Gateway. All PCs in the home are on the same subnet and bridging between networks in the home is performed by the Home Gateway.

Before any ISPs are provisioned NAT has one public IP address and it can only be used to access the provisioning server. The end user can use a browser on any PC to access the ISP provisioning server and provision ISP.

Once an ISP is provisioned, the Home Gateway needs to get a new IP address. The preferred method of doing this is to have the provisioning server send an SNMP command to the Home Gateway that would force the Home Gateway to release the current address and then requested a new one. When the DHCP server renews the address, it will provide one for the provisioned ISP. The proper address is obtained because the provisioning process configured the DHCP server with the MAC address from the NAT function.

At this point all the PCs in the home access the Internet through the same ISP.

If another ISP needs to be configured for the home the user opens up a management window to the Home Gateway and requests a new ISP. NAT can use a single MAC address and share it across multiple IP addresses or NAT can allocate another MAC address. The choice here depends on how the DHCP server works and whether it can handle one option or the other. The preferred is to minimize the use of MAC addresses that NAT needs.

Either way, NAT binds the private address of the PC that made the request to a public address. At this point, that PC can only access the provisioning server.

After the new ISP is provisioned NAT gets another public IP address and now all the PCs in the home except the one that provisioned the new ISP are connected to the old ISP. The PC that was used to provision the new ISP is connected to the new ISP.

A management window to the Home Gateway can be used by the end user to configure which PCs in the home are connected to which of the ISPs.

Additional ISPs can be configured in the same way.

RFCs

IP Network Address Translator (NAT) Terminology and Considerations (RFC 2663)

DNS extensions to Network Address Translators (DNS_ALG) (RFC 2694)

Security Model with Tunnel-mode IPsec for NAT Domains (RFC 2709)

An SNMP Application Level Gateway for Payload Address Translation (RFC 2962)

Traditional IP Network Address Translator (Traditional NAT) (RFC 3022)

Protocol Complications with the IP Network Address Translator (NAT) (RFC 3027)

Network Address Translation - Protocol Translation (NAT-PT) (RFC 2766)

ABOUT THE AUTHORS

Jed Johnson

Jed.Johnson@motorola.com

Jed Johnson is the Director of Engineering for Home Networking products in Motorola's Broadband Communications Sector (formerly General Instrument). Jed has been with Motorola for 17 years and has held management and technical positions in various organizations including Home Networking, Streaming Media, Cable Data Products, Transmission Products and Network Management. Previously Jed worked at The Foxboro Company for 7 years as both a Software Engineer and Hardware Engineer. He began his career at Data General.

Jed has a BET degree from Northeastern University and an MBA from the University of Massachusetts at Dartmouth.

Art Harvey

Arthur.Harvey@motorola.com

Arthur Harvey is the Director of the Broadband Networks Research Laboratory within Motorola Labs and has been with Motorola for two years. He worked previously at the Open Software Foundation as Director of the Architecture Group. Prior to that, Arthur was a Network Architect at Digital Equipment Corporation (now Compaq).

Arthur has Ph.D. in Biophysics from the University of Virginia .