NAT Firewall – Intel Wireless Gateway

Introduction:

By definition, a "firewall" provides a filter that incoming or outgoing data must pass through. NAT stands for **N**etwork **A**ddress **T**ranslation. The Intel Wireless Gateway includes firewall protection via its NAT capabilities. This offers protection for computers on your LAN from intrusions via the Internet.

How does the Intel Wireless Gateway do this? Provided in this document is an overview of NAT, as well as a more technical discussion for those inclined to read on.

Overview:

NAT makes the machines on the local network behind the Intel Wireless Gateway machine more secure essentially because the client computers on the local network use IP addresses that are reserved for use on internal networks only. Those IP addresses will not show up on the Internet – only one publicly available address is available to the Internet. NAT does address/port mapping, and keeps state information that prevents incoming connections. This provides the same protection as stateful packet filtering.

There are many variations and implementations of NAT. The Intel Wireless Gateway's goal is to provide for a trusted network with a reasonable amount of security features, while remaining easy to maintain with little or no technical knowledge.

Features of the Intel Wireless Gateway's NAT capabilities

- Automatic operation – no configuration required
- Works with both wired and wireless clients
- Seamless operation with the built-in DHCP server
- Support for VPN and IPsec pass through
- Supports to 64 clients accessing the Internet simultaneously
- Virtual Server capability – allows for the following services on up to 7 internal addresses.
    - HTTP, SMTP, POP3, FTP, Telnet, and IRC

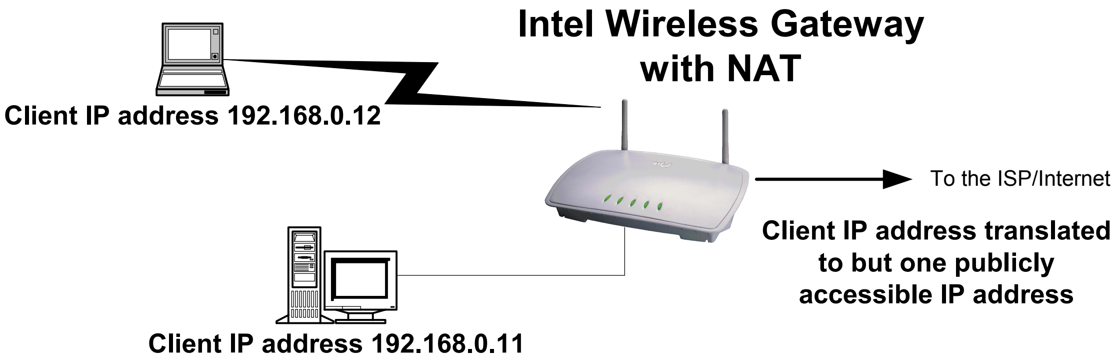
Intel Wireless Gateway
with NAT
Client IP address 192.168.0.12
To the ISP/Internet
Client IP address translated to but one publicly accessible IP address
Client IP address 192.168.0.11

Security Considerations:

Many people view the Internet as a "one-way street"; they forget that while their computer is connected to the Internet, the Internet is also connected to their computer. This is especially true with “always on” Internet connections such as xDSL and Cable Modem connections. The security implications of this are very serious – valuable data may be at risk. To counter the security problem, firewall products are available. Firewall’s are placed between the user’s computer and the Internet and verify all traffic before allowing it to pass through. This means, for example, that no unauthorized user would be allowed to access the company's file or email server. The problem with most firewall solutions is that they are expensive and difficult to set up and maintain, putting them out of reach for home and small business users.

NAT automatically provides firewall-style protection without any special set-up. That is because it only allows connections that are originated on the inside network. This means, for example, that an internal client can connect to an outside web server, but an outside client will not be able to connect to an internal computer because it would have to originate the connection, and the Intel Wireless Gateway’s NAT will not allow that.

It is still possible to make some internal servers available to the outside world via inbound mapping, which maps certain well know TCP ports (e.g.. 21 for FTP) to specific internal addresses, thus making services such as FTP or Web available in a controlled way – the Intel Wireless Gateway provides for this via it’s Virtual Server capabilities.

Technical Description:

The basic purpose of the Intel Wireless Gateway’s NAT function is to multiplex traffic from the internal network and present it to the Internet as if it was coming from a single computer having only one IP address. The TCP/IP protocols include a multiplexing facility so that any computer can maintain multiple simultaneous connections with a remote computer. It is this multiplexing facility that is the key to single address NAT.

To multiplex several connections to a single destination, client computers label all packets with unique "port numbers". Each IP packet starts with a header containing the source and destination addresses and port numbers:

Source address	Source port	Destination address	Destination port

This combination of numbers completely defines a single TCP/IP connection. The addresses specify the two machines at each end, and the two port numbers ensure that each connection between this pair of machines can be uniquely identified. Each separate connection is originated from a unique source port number in the client, and all reply packets from the remote server for this connection contain the same number as their destination port, so that the client can relate them back to its correct connection. In this way, for example, it is possible for a web browser to ask a web server for several images at once and to know how to put all the parts of all the responses back together.

A modern NAT gateway must change the Source address on every outgoing packet to be its single public address. It therefore also renumbers the Source Ports to be unique, so that it can keep track of each client connection. The Intel Wireless Gateway's NAT uses a port-mapping table to remember how it renumbered the ports for each client's outgoing packets. The port-mapping table relates the client's real local IP address and source port plus its translated source port number to a destination address and port. The Intel Wireless Gateway's NAT can therefore reverse the process for returning packets and route them back to the correct clients.

When any remote server responds to an Intel Wireless Gateway client, incoming packets arriving at the Intel Wireless Gateway will all have the same Destination address, but the destination Port number will be the unique Source Port number that was assigned by the Intel Wireless Gateway NAT. The Intel Wireless Gateway looks in its port mapping table to determine which "real" client address and port number a packet is destined for, and replaces these numbers before passing the packet on to the local client. This process is completely dynamic. When a packet is received from an internal client, the Intel Wireless Gateway looks for the matching source address and port in the port-mapping table. If the entry is not found, a new one is created, and a new mapping port allocated to the client:

- Incoming packet received on non-NAT port
- Look for source address, port in the mapping table
- If found, replace source port with previously allocated mapping port
- If not found - allocate a new mapping port
- Replace source address with NAT address, source port with mapping port
- Packets received on the NAT port undergo a reverse translation process:
- Incoming packet received on NAT port
- Look up destination port number in port mapping table
- If found, replace destination address and port with entries from the mapping table
- If not found, the packet is not for us and should be rejected

Many higher-level TCP/IP protocols embed client addressing information in the packets. For example, during an "active" FTP transfer the client informs the server of its IP address & port number, and then waits for the server to open a connection to that address. The Intel Wireless Gateway NAT has to monitor these packets and modify them on the fly to replace the client's IP address (which is on the internal network) with the NAT address. Since this changes the length of the packet, the TCP sequence/acknowledge numbers must be modified as well.

Because the port-mapping table relates complete connection information - source and destination address and port numbers - it is possible for the Intel Wireless Gateway to validate any or all of this information before passing incoming packets back to the client. This checking helps to provide effective firewall protection against Internet-launched attacks on the private LAN. Each IP packet also contains checksums that are calculated by the originator. They are recalculated and compared by the recipient to see if the packet has been corrupted in transit. The checksums depend on the contents of the packet. Since the NAT must modify the packet addresses and port numbers, it must also recalculate and replace the checksums. Before doing so it must check for, and discard, any corrupt packets to avoid converting a bad packet into a good one.