

# Understanding NAT

*As interest in using the Internet grows and grows, Internet address space comes under greater demand, and the threat looms that it could run out. Network Address Translation is one method of conserving Internet address space which you can put to use on your network.*

By Neil Briscoe

**N**etwork Address Translation is commonly known as NAT. In this article I'll explain what NAT is and give a brief technical explanation of how it works.

Why do we need it? Well, it is one attempt at preserving the address space available for use on the Internet.

Internet Service Providers no longer use the A, B, and C classes of address. They allocate addresses in a classless manner. What that means is that, from their blocks of allocated addresses, they will allocate you just as many addresses as you need, and no more. This method is known as CIDR (pronounced Cider) which stands for Classless Internet Domain Routing. CIDR is, however, outside the scope of this article. [See this month's *Professional Update* for an item on CIDR - Ed.]

All you need to know about CIDR for the purposes of this article is that it means that, if you change ISPs, you're going to have to change all of your IP addresses, unless you adhere to RFC 1918 in the first place. Believe you me, if you have even a medium-sized network, changing all the IP addresses on all of your machines requires a considerable investment in time and money, not to mention the fact that you will have to reboot many machines, which can mean things fail that might otherwise have held on a little longer.

On the other hand, if you have adhered to RFC 1918 you will only need to change the address on the external side of your router, and possibly any static HAT mappings you've used - I'll explain that later.

## IP Connections

Before we delve further into NAT we need to talk about how IP connections normally work.

There are 65,535 ports - you can think of them as channels - for IP to choose from. (The number of ports available is determined by the maximum value you can get from the 16 bits allocated to the port number in the IP packet.)

The first 1023 of them are reserved. They're deemed to be privileged ports and are reserved for what are known as the "Well Known Services" (WKS) such as telnet, ftp, gopher, www and so on. So they're normally reserved for server processes for those types of protocol, and can't be used by client processes.

The remaining ports can be used for anything. Some of these are anyway used by various server processes - and some are becoming "Well Known" - but until IANA (see box) grants them a number that use is merely an unofficial convention.

## IANA

The Internet Assigned Numbers Authority ([www.iana.org](http://www.iana.org)) is responsible for allocating port numbers to Well Known Services, allocating blocks of addresses to Internet Service providers (which they do via delegation), and assigning Autonomous System Numbers - amongst other things. Even in an almost unregulated environment like the Internet, responsibility for certain issues has to be centralised otherwise it would all fall into chaos.

## A PC/Server Connection

Leaving NAT aside for the moment, let's consider a connection from your PC to a server. What happens when you type "telnet some.server.com" is that your operating system picks a port above 1023 and assigns it to your session. Let's assume it picks port 1025 to use.

It connects to some.server.com on port 23 - because port 23 is the port reserved by IANA for telnet sessions. But the source information in the packets that reach the server tell it that your machine is using port 1025 to service its replies.

When the packets are returned from the server, therefore, they are aimed at the port number that was indicated in the original packets - in our example, port 1025.

If you've ever wondered how it's possible to talk to two separate servers at once - especially if you use a dial-up connection to the Internet, which makes the whole procedure appear to take place by magic - it is simply because your operating system assigns, say, port 1025 to the first session, and port 1026, say, to the second one. Each server sends its responses back to your machine, but each using a separate port. So, when received, your computer knows which session they were destined for and doesn't mix the sessions up.

## The Need For NAT

The reason why NAT is so important is that address space under IPv4 (in the current version) is limited.

There are four octets, totalling 32 bits of address space. We've tried class-

ing the address space, we've tried using CIDR, but nevertheless address space will eventually run out. Our "stop gap" measures have, however, succeeded in preserving the address space for longer than we all thought possible.

RFC 1918 provides the rules for using a set of reserved numbers guaranteed never to be in use on the Internet. The beauty of it is that every company that connects can pick from the same set of reserved numbers, and it doesn't matter.

However, as these numbers are not routed around the Internet at all we have to have a method of transmitting packets around the Internet, and receiving the replies, and then sending the packets inwards to our network to the correct computer on our LAN, and the correct session on that computer. This is where NAT comes in.

To use NAT, the router which connects your LAN to the Internet will have two addresses. On the LAN side, it will have an address from the particular RFC 1918 address range you chose to use and, on the Internet side, it will have an address assigned to you by your current ISP.

Now let's look at that example again. Your machine sends out a packet aimed at `some.server.com`. The source IP and source port are in the packet, just as before, together with the destination port and IP address.

When it arrives at the router, the router will de-encapsulate the packet, and re-write it. The packet it sends out onto the Internet will contain the router's public IP address, a source port allocated from the router's list of available ports, and the same destination IP address and port number that your machine generated.

The router will also add an entry into a table it keeps, which maps the internal address and source port number your machine generated against the port number it allocated to this session. Therefore, when the machine `some.server.com` sends a reply packet to the router, the router can quickly work out how it needs to re-write the packet before transmitting it back on the LAN.

This works well because most LANs have many computers on them

- certainly more than one. Even if two separate computers happen to pick port 1025 for their respective stacks to use to start a session (quite likely when everyone fires their computers up in the morning) the router just keeps a mapping for each session it sees.

The advantages of NAT are that it works, with RFC 1918, to conserve the address space. The disadvantages are that it slows down the process of transmission, and limits the total number of sessions to the router to slightly less than 65,000 at any one time. That's not really a disadvantage, however, since unless you had an absolutely vast external pipe you wouldn't want to be running that many sessions through your router anyway.

### ***Dynamic And Static NAT***

Finally a quick word on the difference between dynamic NAT and static NAT. What has been defined so far is known as dynamic NAT - all packets leaving your LAN for the Internet contain the same source IP address, which is the public one assigned to your router.

There is one drawback with this. If the router is the only device with a public address, then there is no way for you to provide information services on any computer on your network.

Supposing that, despite being a good network citizen and using RFC 1918 addressing, you nevertheless wish to provide a Web server, or an ftp repository, for example. You can't, because no one on the Internet has any way of specifying that they want to connect to the specific computer containing the server.

However, static NAT allows this to happen. First, you will need to get your ISP to allocate you a block of public addresses. Using CIDR, most ISPs will allocate you a block of eight addresses. Because of certain rules which we won't go into here, you will have five of these eight available for use for static mapping.

The computer on your LAN with the service you wish to make available still gets assigned an RFC 1918 address - otherwise, you wouldn't be able to connect to it from your own computers. Next, you configure the router

with a static NAT mapping rule. You tell it the internal number in use on your LAN for the relevant computer, and you tell it the public address from your ISP-assigned address block that relates to it.

Now that this is properly set up, if someone ftps to the public IP address listed in the static NAT mapping you made, your router will re-write the packets and transmit them inside to the correct machine on your LAN.

To keep things relatively simple, in describing static NAT I have spoken of you providing a computer on your network, or on your LAN, as a server you allow unknown outsiders to connect to via the public Internet. From a security perspective this is a very unwise thing to do. The best advice is to get a router with two Ethernet ports, put all your company private machines on one segment, and put the machines you wish to be public on the other segment. Configure the router with suitable ACLs.

**PCNA**

### **The Author**

Neil Briscoe is a network consultant and can be contacted as [neil.briscoe@itp-journals.com](mailto:neil.briscoe@itp-journals.com).

## New Reviews from [Tech Support Alert](#)

### [Anti-Trojan Software Reviews](#)

A detailed review of six of the best anti trojan software programs. Two products were impressive with a clear gap between these and other contenders in their ability to detect and remove dangerous modern trojans.

### [Inkjet Printer Cartridge Suppliers](#)

Everyone gets inundated by hundreds of ads for inkjet printer cartridges, all claiming to be the cheapest or best. But which vendor do you believe? Our editors decided to put them to the test by anonymously buying printer cartridges and testing them in our office inkjet printers. Many suppliers disappointed but we came up with several web sites that offer good quality cheap inkjet cartridges with impressive customer service.

### [Windows Backup Software](#)

In this review we looked at 18 different backup software products for home or SOHO use. In the end we could only recommend six though only two were good enough to get our "Editor's Choice" award

### [The 46 Best Freeware Programs](#)

There are many free utilities that perform as well or better than expensive commercial products. Our Editor Ian Richards picks out his selection of the very best freeware programs and he comes up with some real gems.