NAT & IP Masquerade

INTRODUCTION

Pre-requisites

TCP/IP

IP Address Space

Internet Protocol version 4 uses a 32 bit IP address. In theory, a 32 bit address space should provide addresses for more than four billion computers, but inefficiencies in address allocation mean that less than half of the addresses are used. The result is a so-called "internet address crisis" - there are more computers than usable addresses. Therefore permanent IP addressses have become expensive – IP addresses are usually assigned dynamically and are often shared among many computers.

Private Networks

Many organizations run private networks using the internet protocol. The IANA has assigned three IP networks for private networking: 10.0.0.0/8, 172.16.0.0/16, and 192.168.0.0/16. These IP networks may be used by anyone provided that the network is not connected to the globally connected Internet. If an internet router receives a packet from or to one of these three networks, the router should discard the packet.

Firewall

Most private networks are connected to the Internet through a firewall. Internal computer are computers on the private network. Computers on the Internet are defined as external computers. External and internal computers communicate with each other through the firewall. But the firewall may not transmit the private IP addresses on the Internet. The firewall must substitute a valid Internet IP address. There are two techniques to accomplish this substitution: Network Address Translation and IP Masquerade.

NETWORK ADDRESS TRANSLATION

Network Address Translation (NAT) is a method of sharing Internet IP addresses among many computers on a private network. NAT works completely on the network level, meaning that NAT operates by changing IP packets. NAT does not modify higher level protocol information, such as TCP, UDP, or ICMP.

How NAT works

Suppose a computer on a private network has IP address 192.168.0.25. The computer sends a packet to a computer, IP address 213.65.37.242, on the Internet (figure 1a). The firewall's internal address (192.168.0.51) is defined as the default gateway, so the packet arrives at the firewall.

The firewall cannot forward the packet to the Internet with a source address of 192.168.0.25. Instead, the firewall has been assigned a pool of valid Internet IP addresses (65.93.135.0/24). The firewall's NAT module replaces the original source address with one of the IP addresses from the pool, records the original and replacement IP addresses in a translation table, then forwards the modified packet to the Internet (see figure 1b).

When the firewall receives packets from the internet (figure 1c) its NAT module compares the arriving packets to the translation table and replaces the destination IP addresses with the appropriate private IP address. The firewall then forwards the packet to the private network (figure 1d).

Problems with NAT

NAT is a one-to-one system, meaning that each private IP address must be mapped to exactly one Internet IP address. Once the Internet IP address is used, it is removed from the pool until the communication is complete (the sysadmin will set a timeout value – if the translation has been idle for that period of time, then the address will be returned to the pool). IP addresses are expensive, so there will typically be many more computers in the private network than there are addresses in the pool. When all the addresses are used up, no more computers may connect to the Internet.

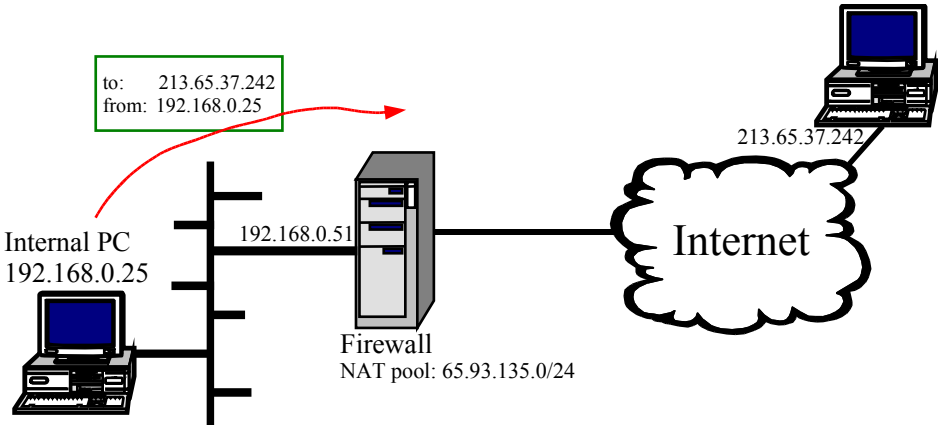

Figure 1a: NAT

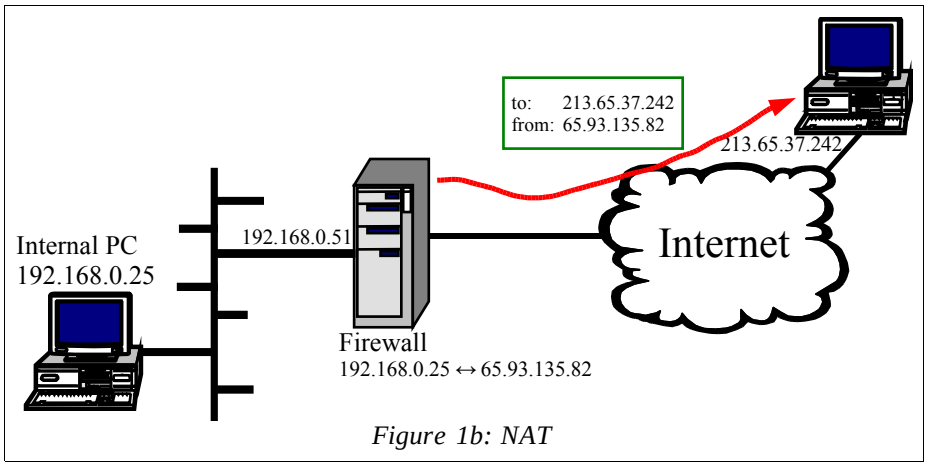

Figure 1b: NAT

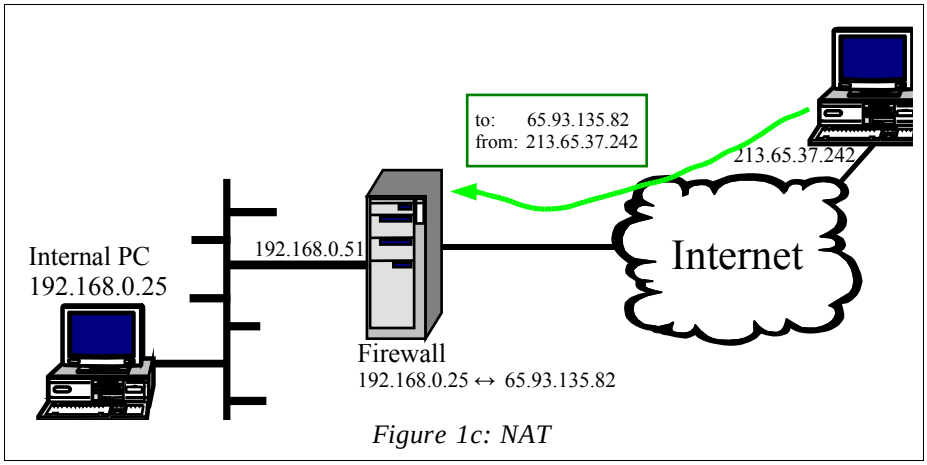

Figure 1c: NAT

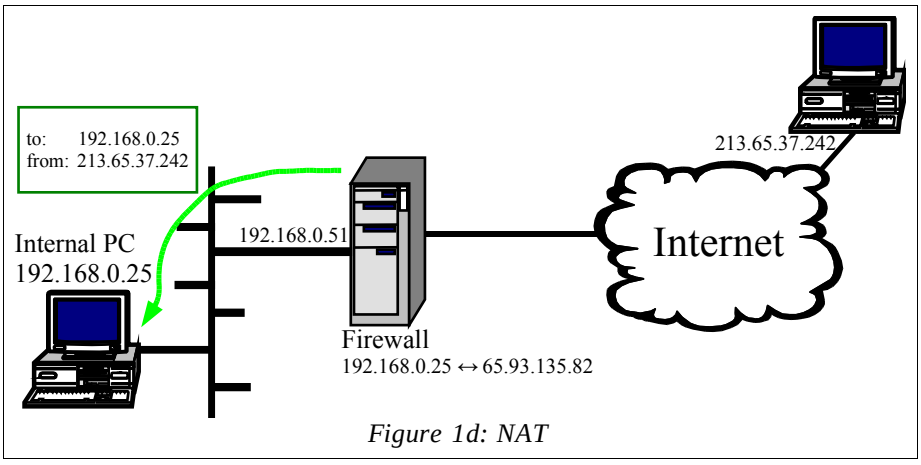

Figure 1d: NAT

IP MASQUERADE

IP Masquerade is a many-to-one translation technique: it allows many private IP addresses to share one Internet IP address simultaneously. IP Masquerade operates at the network and transport level: it changes both IP and TCP (and UDP) header information. on both the IP level and the TCP (or UDP) levels.

How IP Masquerade Works

In an IP Masquerade system, many computers on a private network access the internet through an IP masquerading firewall. When the internal computer sends a packet to the internet, the firewall receives the packet on it's internal interface.

A TCP connection is defined by two sockets – one socket at the source computer and another socket at the destination computer. A socket consists of the computers IP address and the port number (for instance, port 80 for

http). When the internal PC sends a segment to the computer in the internet, it specifies the destination socket 213.65.37.242:80 and its own socket, 192.168.0.25:80.

The firewall cannot forward this TCP segment to the Internet because the IP address 192.168.0.25 is a private network address. The firewall must replace the private IP address with its own Internet IP address, 65.93.135.82, similar to the NAT procedure. But IP Masquerade takes an extra step: it changes the port number. The IP Masquerade module replaces the original port number with a higher number, such as 10000. The IP masquerade module creates a translation table entry listing the original socket number and the modified socket number, then forwards the modified TCP segment to the Internet (figure 2b). Later, if the firewall receives TCP segments from the Internet with destination socket 65.93.135.82:10000, it will translate the socket number back to the original value and forward the segment on the private network.

Now, suppose another computer on the private network sends a TCP segment to the Internet (figure 2c). The segment arrives at the firewall. The IP masquerade module modifies the socket number: it replaces the private IP

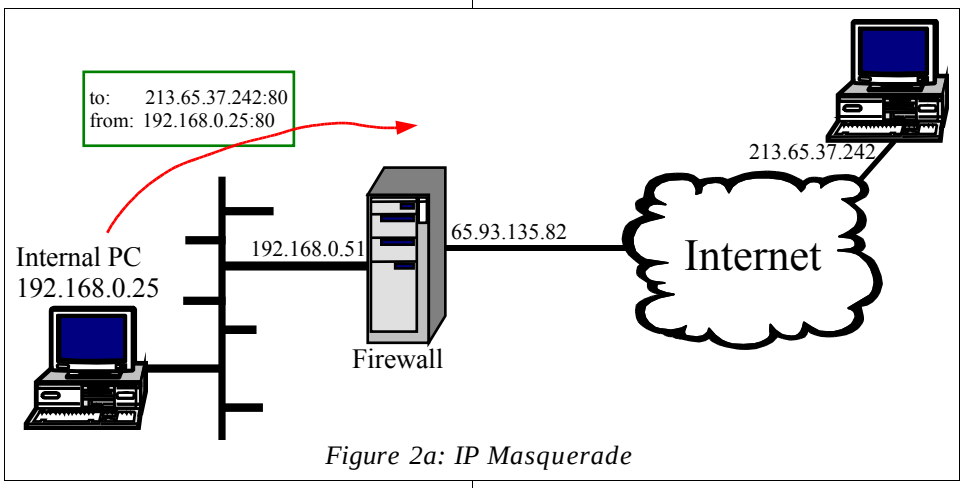

Figure 2a: IP Masquerade

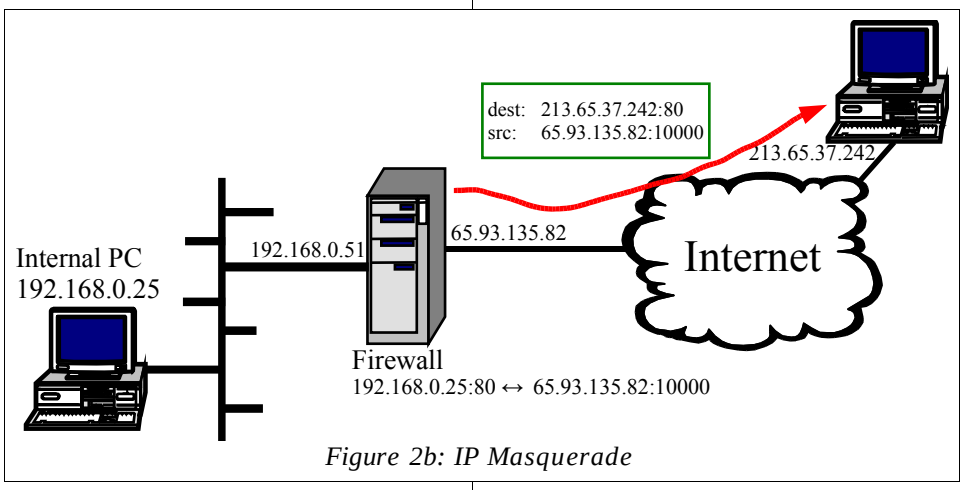

Figure 2b: IP Masquerade

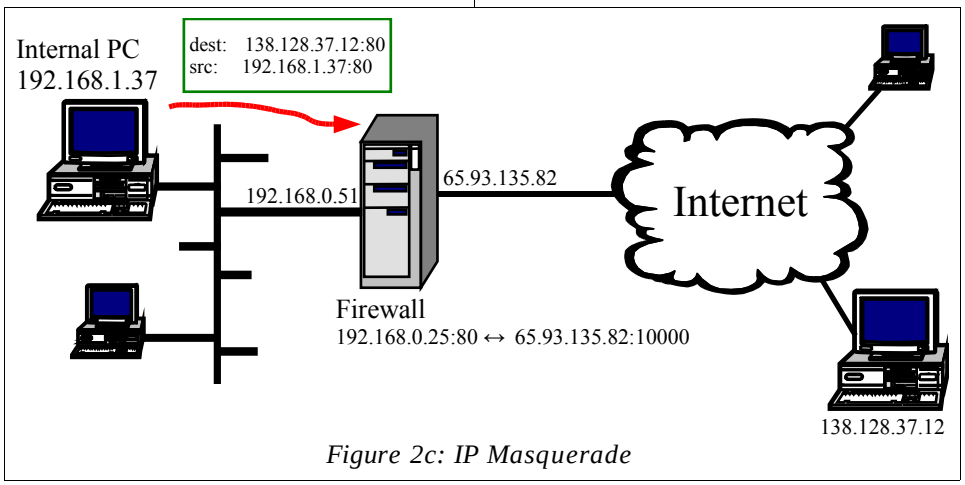

Figure 2c: IP Masquerade

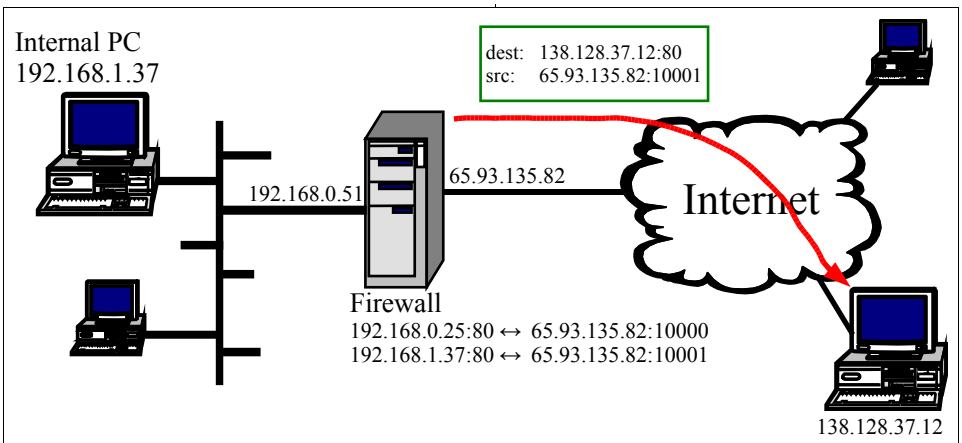

Figure 2d: IP Masquerade

address with the firewall's Internet address, and replaces the original port number with a unique port number from the translation pool. Then the IP masquerade module adds the new entry to the translation table and forwards the modified segment to the Internet (figure 2d). Later, if the firewall receives a segment from the Internet with destination socket 65.93.135.82:10001, then the IP masquerade module will replace the destination socket number with the original socket number from the translation table and transmit the corrected segment on the private network. There is a timeout value associated with the entries in the translation table: if a socket listed in the translation table goes unused for a time period longer than the time out, then the entry is removed from the translation table.

Problems with IP Masquerade

IP Masquerade uses a unique port number for each connection from the private network. There are 65,536 possible port numbers, so it is unlikely that IP masquerade will run out numbers, but under certain circumstances it can happen. For instance, Game Spy software opens thousands of short lived connections. If multiple hosts on the private network run Game Spy, then the firewall can run out of port numbers.

IP masquerade works well when computers on the private network originate all connections. IP masquerade doesn't work when computers on the Internet try to open connections to computers on the private network. For instance, some Internet game servers use multiple connections. The computer on the private network opens the first connection to the game server, and that works fine with IP masquerade. But then the game server tries to open a second connection on a different port to the client on the private network. This is a problem, because when the firewall sees the new incoming connection request the new socket won't be in the translation table, so the IP masquerade module won't know where to send the segment.

Port Forwarding

Suppose you want to run a web server on one of the computers in your private network. This would be a problem, because computers from the Internet cannot send IP datagrams to a private IP address. IP masquerade overcomes this problem using a technique called port forwarding. Using port fowarding, you may program the IP masquerade module to forward any segments destined for a specific port on the firewall to a port on a computer in the private network.

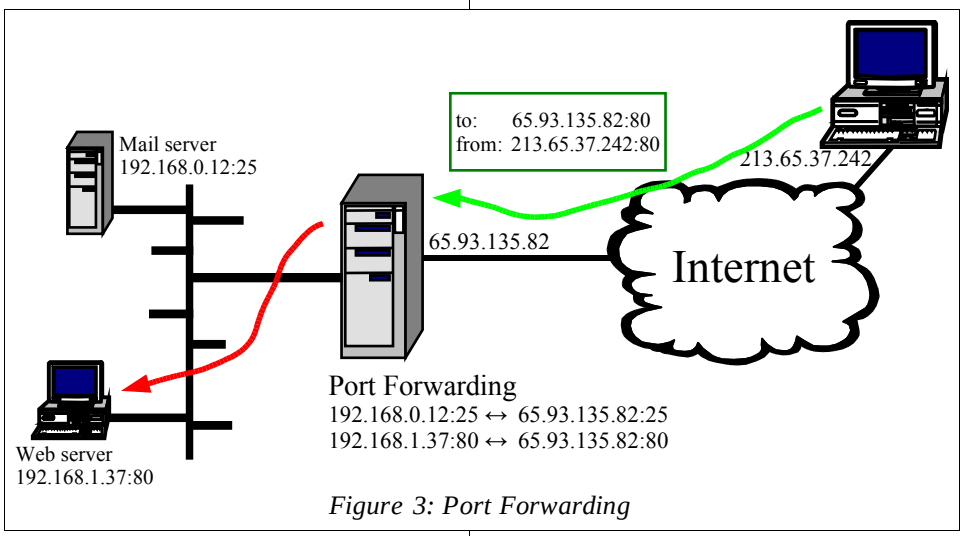

Figure 3: Port Forwarding

For instance, consider the situation depicted in figure 3. The sysadmin wants to operate a mail server and a web server on the private network. To do this, the sysadmin creates two translations in the firewall's IP masquerading table. The first entry translates any segments arriving on the firewall's mail server socket to the private network's mail server socket. The second entry does the same thing for the web server. Unlike normal NAT and IP masquerade, port forwarding translations have no timeout period. The translations stay in the translation table until they are manually removed or the entire table is reset.

DETECTING NAT AND IP MASQUERADE

Some ISPs specify in their user agreements that you can only connect one computer. So you might ask, “Can my ISP tell that I am running NAT or IP Masquerade?” The short answer is Yes.

When a computer originates an IP packet it assigns a value to the Time To Live (TTL) field. Usually this value is 128, sometimes it is 255. If your computer is directly connected to the ISP, then the ISP's edge router will always see values of 128 (or 255) in the TTL field of your packets. But each IP forwarding device reduces the value of the TTL field by 1. So if there is a NAT (or IP masquerade) box between your computer and the ISP edge router, then the ISP will see values of 127 (or 254) in your packets' TTL fields. Therefore the ISP can detect that you are using a NAT box. Obviously, the NAT box can be fixed to modify the TTL field so that the TTL value is always 128 – let's call it “stealth NAT”. But as of this date (April 2003) NAT boxes do not implement stealth NAT.

It is even possible to estimate the number of computers sharing a single Internet IP address through an IP masquerade box. When a computer originates an IP packet, the originating computer creates a 16 bit ID number which goes in the ID field of the IP packet header. Typically ID numbers from a single computer are assigned in sequence, so one packet might have an ID number of 43,321, the next packet might have an ID number of 43,322, and so on. When these packets go through the IP masquerade module, the ID numbers are the same, so the ISP's edge router will see ID number 43,321, ID number 43,322, and so on. But if a second computer is also using IP masquerade, it will probably be using a different series of ID numbers, suppose it is generating packets with ID numbers 23,327, 23,328, etc. And a third computer might be sending packets with IDs of 7,000, 7,001, 7,002, etc. These packets arrive at the edge router in any particular order, so the edge router might see packets arriving with the following ID numbers: 7000, 23327, 23328, 7001, 43321, 7002, 43322, 23329, 7003, 43323. The edge router can easily determine that there are three different sequences of ID numbers, and therefore that there are three different hosts behind the IP masquerade box.

NAT/IP MASQUERADE CONFUSION

Just to make matters interesting, many authors confuse the terms NAT and IP masquerade. Some authors refer toboth NAT and IP Masquerade as NAT. You might also see terms like “one-to-many NAT” or “multipoint NAT” used instead of IP masquerade. Like so much of the computer world, jargon is poorly defined and the few good definitions are rarely enforced. You can help by using NAT and IP Masquerade properly.