
Section 4 -- NetNAT Operations

Starting and Stopping NetNAT

Starting the NetNAT

Most NetNAT installations start the NetNAT software from their autoexec.bat file in one way or another. The command line options for the NetNAT are described here:

Command Syntax

```
nat -d . -a [sernbr],[authkey] {startup file}
```

Where:

nat	Nat.exe is the traditional name of the NetNAT executable module.
-d .	Tells the NetNAT code that it should consider the current directory to be its “runtime” directory, where its files are to be found.
-a [sernbr],[authkey]	Gives the NetNAT its serial number and the authentication key that confirms its accuracy. Absence of this field, or an incorrect value places it in a very limited 2-user mode.
startup file	This is the name of an alternate startup file for the NetNAT other than the default “startup.htm” file.

Additional Information

The 32-bit NetNAT needs DPMI (DOS Protected Mode Interface) services to be able to use all of the available memory in the machine. A driver that provides these services is included in the NetNAT archive, and is named “cwsdpmi.exe.” It must be run before the NetNAT program is started.

Stopping the NetNAT

The “exit” NetNAT console command causes the program to free all memory, close all files and exit the NetNAT program. This returns you to a DOS prompt.

Limitations of Shareware Mode

Shareware mode permits 2 simultaneous users, and prohibits all forms of Service Mapping and VPN tunneling. With a shareware mode NetNAT you have no access to NSC support.

Changing the Setup Parameters

The default system configuration is in the runtime directory in a file called “startup.htm.” The actual file may be specified on the NetNAT startup command line as the last parameter. At present, the file is a normal text file with an HTML wrapper, to make it viewable from a web browser. The normal procedure for startup file update is to fetch it to the administrator workstation using ftp, edit it there, and return it via ftp.

The NetNAT can use an encrypted configuration file. The encryption of this file prevents unauthorized viewing of the setup. The license serial number may specify that the configuration file **must** be encrypted. This prevents unauthorized modification of the file, except by the administrator or his or her workstation only. When using the encrypted configuration file, remember to use **binary mode** for ftp transfer of the file to and from the NetNAT.

Changing the Administrative Password

The administrative or “root” password is kept in encrypted form in the file c:/safe/passwd32. Please see section 5 for discussion of this file and directory.

Using the Web Server

Use the “start www” command to enable the built-in web server. This provides the ability to view the startup configuration file, as well as online documentation in web pages. Use the console command web form (command.htm) to issue simple commands to the NetNAT.

Viewing the Online Documentation

The web page nncfg.htm is the top of the online documentation tree.

Console Operations

The NetNAT local console is a powerful tool for the administrator. In addition to normal functions, the console task handles multiple console windows and window scrolling. The function keys that control these functions are described here.

Multiple Console Windows

A number of commands must be run as a process separate from the main task. When one of these commands starts, a new console window is created for it. The console display switches to that new window until the command completes. The displayed window may be manually switched by keyboard function keys. The left-most field of the console status line shows the window number being displayed.

F8 – Rotate Forward Through Console Windows

The F8 function key switches from console window to console window in the order that they were created. The window being displayed is named on the status line at the bottom of the screen. For example, the main console window will say “0: command interpreter” at the left end of the status line. If you run the command “repeat 1000 pr” from the console, another window (window 1) will be created. Cycling between these with the F8 key, the repeat window will say “1: repeat 1000 pr” at the left end of the status line. As you close a window (see the command descriptions in section 6 for the appropriate method of exiting each command), that window will drop from the rotation.

F7 – Rotate Backward Through Console Windows

The F7 function key switches from console window to console window in the opposite order that they were created. This is the reverse of the operation of the F8 key.

F10 – Shortcut to Console Command Window

The F10 function key switches directly to the console command window from any other window.

F6 – Enable/Disable Window Scrolling

The console command window maintains the most recent inputs and outputs in a scroll buffer. The size of the scroll buffer defaults to 1000 lines, a value that may be displayed or modified with the `scrollback` command (see section 6). When in scroll mode, the status line (at the bottom of the console display) will say “Scroll:nn” where nn is the offset backwards from the most-recent display. Use the up and down arrows, Home and End, and the PgUp and PgDown keys to navigate through the scroll buffer. Press F6 to return to normal window operation.

Alt-Fx – Direct Switch to a Specific Window

Alt-Fx (where x is 1 through 8) allows direct access to a specific console window. F10 always provides direct access to the console command window.

Command Abbreviation

The commands may be abbreviated to the shortest representation that is unique. See section 6.

Configuring a New Interface

The “attach” command creates the relationship between a NetNAT interface object and the associated network card driver. Once attached, the interface is configured using the “ifconfig” command. See section 6 for details on the various commands. Here’s an example of creating a public interface named “en0.” This sequence of commands is usually found in the startup.htm file, but may be run by hand when testing a new interface.


```
#
# Packet driver installed on software int 0x60
#
attach packet 0x60 en0 64 1500
#
# configure the IP Address for the new interface
ifconfig en0 ipaddress 192.168.31.254
#
# set the netmask to 255.255.255.0
ifconfig en0 netmask 255.255.255.0
#
# set the interface broadcast address to the subnet broadcast
ifconfig en0 broadcast 192.168.31.255
#
# private 0 says this is a public interface
ifconfig en0 private 0
```

Viewing the Attached Interfaces

The “ifconfig” console command is used to view the attached interfaces as well as configure them. “Ifconfig” issued with no parameters will display all attached interfaces and the current statistics for them. Before any network interfaces have been attached, there is a single interface created automatically by the NetNAT: the “loopback” interface. This is a required interface on any network device, and is used internally. The ethernet interfaces need to be attached before they will display with ifconfig. See section 6 for examples and interpretation of the ifconfig interface display.

Setting a Route to another IP Network

In most simple installations, the only route the NetNAT needs is the default route to the Internet (via the Internet access router). The NetNAT is capable of handling routes to a large number of additional networks when necessary. An example is shown here. This is a very common case, where there is a substantial corporate Intranet connected to the local network by way of a router.



The NetNAT will have a default route to the Internet via the Internet access router, and one or more routes to subnets in the corporate Intranet by way of the corporate WAN router. The local clients and hosts will have the NetNAT as their default route, and the same set of Intranet routes through the WAN router. Where it's inconvenient to give each client the full set of routes, giving them the NetNAT as a default may be sufficient, since the NetNAT will redirect them to the corporate WAN router where appropriate. This doesn't work where the client IP stack ignores ICMP "host redirect" messages. Here's a route configuration example for this network, assuming that the Intranet uses network 10, the local LAN is 10.1.1.x/24, the NetNAT is 10.1.1.254, the WAN router is 10.1.1.253, and the Internet access router is 204.86.17.12:

```
nscnat> route add default en0 204.86.17.12
nscnat> route add 10/8 en1 10.1.1.253
nscnat> _
```

This sets the NetNAT's default route to the access router's address, and the route to all of network 10 (that isn't local) to the WAN router's address. The NetNAT can handle as many of these "static" routing entries as you need to use.

Viewing the Routing Table

The “route” command with no parameter will display the routing table. There will be more routes than expected, since a number of them are created automatically by the ifconfig commands. In particular, there will be routes to the local networks (those that the NetNAT is connected to) and to the various broadcast addresses.

This is the routing table display for the network pictured in the previous example:

```
nscnat> route
Dest          Len Interface Gateway      Metric P Timer Use
10.1.1.255    32  en1              1      P 0    0
204.86.17.255 32  en0              1      P 0    0
10.1.1.0      24  en1              0      0 159622
204.86.17.0   24  en0              0      0 209331
10.0.0.0      8   en1      10.1.1.253 1      0 87311
default       0   en0      204.86.17.12 1      0
178334
nscnat> _
```

Note the two routes added by explicit “route add” commands and the four others that were created by “ifconfig.” Refer to “route” in section 6 for interpretation of the display information.

Controlling Which Users Get Internet Access

The NetNAT's access control features give you complete control over which internal users may have **direct** access to the Internet or to any network on the public side of the NetNAT. By default, the NetNAT permits all internal users unlimited **direct** access to any network on the public side of the NetNAT. You may add rules that restrict this access by using the "access" command (see section 6). "Direct access" means that the internal user may communicate directly with an external computer without restriction. Denying this **direct** access does not necessarily deny that user access to the Internet. For example, the user computer may use a web or ftp proxy server to access the Internet, and so is not accessing the Internet "directly." In this example, the proxy server has direct access, and is acting on behalf of the internal user.

An Example

Suppose you are a school district that uses the 10-net for all internal devices. You have a filtering web-and-ftp proxy that you want all student and teacher workstations to use, and you wish to give the district office unlimited access to the Internet. You would first deny all direct access from the entire 10-net (see "access deny" in section 6), grant unlimited access to the proxy server host (see "access permit" in section 6), and grant access to the district office subnet (see "access permit" again). Students and teachers would be able to access the Internet only by using the proxy server. If anyone modified the setup of their Web browser to attempt to circumvent the proxy server, the access attempts would be rejected by the NetNAT.

Enabling Internal Services for External Use

One of the strengths of the NetNAT is the flexibility in configuring internal services for external use. Careful thought should be given to what services should be allowed, and to the security impact they have.

Setting Fixed Service Mapping

The “fixed service mapping” is very flexible, but should be used only when absolutely necessary, since it maps all services for a host into view from the public side. Be extra careful in preparing your host to make sure it is secure.

Setting Local Service Mapping

Local service mapping allows a NetNAT server to be used from the public side. For example, you may enable the NetNAT’s ftp server for use from the outside with the following command:

```
service en0 local * 21 tcp
```

This instructs the NetNAT to accept FTP control connections on the natural IP Address of interface en0 (specified by the asterisk). The “natural IP Address” is the one configured with the ifconfig ipaddress command.

In the same way (while specifying the public IP Address), the local WWW server may be enabled:

```
service en0 local 204.87.17.12 80 tcp
```

Setting Port Service Mapping

Port service mapping is the one most used by NetNAT administrators. It allows the mapping of specific services into the public address space. This example illustrates an actual network that we installed using the NetNAT Firewall. The network consists of:

- ◆ A unix host with an internal Web server bound to 10.1.1.1 and an external Web server bound to 10.1.1.2.
- ◆ The same unix host with an internal DNS server on 10.1.1.1 port 53 and an external DNS process on 10.1.1.1 port 54.
- ◆ SMTP and POP3 email services, FTP and Telnet on the unix box at 10.1.1.1.
- ◆ PPTP services on an internal Windows NT Server at 10.1.1.10.
- ◆ Public IP Address of 204.86.17.12 (mythical).

The commands to configure the port service mapping described above are:

```
# Public Web Service
service port en0 204.86.17.12 80 tcp 10.1.1.2 80
# Inbound Email and POP3 access from home
service port en0 204.86.17.12 25 tcp 10.1.1.1 25
service port en0 204.86.17.12 110 tcp 10.1.1.1 110
# External DNS (for our public image)
service port en0 204.86.17.12 53 udp 10.1.1.1 54
service port en0 204.86.17.12 53 tcp 10.1.1.1 54
# FTP and Telnet from the outside
service port en0 204.86.17.12 21 tcp 10.1.1.1 21
service port en0 204.86.17.12 23 tcp 10.1.1.1 23
# PPTP connections from work-at-home users
service port en0 204.86.17.12 1723 tcp 10.1.1.10 1723 p
```

If these are the only service mapping commands in your configuration, then no other services will be visible to the outside.

Setting Pool Service Mapping

Pool service mapping is used by many medical sites, as their telnet-style servers will not accept more than one connection from a single IP Address. A second connection causes any earlier one to be terminated. In those installations, the NetNAT is the Firewall between the hospital and the service bureau, and a pool of “public” IP Addresses is allocated for the clients to use. The pool needs to be large enough to handle the simultaneous users, but might not need to be as large as the client population. The NetNAT can handle up to four pools of 512 addresses each, for a total of 2048 possible concurrent users.

To create a public pool of 128 addresses in the range of 204.86.17.1 through 204.86.17.129 with an address lease time of one hour, and another at 204.86.18.1, use the following pool service mapping commands:

```
service en0 pool 204.86.17.1 128 360  
service en0 pool 204.86.18.1 128 360
```

Refer to the service ... pool command in section 6 for an explanation of the parameters.

Setting Subnet Service Mapping

The subnet service mapping facility has the most options and modes. Refer to the individual command definitions in section 6 for details.

Subnet Mapping Mode 0

Mode 0 allows the administrator to define an internal subnet and assign a single public IP Address to be shared by all members of that subnet. The NetNAT's port address translation allows these computers to share the public address simultaneously. No inbound service requests will be passed into the subnet (unless explicitly enabled by other service mapping commands), and the NetNAT responds to inbound pings on behalf of the subnet.

Subnet Mapping Mode 1

Mode 1 creates a correspondence between an internal subnet and a same-sized public subnet, and translates IP Addresses in messages passing between the internal and public sides. Since there is a one-to-one relationship between computers on the inside and public IP Addresses, there is no port translation necessary. No inbound service requests will be passed into the subnet (unless explicitly enabled by other service mapping commands), and the NetNAT responds to inbound pings on behalf of the subnet.

Subnet Mapping Mode 2

Like mode 1, mode 2 creates a correspondence between an internal subnet and a same-sized public subnet, and translates IP Addresses in messages passing between the internal and public sides. Since there is a one-to-one relationship between computers on the inside and public IP Addresses, there is no port translation necessary. No inbound service requests will be passed into the subnet (unless explicitly enabled by other service mapping commands), except that the NetNAT passes ping requests through to the private side to permit the internal computers to reply for themselves.

Subnet Mapping Mode 3

Like mode 2, mode 3 creates a correspondence between an internal subnet and a same-sized public subnet, translates IP Addresses in messages passing between the internal and public sides, and passes pings into the private subnet. In addition, the NetNAT passes and translates SNMP requests, replies and traps, to permit management of the private network by an external entity. Other than ping and SNMP requests, no inbound service requests will be passed into the subnet (unless explicitly enabled by other service mapping commands).

Virtual Private Networks

Because the Internet is a very inexpensive way to communicate between sites, it is being used in place of dedicated circuits between company locations. These “Virtual Private Networks” use various technologies to create “tunnels” between sites through the Internet. It is fortunate that the speed and dependability of the Internet are improving almost daily. Henceforth, the acronym “VPN” will be used to denote Virtual Private Network.

Security

Security must be considered when creating a VPN. The network at the other end will be treated as a trusted network: their IP Addresses will be visible to you, and your IP Addresses will be visible to them. Of course, this is exactly what we intend when we create a VPN, so we shouldn't be surprised.

PPTP

PPTP is the “Point-to-Point Tunneling Protocol” developed by an Internet Engineering Task Force composed of industry representatives. Initial implementations of PPTP were on Microsoft Windows 95 and NT. Using the Microsoft implementation, it is possible to create an authenticated and encrypted tunnel between a remote Windows 95 workstation and a Windows NT server. Sensitive information may be transferred through this tunnel without fear of eavesdropping. Our support for PPTP will be described on a subsequent page.

Tunneling via IP Encapsulation

One of the VPN technologies offered by the NetNAT uses IP Encapsulation within IP as defined in Internet standard RFC 2003. In this mechanism, an IP datagram that is to be sent to a remote site is encapsulated entirely within another IP datagram. This message-within-a-message is sent through the Internet to a compatible gateway device at the remote site that strips the outer wrapper away and delivers the message contained within. IP Encapsulation is enabled by creating a tunnel route to the remote with IP-in-IP mode specified. Please see the “route tunnel ... i” command in the commands section.

Tunneling via Minimal IP Encapsulation

Another VPN technology provided by the NetNAT is called “Minimal Encapsulation within IP” and is defined in Internet standard RFC 2004. This differs from IP Encapsulation in that less data is added to the message to be forwarded, resulting in a slightly smaller message. As in IP Encapsulation, a compatible gateway device at the remote site is required to restore the original packet for delivery. Minimal Encapsulation is enabled by creating a route to the remote site with the Minimal Encapsulation mode specified. Please see the “route tunnel ... m” command in the commands section.

Tunneling via IPSEC

A growing standard for VPN technology provided by the NetNAT is called “IP Security” and is defined in Internet standards RFC 1825, 1826, 1827 and 1828. This differs from the other encapsulation modes in a number of ways, including authentication of the sender (making sure the incoming packets are from who you think they are), and encryption of the data (to keep prying eyes from reading your messages). As in the other forms of VPN, a compatible gateway device at the remote site is required to restore the original packet for delivery. IPSEC is enabled by creating a route to the remote site with the ESP Encapsulation mode specified. Please see the “route tunnel ... e” command in the commands section. This causes all data messages to be encrypted with the DES algorithm for privacy. In addition, you may specify that the Authentication Header be added to each packet to be more certain of the originator. It is generally believed that the ESP encryption is adequate, since decrypted packets are validated by the NetNAT before being forwarded, and valid decrypted packets could only be created by someone with the encryption key.

National Security Issues

Because the DES encryption is very difficult to decrypt without the key, its export from the United States is controlled by the Federal Government to prevent foreign nationals from using this technology. To comply with the intent of this regulation, the IPSEC feature of the NetNAT is disabled in the standard package. To enable it, the customer must submit documents to Network Safety certifying that the IPSEC features of the NetNAT will not be used outside of the United States. Please contact us to get the required forms.

Using PPTP Through a NetNAT

PPTP is an emerging standard and mechanism to secure tunneling of data through a public network. Typical uses are for gaining secure access to an office LAN from home or on the road. The process includes strong user authentication and encryption of the session data, so it is a very secure mechanism.

The PPTP Mechanism

A PPTP session begins with a TCP connection from the client to the desired server. This connection uses TCP port 1723, and will be the facility through which the authentication is performed. The connection stays in place for the duration of the PPTP session. Once the TCP control connection is running and the user has been authenticated, the actual data passes over a different “connection” using the GRE protocol.

Client Behind NetNAT to Server on the Outside

The NetNAT detects an outbound connection to the standard PPTP control port (1723) and creates the PPTP data session (for the GRE protocol). The PPTP data session stays usable as long as the control session remains established. Due to the nature of the PPTP data session object, there may be no more than one PPTP connection to the same external server from a single NetNAT public IP Address. If more than one connection to an external server is required, the NetNAT configuration should be manipulated to provide for more public IP Addresses. See the “subnet service mapping” command in section 6.

Client on the Outside to Server on the Inside

In order for this to be permitted, a “port service mapping” for the PPTP control session must be added to the NetNAT configuration, including the optional “p” parameter (that specifies that this session is a PPTP control session). Then, when the NetNAT detects an inbound connection to that service, it creates a PPTP data session (for the GRE protocol). The PPTP data session stays usable as long as the control session remains established. Due to the nature of the PPTP data session object, there may be no more than one PPTP connection from the same external client to a single NetNAT public IP Address.

Monitoring PPTP Sessions

Use the PPTP command to display existing PPTP data session objects. See section 6 for details.