

8

CONFIGURING NETWORK ADDRESS TRANSLATION

This chapter describes the Network Address Translation (NAT) feature and how to configure network addresses for translation. This feature allows the IP addresses of a network to be translated into addresses that can be used outside the network.



For conceptual information, refer to “How NAT Works” on page 8-3.

Configuring Network Address Translation

The procedures in this section describe how to configure NAT.

Prerequisites

Before beginning this procedure, complete the following tasks:

- Log on to the system with Network Manager privilege.
- Set up the ports and paths of your bridge/router using the procedure in Chapter 1.

Enabling NAT Ports

To enable NAT on a port connected to an outside network, use:

```
SETDefault !<port> -NAT CONTROL = Enable
```

Defining the Address Mapping

Address mapping allows you to map one set of IP addresses to another.

To configure the address map, use:

```
ADD !<port> -NAT AddressMap <LHS Address(es)> <RHS Address(es)>  
  [InBound | OutBound | BiDirectional | LoadShare] [Log [0..7]]  
  (<LHS and RHS Address(es)> = <IPAddr/mask> | <IPAddr> |  
  <IPAddr>-<IPAddr> [ ,<IPAddr/mask> | <IPAddr> |  
  <IPAddr>-<IPAddr>.. ])
```

Where LHS Address(es) and RHS Address(es) is a single address or a range of addresses.



If you specify BiDirectional, you can use only a one-to-one map.

Address maps take effect immediately. When you remove a map using the DELETE command, all session table entries based on the map are deleted.

If you have static and dynamic maps, the bridge/router checks the static map first and then the dynamic map. Refer to “Address Mapping” on page 8-6 for more information about static and dynamic maps.



You can also enter address maps directly into the `natmap` file. The `natmap` file in your configuration file directory is in ASCII text format that can be edited with any text editor.

If you choose to manually edit the `natmap` file, you must enter the `NATReStart` command or reboot the bridge/router to load the maps listed in the file.

The `NATReStart` command examines the `natmap` file for syntax errors. If a syntax error is encountered in the `natmap` file, processing stops, an error message is displayed, and no further maps are initialized.

One-to-one example

The following command establishes a one-to-one map for BiDirectional connections on port 1. The LHS address can be inside or outside.

```
ADD !1 -NAT AddressMap 144.195.48.20 144.195.40.17 BiDirectional
```

Many-to-one example

The following command establishes a many-to-one map for OutBound connections on port 1. The LHS addresses are inside, and the RHS address is outside.

```
ADD !1 -NAT AddressMap 192.168.0.0/16 144.195.18.4 OutBound
```

Many-to-many example

The following command establishes many-to-many map for OutBound connections on port 1. The command maps four inside addresses to 16 outside addresses.

```
ADD !1 -NAT AddressMap 192.168.0.1-192.168.0.3, 192.168.10.17  
144.195.40.0/28 OutBound
```

One-to-many example

The following command establishes a one-to-many map for LoadShare connections on port 1. The LHS address is outside, and the RHS addresses are inside.

```
ADD !1 -NAT AddressMap 144.195.18.4 192.168.0.0/16 LoadShare
```

Defining TCP/UDP Port Mapping

To map an address and TCP/UDP port to another address and TCP/UDP port, use:

```
ADD !<port> -NAT TcpUdpPortMap <IPaddr>, <TCP/UDP port#> <IPaddr>  
[,<TCP/UDP port#>] [InBound | Outbound | BiDirectional] [Log[0..7]]
```

If you do not specify the second TCP/UDP port number, the software uses the same port number specified on the first address.

For example, to allow inbound telnet traffic to host 10.0.0.1 on port 1, enter:

```
ADD !1 -NAT TcpUdpPortMap 144.195.48.20,23 10.0.0.1 Inbound
```



The software cannot differentiate between TCP and UDP ports.

If you have address maps and TCP/UDP maps defined, the bridge/router checks the TCP/UDP map first and then the address maps.

Logging Messages

To specify whether messages are logged to syslog, the local console, or both, use:

```
SETDefault -NAT Log = [Syslog | NoSyslog] [Console | NoConsole]  
[SessionFail | NoSessionFail] [SessionSuccess | NoSessionSuccess]  
[LogDetail | NoLogDetail]
```

Only the start-of-connection packets are logged, to avoid flooding logging messages. No more than 10 log messages per second are generated. All messages over 10 are suppressed and the next log message (generated after the one second window expires) contains a counter of how many previous messages were suppressed.

Session Information

You can display the address translations, usage, and idle time statistics for active NAT sessions. Each translation is identified by a session ID. When no option is specified, all active NAT sessions are displayed. The FTP session display can have two TCP connections active at the same time. One is the control channel, for passing commands and responses between client and server. The other is the data channel for the actual data transfer.

To display the NAT session information, use:

```
SHoW [!<port> | !*] -NAT SESSions [TCP | UDP | FTP | Others]
```

To determine the maximum timeout period allowed for NAT sessions to remain idle before the session is terminated, use:

```
SETDefault !<port> -NAT SessionTimeout [TCP | Others] <minutes>
(0-99999) [:<seconds>(0-59)]
```

Translation failure actions

To specify how to handle a packet if the address translation fails at the start of a session, use:

```
SETDefault !<port> -NAT XlateFailAction = PassThrough | Drop |
GenerateICMP
```

Translation failure can occur in the following cases:

- The bridge/router ran out of mappable IP addresses, TCP/UDP port resources, or other internal resources such as memory.
- The session direction was not permitted, even though there was a match for the address in the map.

How NAT Works

This section uses the following terms:

- *Inside network* — the network that includes addresses you want to map.
- *Outside network* — the network, such as the Internet, that you want to connect the inside network to.
- *Inside addresses* — the untranslated addresses of the inside network.
- *Outside addresses* — the addresses that are mapped to the inside addresses.

When to Use NAT

Use NAT for the following purposes:

- **Private address space** — You want to connect to the Internet, but your network does not use globally routable IP addresses. If your network uses private addresses, you can use NAT to translate them to access the outside network.
- **Load sharing** — You want to do load sharing of incoming TCP traffic. For example, traffic destined for a web server identified by one IP address can be redirected to multiple servers with duplicate websites.

- Address migration — You must change your inside addresses. You may change Internet service providers (ISPs), for example, and have to change your numbering scheme. Instead of changing IP addresses on every device in your network, you can use NAT to translate the current addresses into new addresses.
- Address redirection — You want to redirect traffic from one host to another. For example, you want to stop telnet traffic to server A but want to allow traffic to a new server B. Translate the address of server A to that of server B to transparently redirect users to the new server.

Guidelines Refer to the following guidelines before configuring NAT:

- Traffic that has embedded IP addresses, such as DNS requests and responses, will not be translated by NAT.
- You should use private IP address space on your inside network recommended by the Internet Assigned Numbers Authority (IANA). The following address ranges are designated as private networks that should not be advertised:

10.0.0.0 — 10.255.255.255 (10.0.0.0/8)

172.16.0.0 — 172.31.255.255 (172.16.0.0/12)

192.168.0.0 — 192.168.255.255 (192.168.0.0/16)

If you use a different address range that can be validly assigned to someone else's network, you will not be able to communicate with that network. For example, if you FTP from the inside host 1.1.1.1 to a valid outside network host with the IP address 1.1.1.7, the bridge/router will not forward the FTP request outside the network.

- NAT may not be practical if large numbers of hosts in the inside network communicate with the outside network, because NAT is slower than untranslated addresses. Most traffic should originate from or go to hosts within the domain. Because most hosts never communicate with an outside network, only a subset of inside addresses need to be translated into outside addresses.

Basic NAT Operation

A bridge/router using NAT has at least one port connected to the inside network and one port connected to the outside network. Enable NAT on the port that is directly connected to the outside network. If you have multiple connections to outside networks, enable NAT on each outside port. You must have separate address maps for each NAT port. The inside network is always the network that contains the address you need to map.

When configuring address maps using the `ADD !<port> -NAT AddressMap` command, you specify the source address (left-hand side (LHS) address) and the translated address (right-hand side (RHS) address). You also specify the direction of the translation: outbound, inbound, bidirectional, or load sharing.

The bridge/router identifies the first packet and the direction of a TCP session and creates a NAT session if a map exists for that direction. However, the bridge/router cannot identify the first packet of a UDP session, so every UDP packet is considered the first packet of a session, which results in a new NAT session for every UDP packet.

Specifying Direction Depending on which direction you specify, the LHS addresses can be either the inside or the outside addresses.

When you specify outbound, the inside addresses are the LHS addresses, which are translated into the outside RHS addresses. Outbound traffic is defined as traffic that leaves the bridge/router through the NAT port.

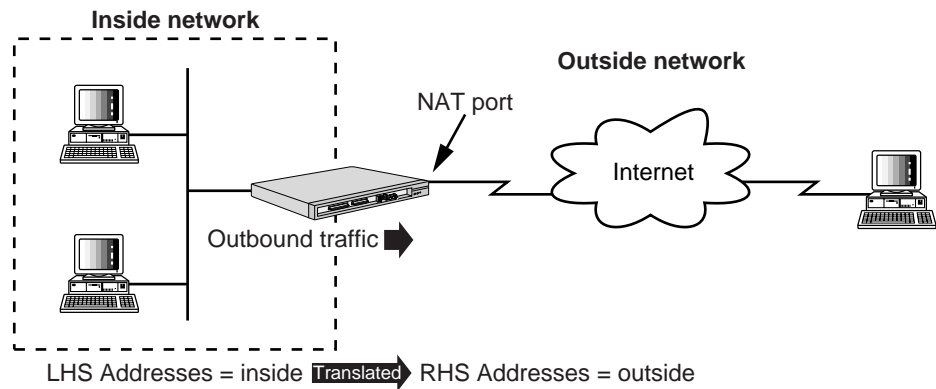


Figure 8-1 Outbound Traffic

When you specify inbound or load sharing, the LHS addresses are the outside addresses, which are translated into the inside RHS addresses. Inbound or load sharing traffic is defined as traffic that enters the bridge/router through the NAT port.

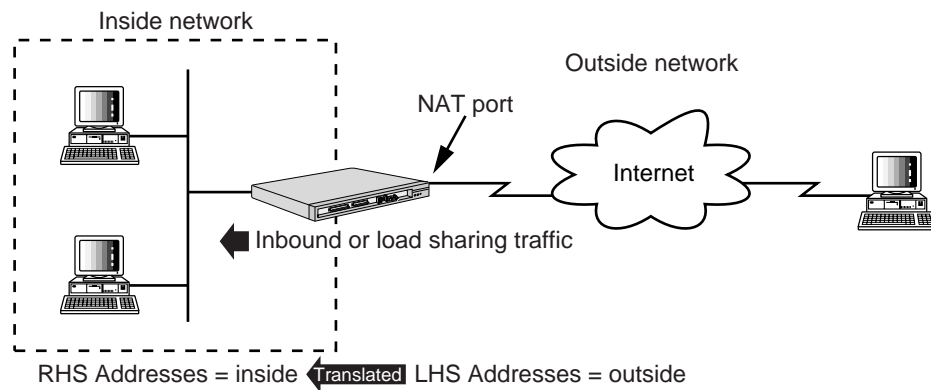


Figure 8-2 Inbound or Load Sharing Traffic

When you specify `bidirectional`, you can use either the LHS or RHS address for the inside address.

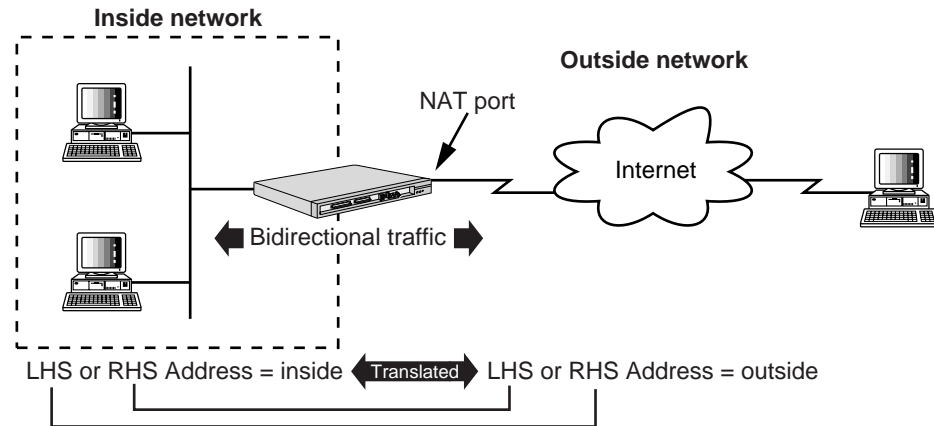


Figure 8-3 Bidirectional Traffic

Address Mapping

Static mapping establishes a **one-to-one** map between two addresses. Static mapping is useful when a host on the inside network must be reachable by a specific IP address, for example, a web server. Use static mapping with outbound, inbound, or bidirectional connections. Bidirectional translation can use only static mapping.

Dynamic mapping establishes a many-to-one, a many-to-many, or a one-to-many map.

A **many-to-one** map, used by outbound or inbound connections, translates multiple addresses into a single address. In an outbound connection, multiple inside hosts can connect using the same outside address because the bridge/router appends a TCP or UDP port number to the outside address for each connection. For example, three hosts are connected at the same time:

Inside Address	Inside Outside Address:Port
10.0.0.5	144.195.23.10:1031
10.0.0.18	144.195.23.10:1032
10.0.0.25	144.195.23.10:1033

In an inbound connection, multiple outside addresses are translated into a single inside address.

A **many-to-many** map, used by outbound or load sharing connections, translates multiple addresses by assigning addresses from a pool. Each subsequent connection is assigned the next available address from the pool. If the bridge/router cannot assign an address because it has run out of addresses in the pool, it drops the packet. Make sure you have enough addresses in the address range if you want to use a many-to-many map.

A **one-to-many** map, used by load sharing or outbound connections, translates a single address to one of many addresses. In a load sharing connection, each subsequent connection to the outside address is translated into the next available address from the inside address pool. In an outbound connection, each subsequent connection from a single inside host is assigned a new IP address, making it appear as though a single user is multiple users.

Using a Mask

When using the ADD !<port> -NAT AddressMap command, you can specify an address block using a mask: <IPaddr/mask>. <mask> is a number in the range of 0-32, which indicates the number of bits in the IP address that remain unchanged for the IP addresses in that block. The remaining bits in the IP address should be all 0s. The address block includes all addresses except for the first address and the last (x.x.x.255) address.

For example:

144.195.0.0/16 All the addresses in the range from 144.195.0.1 to 144.195.255.254

144.195.1.2/32 The host itself 144.195.1.2

0.0.0.0/0 All the IP addresses in your network

224.0.0.0/4 All the class D multicast addresses, from 224.0.0.1 through 239.255.255.254

TCP/UDP Port Mapping

Mapping an address and TCP/UDP port to another address and TCP/UDP port allows more control of the type of traffic NAT translates.

Most TCP and UDP servers use the same port number range, 0-1023, to listen for incoming connections. Most servers use a fixed, well-known port number for listening to a particular service. The major services and their port numbers are listed in Table 8-1. For a detailed list of reserved services and port numbers, refer to RFC 1700.

Table 8-1 TCP/UDP Port Numbers and Services

Service	TCP/UDP Port Numbers	Service	TCP/UDP Port Numbers
DNS	53	SMTP	25
finger	79	SNMP	161, 162
FTP	20, 21	syslog	514
Gopher	70	talk	517, 518
HTTP	80	Telnet	23
NNTP	119	TFTP	69
NTP	123	UUCP	9540
POP	109, 110	WAIS	210
RIP	520	whois	43



TCP/UDP clients use the port number range 1024 - 65535.

NAT Scenarios

This section contains the following scenarios:

- Private Address Space
- Load Sharing
- Address Migration
- Address Redirection

Private Address Space

The following two examples show the same network. Example 1 has the same network defined as the inside network for all steps. Example 2 has opposite networks defined as the inside network for different steps in the same procedure. The inside network is always the network whose addresses need to be mapped.

Example 1 The Human Resources (HR) network is an isolated network using the address block 10.0.0.0/8, which is a recommended private address block that is not advertised outside the domain.

See Figure 8-4 and follow these steps to:

- Enable any host in the HR network to connect to the corporate network.
- Allow access to an HR server from the corporate network.

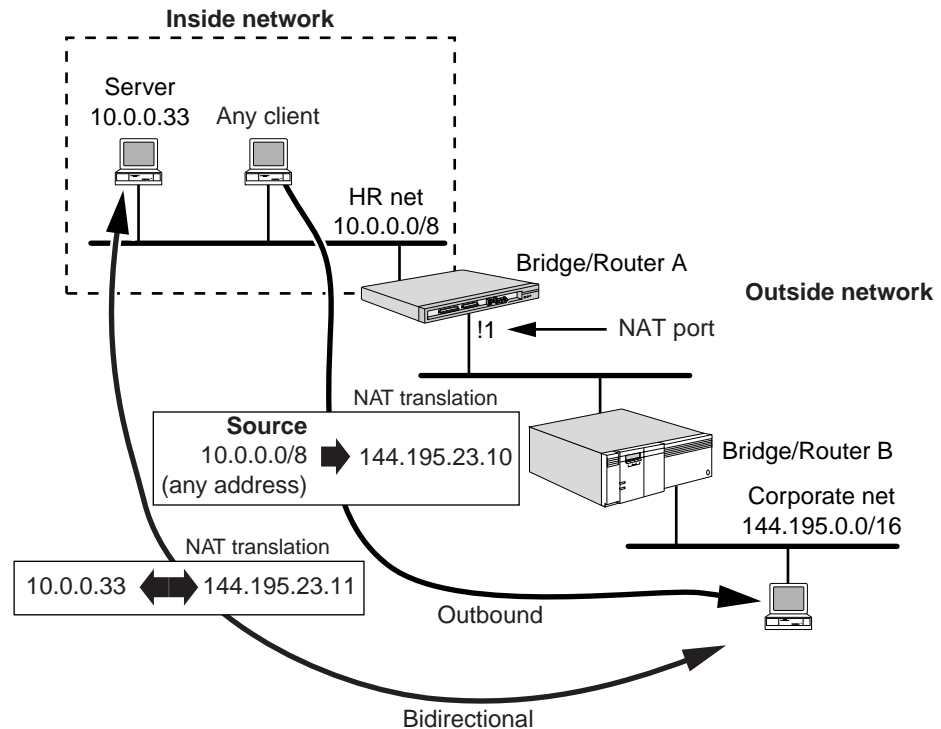


Figure 8-4 Private Network Connecting to the Corporate Network

- 1 Enable NAT on port 1 of bridge/router A by entering:

```
SETDefault !1 -NAT CONTROL = Enable
```

The NAT port must be connected to the outside network.

- 2 To access the corporate network from the HR network, map the HR address block 10.0.0.0/8 to a single outside address, 144.195.23.10, and specify outbound by entering:

```
ADD !1 -NAT AddressMap 10.0.0.0/8 144.195.23.10 OutBound
```

Any address from the HR network will be translated into the outside address 144.195.23.10.

- 3 To make the HR server address 10.0.0.33 available to the corporate network, map it to the outside address 144.195.23.11, and specify bidirectional by entering:

```
ADD !1 -NAT 10.0.0.33 144.195.23.11 BiDirectional
```

The bidirectional option allows the server to access the corporate network as well.

Example 2 The Human Resources (HR) network is an isolated network using the address block 10.0.0.0/8, which is a recommended private address block that is not advertised outside the domain.

See Figure 8-5 and follow these steps to:

- Enable any host in the HR network to connect to the corporate network.
- Allow access to an HR server from the corporate network.
- Limit access to the corporate network to only one server.

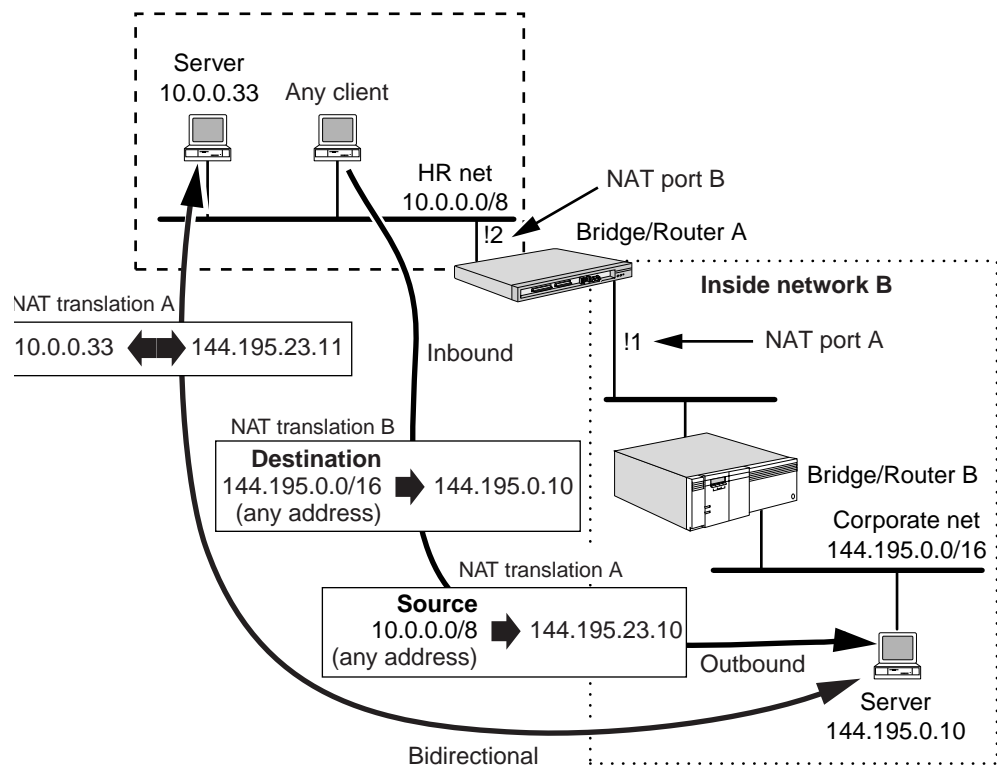


Figure 8-5 Relative Inside Networks

- 1 Enable NAT on port 1 of bridge/router A by entering:

```
SETDefault !1 -NAT CONTROL = Enable
```

This NAT port is connected to the outside network with regard to inside network A.

- 2 To access the corporate network from the HR network, map the HR address block 10.0.0.0/8 to a single outside address, 144.195.23.10, and specify outbound by entering:

```
ADD !1 -NAT AddressMap 10.0.0.0/8 144.195.23.10 OutBound
```

Any address from the HR network will be translated into the outside address 144.195.23.10.

- To make the HR server address 10.0.0.33 available to the corporate network, map it to the outside address 144.195.23.11, and specify bidirectional by entering:

```
ADD !1 -NAT 10.0.0.33 144.195.23.11 BiDirectional
```

The bidirectional option allows the server to access the corporate network as well.

- Enable NAT on port 2 of bridge/router A by entering:

```
SETDefault !2 -NAT CONTROL = Enable
```

This NAT port is connected to the outside network with regard to inside network B.

- To limit access from the HR network to only server 144.195.0.10, map all destination addresses in the corporate network to the server address, and specify InBound by entering:

```
ADD !2 -NAT 144.195.0.0/16 144.195.0.10 InBound
```

Load Sharing

The 3Com web server is replicated on five different servers. The URL `www.3com.com` is accessed thousands of times a day. The domain name server (DNS) advertises the address 192.156.136.22 for `www.3com.com`, even though there is not an actual server at that address.

To forward traffic directed to the virtual server 192.156.136.22 evenly to five web servers, follow these steps (see Figure 8-6):

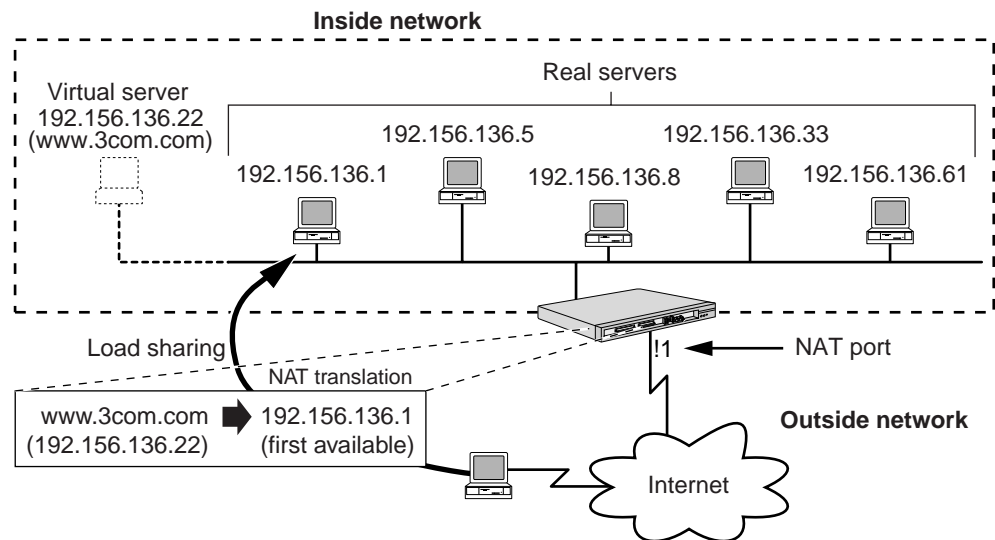


Figure 8-6 Load Sharing

- Enable NAT on port 1 of the bridge/router by entering:

```
SETDefault !1 -NAT CONTROL = Enable
```

The NAT port must be connected to the outside network.

- To distribute traffic evenly to the five web servers, map the virtual server address 192.156.136.22 to all five server addresses, and specify load sharing by entering:

```
ADD !1 -NAT AddressMap 192.156.136.22 192.156.136.1, 192.156.136.5,  
192.156.136.8, 192.156.136.33, 192.156.136.61 LoadShare
```

Load sharing only translates inbound connections. No translation is required for outbound sessions originating from the servers.

The bridge/router does not detect if a load sharing host is down. However, a log message indicating the possibility that a host may be down is logged. If a host is down, and you want to remove the host from the list, you can delete the original load sharing map (this causes all active sessions to be flushed) and create a new load sharing map without the host.

Address Migration

Your company has changed ISPs and the new ISP must reassign your IP addresses to work with their network. Reconfiguring every host with a new address requires extensive effort, time, and interruption to users. Moreover, if you change service providers frequently, the process would have to be repeated every time.

To translate each old IP address into a new address, follow these steps (see Figure 8-7):

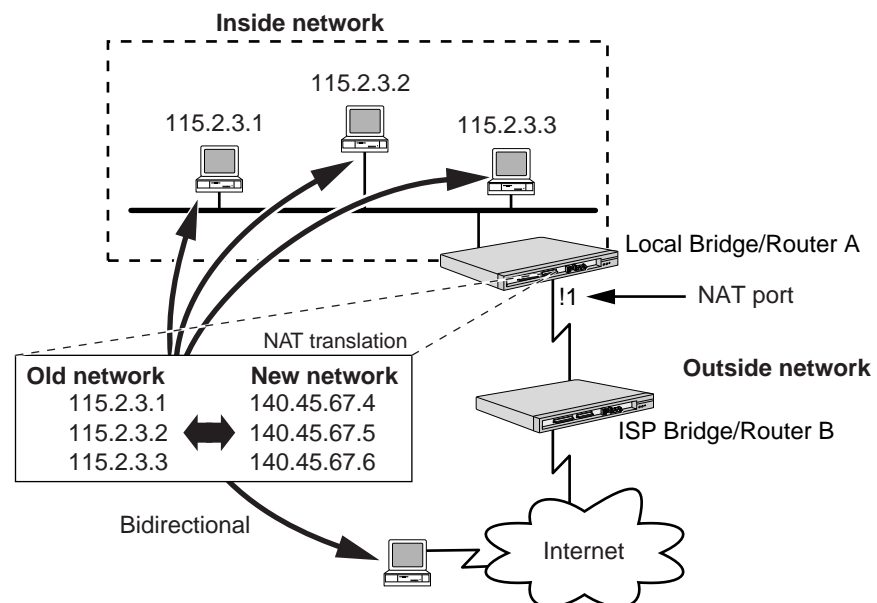


Figure 8-7 Address Migration

- 1 Enable NAT on port 1 of bridge/router A by entering:

```
SETDefault !1 -NAT CONTROL = Enable
```

The NAT port must be connected to the outside network.

- 2 Map each old address to a new address separately, and specify bidirectional by entering:

```
ADD !1 -NAT AddressMap 115.2.3.1 140.45.67.4 BiDirectional
ADD !1 -NAT AddressMap 115.2.3.2 140.45.67.5 BiDirectional
ADD !1 -NAT AddressMap 115.2.3.3 140.45.67.6 BiDirectional
```

If you do not need bidirectional access, you can map a range of addresses to a new range using the OutBound option.

Address Redirection

The new, fast FTP server `butterfly.isp.com` was bought to replace the old FTP server, `tarantula.isp.com`.

To transparently redirect traffic from `tarantula` to `butterfly`, follow these steps (see Figure 8-8):

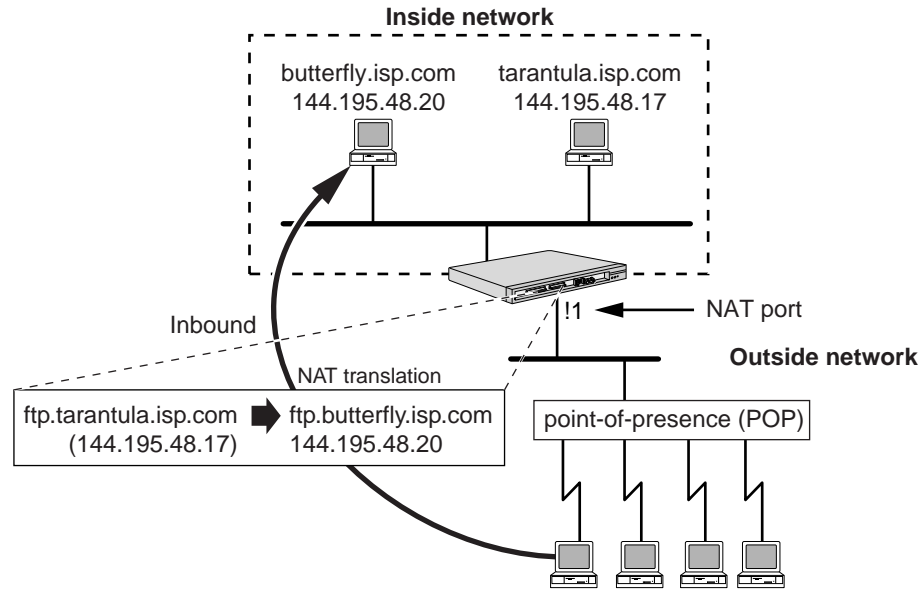


Figure 8-8 Address Redirection

- 1 Enable NAT on port 1 of bridge/router A by entering:

```
SETDefault !1 -NAT Control = Enable
```

The NAT port must be connected to the outside network.

- 2 Map `tarantula` (144.195.48.20) to `butterfly` (144.195.40.17), and specify inbound by entering.

```
ADD !1 -NAT AddressMap 144.195.48.20 144.195.40.17 InBound
```