No. 107

Technology White Paper

NAT: Network Address Translation

Abstract

With the exponential growth of the Internet, a shortage of IP addresses is becoming a problem. One method to help preserve the limited number of legally registered addresses is Network Address Translation (NAT). This paper discusses the concept of NAT and explains how to configure NAT in different scenarios. Riverstone Networks' NAT is a wire-speed solution, which means it is realized in hardware rather than in software. This guarantees the performance required by today's Service Provider (SP) networks. NAT is supported on all Riverstone platforms.

Overview

NAT operates on a router between an inside (local) and outside, public (global) network and helps conserve IP addresses. NAT is very often used with a special group of IP addresses, shown below in Table 1, although it works with any IP address scheme. These addresses were defined (reserved) by the Internet Assigned Numbers Authority (IANA) to be used in private (local) intranets rather than in the public (global) Internet, and are non Internet-routable IP addresses.

Address Class Range	Network Address Range
A	10.0.0.0 – 10.255.255.255
B	172.16.0.0 - 172.31.255.255
C	192.168.0.0 – 192.168.255.255

Table 1: IANA-allocated, non Internet-routable, IP address schemes (RFC 1918).

Service Providers especially benefit from the NAT solution as it saves them the expense of registering large numbers of addresses with IANA. Enterprises can utilize their own internal IP address schemes and only use a limited number of IP addresses for Internet access. This not only reduces costs, but also simplifies changes since the entire internal address scheme does not have to change.

River STONE NETWORKS™

How NAT Works

Basically, NAT does a one-to-one or a many-to-one IP address translation. An inside (local) IP address is mapped to an outside (global) IP address, meaning that an inside IP address is replaced by the appropriate outside IP address, and vice versa.

Network Address Translation includes the following steps:

1. The IP address in the header is replaced with the new inside or outside IP address. The port numbers in the TCP/UDP header are replaced by the new port if port translation is enabled.
2. The checksum for the IP packet is recalculated and checked for integrity.
3. The TCP header checksum must also be recalculated since this checksum is calculated using the new inside or outside IP address, new port (if applicable), and the payload (if applicable).

There are two types of NAT—static and dynamic. These can be used simultaneously.

Static NAT

As the name implies, static NAT defines a fixed address translation from the inside (local) network to the outside (global) network. The (simplified) CLI command for Riverstone Networks' RS Router is:

nat create static ip <local IP address> <global IP address>

Static NAT with Port Address Translation (PAT)

It is also possible to include the port numbers. This is known as Port Address Translation (PAT) when TCP or UDP is used, and it works in both static and dynamic NAT. The (simplified) RS Switch Router command would then be:

nat create static protocol <ip|tcp|udp> local-ip <IP address> global-ip <IP address> local-port <port number> global-port <port number>

To make the configuration faster, it is possible to map a range of IP addresses. For details, see the Appendix "NAT CLI Commands."

Static NAT Example

The following command shows how to use static NAT for local address 10.1.1.13 and global address 136.1.1.13:

nat create static local-ip 10.1.1.13 global-ip 136.1.1.13 protocol ip (see Figure 1)

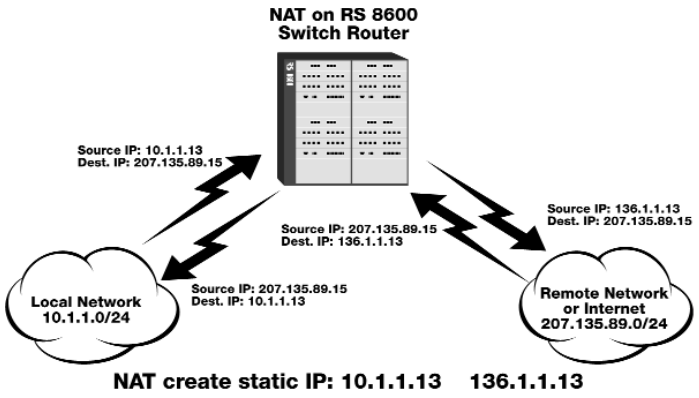

Figure 1: Example of static NAT

(Source is translated from 10.1.1.13 to 136.1.1.13 and destination is translated from 136.1.1.13 to 10.1.1.13.)

Dynamic NAT

Dynamic NAT, in contrast, translates from a pool of local IP addresses to a pool of global IP addresses. The user must define both pools. The address assignment is done automatically by the NAT-enabled router, and the user has no influence over which IP address is picked from the address pool.

Dynamic NAT Example

The following commands show how to use dynamic NAT for a local address pool, 10.1.1.0 to 10.1.1.254, and a global address pool, 136.1.1.0 to 136.1.1.254:

```
acl Local permit ip 10.1.1.0/24
nat create dynamic local-acl-pool local global-pool 136.1.1.0/24
```

In the example below, the two ACLs were named "Local" and "Global." Of course any name can be chosen.

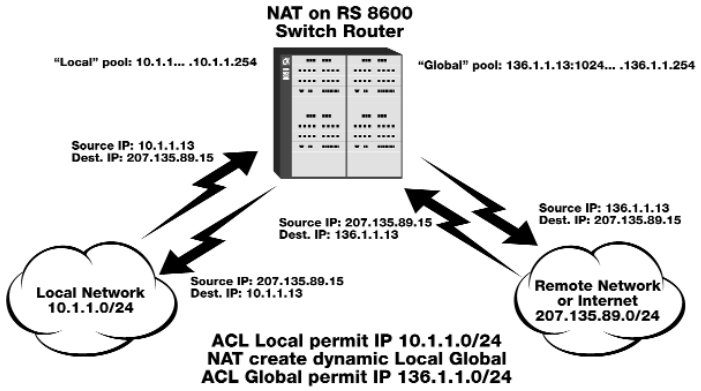

Figure 2: Example of dynamic NAT between local and global (remote) IP address pools.

Dynamic NAT with IP Overload

Port Address Translation (PAT) allows many-to-one address mapping, since many inside IP addresses can be mapped to one outside IP address. The port number (TCP or UDP) is sufficient to ensure that packets are delivered properly. To enable this feature, the keyword "enable-port-overload" must be added to the "nat create dynamic" command. "Overload" refers to a situation in which no more free IP addresses are available from the pool and ports have to be assigned to distinguish two different connections.

The range of ports that the RS Router uses is 1024 to 4999; this is a total of approximately 4,000 ports.

Dynamic NAT with IP Overload Example

The following commands show how to use Dynamic NAT with IP overload for local address pool 10.1.1.1 to 10.1.1.254 and global address 136.1.1.13:

```
acl local permit tcp 10.1.1.0/24
nat create dynamic local-acl-pool local global-pool 136.1.1.13 enable-ip-overload
```

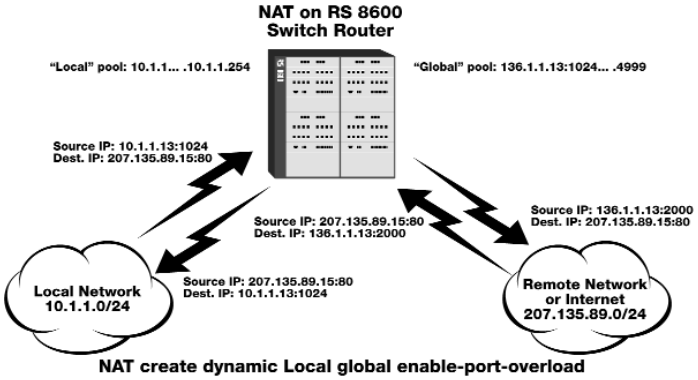

Figure 3: Dynamic NAT with the port overloading feature enabled.

Dynamic NAT with Outside Interface Redundancy

If you have redundant connections to the remote network via two different interfaces, you can use NAT for translating the local address to different global pools specified for the two connections. This case is possible when you have two ISPs connected on two different interfaces to the Internet through a routing protocol; some routes will result in traffic going out of one interface and others going out of the other interface. NAT will check which interface the packet is going out of before selecting a global pool. Hence you can specify two different global pools with the same local ACL pool on two different interfaces.

Dynamic NAT with Outside Interface Redundancy Example

The following commands show how to use Dynamic NAT with redundant interfaces for local address pool 10.1.1. to 10.1.1.254 and global address pool 136.1.1.13 and 208.1.1.13 connected to two different interfaces:

```
acl local permit tcp 10.1.1.0/24

nat create dynamic local-acl-pool local global-pool 136.1.1.13 enable-ip-overload
matching-interface 136net

nat create dynamic local-acl-pool local global-pool 208.1.1.13 enable-ip-overload
matching-interface 208net
```

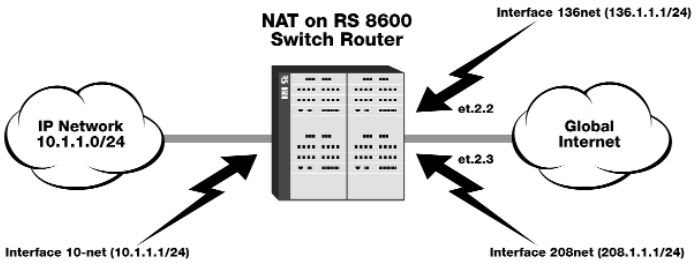

Figure 4: Dynamic NAT with port-overloading feature enabled.

NAT of Inside Local Addresses

In this scenario, a private network was set up with its own private IP address scheme, and NAT is configured on an RS Switch Router connecting the inside (local, 10.1.1.0) network to an outside (global, 192.50.20.2) network, using the public Internet addressing scheme. NAT can be configured in a static or dynamic way. Since the NAT feature in the RS Switch Router maps addresses between both networks, the end stations do not know the real IP addresses of the station they are communicating with on the other side. Figure 4 illustrates this scenario. When station A wants to access the outside network, it uses one of the IP addresses from the pool defined in the router (192.50.20.1 to 192.50.20.254). Station B knows only about station A's router-assigned IP address; it is not aware that station A's real IP address is 10.1.1.2.

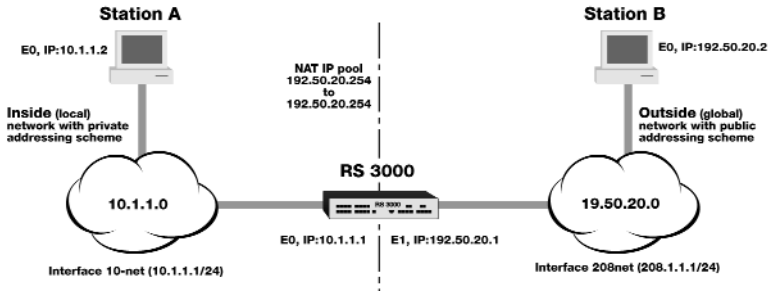

Figure 5: NAT of inside local addresses.

Dynamic NAT with Port Address Translation of Global Addressing

In this scenario, there are more inside IP hosts than NAT pool-allocated addresses. The PAT feature of NAT is able to handle this situation using socket (socket = address: port) mappings. To ensure that multiple applications can use the same IP address from the limited IP pool, the "overload" feature has to be enabled. Once this feature is enabled, NAT requires only unique ports and not unique IP addresses to establish a connection between two hosts. Figure 5 illustrates this scenario. There is only one IP address in the NAT pool (192.50.20.1). Station A establishes a TCP connection via port 1024 (10.1.1.2:1024) with station B (192.50.20.2:23) by mapping the IP address from the NAT pool. If station C now wants to connect to station D, it would fail since no more IP addresses are available from the NAT pool. With the overload feature enabled, this would work, however. Table 2 shows how NAT maps the inside addresses to the outside by using the NAT pool IP address.

Note that some applications require unique IP addresses for a client-server connection and therefore cannot work with PAT.

NAT Handling of Overlapping Networks

This scenario occurs when inside, user-defined IP network addresses (not the IANA-assigned IP addresses shown in Table 1) coexist with legally registered outside networks. Other networks may have the same addresses, whether legally registered or picked without using the IANA-allocated IP addresses.

NAT Support for TCP Load Sharing (Distribution)

Load sharing or load distribution is a separate feature in Riverstone Networks' RS Routers called LS-NAT.

Note though that NAT and LS-NAT will not work for the same IP address on the same RS Switch Router.

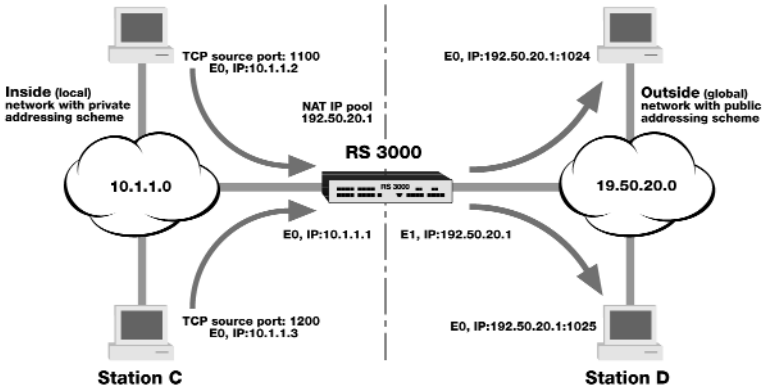

Figure 6: NAT with Port Address Translation (PAT) of global addressing.

Protocol (connection)	Inside Source Address : Port	Outside Source Address : Port	Outside Destination Address
TCP (A to B)	**10.1.1.2:1100**	**192.50.20.1:1024**	**192.50.20.2**
TCP (C to D)	**10.1.1.3:1200**	**192.50.20.1:1025**	**192.50.20.3**

Table 2: NAT table for example shown in Figure 6.

NAT Configuration

Configuring NAT of Inside Local Addresses

1. Configure the RS Switch Router outside and inside interfaces for NAT
   rs(config) # interface create ip 10-net address-netmask 10.1.1.1/24 port et.2.1
   rs(config) # nat set interface 10-net inside

   rs(config) # interface create ip 192-net address-netmask 192.50.20.1/24 port et.2.2
   rs(config) # nat set interface 192-net outside

2. Configure NAT static and dynamic
   rs(config) # nat create static ip 10.1.1.2 192.50.20.2
   rs(config) # acl local permit ip 10.1.1.0/24
   rs(config) # acl global permit ip 192.170.2.0/24

   rs(config) # nat create dynamic local global

In this configuration example, it is not necessary to separately define ACL lists and NAT pools and then combine them. This can all be done with one command.

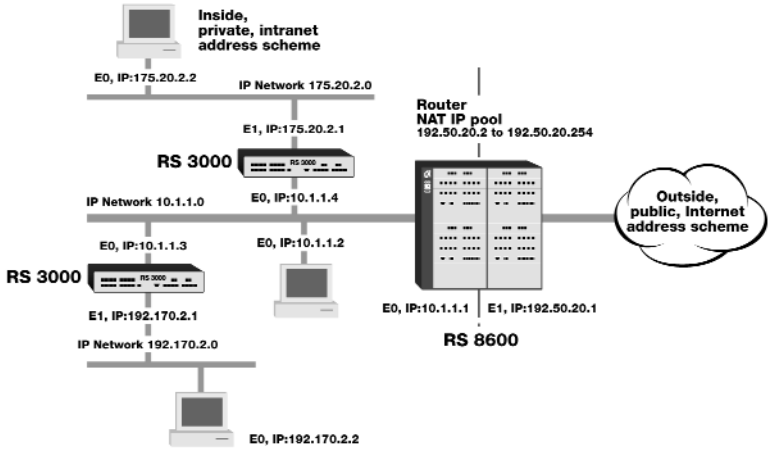

Figure 7: Configuring NAT of inside local addresses

Dynamic NAT with Port Address Translation of Global Addressing

In this example, the address pool consists of only one IP address. In order to make multiple connections possible, Port Address Translation has to be enabled. This is achieved in a very simple way, as shown next.

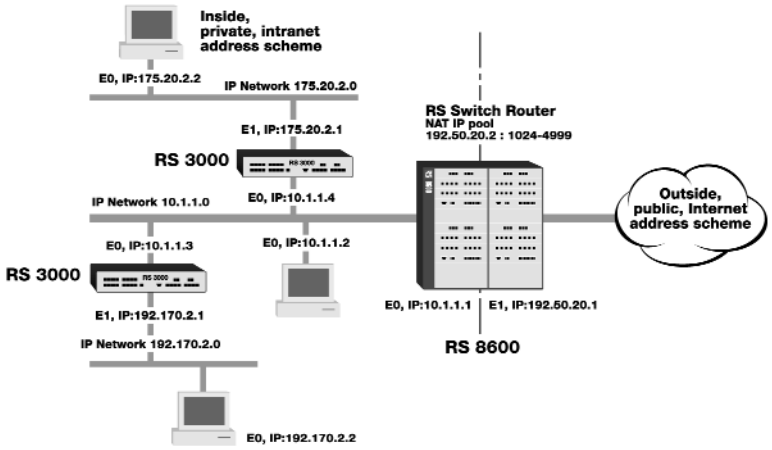

Figure 8: Configuring NAT with Port Address Translation of global addressing.

Configure the router outside and inside interfaces for NAT

rs(config) # interface create ip 10-net address-netmask 10.1.1.1/24 port et.2.1
rs(config) # nat set interface 10-net inside

rs(config) # interface create ip 192-net address-netmask 192.50.20.1/24 port et.2.2
rs(config) # nat set interface 192-net outside

Configure NAT static and dynamic

rs(config) # nat create static ip 10.1.1.2 192.50.20.2

rs(config) # acl local permit ip 10.1.1.0/24
rs(config) # acl global permit ip 192.170.2.0/24

rs(config) # nat create dynamic local global enable-port-overload

The key word here is "enable-port-overload," which enables the one-to-many address mapping by using ports.

The following table explains other NAT commands.

Command:	**rs(config) # nat set ftp-control-port 100**
Result:	**Changes the FTP control port to 100 from default 21**
Command:	**rs(config) # nat set dynamic-binding-timeout 100**
Result:	**Changes the NAT dynamic timeout in minutes from default 24h**
Command:	**rs(config) # nat set ftp-session-timeout 10**
Result:	**Changes the NAT FTP session timeout to 10 minutes from default 30 minutes**
Command:	**rs # nat show translations all|type|owner|local-filter-in |verbose**
Result:	**Shows the currently active translations**
Command:	**rs # nat show timeouts**
Result:	**Shows the current set of timeouts**
Command:	**rs # nat show statistics**
Result:	**Shows the NAT statistics**
Command:	**rs # nat flush-dynamic-binding pool-specified|type-specified |owner-specified|all**
Result:	**Deletes all the dynamic NAT bindings for the local and global settings**

River STONE NETWORKS™

Summary

NAT Pros

- Conserves the legally registered addressing scheme. Enterprises and ISPs benefit since they can reduce the number of legally registered IP addresses.
- Network design is simplified by now-limitless availability of addressing schemes.
- Merging and changing of networks is simplified. An enterprise can change its ISP vendor without having to completely renumber the network.

NAT Cons

- Loss of end-to-end IP trace ability; the command "traceroute <hostname>" will not be of great help anymore.
- Applications that send IP addressing information within their data require special handling (e.g., FTP).

Overall, NAT is a great solution to the problem of limited IP addresses and of having private (local) address schemes connected to the outside (global) Internet. Riverstone has implemented NAT in hardware rather than in software. This guarantees the performance required by today's Service Providers. Riverstone Networks provides its customers with enabling service-provider infrastructure that yields increased network reliability, improved performance and enhanced services.

Appendices

NAT CLI Commands

- nat set interface <interface-name> <inside|outside>

  This command is used to define the interface as being inside or outside. When NAT is enabled, static or dynamic, it will only be applied to the interfaces to which NAT is defined. You can apply NAT to all interfaces by using the keyword "all." Here it is determined whether that interface's channels are all NAT compliant. If even one channel is not NAT compliant, then the command will return an error.

- nat create static protocol <ip|tcp|udp> local-ip <ip-address|ip-range> global-ip <ip-address|ip-range> local-port <port-number|[any]> global-port <global-port|[any]>

  This command is used to create static NAT bindings between local IP addresses and global IP addresses. If the protocol is UDP or TCP, then you can also specify Port Address Translation by giving the ports to be translated local-port and global-port options.

  When a packet arrives at an inside interface, an ACL check is done on that packet prior to NAT processing. If the ACL check is positive, then the local IP address is translated to a global IP address and forwarded to the destination.

- nat create dynamic local-acl-pool <acl-name> global-pool <ip-address|ip rangel ip-address|netmask|ip-list> matching-interface <interface-name> enable-ip-overload

River STONE NETWORKS™

This command is used to create dynamic NAT bindings between a local ACL list and a global ACL list. The keyword "enable-port-overload" allows for Port Address Translation, if no more global IP addresses are left over from the list.

Every time a packet arrives from an inside interface, an ACL check is done on that packet prior to NAT processing. If the ACL check is positive, then a local IP address is assigned to a global IP address by picking up a free address from the global list. A binding is assigned for that address and the flow is installed in hardware for that binding. If port overloading is enabled, Port Address Translation will be performed if all the addresses are used up and no free global addresses are available from the list.

Glossary

Global IP address. IP address used within the pubic Internet. These are assigned by the Internet Assigned Numbers Authority (IANA).

Internet Assigned Numbers Authority (IANA). This organization administers the distribution of IP addresses.

IP address pool. As the name suggests, a list of user-defined IP addresses used by dynamic NAT.

Local IP address. IP address used within private organizations only.

Port Address Translation (PAT). In addition to NAT, PAT uses port addresses to map inside (local) and outside (global) IP addresses. This allows a many-to-one IP address mapping from inside (local) to outside (global) addresses.

Transmission Control Protocol (TCP). This connection-oriented transport protocol uses port numbers to establish a reliable connection between hosts. TCP uses various methods to guarantee the reliable delivery of data segments, such as sliding windows and three-way handshake to establish and terminate a connection. This is required since TCP uses IP, a best-effort protocol.

UDP. This connectionless transport protocol uses port numbers to establish an unreliable connection between hosts. UDP has less overhead than TCP.

Riverstone Networks, Inc.
5200 Great America Parkway, Santa Clara, CA 95054 USA

408 / 878-6500 *or* ***www.riverstonenet.com***

Printed in the USA *v 1.5 9/00*