Global Knowledge®

Expert Reference Series of White Papers

Network Address Translation (NAT) and Port Address Translation (PAT)

Network Address Translation (NAT) and Port Address Translation (PAT)

Al Friebe, Global Knowledge Instructor, CCDA, CCDP, CCNA, CCNP, A+, CCSI

Introduction

It's common today to use private addressing within an Autonomous System (AS), a collection of routers and subnets under a common administrative domain. The problem is that per RFC 1918 (Address Allocation for Private Internets), advertising these address spaces to the public Internet is not allowed. This means that if you send a packet with a "private" source address to the Internet, the destination will not be able to reply to you because the routers on the Internet backbone won't know where you are. The solution to this problem is Network Address Translation (NAT).

Static NAT

In the Static NAT method, you build the translation table by hand. For example, let's say that we want to translate addresses on the 10.1.2.0/24 subnet (private address space) to addresses on the 200.1.2.0/24 network (public). We could translate the first address like this:

- Router(config)#ip nat inside source static 10.1.2.1 200.1.2.1

The translation tells the router that if a packet with the specified source address (10.1.2.1) hits the inside interface and is bound for the outside interface; translate the source address statically to the second address (200.1.2.1). You can have multiple translation lines, as many as you need.

The next thing to do is to tell the router which interface (or subinterface) is the "inside" and which is the "outside." For our example, let's assume that the FastEthernet0/0 interface connects to our LAN, and the Serial0/0 interface leads to our Internet Service Provider (ISP).

- Router(config)#interface fa0/0
- Router(config-if)#ip nat inside
- Router(config-if)#int s0/0
- Router(config-if)#ip nat outside

Notice that although we only specified the translation of the source address as the packet transited from the inside to outside interface, the router will automatically translate the destination addresses of packets traversing the router from the outside to inside interface. The beauty of it is that the translation is invisible to all devices, not just the one device performing the translation.

Copyright ©2010 Global Knowledge Training LLC. All rights reserved.

You can view the translation table with the command show ip nat translations (or shortened to sh ip nat trans), and which interfaces are the inside and outside (along with other info) with show ip nat statistics.

When you display the translation table, you'll notice that it specifies "inside local" and "inside global" addresses.

- Inside refers to where the addressed device physically resides (inboard of the inside interface, that is, on our side of the router).
- Local means "as seen from the inside"
- Global means "as seen from the outside"

In other words, the "inside local address" is our host's untranslated (actual) address, and the "inside global address" is the translated address as seen by those outboard of the outside interface.

Dynamic NAT

While static NAT works, since it uses manually constructed "one-to-one" translations, it's not scalable. For example, translating all of the legal host addresses on the 10.1.2.0/24 subnet would require 254 lines. And if we were dealing with the entire 10.0.0.0/8 network, covering all possible addresses would require over sixteen million lines! The solution is dynamic NAT.

In dynamic NAT, instead of specifying the translations one-by-one, you give the NAT device some rules that specify how to translate addresses. With a Cisco router, the addresses to be translated are specified by an access control list (ACL), and the addresses to which they are translated are specified by a pool.

For example, to translate any address on the 10.1.2.0/24 subnet (those permitted by ACL 1) to an address on the 200.1.2.0/24 network (as specified by the pool named "Test"), you could do the following.

- Router(config)#ip nat inside source list 1 pool Test

The translation tells the router that if a packet with source address matching a permit in ACL 1 hits the inside interface, and it is bound for the outside interface, translate the source address to an available address in the pool named "Test."

Obviously, you also need to create ACL 1 and the pool Test. Let's create the ACL first.

- Router(config)#access-list 1 permit 10.1.2.0 0.0.0.255

As is the usual case with a standard IP ACL, this list specifies the source address. Remember that ACLs use a wildcard (inverse) mask.

Now, let's create the pool named Test (remember pool names are case-sensitive).

- Router(config)#ip nat pool Test 200.1.2.0 200.1.2.127 netmask 255.255.255.0

The netmask specified for the pool is the subnet mask of the network or subnet containing the translated addresses, and this is not a wildcard (inverse) mask. If you prefer, you can specify the pool's mask using "slash" ("bitcount", "CIDR") notation by using the "prefix-length" option.

- Router(config)#ip nat pool Test 200.1.2.0 200.1.2.127 prefix-length 24

Notice that while the ACL covers 254 addresses, the pool only specifies 127 addresses since the size of the pool only needs to cover the number of hosts that simultaneously require translation. It's unlikely that all 254 host addresses are actually in use, and even if they are, it's unlikely that they are simultaneously trying to access the Internet.

Also, if the public addresses are being rented from a provider (typically the case), conserving public IP addresses can save money.

Finally, if it hasn't already been done, the inside and outside interfaces must be assigned, just as with static NAT. Let's assume that FastEthernet0/1 is on the inside, and Serial1/2 is on the outside.

- Router(config)#interface fa0/1
- Router(config-if)#ip nat inside
- Router(config-if)#int s1/2
- Router(config-if)#ip nat outside

If we view the translation table at this point, we would see no entries, because no traffic matching the ACL has attempted to traverse the router from inside to outside. To trigger a translation, we generate traffic from an inside host that's destined for an outside host.

For TCP, entries are placed in the table when a session is built (when the NAT device sees the SYN marking the start of a three-way handshake) and removed when the session is terminated. For UDP and ICMP, the translation table entries are created with the first packet in a particular data stream, and the entries are removed when an inactivity timer expires.

Port Address Translation (PAT)

Dynamic NAT, which could allow several hosts to use the same public IP address at different times of the day, still translates on a "one-to-one" address basis. That is, each inside local address (usually private) being actively translated requires one global address (usually public).

In PAT (also known as "overloading"), many inside local addresses are simultaneously mapped to one inside global address (that is, the global address is "overloaded"). Thus, PAT is a "many-to-one" translation scheme. To configure PAT, the syntax is:

- Router(config)#ip nat inside source list 1 interface serial 0/0 overload

The translation tells the router that if a packet with source address matching a permit in ACL 1 hits the inside interface, and it is bound for the outside interface, translate the source address to the address of the Serial0/0 interface. Thus, all translated traffic has the same source address, and no pool is required.

You may ask what happens when the return traffic hits the router if it's all destined for the same address? The key to PAT is that the port numbers are also tracked and, if necessary, manipulated.

Remember that when an application using TCP or UDP starts, it is assigned a port number by the IP stack. Specifically, server-side apps are assigned well-known ports below 1024 (for example, TCP 23 for Telnet, TCP 80 for HTTP, UDP 69 for TFTP). Client-side apps are assigned random port numbers in the range 1024 and above.

Let's say that host 10.0.0.1 (the client) initiates a Telnet session with a host at address 200.1.2.3 (the server). The client process on host 10.0.0.1 will be assigned a TCP port by host 10.0.0.1's IP stack, which we'll assume is 2000 (and, of course, the Telnet server at 200.1.2.3 is using TCP port 23). When the traffic from 10.0.0.1 hits the inside interface and is bound for the outside interface, it's checked against the ACL.

Since no corresponding entry yet exists in the translation table, the inside local address and port number will be entered (10.0.0.1:2000). For the corresponding inside global address, the Serial 0/0 address will be used (let's assume it's 123.4.5.6), and the port number will stay the same (123.4.5.6:2000). If the port number is already in use by another host, the port will be changed to a value that is not already in use (the algorithm for this is implementation-specific).

Timeouts are used to free up unused translation entries for ICMP (60 seconds by default) and UDP (300 seconds by default), and TCP has a fail-safe timeout of 24 hours.

Finally, what if you want to use a pool (dynamic NAT), but switch to PAT if the addresses in the available pool addresses are exhausted? In this case, you combine the pool and overload options.

- Router(config)#ip nat inside source list 1 pool test overload

The effect of this is implementation-specific, but in my experience, a Cisco router will allocate the pool addresses in ascending order and then overload on the last address, if necessary.

Summary

This has been a brief overview of three address translation options: Static NAT, Dynamic NAT, and PAT. Each has its own strengths and weaknesses, and can perform admirably in the right circumstances. Hopefully this white paper will help you decide when to use each method in your networks.

Copyright ©2010 Global Knowledge Training LLC. All rights reserved.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge. Check out the following Global Knowledge course(s):

CCNA Boot Camp v2.0

ICND2 – Interconnecting Cisco Network Devices 2

TCP/IP Networking

For more information or to register, visit **www.globalknowledge.com** or call **1-800-COURSES** to speak with a sales representative.

Our courses and enhanced, hands-on labs and exercises offer practical skills and tips that you can immediately put to use. Our expert instructors draw upon their experiences to help you understand key concepts and how to apply them to your specific work situation. Choose from our more than 1,200 courses, delivered through Classrooms, e-Learning, and On-site sessions, to meet your IT and business training needs.

About the Author

Al Friebe (CCDA, CCDP, CCNA, CCNP, A+, CCSI) has taught networking classes since 1995. He previously served as Global Knowledge's Course Director for BGP and BSCI, and he is the author of our current ICND2 labs. His previous experience includes instructor duty in the U.S. Navy's Nuclear Power School, radio-chemistry, software engineering, and network management.

Copyright ©2010 Global Knowledge Training LLC. All rights reserved.