

Network Address Translation (NAT) FAQ

Document ID: 26704

Questions

Introduction

Generic NAT

Voice-NAT

NAT with VRF/MPLS

NAT NVI

SNAT

NAT-PT (v6 to v4)

Platform-Dependent Cisco 7300/7600/6k

Platform-Dependent Cisco 850

NAT Deployment

NAT Best Practices

Related Information

Introduction

This document provides answers to frequently asked questions about Network Address Translation (NAT).

Generic NAT

Q. What is NAT?

A. Network Address Translation (NAT) is designed for IP address conservation. It enables private IP networks that use unregistered IP addresses to connect to the Internet. NAT operates on a router, usually connecting two networks together, and translates the private (not globally unique) addresses in the internal network into legal addresses, before packets are forwarded to another network.

As part of this capability, NAT can be configured to advertise only one address for the entire network to the outside world. This provides additional security by effectively hiding the entire internal network behind that address. NAT offers the dual functions of security and address conservation and is typically implemented in remote-access environments.

Q. How does NAT work?

A. Basically, NAT allows a single device, such as a router, to act as an agent between the Internet (or public network) and a local network (or private network), which means that only a single unique IP address is required to represent an entire group of computers to anything outside their network. Refer to How NAT Works for more information.

Q. How do I configure NAT?

A. In order to configure traditional NAT, you need to make at least one interface on a router (NAT outside) and another interface on the router (NAT inside) and a set of rules for translating the IP addresses in the packet headers (and payloads if desired) need to be

configured. In order to configure Nat Virtual Interface (NVI), you need at least one interface configured with NAT enable along with the same set of rules as mentioned above.

For more information, refer to Cisco IOS IP Addressing Services Configuration Guide or Configuring the NAT Virtual Interface.

Q. What are the main differences between the Cisco IOS[®] Software and Cisco PIX[®] Security Appliance implementations of NAT?

A. Cisco IOS software–based NAT is not fundamentally different from the NAT function in the Cisco PIX Security Appliance. The main differences include the different traffic types supported in the implementations. Refer to Cisco PIX 500 Series Security Appliances and NAT Configuration Examples for more information on the configuration of NAT on Cisco PIX devices (includes the traffic types supported).

Q. On which Cisco routing hardware is Cisco IOS NAT available? How can the hardware be ordered?

A. The Cisco Feature Navigator tool allows customers to identify a feature (NAT) and find on which release and hardware version this Cisco IOS Software feature is available. Refer to Cisco Feature Navigator in order to use this tool.

Q. Does NAT occur before or after routing?

A. The order in which the transactions are processed using NAT is based on whether a packet is going from the inside network to the outside network or from the outside network to the inside network. Inside to outside translation occurs after routing, and outside to inside translation occurs before routing. Refer to NAT Order of Operation for more information.

Q. Can NAT be deployed in a public wireless LAN environment?

A. Yes. The NAT – Static IP Support feature provides support for users with static IP addresses, enabling those users to establish an IP session in a public wireless LAN environment. Refer to NAT – Static IP Support for more information about this feature.

Q. Does NAT do TCP load–balancing for Servers on the internal network?

A. Yes. Using NAT, you can establish a virtual host on the inside network that coordinates load sharing among real hosts. Refer to Avoiding Server Overload Using TCP Load Balancing for more information.

Q. Can I rate limit the number of NAT translations?

A. Yes. The Rate–Limiting NAT Translation feature provides the ability to limit the maximum number of concurrent NAT operations on a router. In addition to giving users more control over how NAT addresses are used, the Rate–Limiting NAT Translation feature can be used to limit the effects of viruses, worms, and denial–of–service attacks. Refer to Rate Limiting NAT Translation for more information.

Q. How is routing learned or propagated for IP subnets or addresses that are used by NAT?

A. Routing for IP addresses created by NAT is learned if:

- ◆ The inside global address pool is derived from the subnet of a next-hop router.
- ◆ Static route entry is configured in the next-hop router and redistributed within the routing network.

When the inside global address is matched with the local interface, NAT installs an IP alias and an ARP entry, in which case the router will **proxy-arp** for these addresses. If this behavior is not wanted, use the *no-alias* keyword.

When a NAT pool is configured, the *add-route* option can be used for automatic route injection.

Q. How many concurrent NAT sessions are supported in Cisco IOS NAT?

A. The NAT session limit is bounded by the amount of available DRAM in the router. Each NAT translation consumes about 312 bytes in DRAM. As a result, 10,000 translations (more than would generally be handled on a single router) consume about 3 MB. Therefore, typical routing hardware has more than enough memory to support thousands of NAT translations.

Q. What kind of routing performance can be expected when using Cisco IOS NAT?

A. Cisco IOS NAT supports Cisco Express Forwarding switching, fast switching, and process switching. For 12.4T release and later, fast-switching path is no longer supported. For Cat6k platform, the switching order is Netflow (HW switching path), CEF, process path.

Performance depends on several factors:

- ◆ The type of application and its type of traffic
- ◆ Whether IP addresses are embedded
- ◆ Exchange and inspection of multiple messages
- ◆ Source port required
- ◆ The number of translations
- ◆ Other applications running at the time
- ◆ The type of hardware and processor

Q. Can Cisco IOS NAT be applied to subinterfaces?

A. Yes. Source and/or destination NAT translations can be applied to any interface or subinterfaces having an IP address (including dialer interfaces). NAT cannot be configured with Wireless Virtual Interface. Wireless Virtual Interface does not exist at the time of writing to NVRAM. Thus, after reboot, the router loses NAT configuration on the Wireless Virtual Interface.

Q. Can Cisco IOS NAT be used with Hot Standby Router Protocol (HSRP) to provide redundant links to an ISP?

A. Yes. NAT does provide HSRP redundant. However, it is different from SNAT (Stateful NAT). NAT with HSRP is a stateless system. The current session is not maintained when failure takes place. During static NAT configuration (when a packet does not match any STATIC rule configuration), the packet is sent through without any translation. Refer to NAT – Static Mapping Support with HSRP for High Availability for more information about this feature.

Q. Does Cisco IOS NAT support inbound translations on a Frame Relay interface? Does it support outbound translations on the Ethernet side?

A. Yes. Encapsulation does not matter for NAT. NAT can be done where there is an IP address on an interface and the interface is NAT inside or NAT outside. There must be an inside and an outside for NAT to function. If you use NVI, there must be at least one NAT enabled interface. See How do I configure NAT? for more details.

Q. Can a single NAT-enabled router allow some users to use NAT and other users on the same Ethernet interface to continue to use their own IP addresses?

A. Yes. This can be accomplished through the use of an access list describing the set of hosts or networks that require NAT. All sessions on the same host will be either translated or will pass through the router and not be translated.

Access lists, extended access lists, and route maps can be used to define *rules* by which IP devices get translated. The network address and appropriate subnet mask should always be specified. The keyword *any* should not be used in place of the network address or subnet mask (see NAT FAQ, Best Practices and Deployment Guide for more detail). With Static NAT configuration, when packet doesn't match with any STATIC rule configuration, packet will be sent through without any translation.

Q. When configuring for PAT (overloading), what is the maximum number of translations that can be created per inside global IP address?

A. PAT (overloading) divides the available ports per global IP address into three ranges: 0–511, 512–1023, and 1024–65535. PAT assigns a unique source port for each UDP or TCP session. It attempts to assign the same port value of the original request, but if the original source port has already been used, it starts scanning from the beginning of the particular port range to find the first available port and assigns it to the conversation. There is an exception for 12.2S code base. 12.2S code base uses different port logic, and there is no port reservation.

Q. How does PAT work?

A. PAT works with either one global IP address or multiple addresses.

PAT with One IP Address

Condition	Description
1	NAT/PAT inspects traffic and matches it to a translation rule.

2	Rule matches to a PAT configuration.
3	If PAT knows about the traffic type and if that traffic type has "a set of specific ports or ports it negotiates" that it will use, PAT sets them aside and does not allocate them as unique identifiers.
4	If a session with no special port requirements attempts to connect out, PAT translates the IP source address and checks availability of the originated source port (433, for example). Note: For Transmission Control Protocol (TCP) and User Datagram Protocol (UDP), the ranges are: 1–511, 512–1023, 1024–65535. For Internet Control Message Protocol (ICMP), the first group starts at 0.
5	If the requested source port is available, PAT assigns the source port, and the session continues.
6	If the requested source port is not available, PAT starts searching from the beginning of the relevant group (starting at 1 for TCP or UDP applications, and from 0 for ICMP).
7	If a port is available it is assigned, and the session continues.
8	If no ports are available, the packet is dropped.

PAT with Multiple IP Addresses

Condition	Description
1–7	The first seven conditions are the same as with a single IP address.
8	If no ports are available in the relevant group on the first IP address, NAT moves on to the next IP address in the pool and tries to allocate the original source port requested.
9	If the requested source port is available, NAT assigns the source port and the session continues.
10	If the requested source port is not available, NAT starts searching from the beginning of the relevant group (starting at 1 for TCP or UDP applications, and

	from 0 for ICMP).
11	If a port is available, it is assigned and the session continues.
12	If no ports are available, the packet is dropped, unless another IP address is available in the pool.

Q. What are NAT IP pools?

A. NAT IP pools are a range of IP addresses that are allocated for NAT translation as needed. To define a pool, the configuration command is used:

```
ip nat pool <name> <start-ip> <end-ip>
    {netmask <netmask> | prefix-length <prefix-length>}
    [type {rotary}]
```

Example 1

The following example translates between inside hosts addressed from either the 192.168.1.0 or 192.168.2.0 network to the globally unique 10.69.233.208/28 network:

```
ip nat pool net-208 10.69.233.208 10.69.233.223 prefix-length 28
ip nat inside source list 1 pool net-208
!
interface ethernet 0
ip address 10.69.232.182 255.255.255.240
ip nat outside
!
interface ethernet 1
ip address 192.168.1.94 255.255.255.0
ip nat inside
!
access-list 1 permit 192.168.1.0 0.0.0.255
access-list 1 permit 192.168.2.0 0.0.0.255
```

Example 2

In the following example, the goal is to define a virtual address, connections to which are distributed among a set of real hosts. The pool defines the addresses of the real hosts. The access list defines the virtual address. If a translation does not already exist, TCP packets from serial interface 0 (the outside interface) whose destination matches the access list are translated to an address from the pool.

```
ip nat pool real-hosts 192.168.15.2 192.168.15.15 prefix-length 28 type rotary
ip nat inside destination list 2 pool real-hosts
!
interface serial 0
ip address 192.168.15.129 255.255.255.240
ip nat outside
!
interface ethernet 0
ip address 192.168.15.17 255.255.255.240
ip nat inside
!
access-list 2 permit 192.168.15.1
```

Q. What is the maximum number of configurable NAT IP pools (ip nat pool "name")?

A. In practical use, the maximum number of configurable IP pools is limited by the amount of available DRAM in the particular router. (Cisco recommends that you configure a pool size of 255.) Each pool should be no more than 16 bits. In 12.4(11)T and later, IOS introduces CCE (Common Classification Engine). This has limited NAT to only have a maximum of 255 pools. In 12.2S code base, there is no maximum pools restriction.

Q. What is the advantage of using route-map vs ACL on a NAT pool?

A. A route-map is protecting unwanted outside users to reach to the inside users/servers. It also has the capability to map a single inside IP address to different Inside Global addresses based on the rule. Refer to NAT Support for Multiple Pools Using Route Maps for more information.

Q. What is IP address "overlapping" within the context of NAT?

A. IP address overlapping refers to a situation where two locations that want to interconnect are both using the same IP address scheme. This is not an unusual occurrence; it often happens when companies merge or are acquired. Without special support, the two locations will not be able to connect and establish sessions. The overlapped IP address can be a public address assigned to another company, a private address assigned to another company, or can come from the range of private addresses as defined in RFC 1918.

Private IP addresses are unroutable and require NAT translations to allow connections to the outside world. The solution involves intercepting Domain Name System (DNS) name-query responses from the outside to the inside, setting up a translation for the outside address, and fixing up the DNS response before forwarding it to the inside host. A DNS server is required to be involved on both sides of the NAT device to resolve users wanting to have connection between both networks.

NAT is able to inspect and perform address translation on the contents of DNS *A* and *PTR* records, as shown in Using NAT in Overlapping Networks.

Q. What are static NAT translations?

A. Static NAT translations have one-to-one mapping between local and global addresses. Users can also configure static address translations to the port level, and use the remainder of the IP address for other translations. This typically occurs where you are performing Port Address Translation (PAT).

The following example shows how to configure routemap to allow outside-to-inside translation for static NAT:

```
ip nat inside source static 1.1.1.1 2.2.2.2 route-map R1 reversible
!
ip access-list extended ACL-A
permit ip any 30.1.10.128 0.0.0.127'
route-map R1 permit 10
match ip address ACL-A
```

Q. What is meant by the term NAT *overloading*; is this PAT?

A. Yes. NAT overloading is PAT, which involves using a pool with a range of one or more addresses or using an interface IP address in combination with the port. When you overload, you create a fully extended translation. This is a translation table entry containing IP address and source/destination port information, which is commonly called PAT or overloading.

PAT (or overloading) is a feature of Cisco IOS NAT that is used to translate *internal* (inside local) private addresses to one or more *outside* (inside global, usually registered) IP addresses. Unique source port numbers on each translation are used to distinguish between the conversations.

Q. What are dynamic NAT translations?

A. In dynamic NAT translations, the users can establish dynamic mapping between local and global addresses. Dynamic mapping is accomplished by defining the local addresses to be translated and the pool of addresses or interface IP address from which to allocate global addresses and associating the two.

Q. What is ALG?

A. ALG is an Application Layer Gateway (ALG). NAT performs translation service on any Transmission Control Protocol/User Datagram Protocol (TCP/UDP) traffic that does not carry source and/or destination IP addresses in the application data stream.

These protocols include FTP, HTTP, SKINNY, H232, DNS, RAS, SIP, TFTP, telnet,archie, finger, NTP, NFS, rlogin, rsh, rcp. Specific protocols that embed IP address information within the payload require support of an Application Level Gateway (ALG).

Refer to Using Application Level Gateways with NAT for more information.

Q. Is it possible to build a configuration with both static and dynamic NAT translations?

A. Yes. However, the same IP address cannot be used for the NAT static configuration or in the pool for NAT dynamic configuration. All the public IP addresses need to be unique. Note that the global addresses used in static translations are not automatically excluded with dynamic pools containing those same global addresses. Dynamic pools must be created to exclude addresses assigned by static entries. For more information, refer to Configuring Static and Dynamic NAT Simultaneously.

Q. When a traceroute is done through a NAT router, should traceroute show the NAT-Global address or should it leak the NAT-Local address?

A. Traceroute from outside should always return the global address.

Q. How does PAT allocate port?

A. NAT introduces additional port features: full-range and port-map.

- ◆ Full-range allows NAT to use all ports regardless of its default port range.
- ◆ Port-map allows NAT to reserve a user define port range for specific application.

Refer to User Defined Source Port Ranges for PAT for more information.

In 12.4(20)T2 onward, NAT introduces port randomization for L3/L4 and symmetric-port.

- ◆ Port randomization allows NAT to randomly select any global port for the source port request.
- ◆ Symmetric-port allows NAT to support *endpoint independent*.

Refer to Anatomy: A Look Inside Network Address Translators for more information.

Q. What is the difference between IP fragmentation and TCP segmentation?

A. IP fragmentation occurs at Layer 3 (IP); TCP segmentation occurs at Layer 4 (TCP). IP fragmentation takes place when packets that are larger than the Maximum Transmission Unit (MTU) of an interface are sent out of this interface. These packets will have to be either fragmented or discarded when they are sent out the interface. If the Don't Fragment (DF) bit is not set in the IP header of the packet, the packet will be fragmented. If the DF bit is set in the IP header of the packet, the packet is dropped and an ICMP error message indicating the next-hop MTU value will be returned to the sender. All the fragments of an IP packet carry the same Ident in the IP header, which allows the final receiver to reassemble the fragments into the original IP packet. Refer to Resolve IP Fragmentation, MTU, MSS, and PMTUD Issues with GRE and IPsec for more information.

TCP segmentation takes place when an application on an end station is sending data. The application data is broken into what TCP considers the best-sized chunks to send. This unit of data passed from TCP to IP is called a segment. TCP segments are sent in IP datagrams. These IP datagrams can then become IP fragments as they pass through the network and encounter lower MTU links than they can fit through.

TCP will first segment this data into TCP segments (based on TCP MSS value) and will add the TCP header and pass this TCP segment to IP. Then IP will add an IP header to send the packet to the remote end host. If the IP packet with the TCP segment is larger than the IP MTU on an outgoing interface on the path between the TCP hosts then IP will fragment the IP/TCP packet in order to fit. These IP packet fragments will be reassembled on the remote host by the IP layer and the complete TCP segment (that was originally sent) will be handed to the TCP layer. The TCP layer has no idea that IP had fragmented the packet during transit.

NAT supports IP fragments, but it does not support TCP segments.

Q. Does NAT support out-of-order for IP fragmentation and TCP segmentation?

A. NAT supports only out-of-order IP fragments because of **ip virtual-reassembly**.

Q. How to debug IP fragmentation and TCP segmentation?

A. NAT uses the same debug CLI for both IP fragmentation and TCP segmentation: **debug ip nat frag**.

Q. Is there a supported NAT MIB?

A. No. There is no supported NAT MIB, including CISCO-IETF-NAT-MIB.

Q. What is *TCP timeout*, and how does it relate to the NAT TCP timer?

A. If the three-way handshake is not completed and NAT sees a TCP packet, then NAT will start a 60-second timer. When the three-way handshake is completed, NAT uses a 24-hour timer for a NAT entry by default. If an end host sends a RESET, NAT changes the default timer from 24 hours to 60 seconds. In the case of FIN, NAT changes the default timer from 24 hours to 60 seconds when it receives FIN and FIN-ACK.

Q. Can I change the amount of time it takes for a NAT translation to time out from the NAT translation table?

A. Yes. You can change the NAT timeout values for all entries or for different types of NAT translations (such as udp-timeout, dns-timeout, tcp-timeout, finrst-timeout, icmp-timeout, pptp-timeout, syn-timeout, port-timeout and arp-ping-timeout). For more information about changing the default values for these timeout settings, refer to ip nat translation (timeout).

Q. How do I stop Lightweight Directory Access Protocol (LDAP) from attaching extra bytes to each LDAP reply packet?

A. The LDAP settings add the extra bytes (LDAP search results) while processing messages of type Search-Res-Entry. LDAP attaches 10 bytes of search results to each of the LDAP reply packet. In the event that this 10 extra bytes of data result in the packet exceeding the Maximum Transmission Unit (MTU) in a network, the packet is dropped. In this case, Cisco recommends that you turn off this LDAP behavior using the CLI **no ip nat service append-ldap-search-res** command in order for the packets to be sent and received.

Q. What is the route recommendation for the inside global/outside local IP address on the NAT box ?

A. A route has to be specified on the NAT configured box for the inside global IP address for features such as NAT-NVI. Similarly, a route should also be specified on the NAT box for the outside local IP address. In this case, any packet from an in to out direction using the outside static rule will require this kind of route. In such scenarios, while providing the route for IG/OL, the next hop IP address should also be configured. If the next hop configuration is missing, this is considered a configuration error and will result in undefined behavior.

NVI-NAT is present in the output feature path only. If you have directly connected subnet with NAT-NVI or the outside NAT translation rule configured on the box, then in those scenarios, you need to provide a dummy Next Hop IP address and also an associated ARP for the Next Hop. This is needed for the underlying infrastructure to hand the packet to NAT for the translation.

Q. Does Cisco IOS NAT support ACLs with a "log" keyword?

A. When you configure Cisco IOS NAT for dynamic NAT translation, an ACL is used to identify packets that can be translated. The current NAT architecture does not support ACLs with a "log" keyword.

Voice–NAT

Q. Does NAT support Skinny Client Control Protocol (SCCP) v17 which is shipped with Cisco Unified Communications Manager (CUCM) V7?

A. CUCM 7 and all of the default phone loads for CUCM 7 support SCCPv17. The SCCP version used is determined by the highest common version between CUCM and the phone when the phone registers.

NAT does not yet support SCCP v17. Until NAT support for SCCP v17 is implemented, the firmware must be downgraded to version 8–3–5 or below so that SCCP v16 is negotiated. CUCM6 will not encounter the NAT problem with any phone load as long as it uses SCCP v16. Cisco IOS does not currently support SCCP version 17.

Q. Which CUCM /SCCP/firmware load versions are supported by NAT?

A. NAT supports CUCM version 6.x and earlier releases. These CUCM versions are released with the default 8.3.x (or earlier) phone firmware load that support SCCP v15 (or earlier).

NAT does not support CUCM versions 7.x or later releases. These CUCM version are released with the default 8.4.x phone firmware load that supports SCCP v17 (or later).

If CUCM 7.x or later is used, an older firmware load must be installed on the CUCM TFTP server so that the phones use a firmware load with SCCP v15 or earlier in order to be supported by NAT.

The link below confirms that firmware load 8.3.x contains SCCP v15 or earlier and will work with NAT and that firmware load 8.4.x contains SCCP v17 and will NOT work with NAT.

<http://third-gen-phones.gforge.cisco.com/twiki/prod/bin/view/Thirdgenphones/CCMLoadNumberAndCodeN>

Refer to NAT–Support of IP Phone to Cisco CallManager for more information about NAT and SCCP.

Q. What is Service Provider PAT Port Allocation Enhancement for RTP and RTCP?

A. The Service Provider PAT Port Allocation Enhancement for RTP and RTCP feature ensures that for SIP, H.323, and Skinny voice calls. The port numbers used for RTP streams are even port numbers, and the RTCP streams are the next subsequent odd port number. The port number is translated to a number within the range specified conforming to RFC–1889. A call with a port number within the range will result in a PAT translation to another port number within this range. Likewise, a PAT translation for a port number outside this range will not result in a translation to a number within the given range.

Refer to Service Provider PAT Port Allocation Enhancement for RTP and RTCP for more information.

Q. What is Session Initiation Protocol (SIP) and can SIP packets be NATted?

A. Session Initiation Protocol (SIP) is an ASCII-based, application-layer control protocol that can be used to establish, maintain, and terminate calls between two or more endpoints. SIP is an alternative protocol developed by the Internet Engineering Task Force (IETF) for multimedia conferencing over IP. The Cisco SIP implementation enables supported Cisco platforms to signal the setup of voice and multimedia calls over IP networks. Refer to Overview of SIP for more information.

SIP packets can be NATted. Refer to NAT Support for SIP and NAT Support of H.323 v2 RAS for more information.

Q. What is Hosted NAT Traversal support for Session Border Controller (SBC)?

A. The Cisco IOS Hosted NAT Traversal for SBC feature enables a Cisco IOS NAT SIP Application-Level Gateway (ALG) router to act as a SBC on a Cisco Multiservice IP-to-IP Gateway, which helps to ensure smooth delivery of voice over IP (VoIP) services.

Refer to Configuring Cisco IOS Hosted NAT Traversal for Session Border Controller and SP Hosted NAT Traversal for SIP Calls Using Cisco IOS Session Border Controller for more information.

Q. How many SIP, Skinny, and H323 calls can a routers memory and CPU handle with NAT?

A. The number of calls handled by a NAT router is contingent on the amount of memory available on the box and the processing power of the CPU.

Q. Does a NAT router support TCP segmentation of Skinny and H323 packets?

A. IOS-NAT support TCP segmentation for H323 in 12.4 Mainline and TCP segmentation support for SKINNY from 12.4(6)T onward.

Q. Are there any caveats to watch out for when using a NAT overload configuration in a voice deployment?

A. Yes. When you have NAT overload configs and a voice deployment, you need the registration message to go through NAT and create an association for out->in to reach this inside device. The inside device sends this registration in a periodic fashion and NAT updates this pin-hole/association from the information as in the signalling message.

Q. Are there any known problems caused by issuing the clear ip nat trans * command or the clear ip nat trans forced command in a voice deployment?

A. In voice deployments when you issue a **clear ip nat trans *** command or a **clear ip nat trans forced** command and have dynamic NAT, you will wipe out the pin-hole/association and must wait for the next registration cycle from the inside device to re-establish this. Cisco recommends that you do not use these clear commands in a voice deployment.

Q. Does NAT support voice co-located solution?

A. No. The co-located solution is currently not supported. The following deployment with NAT (on the same box) is considered a co-located solution: CME/DSP-Farm/SCCP/H323.

Q. Does NVI support Skinny ALG, H323 ALG, and TCP SIP ALG?

A. No. Note that UDP SIP ALG (used by most deployments) is not impacted.

NAT with VRF/MPLS

Q. Will a NAT router ever support NATting the same address space in a VRF as is being NATted in a global address space? Currently, I receive this warning: "% similar static entry (1.1.1.1 ----> 22.2.2.2) already exists" when I attempt to configure the following:

```
72UUT(config)#ip nat inside source static 1.1.1.1 22.2.2.2 72UUT(config)#ip nat inside source s
```

A. Legacy NAT supports overlapping address config over different VRFs. You would have to configure overlapping at rule with the *match-in-vrf* option and set up **ip nat inside/outside** in the same VRF for traffic over that specific VRF. The overlapping support does not include the global routing table.

You must add the *match-in-vrf* keyword for the overlapping VRF static NAT entries for different VRFs. However, it is not possible to overlap global and vrf NAT addresses.

```
72UUT(config)#ip nat inside source static 1.1.1.1 22.2.2.2 vrf RED match-in-vrf
72UUT(config)#ip nat inside source static 1.1.1.1 22.2.2.2 vrf BLUE match-in-vrf
```

Q. Does legacy NAT support VRF-Lite (NATting from a VRF to a different VRF)?

A. No. You must use NVI for NATting between different VRFs. Refer to NAT Virtual Interface for more information and examples. You can use legacy NAT to do NAT from VRF to global or NAT within the same VRF.

NAT NVI

Q. What is NAT NVI?

A. NVI stands for NAT Virtual Interface. It allows NAT to translate between two different VRFs. This solution should be used in lieu of Network Address Translation on a Stick. Refer to NAT Virtual Interface for more information.

Q. Should NAT NVI be used when NATting between an interface in global and an interface in a VRF?

A. Cisco recommends that you use legacy NAT for VRF to global NAT (ip nat inside/out) and between interfaces in the same VRF. NVI is used for NAT between different VRFs.

Q. Is TCP segmentation for NAT–NVI supported?

A. There is no support for TCP segmentation for NAT–NVI.

Q. Does NVI support Skinny ALG, H323 ALG, and TCP SIP ALG?

A. No. Note that UDP SIP ALG (used by most deployments) is not impacted.

Q. Does TCP segmentation supported with SNAT?

A. SNAT does not support any TCP ALGs (such as, SIP, SKINNY, H323, or DNS). Therefore, TCP segmentation is not supported. However, UDP SIP and DNS are supported.

SNAT

Q. What is Stateful NAT (SNAT)?

A. SNAT allows two or more network address translators to function as a translation group. One member of the translation group handles traffic requiring translation of IP address information. Additionally, it informs the backup translator of active flows as they occur. The backup translator can then use information from the active translator to prepare duplicate translation table entries. Therefore, if the active translator is hindered by a critical failure, the traffic can rapidly be switched to the backup. The traffic flow continues since the same network address translations are used and the state of those translations has been previously defined. Refer to Enhanced IP Resiliency Using Cisco Stateful NAT for more information.

Q. Is TCP segmentation supported with SNAT?

A. SNAT does not support any TCP ALGs (such as, SIP, SKINNY, H323, or DNS). Therefore, TCP segmentation is not supported. However, UDP SIP and DNS are supported.

Q. Is SNAT support for asymmetric routing?

A. Asymmetric routing supports NAT by enabling as queuing. By default, as–queueing is enable. However, from 12.4(24)T onward, as–queueing is no longer supported. Customers must make sure packets are routed properly and proper delay is added in order for asymmetric routing to work correctly.

NAT–PT (v6 to v4)

Q. What is NAT–PT?

A. NAT–PT is v4 to v6 translation for NAT. Protocol Translation (NAT–PT) is an IPv6–IPv4 translation mechanism, as defined in RFC 2765 and RFC 2766, allowing IPv6–only devices to communicate with IPv4–only devices and vice versa. Refer to Implementing NAT–PT for IPv6 and Cisco IOS NAT for IPv6 for more information about this feature

Q. Is NAT–PT supported in the Cisco Express Forwarding (CEF) path?

A. NAT–PT is not supported in the CEF path.

Q. What ALGs are supported in NAT–PT?

A. NAT–PT supports TFTP/FTP and DNS. There is no support for voice and SNAT in NAT–PT.

Q. Does ASR 1004 support NAT–PT?

A. Aggregation Services Routers (ASR) uses NAT64. For more information on configuring NAT64, refer to *Configuring a Routing Network for Stateless NAT64*.

Platform–Dependent Cisco 7300/7600/6k

Q. Is Stateful NAT (SNAT) available on Catalyst 6500 on the SX train?

A. SNAT is not available on Catalyst 6500 on the SX train.

Q. Is VRF–aware NAT supported in hardware on the 6k?

A. VRF–aware NAT is not supported in hardware on this platform.

Q. Do the 7600 and Cat6000 support VRF–aware NAT?

A. On the 65xx/76xx platform, VRF–aware NAT is not supported, and the CLIs are blocked.

Note: You can implement a design by leveraging a FWSM that runs in virtual context transparent mode.

Platform–Dependent Cisco 850

Q. Does the Cisco 850 support Skinny NAT ALG in release 12.4T?

A. No. There is no support for Skinny NAT ALG in 12.4T on the 850 series.

NAT Deployment

Q. How do I implement NAT?

A. NAT enables private IP internetworks that use nonregistered IP addresses to connect to the Internet. NAT translates the private (RFC1918) address in the internal network into legal routable addresses before packets are forwarded onto another network.

For more information about implementing NAT, refer to *Configuring NAT for IP Address Conservation*.

Q. How do I implement NAT with voice?

A. The NAT support for voice feature allows SIP embedded messages passing through a router configured with Network Address Translation (NAT) to be translated back to the packet. An application layer gateway (ALG) is used with NAT to translate the voice packets.

For more information about implementing NAT with voice, refer to NAT Support for ALGs.

Q. How do I integration NAT with MPLS VPNs?

A. The NAT integration with MPLS VPNs feature allows multiple MPLS VPNs to be configured on a single device to work together. NAT can differentiate from which MPLS VPN it receives IP traffic even if the MPLS VPNs all use the same IP addressing scheme. This enhancement enables multiple MPLS VPN customers to share services while ensuring that each MPLS VPN is completely separate from the other.

For more information, refer to NAT Integration with MPLS VPNs and Integrating NAT with MPLS VPNs.

Q. Does NAT static mapping support HSRP for high availability?

A. When an Address Resolution Protocol (ARP) query is triggered for an address that is configured with Network Address Translation (NAT) static mapping and owned by the router, NAT responds with the BIA MAC address on the interface to which the ARP is pointing. Two routers act as HSRP active and standby. Their NAT inside interfaces must be enabled and configured to belong to a group.

For more information, refer to NAT – Static Mapping Support with HSRP for High Availability.

Q. How do I implement NAT NVI?

A. The NAT virtual interface (NVI) feature removes the requirement to configure an interface as either NAT inside or NAT outside. For more information about NAT NVI, refer to Configuring the NAT Virtual Interface.

Q. How do I implement load balancing with NAT?

A. There are two kinds of load balancing that can be done with NAT: you can load balance inbound to a set of servers to distribute the load on the servers, and you can load balance your user traffic to the Internet over two or more ISPs.

For more information about inbound load balancing, refer to Avoiding Server Overload Using TCP Load Balancing.

For more information about outbound load balancing, refer to IOS NAT Load-Balancing for Two ISP Connections.

Q. How do I implement NAT in conjunction with IPSec?

A. There is support for IP Security (IPSec) Encapsulating Security Payload (ESP) through NAT and IPSec NAT Transparency.

The IPsec ESP through NAT feature provides the ability to support multiple concurrent IPsec ESP tunnels or connections through a Cisco IOS NAT device configured in overload or Port Address Translation (PAT) mode. For more information about this feature, refer to Support for IPsec ESP Through NAT and NAT Support for IPsec ESP – Phase II.

The IPsec NAT transparency feature introduces support for IPsec traffic to travel through NAT or PAT points in the network by addressing many known incompatibilities between NAT and IPsec. For more information about this feature, refer to IPsec NAT Transparency.

Q. How do I implement NAT-PT?

A. NAT-PT (Network Address Translation Protocol Translation) is an IPv6-IPv4 translation mechanism, as defined in RFC 2765 and RFC 2766, that allows IPv6-only devices to communicate with IPv4-only devices and vice versa.

For more information about implementing and configuring NAT-PT, refer to Implementing NAT-PT for IPv6.

Q. How do I implement multicast NAT?

A. It is possible to NAT the source IP for a multicast stream. A route-map can not be used when doing dynamic NAT for multicast, only an access list is supported for this.

For more information, refer to How Does Multicast NAT Work on Cisco Routers. The destination multicast group is NATted using a Multicast Service Reflection solution.

Q. How do I implement stateful NAT (SNAT)?

A. SNAT enables continuous service for dynamically mapped NAT sessions. Sessions that are statically defined receive the benefit of redundancy without the need for SNAT. In the absence of SNAT, sessions that use dynamic NAT mappings would be severed in the event of a critical failure and would have to be reestablished. Only the minimal SNAT configuration is supported. Future deployments should be performed only after talking to your Cisco Account Team in order to validate the design relative to current restrictions.

For more information about implementing SNAT, refer to Configuring NAT for High Availability.

SNAT is recommended for the following scenarios:

- ◆ HSRP mode as described in the SNAT white-paper: Enhanced IP Resiliency Using Cisco Stateful NAT.
- ◆ Primary/backup is not a recommended mode since there are some features missing compared to HSRP.
- ◆ For fail-over scenarios and for 2-router setup. That is, if one router crashes, the other router takes over seamlessly. (SNAT architecture is not designed to handle Interface-flaps.)
- ◆ Non-asymmetric routing scenario is supported. Asymmetric routing can be handled only if the latency in the reply packet is higher than that between 2 SNAT routers to exchange the SNAT messages.

Currently SNAT architecture is not designed to handle robustness; therefore, these tests are not expected to succeed:

- ◆ Clearing NAT entries while there is traffic.
- ◆ Changing interface parameters (like IP address change, shut/no-shut, etc.) while there is traffic.
- ◆ SNAT specific **clear** or **show** commands are not expected to execute properly and not recommended.

Some of the SNAT related **clear** and **show** commands are as follows:

```
clear ip snat sessions *
clear ip snat sessions <ip address of the peer>
clear ip snat translation distributed *
clear ip snat translation peer < IP address of SNAT peer>
sh ip snat distributed verbose
sh ip snat peer < IP address of peer>
```

- ◆ If the user wants to clear entries, **clear ip nat trans forced** or **clear ip nat trans *** commands can be used.

If the user wants to view entries, **show ip nat translation**, **show ip nat translations verbose**, and **show ip nat stats** commands can be used. If *service internal* is configured, it will show SNAT specific information as well.

- ◆ Clearing NAT translations at the back up router is not recommended. Always clear the NAT entries on the primary SNAT router.
- ◆ SNAT is not HA; therefore, configurations on both routers should be the same. Both routers should have the same image running. Also make sure that the underlying platform used for both the SNAT routers are the same.

NAT Best Practices

Q. Are there any NAT best practices?

A. Yes. These are the NAT best practices:

1. When using both dynamic and static NAT, the ACL that sets the rule for dynamic NAT should exclude the static local hosts so there is no overlap.
2. Beware of using ACL for NAT with **permit ip any any** as you can get unpredictable results. After 12.4(20)T NAT will translate locally generated HSRP and routing protocol packets if they are sent out the outside interface, as well as locally encrypted packets matching the NAT rule.
3. When you have overlapping networks for NAT, use the **match-in-vrf** keyword.

You must add the **match-in-vrf** keyword for the overlapping VRF static NAT entries for different VRFs, but it is not possible to overlap global and vrf NAT addresses.

```
Router(config)#ip nat inside source static 1.1.1.1 22.2.2.2 vrf RED match-in-
```

```
Router(config)#ip nat inside source static 1.1.1.1 22.2.2.2 vrf BLUE match-in-
```

4. NAT pools with same address range can not be used in different VRFs unless the **match-in-vrf** keyword is used.

For example:

```
ip nat pool poolA 171.1.1.1 171.1.1.10 prefix-length 24
ip nat pool poolB 171.1.1.1 171.1.1.10 prefix-length 24
ip nat inside source list 1 poolA vrf A match-in-vrf
ip nat inside source list 2 poolB vrf B match-in-vrf
```

Note: Wven though CLI configuration is valid, without the **match-in-vrf** keyword the configuration is not supported.

5. When deploying ISPs load balancing with NAT interface overload, the best practice is to use route-map with interface match over ACL matching.
6. When using pool mapping, you should not use two different mapping (ACL or route-map) to share the same NAT pool address.
7. When deploying the same NAT rules on two different routers in the failover scenario, you should use HSRP redundancy.

Related Information

- [IP Routing Technology Support](#)
 - [Technical Support & Documentation – Cisco Systems](#)
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2010 – 2011 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Feb 17, 2010

Document ID: 26704
