Cisco Network Address Translation (NAT)

Introduction

IETF NGTrans working group defined several translation mechanisms to enable communications between IPv6-only and IPv4-only hosts. One such example is Network Address Translator-Protocol Translator (NAT-PT)—RFC 2766; network administrators already familiar with NAT may find it useful to insure the co-existence between hosts when native communication cannot be achieved. The application of each area must be well understood, as the protocol does not represent a generic mechanism that would be universally applicable.

Since IPv6 deployment will be a gradual process, there will be a transitional period, during which IPv6 hosts will need to communicate with the global Internet, which currently has majority of IPv4 hosts. Simply stated, IPv4 and IPv6 nodes will need to coexist and communicate during the lengthy transition. A strong set of flexible IPv4-to-IPv6 transition and coexistence mechanisms will be required during this period. In these environments, NAT-PT is the translator that provides the solution.

NAT-PT is an interoperability solution that does not require any modifications or extra software, such as dual stacks, to be installed on any end user host of either IPv4 or IPv6 networks. It performs the required interoperability functions within a stub network, making interoperability between hosts easier to manage and faster to deploy.

NAT-PT provides IPv4/IPv6 protocol translation. It resides within an IP router, situated at the boundary of an IPv4 network and an IPv6 network. By installing NAT-PT between an IPv4 and IPv6 network, all IPv4 users are given access to the IPv6 network without modification in the local IPv4-hosts (and vice versa). Equally, all hosts on the IPv6 network are given access to the IPv4 hosts without modification to the local IPv6-hosts. This is accomplished with a pool of IPv4 addresses for assignment to IPv6 nodes on a dynamic basis as sessions are initiated across IPv4-IPv6 boundaries. However, NAT-PT requires tracking the sessions it supports and mandates that inbound and outbound datagrams pertaining to a session traverse the same NAT-PT router. In that sense, NAT-PT inherits many NAT restrictions. Suppose that some applications carry network addresses in payloads. NAT-PT can be application unaware, because it does not snoop the payload. NAT-PT requires some Application Level Gateway (ALG), an application specific agent that allows an IPv6 application to communicate with an IPv4 application and vice versa. ALG could work in conjunction with NAT-PT to provide support for such applications.

Deployment Areas

NAT-PT will likely be deployed in these networking environments:

- **Enterprises:** Enterprise customers may maintain IPv4-only hosts for several years before upgrading to IPv6 (e.g. IPv4-only Token-Ring PC communicating with IPv6-only PDA devices), depending on their applications. At the same time, it is expected that IPv6-only hosts will be configured, even on capable of Dual Stack Operating System, to get benefit of the IPv6 auto-configuration, global addressing and simpler management (no dual-protocol).
- **ISPs:** ISP offering IPv4 and IPv6 connectivity and services will look at providing communications between both protocol versions, i.e.: an IPv6-only client accessing an IPv4-only web server. NAT-PT is an added-value capability in an ISP POP site. In an IPv4-only legacy server farm, the addition of NAT-PT routers would open it to the new IPv6 world at a low cost.
- **Homes:** New IPv6-only Internet appliances (i.e.: mobile phones, PDAs, refrigerators, and audio components) will become available in the future. Home users will not focus on the IP version supported by these classes of product. ISPs can seize this opportunity by offering a NAT-PT service, and thus decreasing the complexity for home users.

NAT-PT Architecture

With NAT-PT, sessions can be initiated from IPv4 and IPv6 hosts, because bi-directional connections are established. IPv6 addresses are mapped to IPv4 addresses and vice and versa, statically or dynamically. The source IP address and related fields (i.e.: IP headers and ICMP headers) are translated for those packets leave one realm for another.

In the vast majority of the cases, DNS is used for IPv4 to IPv6 and IPv6 to IPv4 communications, as it is the common way to access an IP host. As a result, DNS requests cross the NAT-PT box. A DNS-ALG must be implemented in NAT-PT routers to facilitate name to address mapping. The DNS-ALG is capable of translating IPv6 addresses in DNS queries and responses into their IPv4 address bindings (and vice versa), as DNS packets traverse between IPv6 and IPv4 realms.

NAT-PT Operations

This section details all the steps of a data communication crossing a NAT-PT router in the following environment (Figure 1).

Figure 1. A Basic NAT-PT Topology

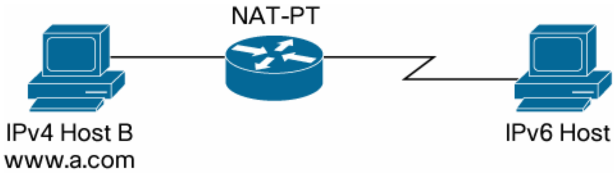

NAT-PT Communications with Static Mapping

IPv6 Initiated Communications

In the case of an IPv6 initiated communication between Hosts A and B, the data communication follows the sequence depicted in Figure 2. Storing static IPv6 to IPv4 address correspondences in the NAT-PT router configuration creates static mapping.

Figure 2. IPv6 Initiated Communication with Static Mapping

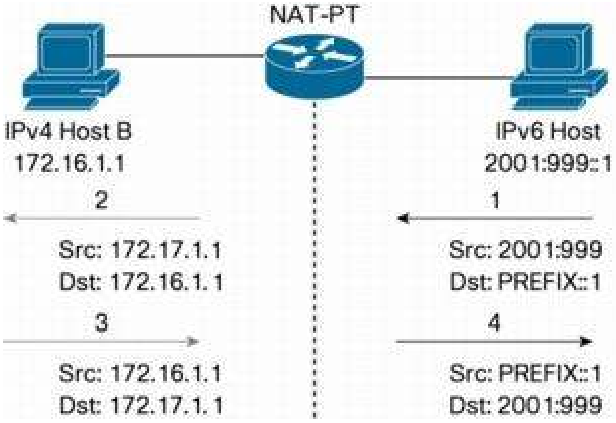

Step 1. A packet leaves Host A for its destination of Host B. The source address of the packet is its actual IPv6 address (2001:9999:: 1). The destination IPv6 address is in the range of a /96 prefix arbitrarily assigned on the NAT-PT router as the destination IPv6 address for IPv6 to IPv4 communications. The prefix length is 96 bits long as it spares 32 bits for addressing the entire IPv4 Internet. In our case, destination address is "PREFIX::1" as it exists a static translation for this IPv6 address in the router.

Step 2. The packet reaches the NAT-PT router and the source and destination address of this IPv6 incoming packet are translated following the static mapping available in the router. An IPv4 packet is issued from the router with the following IPv4 addresses couple (source 172.17.1.1, destination 172.16.1.1).

Step 3. Packets originated from Host B to Host A swap IPv4 source and destination addresses.

Step 4. Using the same IPv6-IPv4 mapping as in Step 1, source and destination IPv4 addresses are translated into IPv6 addresses.

IPv4 Initiated Communications

In the case of an IPv4 initiated communication, the Step 1 (above) is modified:

Step 1. A packet leaves Host A for its destination of Host B. The source address of the packet is its actual IPv4 address (172.16.1.1). The destination IPv4 address is falling into an IPv4 subnet configured on the NAT-PT router. This IPv4 subnet configured on the router is receiving the traffic targeted to the attached IPv6 cloud. In our case destination address is 172.17.1.1 and it exists a static translation to IPv6 for this IPv4 address in the router.

The rest of the process remains unchanged.

NAT-PT Communications with DNS-ALG

This paragraph describes a real-life communication through a NAT-PT router. It will clarify the need for a DNS-ALG, as hosts are commonly accessed by name and not by IP address.

IPv6 Initiated Communications

IPv6 Host A wants to communicate with IPv4 Host B using its name instead of its IP address. The first step of the communication is to resolve the name (www.a.com) of Host B. The DNS requests flow is following the steps depicted in Figure 3.

Figure 3. DNS Requests Flow in a NAT-PT for IPv6 Initiated Communications.

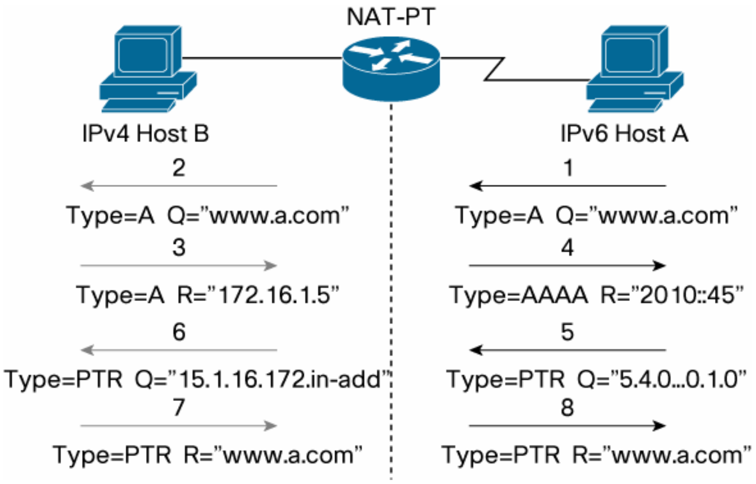

Step 1. An IPv4 server in the IPv4 cloud maintains the name/IPv4 address mapping for www.a.com. Host A sends the AAAA query with the name to resolve (www.a.com) in the query. The destination address of this packet is an IPv6 address part of the /96 prefix setup for IPv4 Internet traffic. The mapped destination IPv6 address to target for this DNS query is statically defined in Host A.

Step 2. In the NAT-PT router, the AAAA request is translated into an AAAA request and an additional A request with the same query string "www.a.com". A static mapping exists for the IPv4 DNS server address. The source IPv6 address of the packet (Host A IPv6 address) is statically or dynamically translated into an IPv4 address.

Step 3. The IPv4 DNS server answers the query with the IPv4 address corresponding to "www.a.com". From the content of the DNS response packet, a dynamic translation is created in the router (if not static mapping exists) between the response IPv4 address (172.16.1.5) and an address of the /96 prefix. This translation will be used later by the NAT-PT router to translate the incoming IPv6 packet to host B into host B IPv4 address.

Step 4. The NAT-PT router changes the type of DNS response from A to AAAA, and modifies the IPv6 address in the payload of the response according the aforementioned dynamic translation. With the result of this DNS request, Host A can send a packet to Host B using a translated IPv6 address. When the IPv6 packet targeted to Host B reaches the router, the previously created dynamic translation will translate this destination IPv6 address into the real Host B IPv4 address. The destination IPv4 address of this DNS response is translated into an IPv4 address following Step 2.

Step 5. For security reasons, servers like FTP servers check that a pointer record exists for the IP address of the incoming connection, and then a NAT-PT router must be able to perform PTR record translation.

In our case, the mapped IPv6 address of "www.a.com" is known. Then to check a PTR record exists for this address; host A creates an IPv6 PTR query for "2010::45".This packet is sent to the same IPv6 address as in Step 1.

Step 6. In the NAT-PT router, the IPv4 to IPv6 translation for Host B (Step 3) translates the IPv6 PTR query into an IPv4 PTR query. If this translation had not already been created, it would have been impossible to translate the IPv6 PTR query into an IPv4 PTR query, because the NAT-PT router would not have known the target IPv4 address. The source IPv6 address of the DNS server is translated as in Step 2.

Step 7. The IPv4 DNS answers the PTR query with "www.a.com".

Step 8. The NAT-PT router translates IPv4 PTR response into an IPv6 PTR response. IPv4 header IP addresses are translated following the dynamic translation created in Step 3.

In this scenario we are making the assumption that DNS servers located inside an IPv4 cloud are maintaining IPv4 reverse zones. It is technically possible to maintain a zone for IPv4 hosts, referencing IPv6 NAT-PTed address, in a DNS hosted in an IPv6 cloud. This would alleviate the need for a DNS-ALG, because all the mappings would reside in the IPv6 zone description.

IPv4 Initiated Communications

In the case of an IPv4 initiated communication, the process is symmetrical, given the modification described above for IPv4 initiated communication in a case of a static mapping.

In this scenario, we are making the assumption that DNS servers located inside an IPv6 cloud are maintaining IPv6 zone and reverse zone. It is technically possible to maintain a zone for IPv6 hosts, referencing IPv4 NAT-PTed address, in a DNS hosted in an IPv4 cloud. This would alleviate the need for a DNS-ALG because all the mappings would reside in the IPv4 zone description.

NAT-PT Communications with Dynamic Mapping

IPv6 Initiated Communications

In the case of an IPv6 initiated communication with dynamic mapping between Host A and Host B, the data communication is following the steps depicted in Figure 2. Unlike static mapping, dynamic mapping does not create a hard-coded one-to-one mapping between an IPv6 address and an IPv4 address. A pool of IPv4 addresses is created and the IPv4 to IPv6 mappings are allocated, as needed, pulling IP address from a pre-defined IPv4 pool. The IPv6 packet crossing the NAT-PT box is following the above steps:

Step 1. A packet leaves Host A for its destination of Host B. The source address of the packet is its actual IPv6 address (2001:9999:: 1). The destination IPv6 address is in the range of a /96 prefix arbitrarily assigned on the NAT-PT router as the destination for IPv6 to IPv4 communications. The prefix length is 96 bits long, as it spares 32 bits for addressing the entire IPv4 Internet.

Step 2. The packet reaches the NAT-PT router. The source address of this IPv6 incoming packet is translated by pulling an IPv4 address from a locally defined IPv4 pool. This dynamic translation mechanism identifies source IPv6 addresses to be translated from a prefix-list, IPv6 access-list or route-map. The destination IPv6 address is translated either statically or from a previously created DNS-ALG dynamically translation. An IPv4 packet is issued from the router.

Step 3. Packets sent from Host B to Host A swap IPv4 source and destination addresses.

Step 4. The destination IPv4 address of the packet is translated into an IPv6 address along the dynamic mapping created in Step 2. As in Step 2, the incoming IPv4 source address is translated using static or previously created DNS-ALG translation.

IPv4 Initiated Communications

In the case of an IPv4 initiated communication, the process is symmetrical. A pool of IPv6 address is formed for translated source address instead. Static mapping or DNS-ALG created translations for destination address are "IPv4 to IPv6".

Auto-Aliasing of NAT-PT IPv4 Pool Addresses

In order to save as many IPv4 addresses as possible (in this time of scarcity!), NAT-PT makes it possible to define a pool of IPv4 address as an attached interface IPv4 subnet. In this case, only the subnet unused IPv4 address will be translated. The NAT-PT router answers ARP requests for those addresses; therefore, the router accepts all packets destined for this IPv4 address and translates them.

NAT-PT Translations and Application Level Gateway

IP Header Translation

From IPv6 to IPv4

The Protocol Translator translates the IPv6 header to an IPv4 header under the following conditions:

- IPv6 packet is received with an IPv4-mapped IPv6 address (i.e. pre-configured /96 prefix)
- Translation exists for the incoming packet

The payload is untouched, with the exception of ICMP and DNS packets. The resulting IPv4 packet is routed into an IPv4 cloud.

If there is no IPv6 fragment header (fragmented packets are not supported in this release), the IPv6 header is translated into an IPv4 header (Table 1):

Table 1. IPv6 to IPv4 Header Translation

IPv4 Field	Operation
TOS	Copied from "Traffic Class"
Total Length	Payload length + IP header length
Identification	Set to 0
Flag	MF=0/DF=1
Fragment Offset	0
Header Checksum	Recalculated after building IPv4 packet
TTL	Copied from "Hop Limit"

From IPv4 to IPv6

When NAT-PT receives a packet addressed to a destination that lies outside of the attached IPv4 realm, the IPv4 header is translated to an IPv6 header.

If there are no IPv4 fragments in the header (fragmented packets are not supported in this release), the IPv4 header is translated into an IPv6 header (Table 2):

Table 2. IPv4 to IPv6 Header Translation

IPv6 Field	Operation
Traffic Class	Copied from "TOS"
Flow Label	Set to 0

IPv6 Field	Operation
Payload Length	Total Length—(IPv4 Header + options)
Next Header	Copied from "Protocol"
Hop Limit	Copied from "TTL"
Header Checksum	Recalculated after building IPv6 packet

ICMP Messages Translation

There are some differences between IPv4 and IPv6 ICMP error message formats [SIIT]. The NAT-PT device must then convert the error messages appropriately. Note that ICMP packets require the contents of ICMP error message to be translated and the ICMP pseudo-header checksum to be re-calculated. ICMP-ALG performs this task as part of NAT-PT.

The following scenario (Figure 4) explains the ICMP message translation in the case of "Destination unreachable".

Figure 4. A Simple NAT-PT Setup where ICMP Packet Translations Occurs

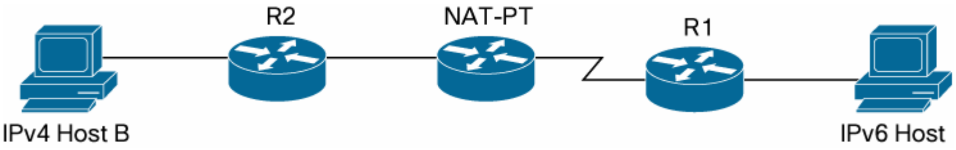

Step 1. A packet is sent from Host B to Host A. For some reason, Router R1 receives the packet and cannot route it to the destination Node A.

Step 2. This will trigger Router R1 to generate an ICMP message (Destination Unreachable) back to the sender.

R1 generates the following IPv6 ICMP error message:

- Type = Destination Unreachable
- Destination = IPv6 translated host B address
- Source = Router R1 IPv6 address
- Contents will contain the IPv6 address of host A that was not reachable

The allocation of IPv6 address for Host B would have happened when the first packet from Host B to Host A goes through the NAT-PT device.

Step 3. The IPv6 ICMP error message reaches the NAT-PT device and is translated.

The following IPv6 ICMP error message is generated:

- Type = Destination Unreachable
- Destination= host B IPv4 address
- Source= Router R1 translated IPv4 address
- Contents will also be changed to contain the IPv4 address of host A

The allocation of IPv4 address for Host A would have happened when the first packet from Host B to Host A goes through the NAT-PT device. The allocation of IPv4 address for Router R1 occurs when the ICMP message crosses the NAT-PT router (if it does not already exist).

Step 4. The IPv4 ICMP error message reaches Host B.

From a more general viewpoint, translations do exist from IPv4 ICMP messages to IPv6 ICMP messages and vice-and-versa [SIIT]. However, not all IPv6 or IPv4 ICMP message can be translated, as a concept present in one protocol may not exist in the other protocol. If a translation is not possible, then the message is silently dropped.

DNS ALG

DNS uses AAAA records for IPv6 name-to-address mappings, which is different from A records in IPv4 DNS. NAT-PT DNS ALG needs to convert the DNS query/reply packets from AAAA records to A records and vice versa.

DNS ALG supports the following translation on the DNS query/reply:

1. For node name to address queries, it will change the query type from A to AAAA and vice versa.
2. For node address to name queries, replace the string IN-ADDR.ARPA with IP6.INT. Replace the IPv4 address octets with the corresponding IPv6 address octets (in reverse order).
3. For name to address replies, replace the IPv6 address with the IPv4 address by allocating an address. NAT-PT sets the TTL values to 0 on all resource records (RRs) passing through NAT-PT, so that clients and servers do not cache the temporarily assigned RRs.
4. When a NAT-PT receives a AAAA query, the DNS-ALG on the NAT-PT forwards the query unchanged and also forms another DNS query with the rules specified as above: sends an A query for the same node, since NAT-PT does not know whether the destination node is an IPv4 or IPv6 node. If the reply is an AAAA record, NAT-PT forwards it unchanged otherwise it converts the A record to AAAA record and forwards the packet back to IPv6 domain.

The following DNS record types are only supported and the embedded IP address contained are translated when necessary:

- A queries from IPv4 hosts only
- AAAA queries from IPv6 hosts only
- A/AAAA replies from IPv4 DNS server
- AAAA replies only from IPv6 DNS server
- PTR, CNAME, ANY queries from both IPv4/IPv6 hosts
- PTR, CNAME, ANY replies from both IPv4/IPv6 DNS server

Command Line Interface

Cisco NAT-PT CLI shares similarities with Cisco IOS NAT configuration. The network administrator should not need intensive training to configure NAT-PT.

Cisco IOS NAT-PT Implementation and Switching Services

Cisco IOS Software Release 12.2(13)T provides a process-switched implementation of the feature. Later revisions of NAT-PT will enable CEF/dCEF switching. The packets that need only the header translations will be processed in CEF/dCEF switching mode. Like the IPv4 NAT, NAT-PT will do all the payload translations in process switching mode.

Global Configuration Commands

IPv6 NAT Prefix

This global command creates a /96 prefix to receive all IPv6 traffic targeted to NAT-PT at the router level. Only a /96 prefix is permitted. It would not make sense to offer more than a /96, because the IPv4 address space is 32 bits. The same command exists at the interface level with an interface scope. Alternatively, when a packet moves from IPv4 to IPv6, NAT-PT first checks if there is such a prefix configured on the box. NAT-PT will not translate the IPv4 packet if no NAT-PT prefix is configured.

```
ipv6 nat prefix 2005::/96
```

Interface Configuration Commands

IPv6 NAT

This command enables the interface for NAT-PT operations, and it must be present on the IPv6 and IPv4 interfaces. Note that, on the IPv4 side, it is possible to configure "ipv6 nat" on a non-IPv6 supported type on interface like Token Ring.

```
interface TokenRing0
ip address 172.16.1.1 255.255.255.0
ipv6 nat
!
interface Ethernet 0
ip address 172.16.2.1 255.255.255.0
ipv6 nat
!
ipv6 nat prefix 2005::/96
```

IPv6 NAT Prefix

This command is set up exclusively on IPv4 interfaces at the interface level. It creates a /96 prefix to receive all IPv6 NAT-PT traffic to this prefix. Only a /96 prefix is permitted. Setting different prefixes on different interfaces enables an exit selection feature to access the IPv4 realm. Based on the destination NATed IPv6 prefix the packet is translated and forwarded on a different interface.

In the following example, all packets with 2005:: /96 as a destination address are forwarded to interface Ethernet 0/0; all the packets with a destination address of 2006::/96 are forwarded to interface Ethernet 0/1.

```
interface Ethernet 0/0
ip address 172.16.1.1 255.255.255.0
ipv6 nat
ipv6 nat prefix 2005::/96
!
interface Ethernet 0/1
ip address 172.16.2.1 255.255.255.0
ipv6 nat
ipv6 nat prefix 2006::/96
!
```

Configuring Static Address Mappings

IPv6 NAT v6v4 Source

With this command, NAT-PT will translate the source address of an IPv6 packet, matching the given IPv6 address to an IPv4 address. An equivalent reverse rule is added by default. Therefore, the IPv4 address will be translated when it appears in the destination in the reverse direction. The keyword "source" indicates that a packet's source address will be translated based on the direction specified before by the "v6v4" keyword (i.e. IPv6 to IPv4).

```
ipv6 nat v6v4 source 2001::5 192.168.1.1
```

IPv6 NAT v4v6 Source

Alternately, this command translates the source address of an IPv4 packet into an IPv6 address. An equivalent reverse rule is added by default. Therefore, the IPv6 address will be translated when it appears in the destination in the reverse direction. The keyword "source" indicates that a packet's source address will be translated based on the direction specified before by the "v4v6" keyword (i.e. IPv4 to IPv6).

```
ipv6 nat v4v6 source 192.168.1.2 2001::6
```

Configuring Dynamic Address Mappings

When the above two static one-to-one mappings are used together, they will translate traffic flowing between single IPv4 node IPv6 nodes. However, the above commands are limited, in that they could be used to represent a single IPv4/IPv6 host communicating with the public IPv4/IPv6 world. To have multiple IPv4 or IPv6 nodes communicate, multiple static mappings must be configured. Alternately, a single dynamic mapping configuration command can be used to accomplish this.

Dynamic Mappings allow NAT-PT to translate IPv4/IPv6 addresses using addresses from an IPv6/IPv4 pool. The addresses are dynamically allocated from a pool.

The following paragraphs describe the sequence of steps to follow in order to configure a dynamic mapping.

Identifying the Translated IPv4/IPv6 Addresses

Cisco IOS Software offers three ways to identify the IPv4/IPv6 address to be translated:

- Prefix lists
- IPv6 access lists
- Route maps

IPv6 Prefix Lists

A "Prefix-list" is a collection of permit/deny conditions that use IPv6 address prefixes for source and/or destination address matching. The following example command defines a list of IPv6 prefixes (2001:1:::/64 and 2001:2::/64) to be used as criteria for selection of IPv6 packets to be translated by NAT-PT.

```
ipv6 prefix-list v6-prefix-list seq 5 permit 2001:1::/64
ipv6 prefix-list v6-prefix-list seq 10 permit 2001:2::/64
```

IPv4/IPv6 Access Lists

An IPv6 standard access list is a sequential collection of permit and denies conditions that apply to IPv6 addresses. Source and/or destination IPv6 prefixes are used for matching operations. The following example selects prefix "2005:: /64" as the source of the IPv6 packets.

```
ipv6 access-list v6-list
   permit 2005::/64 any
```

IPv4 standard or extended access lists can be used to translate IPv4 addresses.

In the following examples, "access list 1" is a standard access list which permits those IPv4 packets whose address matches 192.168.1 in its most significant 24 bits, to be translated. Access list 101 is an extended list, which matches most significant 24 bits of both the source and destination of IPv4 packets.

```
access-list 1 permit 192.168.1.0 0.0.0.255
access-list 101 permit ip 192.168.1.0 0.0.0.255 10.10.10.1 0.0.0.255
```

Route Maps

Route maps can also be used instead of prefix lists or access lists. A route map allows the user to match an access-list, prefix list, or output interface to determine which pool to use. The following route-map selects for translation all the packets leaving the router through interface Ethernet 2.

```
!
route-map map1 permit 10
match interface Ethernet 2
!
```

IPv4/IPv6 Addresses Pool

Once the selection of an IPv6/IPv4 source address to translate is complete, pools of IP addresses have to be created to enact this translation.

IPv6 NAT v6v4 Pool

This command creates a pool of forty IPv4 address for IPv6 to IPv4 translation. As for IPv4 NAT, the prefix length argument enforces that the broadcast or network addresses are not assigned. Note that those IPv4 pool definitions are interchangeable with IPv4 NAT pool definitions.

```
ipv6 nat v6v4 pool pool1 192.168.1.1 192.168.1.40 prefix-length 24
```

IPv6 NAT v4v6 Pool

This command creates a pool of IPv6 address for IPv4 to IPv6 translation. When the pool runs out of IPv6 addresses, NAT-PT translates the addresses by pre-pending the NAT-PT prefix with the IPv4 address so that IPv4 address appears in last 32 bits of the translated IPv6 address.

The following example creates a pool of 100 IPv6 addresses to reach the IPv6 realm for the IPv4 realm.

```
ipv6 nat v4v6 pool pool2 2010::1 2010::100 prefix-length 96
```

Dynamic Mapping Definition

When all the above steps are completed, the following commands create the actual dynamic translations.

IPv6 NAT v6v4 Source

This command is a dynamic mapping defined for IPv6 to IPv4 source address translation. In the following example, IPv6 packets, whose source addresses match the IPv6 standard access list "v6-list", will be translated using IPv4 addresses allocated from the address pool "v4-pool".

```
ipv6 nat v6v4 source list v6-list pool v4-pool
```

IPv6 NAT v4v6 Source List

This command is a dynamic mapping defined for IPv4 to IPv6 source address translation. IPv4 packets whose source addresses match the simple access-list "v4-list" will be translated using IPv6 addresses allocated from the address pool "v6-pool".

```
ipv6 nat v4v6 source list v4-list pool v6-pool
```

IPv6 NAT v6v4 Source Route-Map

This command configures a dynamic mapping using a route map.

The following example uses a route-map "map1" to select IPv6 packets for translation. The IPv4 addresses used for translation are allocated from the named pool "v4-pool". NAT-PT does not support "route-map" selection criteria for IPv4-to-IPv6 dynamic mappings.

```
ipv6 nat v6v4 source route-map map1 pool v4-pool
```

Configure Translation Entry Limit

The following commands sets the maximum NAT-PT translations stored in the router. It is particularly useful when the router memory is limited or the number of translations is important. A translation consumes about 250 bytes of memory.

```
ipv6 nat translation max-entries <n>
```

Configuring timeouts

IPv6 NAT Translation Timeout

By default, dynamic translations time out after 24 hours. This command adjusts that timeout from 1 to 2^32 seconds. The first example (below) adjusts the time to 1,000 seconds. The second example sets the timer to never expire. Two forms exist: 0 second or never.

```
ipv6 nat translation timeout 1000
ipv6 nat translation timeout 0
ipv6 nat translation timeout never
```

IPv6 NAT Translation UDP-Timeout

By default, non-DNS UDP translations time out after 5 minutes. This timeout is adjustable, using the aforementioned method.

```
ipv6 nat translation udp-timeout 1000
ipv6 nat translation udp-timeout 0
ipv6 nat translation udp-timeout never
```

IPv6 NAT Translation DNS-Timeout

In a NAT-PT scenario, DNS translations are transient before the default timeout is set to 1 minute. This timeout is adjustable, using the aforementioned method.

```
ipv6 nat translation dns-timeout 1000
ipv6 nat translation dns-timeout 0
ipv6 nat translation dns-timeout never
```

IPv6 NAT Translation TCP-Timeout

TCP translations time out after 24 hours; however, it times out after one minute if a RST or FIN is seen on the stream. This timeout is adjustable, using the aforementioned method.

```
ipv6 nat translation tcp-timeout 1000
ipv6 nat translation tcp-timeout 0
ipv6 nat translation tcp-timeout never
```

IPv6 NAT Translation RST-Timeout

It is possible to immediately terminate a TCP session with an RST. TCP session timeout after receiving a RST can be adjusted in the same way as the TCP timeout.

```
ipv6 nat translation rst-timeout 1000
ipv6 nat translation rst-timeout 0
ipv6 nat translation rst-timeout never
```

IPv6 NAT Translation FINRST-Timeout

In addition to immediately terminating a TCP session, it is possible to gracefully terminate a TCP session with an FIN. TCP session timeout after receiving a FIN or a RST can be adjusted in the same way as the TCP timeout.

```
ipv6 nat translation finrst-timeout 1000
ipv6 nat translation finrst-timeout 0
ipv6 nat translation finrst-timeout never
```

IPv6 NAT Translation SYN-Timeout

If a SYN flag has been received and is not followed by data belonging to the same session TCP session, the translation associated with this session does timeout. Timeout after receiving a SYN, but no further data can be adjusted in the same way as the TCP timeout.

```
ipv6 nat translation syn-timeout 1000
ipv6 nat translation syn-timeout 0
ipv6 nat translation syn-timeout never
```

IPv6 NAT Translation ICMP-Timeout

This command adjusts the timeout associated with ICMP ALG translations in the aforementioned method. The default timeout is 5 minutes.

```
ipv6 nat translation icmp-timeout 1000
ipv6 nat translation icmp-timeout 0
ipv6 nat translation icmp-timeout never
```

Show/Clear/Debug Commands

Show IPv6 NAT Statistics

This command gives aggregated statistics on the NAT-PT process. Total active translations are broken into static, dynamic and extended (when the port number is recorded or a route-map is used).

When a NAT lookup is completed, the "Hits" counter reports how many times the router was able to find a translation in its table; the "Misses" counter reports how many times it could not find a translation in its table.

```
router#show ipv6 nat statistics
Total active translations: 29 (9 static, 15 dynamic; 5 extended)
NAT-PT interfaces: Ethernet0/0, Ethernet0/1
Hits: 200 Misses: 2
Expired translations: 101
```

Show IPv6 NAT Translation

This command reveals detailed mappings between IPv4/IPv6 addresses. It is possible to restrict the displayed translation to a subset of it (ICMP, TCP, and UDP). The following example displays 2 translations.

```
One from 2010::60 to 192.168.1.200 and another from 2001:2::1 to
192.168.2.1.
Router#show ipv6 nat translations
Pro IPv4 source IPv6 source IPv6 destn IPv4 destn
--- --- --- 2010::60 192.168.1.200
--- 192.168.2.1 2001:2::1
It corresponds to the following router configuration:
ipv6 nat v6v4 source 2001:2::1 192.168.2.1
ipv6 nat v4v6 source 192.168.1.200 2010::60
```

Clear IPv6 NAT Translation

This command clears all the dynamic translation stored on the NAT-PT router.

Clear IPv6 NAT Statistics

This command clears all the NAT-PT statistics stored on the NAT-PT router and displayed by the "show ipv6 nat statistics" command.

Debug IPv6 NAT

This command provides debugging data on NAT operations. The output can be narrowed by an "access-list" or expanded by a "verbose" statement.

```
router# debug ipv6 nat
udp (3FFE:B00:FFFF:1::1) -> (192.168.99.51), dst (3FFE:B00::53) ->
(192.168.31.1)
udp (192.168.31.1) -> (3FFE:B00::53), dst (192.168.99.51) ->
(3FFE:B00:FFFF:1::1)
icmp (3FFE:B00:FFFF:1::1)->(192.168.99.51),dst(3FFE:B00::DB19)->
(198.133.219.25)
icmp (3FFE:B00:FFFF:1::1)->(192.168.99.51),dst (3FFE:B0:DB19) ->
(198.133.219.25)
icmp (198.133.219.25)->(3FFE:B00::DB19),dst (192.168.99.51)-
>(3FFE:B00:FFFF:1::1)
```

Debug IPv6 NAT Detailed

This command adds more debug messages for packets being dropped by NAT-PT or address allocated out of pools.

```
router# debug ipv6 nat detailed
IPv6 NAT: address allocated 172.16.20.1
IPv6 NAT: Dropping v6tov4 packet
```

Deployment Scenarios

Static vs. Dynamic Translation

Some applications count on a certain degree of address stability for their operation. Dynamic address reused by NAT-PT might not be agreeable for these applications. For hosts running such address critical applications, NAT-PT may be configured to provide static address mapping between the host's IPv6 address and a specific IPv4 address. Figure 5 depicts a topology where the IPv4 DNS server uses a static translation.

Figure 5. A NAT-PT Topology with Static Translation for DNS Server

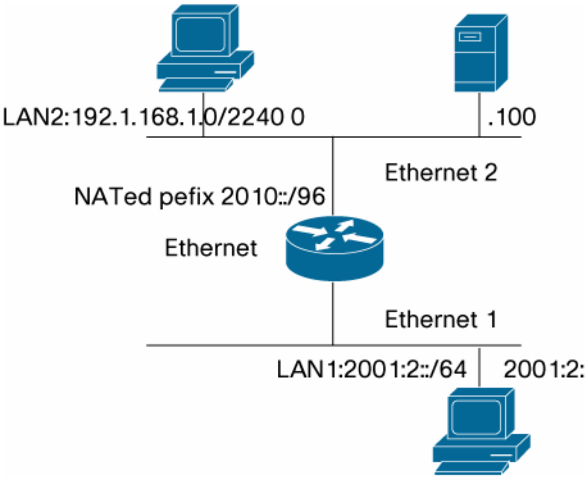

Static Translation Configuration Example

```
interface Ethernet 1
ipv6 address 2001:2::10/64
ipv6 nat
!
interface Ethernet 2
ip address 192.168.1.1 255.255.255.0
ipv6 nat
!
ipv6 nat v6v4 source 2001:2::1 192.168.2.1
ipv6 nat v4v6 source 192.168.1.100 2010::60
!
ipv6 nat prefix 2010::/96
Router1 #show ipv6 nat translations
Pro IPv4 source IPv6 source IPv6 destn IPv4 destn
--- --- --- 2010::60 192.168.1.100
--- 192.168.2.1 2001:2::1 ---
```

Dynamic Translation Configuration Example

There are three ways to select source address undergoing an IPv6 to IPv4 translation. The following section gives a configuration for all the possibilities.

Route-Map

```
interface Ethernet 1
ipv6 address 2001:2::10/64
ipv6 nat
!
interface Ethernet 2
ip address 192.168.1.1 255.255.255.0
ipv6 nat
ipv6 nat prefix 2010::/96
!
ipv6 nat v6v4 source route-map map1 pool v4pool1
ipv6 nat v6v4 pool v4pool1 192.168.2.1 192.168.2.10 prefix-length 24
!
route-map map1 permit 10
```

```
match interface Ethernet 2
```

Access-List

```
interface Ethernet 1
ipv6 address 2001:2::10/64
ipv6 nat
!
interface Ethernet 2
ip address 192.168.1.1 255.255.255.0
ipv6 nat
ipv6 nat prefix 2010::/96
!
ipv6 nat v6v4 source list v6-list pool v4pool1
ipv6 nat v6v4 pool v4pool1 192.168.2.1 192.168.2.10 prefix-length 24
!
ipv6 access-list v6-list
permit 2001:2::/64 any
```

Prefix List

```
interface Ethernet 1
ipv6 address 2001:2::10/64
ipv6 nat
!
interface Ethernet 2
ip address 192.168.1.1 255.255.255.0
ipv6 nat
ipv6 nat prefix 2010::/96
!
ipv6 nat v6v4 source list v6-prefix-list pool v4pool1
ipv6 nat v6v4 pool v4pool1 192.168.2.1 192.168.2.10 prefix-length 24
!
ipv6 prefix-list v6-prefix-list seq 5 permit 2001:2::/64
```

Multiple Exit Points Configuration Example

In the following diagram, a NAT-PT router offers two different exits toward the IPv4 world to an IPv6 cloud. A possible application is the ISP selection based on the destination IPv6 address. Another possibility is the connection of two overlapping IPv4 realms to IPv6 through a single NAT-PT router.

Figure 6 depicts the topology associated with this feature. To reach the IPv4 cloud behind Ethernet 2, the IPv6 packet leaving Host A will be destined to 2010::/96 and to reach the IPv4 cloud behind Ethernet 3, the IPv6 packets leaving Host A will be destined to 2020::/96.

Figure 6. A Multiple Exits NAT-PT Topology

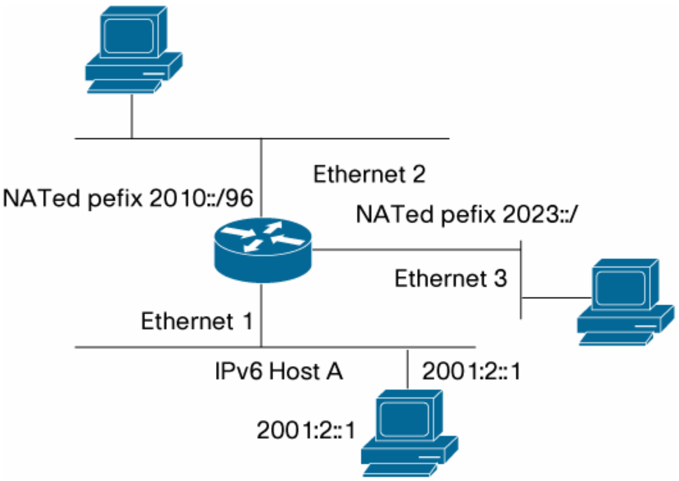

```
interface Ethernet 1
ipv6 address 2001:2::10/64
ipv6 nat
!
interface Ethernet 2
ip address 192.168.1.1 255.255.255.0
ipv6 nat
ipv6 nat prefix 2010::/96
!
interface Ethernet 3
ip address 192.168.2.1 255.255.255.0
ipv6 nat
ipv6 nat prefix 2020::/96
!
ipv6 nat v6v4 source route-map map2 pool v4pool1
ipv6 nat v6v4 source route-map map3 pool v4pool2
!
ipv6 nat v6v4 pool v4pool1 192.168.3.1 192.168.3.10 prefix-length 24
ipv6 nat v6v4 pool v4pool2 192.168.4.1 192.168.4.10 prefix-length 24
!
route-map map2 permit 10
match interface Ethernet 2
!
route-map map3 permit 10
match interface Ethernet 3
```

Topology Limitations

There are limitations to using the NAT-PT translation method. It is mandatory that all requests and responses pertaining to a session be routed via the same NAT-PT router. One way to guarantee this is to have NAT-PT based on a border router that is unique to a stub domain, where all IP packets are either originated from the domain or destined to the domain. This is a generic problem with NAT.

Topology Example with DNS

The DNS deployment in a NAT-PT topology generally looks like Figure 7 (below). A DNS server is deployed in each realm: IPv6 and IPv4. The IPv4 and IPv6 DNS servers are accessed via an IPv4 transport and maintain IPv4 zones and reverse zones.

Host A is configured with a DNS server that points to the IPv6 DNS. Host A resolves names through the IPv6 DNS. DNS resolution occurs normally for zones pointing to IPv6 hosts. If the name to resolve is from a zone of IPv4 hosts, there are 2 solutions:

1. The IPv6 DNS points to an IPv4 DNS server to find the answer to the query. Then there is a need for DNS-ALG, because DNS messages will cross the border between IPv4 and IPv6.
2. IPv6 DNS maintains NAT-PTed AAAA record for host-to-address mapping between the name and the IP addresses. This NAT-PTed address will be returned to Host A, which will send a request to this NAT-PT IPv6 address. This translates into an IPv4 address in the NAT-PT router.

Host A may also point directly to IPv4 DNS for address resolution through a NAT-PTed address (belonging to the /96 prefix). This is the case where Host A accesses only the IPv4 realm, and there is no IPv6 DNS server.

The problem is symmetrical for IPv6 initiated communications. In the IPv4 case there is no /96 prefix but an IPv4 subnet, routed to the NAT-PT router. A static mapping is maintained between IPv4 hosts of this subnet and crucial IPv6 hosts like DNS servers.

Figure 7. DNS Server Deployment in a NAT-PT Topology

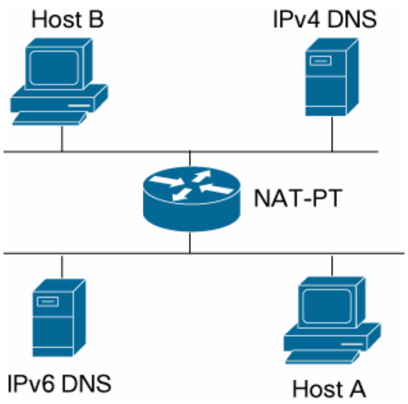

Security Considerations

One of the most important limitations of the NAT-PT is that end-to-end network layer security is impossible. Transport and application layer security may also be impossible for applications that carry IP addresses to the application layer. This is an inherent limitation of the NAT function.

Independent of NAT-PT, end-to-end IPSec security is not possible across different address realms. The two end-nodes that seek IPSec network level security must both support IPv4 or IPv6.

DNS Translation and DNSSEC

The use of NAT-PT in a real environment involves the translation of DNS messages. It is clear that it precludes the utilization of secure DNS (i.e.: an authoritative DNS name server in the IPv6 domain cannot sign replies to queries that originate from the IPv4 realm). As a result, an IPv4 host that demands DNS replies to be signed will reject replies that have been tampered with by NAT-PT.

Compatibility Matrix

NAT-PT offers the following features which are not all supported in 12.2(13)T.

Table 3.

NAT-PT Feature	12.2(13)T Supported
Static Translation	Yes
Dynamic Translation	Yes
Auto Aliasing	Yes
NATPT-PT (Address Overload)	Future
DNS-ALG	Yes
ICMP-ALG	Yes
FTP-ALG	Future
Fragmentation	Future
CEF Switching	Future

Conclusion

In addition to the other transition mechanisms from IPv4 to IPv6, NAT-PT also offers a useful tool to enable an IPv4/IPv6 co-existence. It should be well received from the Network Administrators, as it maintains many IPv4 NAT concepts.

NAT-PT is an attractive solution whenever IPv4 only to IPv6 only communication (or vice versa) is needed. With NAT-PT, legacy IPv4 only hosts will still be able to connect, while new IPv6 only home appliances will interact with the initially predominant IPv4 world. However, it is generally limited to the stub network topology, due to inherited limitations of NAT. Similarly, it does suffer a lack of application transparency in spite of the use of ALG.

References

- [DNS-ALG] RFC 2694 DNS extensions to Network Address Translators.
- [SIIT] RFC 2765 Stateless IP/ICMP Translation Algorithm.
- [NAT-PT] RFC 2766 Network Address Translation—Protocol Translation.
- [V6ADDR]RFC 2373 Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture"

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA

C11-416070-00 7/07