White Paper

Enhanced IP Resiliency Using Cisco Stateful NAT

INTRODUCTION

Stateful Network Address Translation is a Cisco IOS® Software feature allowing two or more network address translators to function as a translation group. A backup NAT provides translation services in the event of failure to the active translator. The result is a more resilient IP network.

The goal is to create a more globally resilient IP network. Networked applications are placing increased demands on the core IP network. Users expect continuous access to servers and data, regardless of location. Although the mean time between failure (MTBF) of hardware components has increased, failures can and do occur. Administrative activities can also cause downtime. A resilient IP network offers continuous service, despite failures that may occur.

The concept of a highly resilient IP network is not new; however, this paper introduces an innovative approach. The intelligent systems approach creates a highly optimized, resilient IP network where individual component features interact and share services. The result is a network that is inherently more intelligent and less labor-intensive in terms of design and management. Cisco IOS Software is evolving into a more intelligent, shared function system that helps reduce support costs and increase the benefit and return on investment in network equipment and services.

NAT has been a core Cisco IOS Software feature since its introduction. It has helped to reduce address depletion and promote Internet growth. NAT has been used to permit interconnection of private networks, regardless of their use of independent addressing schemes, even when these schemes use addresses that conflict. NAT has also been used to effectively hide networks from outside the administrative domain while allowing predetermined connections to occur. NAT fulfills an important role and will likely do so even as IPv6 is deployed.

This enhancement can make NAT even more resilient and allow application connectivity to continue, unaffected by potential failures to links and routers at the NAT border. Cisco Stateful NAT provides this enhanced capability.

In IP networking, “stateful” is defined as applying a more global context to the task of forwarding a particular datagram. There is consideration of not just where to forward the datagram, but also of the application/connection state with regard to this datagram. With this knowledge, devices can take action so that potential failures will have less impact on the flow and on the application that is transmitting data. Multiple NAT routers that share stateful context can work cooperatively and thus increase service availability.

STATEFUL NAT OVERVIEW

Stateful NAT (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.

Only sessions that are already statically defined receive the benefit of redundancy without the need for this feature. 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. SNAT enables the maintenance of continuous service for dynamically mapped NAT sessions. The end result is a more resilient IP network.

Phased Release

Cisco is releasing SNAT in phases. Phase I (Cisco IOS Release 12.2(13)) provides a subset of the intended function. Application-level gateway support is not included in Phase I, so protocols that embed IP address data within the payload of the IP packet will not be able to take advantage of the enhanced redundancy provided by SNAT.

Phase II (Cisco IOS Release 12.3(7)) provides increased application-level gateway and asymmetric routing support in SNAT.

Protocols and applications supported in Phase I are:

Any TCP/UDP traffic that does not carry source or destination addresses in the payload

- Archie
- Finger
- HTTP
- Internet Control Message Protocol (ICMP)
- Ping
- rcp, rlogin, rsh
- TCP
- Telnet

Protocols and applications supported in Phase II are:

- FTP
- H.225, H.245
- Point-to-Point Tunneling Protocol (PPTP)/generic routing encapsulation (GRE)
- NetMeeting Directory (ILS)
- H323 Registration, Admission and Status (RAS)
- Session Initiation Protocol (SIP); both TCP- and UDP-based
- Skinny
- Trivial File Transfer Protocol (TFTP)

Support for additional protocols may be offered in later releases.

There are additional deployment restrictions for SNAT Phase I. It only function properly when the return traffic path traverses the primary SNAT router. In other words, asymmetrical routing should be prevented. To ensure that return traffic follows a single path to the NAT router, the routing path cost must be adjusted or the Border Gateway Protocol (BGP) metric must be set appropriately. Phase II will allow for asymmetric routing, which will remove the restriction.

Phase II includes additional support for the following:

- Support for outside NAT pools, using the configuration command ip nat outside source pool. SNAT Phase I will only permit inside NAT pools.
- Dynamic entries, which are extended out of static definitions.
- Support for ip nat inside destination.

Scalability for Stateful NAT

There is a potential problem for multiple NAT routers that share stateful context: because Phase II SNAT has no control of Hot Standby router Protocol (HSRP), NAT databases are out of sync between the NAT routers and result in connection losses between end applications.

Scalability for SNAT was integrated into Cisco IOS Software Release 12.4(3), allowing users to enable the feature that allows SNAT to control the HSRP state change until the NAT information is completely exchanged at HSRP mode. Cisco IOS Software Release 12.4(10) will enable scalability for SNAT at both HSRP mode and Primary/Backup mode.

Note: It is highly recommended to run the same Cisco IOS image and be configured with the same NAT configuration, including the global address pools for dynamic NAT, static NAT, and NAT timeout values on SNAT peer routers, to ensure stability and compatibility

Scalability for SNAT can disable queuing during asymmetric routing to avoid delays in the data path for the creation of new entries and traffic on special ports (application-layer gateway support).

PLATFORM SUPPORT

SNAT will be supported on all platforms running Cisco IOS Software. Platforms that include hardware acceleration for NAT will benefit, since the mechanism for creating NAT table entries is compatible with the hardware acceleration implementation.

Cisco IOS Software is packaged in feature sets that support specific platforms. To get updated information regarding platform support for this feature, please visit Cisco Feature Navigator at http://www.cisco.com/go/fn/. This application dynamically updates the list of supported platforms as new platform support is added for the feature.

STATEFUL NAT PROTOCOL

SNAT using UDP to communicate NAT table updates between the primary and backup NAT routers was introduced in Cisco IOS Software Release 12.4(3) along with TCP. When UDP mode is used, SNAT will send NAT database exchange information over UDP using proprietary acknowledgement/retransmit mechanism.

Note: SNAT using TCP as the transport mechanism is no longer accepted by Cisco IOS Software Release 12.4(10) and later. TCP configuration will be ignored and replaced by UDP communication by SNAT.

Figure 1 is a SNAT functional diagram. Once configured for SNAT, the UDP session is established between the SNAT peer routers and is used to transmit messages that communicate updates to the NAT tables and maintain session state.

Figure 1. SNAT Functional Diagram

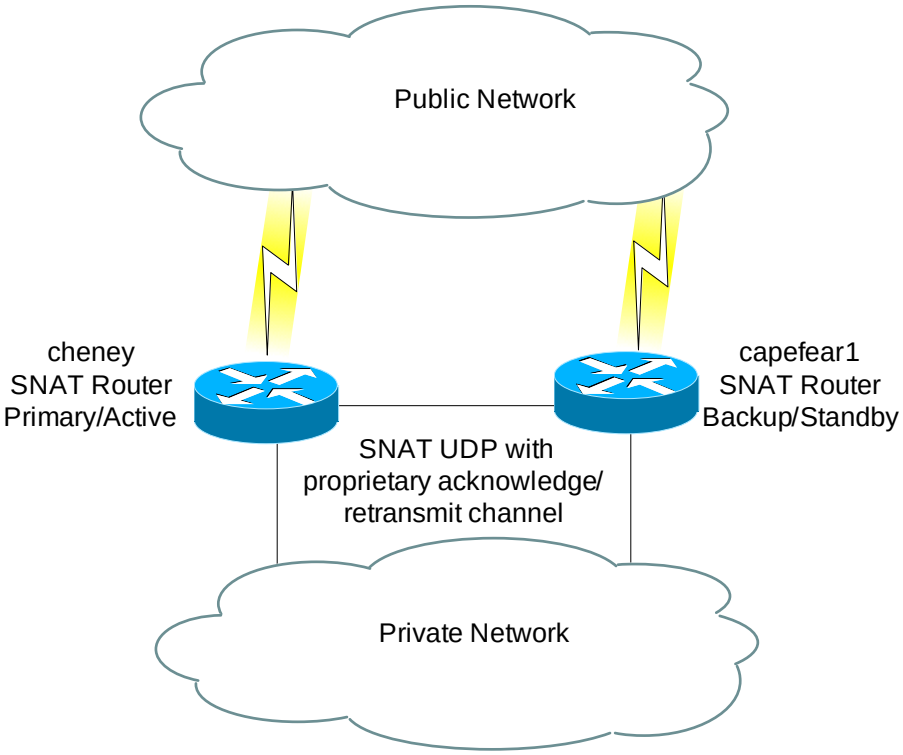

The distributed NAT protocol will ensure that dynamic NAT entries created at the primary or active NAT are duplicated consistently on the backup or standby NAT router. This prepares the backup NAT to take over in the event of a critical failure.

The distributed NAT protocol defines a set of messages that are exchanged between NAT routers:

- **Add message**—Sent to the peer NAT router whenever traffic flow dictates that a dynamic entry be created locally. The action creates an entry at the recipient's database, based on information in the message. This is also discussed in the Mapping ID section.
- **Delete message**—Sent to the peer when a dynamic entry is deleted from the local database. The action deletes the corresponding entry at the recipient's database. In SNAT Phase II, the Delete message will be extended to include three types of delete operations:
  - **Forced-Delete:** The recipient will delete the entry.
  - **Delete-Query:** Upon entry-timeout, the Active/Primary that timed out the entry will query the other router as to whether it has received packets later than the NAT router, which is actually running the timer on the entry. In other words, the query permits adjustment of the timer so an entry is not prematurely deleted due to asymmetric flow of traffic.
  - **Delete-Response:** This is sent in response to the Delete-Query. A time-to-restart value is included to adjust the timer on the entry at the Active/Primary that is handling the timers for this entry. A value of 0 in the time-to-restart field will indicate that the recipient has not received packets for this flow later than the Active/Primary.
- **Dump-Request message**—This message is sent whenever the router comes up asking for the snapshot of the NAT database from the peer NAT router.

- **Dump-Reply message**—This message is sent in response to the **Dump-Request**. The message will include the previously learnt dynamic entries from the router that issued the Dump-Request plus the dynamic entries created locally. This is also discussed in the Mapping ID section.
- **Update message**—Distributes application-specific information (valid only in SNAT Phase II).
- **Sync message**—Informs the peer of the local SNAT ID number. After the HSRP connection is established between the SNAT peer routers, SNAT starts to send the Sync message. This informs every NAT peer router about the configured SNAT ID number at the peer.

A consistent set of NAT entries is maintained through the exchange of the aforementioned messages. When a SNAT router fails or reloads, it will request a dump of the current NAT entries from the currently active **SNAT** router upon restoration, and will assume its role in the SNAT group.

It is not recommended to perform dynamic translation clearing on the Standby/Backup router; doing so will cause NAT tables out of sync between SNAT peer routers. If needed, it should be done at the Active/Primary router, and the Active/Primary router will propagate the updates to the Standby/Backup routers automatically.

Mapping ID

The **mapping id** command is used to specify whether the local SNAT router will distribute a particular set of locally created entries to a peer SNAT router.

The logic used for distributing the entries created locally to the peer is as follows:

Each dynamically created entry inherits a mapping ID number based on the configuration setting at the point of creation. This comes from the mapping defined on the NAT rule. For example, entries created using rule `ip nat inside source route-map rm-101 pool SNATPOOL1 mapping-id 10 overload` will have ID 10 associated with them.

For each SNAT router, a mapping list may also be defined using the command **mapping-id** within the SNAT configuration as shown below:

```
ip nat Stateful id 1
   redundancy SNATHSRP
      mapping-id 10
      mapping-id 11
```

Multiple mapping ID statements can be used to form a mapping list. The list specifies which of the entries will be forwarded to peers in that group. It provides a way to specify that entries from particular NAT rules should be forwarded.

Show Commands

Use the command `show ip snat distributed verbose` to get status information about the SNAT processes. In Example 1: **show ip snat distributed verbose** shows a router that is configured for HSRP mode and is currently in STANDBY due to the corresponding HSRP group being in STANDBY state. This is because the tracked interface (FastEthernet 0/1) is down; the HSRP group priority is decreased from 105 to 95, its peer router (10.88.194.18) with higher HSRP group priority 100 goes active.

Example 1 show ip snat distributed verbose

```
cheney#show ip snat distributed verbose
Stateful NAT Connected Peers
SNAT: Mode IP-REDUNDANCY :: STANDBY
```



```
    : State READY
    : Local Address 10.88.194.17
    : Local NAT id 1
    : Peer Address 10.88.194.18
    : Peer NAT id 2
    : Mapping List 10
    : InMsgs 5210, OutMsgs 5212, tcb  0x82C2DC28, listener 0x826D1790

Cheney#show standy
FastEthernet0/0 - Group 0
  State is Standby
    4 state changes, last state change 1d04h
  Virtual IP address is 10.88.194.20
  Active virtual MAC address is 0000.0c07.ac00
    Local virtual MAC address is 0000.0c07.ac00 (default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.372 secs
  Preemption enabled, delay min 20 secs
  Active router is 10.88.194.18, priority 100 (expires in 9.796 sec)
```

The TCP Control Block (TCB) value for the Stateful NAT using UDP communication mechanism displayed here is a dummy value. This value is consistent with the output of Stateful NAT using the TCP communication mechanism, which is no longer accepted by Cisco IOS Software Release 12.4(10) and later.

```
  Standby router is local
  Priority 95 (configured 105)
    Track interface FastEthernet0/1 state Down decrement 10
  IP redundancy name is "SNATHSRP" (cfgd)

cheney#show ip int brief
Interface                  IP-Address      OK? Method Status                Protocol
FastEthernet0/0            10.88.194.17    YES NVRAM  up                    up
FastEthernet0/1            10.88.161.6     YES NVRAM  up                    down
Loopback0                  10.88.194.5     YES NVRAM  up                    up
Virtual-Access1            unassigned      YES unset  up                    up
cheney#
```


This illustrates how SNAT works together with HSRP to achieve improved redundancy.

The current NAT entries can be displayed using the command `show ip nat translation`. Additional information is shown when the **verbose** option is included.

NAT entries have been extended to include information about which of the SNAT routers created them, and which router is responsible for the state and timing of that particular entry. The combination of the entry id-number and the SNAT router id-number make each entry unique within the group.

In Example 2: **show IP NAT translations**, SNAT router "cheney" has two entries numbered 1173 and 1174 that have "left" values counting down from 00:00:35. These entries are timing out. The active SNAT router is responsible for timing out the entries. All three entries are duplicated on a standby SNAT router (capefear1) and are flagged "created-by-remote". This indicates that this router is a backup for these entries, as they are not being timed locally.

Example 2 show IP NAT translations

```
cheney#show ip nat translations
Pro Inside global   Inside local       Outside local    Outside global
tcp 11.1.1.3:1173   10.88.194.22:1173  88.1.88.8:10115  88.1.88.8:10115
tcp 11.1.1.3:1174   10.88.194.22:1174  88.1.88.8:10115  88.1.88.8:10115
tcp 11.1.1.3:1175   10.88.194.22:1175  88.1.88.8:47950  88.1.88.8:47950

cheney#show ip nat translations verbose
Pro Inside global   Inside local       Outside local    Outside global
tcp 11.1.1.3:1173   10.88.194.22:1173  88.1.88.8:10115  88.1.88.8:10115
   create 00:00:24, use 00:00:24, left 00:00:35, Map-Id(In): 1,
   flags: extended, timing-out, use_count: 0 nat_id: 1 nat_entry_num: 3
nat_mapping_id: 10

tcp 11.1.1.3:1174   10.88.194.22:1174  88.1.88.8:10115  88.1.88.8:10115
    create 00:00:24, use 00:00:24, left 00:00:35, Map-Id(In): 1,
    flags: extended, timing-out, use_count: 0 nat_id: 1 nat_entry_num: 4 nat_mapping_id: 10

tcp 11.1.1.3:1175   10.88.194.22:1175  88.1.88.8:47950  88.1.88.8:47950
    create 00:00:24, use 00:00:00, left 1d00h, Map-Id(In): 1,
    flags: extended, use_count: 0 nat_id: 1 nat_entry_num: 5 nat_mapping_id: 10

capefear1#show ip nat translations
Pro Inside global   Inside local       Outside local    Outside global
```



```
tcp 11.1.1.3:1173   10.88.194.22:1173  88.1.88.8:10115  88.1.88.8:10115
tcp 11.1.1.3:1174   10.88.194.22:1174  88.1.88.8:10115  88.1.88.8:10115
tcp 11.1.1.3:1175   10.88.194.22:1175  88.1.88.8:47950  88.1.88.8:47950

capefear1#sh ip nat translations verbose
Pro Inside global   Inside local       Outside local    Outside global
tcp 11.1.1.3:1173   10.88.194.22:1173  88.1.88.8:10115  88.1.88.8:10115
    create 00:00:47, use 00:00:47, Map-Id(In): 2,
    flags: extended, created-by-remote, use_count: 0

tcp 11.1.1.3:1174   10.88.194.22:1174  88.1.88.8:10115  88.1.88.8:10115
    create 00:00:46, use 00:00:46,
    flags: extended, created-by-remote, use_count: 0

tcp 11.1.1.3:1175   10.88.194.22:1175  88.1.88.8:47950  88.1.88.8:47950
    create 00:00:45, use 00:00:45,
    flags: extended, created-by-remote, use_count: 0
```

CONFIGURATION

Configuration for SNAT is the same as regular NAT, with some simple additional commands. The first step in defining SNAT is to determine the method of redundancy. SNAT can be configured to work with HSRP by using the IP Redundancy API built into Cisco IOS Software. When HSRP mode is set, the primary and backup NAT routers are elected according to the HSRP standby state. Alternatively, SNAT can be manually defined as primary or backup.

Stateful NAT Interaction with HSRP

The Active and Standby routers are determined from the IP Redundancy API and do not need to be explicitly defined. An example of configuration using IP Redundancy mode is depicted in Example 3. Merely coding redundancy SNATHSRP causes SNAT to make use of the IP Redundancy API. The name is the same as that used in the command standby name SNATHSRP.

Example 3 HSRP Example

```
CHENEY                          CAPEFEAR1

ip nat Stateful id 1            ip nat Stateful id 2
  redundancy SNATHSRP             redundancy SNATHSRP
  as-queuing disable              as-queuing disable
  mapping-id 10                   mapping-id 10
```


The two routers, CHENEY and CAPEFEAR1, form a NAT group. They are designated members of the group by coding the command:

```
ip nat stateful id <id-number>
```

Note: id-number is a unique number given to each router in the stateful translation group. Each SNAT router should have a unique ID number.

Establish HSRP as the method of redundancy by coding the command:

```
redundancy <name>
```

Note: SNAT can only listen to one HSRP group. If it is necessary to listen to multiple groups, you need to tie multiple HSRP groups to one group.

Disable asymmetric routing during queuing in HSRP mode by coding the command:

```
as-queuing disable
```

Note: For most of network topologies, the asymmetric routing can be handled by a proper routing configuration. It is recommended to disable asymmetric routing if the router can handle asymmetric routing to improve CPU performance. The asymmetric process is enabled by default.

The Mapping ID section offers more information on how the mapping-id command is used.

The NAT configuration must be configured identically on SNAT peer routers. If the NAT configuration comprises both dynamic NAT and static NAT, the translated IP addresses in the global address pools for the dynamic NAT should not overlap with the translated IP addresses for the static NAT. It is important to not overlap the designated address for the Hot Standby group with the translated IP addresses for both dynamic NAT and static NAT.

Note: For SNAT configuration, the router's interface addresses cannot be used as the translated IP addresses for both dynamic NAT and static NAT.

To learn more about translated IP addresses, read the Cisco IOS NAT Overview at http://www.cisco.com/en/US/products/ps6640/products_white_paper09186a0080091cb9.shtml.

For more information on configuring HSRP, refer to the HSRP documentation at http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_configuration_guide_chapter09186a00800b3e13.html#wp1001531.

Stateful NAT Primary/Backup

Primary/Backup mode allows explicit configuration of the primary SNAT router and the backup SNAT router. Each router is defined explicitly, and the IP address of the peer router is specified (Example 4).

Example 4 Primary/Backup Example


```
PRIMARY                          BACKUP

ip nat Stateful id 1             ip nat Stateful id 2
  primary 10.88.194.17             backup 10.88.194.18
  peer 10.88.194.18                peer 10.88.194.17
  mapping-id 10                    mapping-id 10
```

The primary command identifies an interface and IP address that the primary SNAT will use as the source for communicating with the backup SNAT router (for sending SNAT protocol messages). Likewise, the backup command does the same for the backup SNAT router. The peer command defines the destination IP address to use for communicating with the peer.

Verification

The status of the SNAT configuration can be examined by using the command: `show ip snat distributed verbose`

If SNAT routers have an entry with the correct peer router ID, please see Example 5. cheney, the active SNAT/HSRP router, has peer address 10.88.194.18 and peer NAT ID 2, which matches its peer standby router capefear1's local address 10.88.194.18 and local NAT ID 2. The communication is established between the SNAT/HSRP routers.

Example 5 Show IP SNAT Commands

```
cheney#sh ip snat distributed verbose
Stateful NAT Connected Peers
SNAT: Mode IP-REDUNDANCY :: ACTIVE
    : State READY
    : Local Address 10.88.194.17
    : Local NAT id 1
    : Peer Address 10.88.194.18
    : Peer NAT id 2
    : Mapping List 10
    : InMsgs 3246, OutMsgs 3247, tcb 0x82BF8BFC, listener 0x0


capefear1#sh ip snat distributed verbose

Stateful NAT Connected Peers

SNAT: Mode IP-REDUNDANCY :: STANDBY
    : State READY
```



```
: Local Address 10.88.194.18
: Local NAT id 2
: Peer Address 10.88.194.17
: Peer NAT id 1
: Mapping List 10
: InMsgs 3249, OutMsgs 3248, tcb 0x82BB4D0C, listener 0x826D1790
```

CONFIGURATION EXAMPLE

The network example shown in Figure 2 is from a customer test case for deploying SNAT. It was configured within a test lab. In this case, a shared test FTP server is in the Internet that provides services to the test client inside the customer's private network. The customer deploys SNAT to ensure the client can receive the Internet service 24x7.

The FTP server is statically NATed to 192.168.241.12 to hide its private subnet, 182.16.200.0, and provide service to clients in the Internet. The FTP clients in private subnet 12.240.0.0 are dynamically PATed to a pool of addresses ranging from 172.16.201.1 to 172.16.201.10 to gain the Internet access.

The network diagram in Figure 2 shows two SNAT routers, DNS1-INS and DNS2-INS:

- DNS1-INS is the active router and tracks the FastEthernet0 and FastEthernet1 interface states. When DNS1-INS is the active router, the traffic from the hosts (test client PC) to the test FTP server is routed through DNS1-INS.
- DNS2-INS is the standby router and tracks the FastEthernet0/0 and FastEthernet0/1 interface states.

Two aggregate routers, OUT182 and IN12 are connected to SNAT routers and provide LAN connectivity.

Figure 2. SNAT Test Network

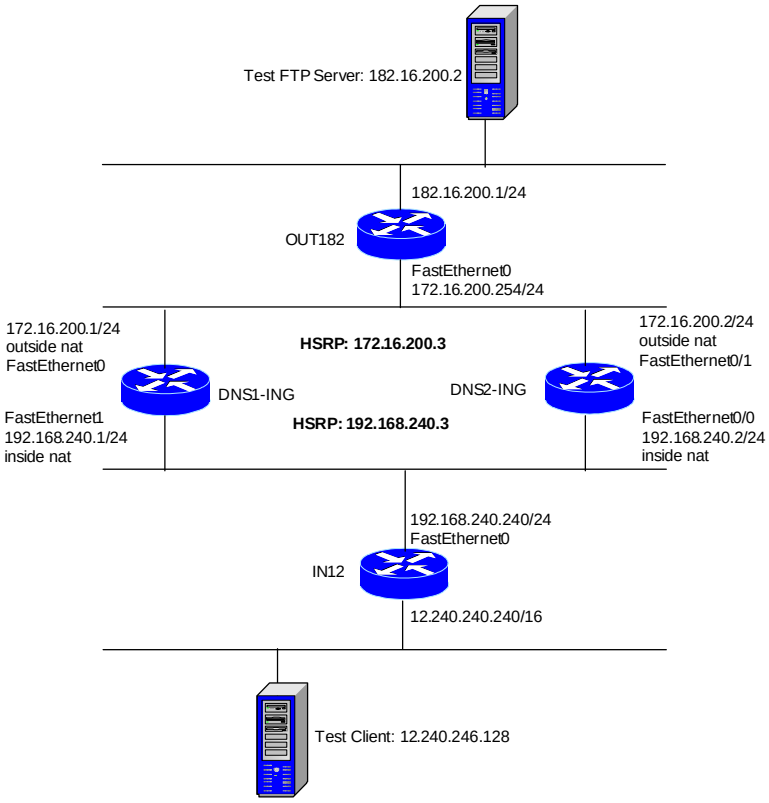

The configuration examples include commands required for SNAT and HSRP. These are shown in Examples 6–9.

Example 6 DNS1-INS

```
hostname DNS1-INS
!
interface FastEthernet0
   ip address 192.168.240.1 255.255.255.0
   ip nat inside

! --- Configure the delay period before the initialization of HSRP groups
! --- after the router has reloaded. The feature helps prevent HSRP state flapping
   standby delay minimum 60 reload 60

! --- Activate HSRP on the interface
! --- Assign a standby group (1 in this case) and the designated IP address for
```



```
! --- the Hot Standby group (192.168.240.3 in this case)
   standby 1 ip 192.168.240.3

! --- Assign a priority (105 in this case) to the router interface FE0
! --- for a particular group number (1)
   standby 1 priority 105

! --- The HSRP preempt feature enables a router with  highest priority to immediately
! --- become the active router at any time.
! --- The preempt delay feature allows preemption to be delayed for a configurable
! --- time period, allowing the router to populate its routing table before becoming
! --- the active router
! ---  minimum causes the router to postpone taking over the active role for
! --- a minimum of seconds since the router was last restarted
! --- reload specifies the preemption delay after a reload only
! --- sync specifies the maximum number of seconds to allow IP redundancy clients
! --- to prevent preemption
   standby 1 preempt delay minimum 60 reload 60 sync 60

! --- Sets HSRP group name (in this case group 1 is named HSRP_IN)
   standby 1 name HSRP_IN

! --- The tracking feature specifies another interfaces (FastEthernet1 in this case)
! --- on the router for the HSRP process to monitor in order to alter the HSRP priority
! ---  The priority decreases (20 in this case) when the interface goes down.
   standby 1 track FastEthernet1 20
!
interface FastEthernet1
   ip address 172.16.200.1 255.255.255.0
   ip nat outside
   standby delay reload 60
   standby 2 ip 172.16.200.3
   standby 2 priority 105
   standby 2 preempt delay minimum 60 reload 60 sync 60
   standby 2 name HSRP_OUT
```



```
   standby 2 track FastEthernet0 20
!
ip route 0.0.0.0 0.0.0.0 192.168.240.240
ip route 192.168.241.0 255.255.255.0 172.16.200.254
!
! --- Enable SNAT on routers configured for HSRP using the keyword redundancy
ip nat Stateful id 1
      redundancy HSRP_IN
            mapping-id 1
            as-queuing disable
            protocol udp
!
ip nat translation udp-timeout 120
ip nat translation dns-timeout 120
ip nat pool VIRPOOL 172.16.201.1 172.16.201.10 netmask 255.255.255.0
ip nat pool INGPOOL 192.168.241.100 192.168.241.110 netmask 255.255.255.0

! --- Enable NAT of the inside source address, in this case
! -- - A packet from inside interface with the source address defined by access-list 100
! --- is NATed to global IP addresses defined by VIRPOOL dynamically
ip nat inside source list 100 pool VIRPOOL mapping-id 1 overload

! --- Enable NAT of the outside source address, in this case
! --- A packet from outside interface with the source address defined by access-list 110
! --- is NATed to global IP address defined by INGPOOL dynamically
ip nat outside source list 110 pool INGPOOL mapping-id 1

! --- Enable SNAT for the HSRP translation group:
! --- In this case, A packet from outside interface with source address 182.16.200.1
! --- is statically NATed to 192.168.241.11 with NAT redundancy operation enabled
ip nat outside source static 182.16.200.1 192.168.241.11 redundancy HSRP_IN mapping-id 1

! --- In this case, the Test FTP server is NATed to a global IP address 192.168.241.12
! --- The private address 182.16.200.2 is hidden
ip nat outside source static 182.16.200.2 192.168.241.12 redundancy HSRP_IN mapping-id 1
```



```
ip nat outside source static 182.16.200.3 192.168.241.13 redundancy HSRP_IN mapping-id 1
!
access-list 100 permit ip 192.168.240.0 0.0.0.255 any
access-list 100 permit ip 12.240.0.0 0.0.255.255 any
access-list 110 permit ip 182.16.200.0 0.0.0.255 any
```

Example 7 DNS2-INS

```
hostname DNS2-INS
!
interface FastEthernet0/0
   ip address 192.168.240.2 255.255.255.0
   ip nat inside
   standby delay minimum 60 reload 60
   standby 1 ip 192.168.240.3
   standby 1 priority 95
   standby 1 preempt delay minimum 60 reload 60 sync 60
   standby 1 name HSRP_IN
   standby 1 track FastEthernet0/1 20
!
interface FastEthernet0/1
   ip address 172.16.200.2 255.255.255.0
   ip nat outside
   standby delay reload 60
   standby 2 ip 172.16.200.3
   standby 2 priority 95
   standby 2 preempt delay minimum 60 reload 60 sync 60
   standby 2 name HSRP_OUT
   standby 2 track FastEthernet0 20
!
ip route 0.0.0.0 0.0.0.0 192.168.240.240
ip route 192.168.241.0 255.255.255.0 172.16.200.254
!
ip nat Stateful id 2
      redundancy HSRP_IN
```



```
            mapping-id 1
            as-queuing disable
            protocol udp
!
ip nat translation udp-timeout 120
ip nat translation dns-timeout 120
ip nat pool VIRPOOL 172.16.201.1 172.16.201.10 netmask 255.255.255.0
ip nat pool INGPOOL 192.168.241.100 192.168.241.110 netmask 255.255.255.0
ip nat inside source list 100 pool VIRPOOL mapping-id 1 overload
ip nat outside source list 110 pool INGPOOL mapping-id 1
ip nat outside source static 182.16.200.1 192.168.241.11 redundancy HSRP_IN mapping-id 1
ip nat outside source static 182.16.200.2 192.168.241.12 redundancy HSRP_IN mapping-id 1
ip nat outside source static 182.16.200.3 192.168.241.13 redundancy HSRP_IN mapping-id 1
!
access-list 100 permit ip 192.168.240.0 0.0.0.255 any
access-list 100 permit ip 12.240.0.0 0.0.255.255 any
access-list 110 permit ip 182.16.200.0 0.0.0.255 any
```

Example 8 IN12

```
hostname IN12
!
interface FastEthernet0
   ip address 192.168.240.240 255.255.255.0
!
interface Vlan1
   ip address 12.240.240.240 255.255.0.0
!
Ip route 0.0.0.0 0.0.0.0 FastEthernet0
```

Example 9 OUT182

```
hostname OUT182
!
interface FastEthernet4
   ip address 172.16.200.254 255.255.255.0
```



```
!
interface Vlan1
   ip address 182.16.200.1 255.255.255.0
!
Ip route 0.0.0.0 0.0.0.0 FastEthernet4
```

To display HSRP information, use the show standby command. In this case, the DNS1-INS [[ING?]]is configured with higher router priority and is elected as the Active router (Examples 10 and 11).

Example 10 DNS1-ING:HSRP information

```
DNS1-ING#show standby
FastEthernet0 - Group 1
  State is Active
    32 state changes, last state change 03:47:43
  Virtual IP address is 192.168.240.3
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.904 secs
  Preemption enabled, delay min 60 secs, reload 60 secs, sync 60 secs
  Active router is local
  Standby router is 192.168.240.2, priority 95 (expires in 8.952 sec)
  Priority 105 (configured 105)
    Track interface FastEthernet1 state Up decrement 20
  IP redundancy name is "HSRP_IN" (cfgd)
FastEthernet1 - Group 2
  State is Active
    32 state changes, last state change 03:48:23
  Virtual IP address is 172.16.200.3
  Active virtual MAC address is 0000.0c07.ac02
    Local virtual MAC address is 0000.0c07.ac02 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.492 secs
  Preemption enabled, delay min 60 secs, reload 60 secs, sync 60 secs
  Active router is local
```



```
  Standby router is 172.16.200.2, priority 95 (expires in 8.276 sec)
  Priority 105 (configured 105)
    Track interface FastEthernet0 state Up decrement 20
  IP redundancy name is "HSRP_OUT" (cfgd)
```

Example 11 DNS2-ING:HSRP information

```
DNS2-ING#show standby
FastEthernet0/0 - Group 1
  State is Standby
    31 state changes, last state change 03:48:52
  Virtual IP address is 192.168.240.3
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.972 secs
  Preemption enabled, delay min 60 secs, reload 60 secs, sync 60 secs
  Active router is 192.168.240.1, priority 105 (expires in 9.904 sec)
  Standby router is local
  Priority 95 (configured 95)
    Track interface FastEthernet0/1 state Up decrement 20
  IP redundancy name is "HSRP_IN" (cfgd)
FastEthernet0/1 - Group 2
  State is Standby
    31 state changes, last state change 03:49:25
  Virtual IP address is 172.16.200.3
  Active virtual MAC address is 0000.0c07.ac02
    Local virtual MAC address is 0000.0c07.ac02 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.560 secs
  Preemption enabled, delay min 60 secs, reload 60 secs, sync 60 secs
  Active router is 172.16.200.1, priority 105 (expires in 8.193 sec)
  Standby router is local
  Priority 95 (configured 95)
    Track interface FastEthernet0/0 state Up decrement 20
```



```
  IP redundancy name is "HSRP_OUT" (cfgd)
```

Use show ip snat distributed verbose to display SNAT status (Examples 12 and 13). The peer address and peer NAT ID on DNS1-ING SNAT router should match the local address and local NAT ID on DNS2-ING SNAT router. The converse is also true.

Example 12 DNS1-ING:Stateful NAT Status

```
DNS1-ING#show ip snat distributed verbose

Stateful NAT Connected Peers

SNAT: Mode IP-REDUNDANCY :: ACTIVE
    : State READY
    : Local Address 192.168.240.1
    : Local NAT id 1
    : Peer Address 192.168.240.2
    : Peer NAT id 2
    : Mapping List 1
    : InMsgs 670, OutMsgs 0, tcb 0x832CD730, listener 0x0
```

Example 13 DNS2-ING:Stateful NAT Status

```
DNS2-ING#show ip snat distributed verbose

Stateful NAT Connected Peers

SNAT: Mode IP-REDUNDANCY :: STANDBY
    : State READY
    : Local Address 192.168.240.2
    : Local NAT id 2
    : Peer Address 192.168.240.1
    : Peer NAT id 1
    : Mapping List 1
    : InMsgs 850, OutMsgs 0, tcb 0x82F91C08, listener 0x0
```

A test client PC attached to the LAN in subnet 12.240.0.0 will FTP to the shared test FTP server. Use `show ip nat translation` to show the content of the translation table (Example 14). The table shows the test client (12.240.246.128) establishing TCP sessions with the

test FTP server (182.16.200.2). The active SNAT router, DNS1-ING, translates 12.240.126.128 to 172.16.201.1 dynamically, and translates 182.16.200.2 to 192.168.241.12 statically.

Example 14 DNS1-ING:translation table

```
DNS1-ING#show ip nat translations
Pro Inside global   Inside local       Outside local      Outside global
--- ---             ---                192.168.241.11     182.16.200.1
--- ---             ---                192.168.241.12     182.16.200.2
--- ---             ---                192.168.241.13     182.16.200.3
icmp 172.16.201.1:768  12.240.246.127:768 192.168.241.11:768 182.16.200.1:768
tcp 172.16.201.1:1807  12.240.246.128:1807 192.168.241.12:21 182.16.200.2:21
tcp 172.16.201.1:1810  12.240.246.128:1810 192.168.241.12:21 182.16.200.2:21
tcp 172.16.201.1:1811  12.240.246.128:1811 192.168.241.12:20 182.16.200.2:20
tcp 172.16.201.1:1813  12.240.246.128:1813 192.168.241.12:20 182.16.200.2:20
```

The same translation table should be found in the Standby SNAT router, DNS2-ING.

If it is necessary to clear peer SNAT translations from the translation table at the Standby/Backup router, use the `clear ip snat translation peer ip-addrees-active-router refresh` command in EXEC mode. The key word `refresh` provides a fresh dump of the NAT table from the Active/Primary router to ensure NAT tables on SNAT peer routers are in sync.

Time-sensitive applications, such as client-server-based applications, might experience timeout and traffic drop. Tuning the HSRP timers is recommended. In the Cisco lab, FTP sessions experience timeout, with `standby delay reload 60` and `standby preempt delay minimum 60 reload 60 sync 60` configuration. To improve network convergence, the HSRP timers are changed to lower values as shown in Examples 15 and 16.

Example 15 DNS1-INS

```
hostname DNS1-INS
!
interface FastEthernet0
   ip address 192.168.240.1 255.255.255.0
   ip nat inside

   standby delay minimum 20 reload 20
   standby 1 ip 192.168.240.3
   standby 1 priority 105
   standby 1 preempt delay minimum 20 reload 20 sync 10
   standby 1 name HSRP_IN
   standby 1 track FastEthernet1 20
```



```
!
interface FastEthernet1
   ip address 172.16.200.1 255.255.255.0
   ip nat outside
   standby delay minimum 20 reload 20
   standby 2 ip 172.16.200.3
   standby 2 priority 105
   standby 2 preempt delay minimum 20 reload 20 sync 10
   standby 2 name HSRP_OUT
   standby 2 track FastEthernet0 20
```

Example 16 DNS2-INS

```
Hostname DNS2-INS
!
interface FastEthernet0/0
   ip address 192.168.240.2 255.255.255.0
   ip nat inside
   standby delay minimum 20 reload 20
   standby 1 ip 192.168.240.3
   standby 1 priority 95
   standby 1 preempt delay minimum 20 reload 20 sync 10
   standby 1 name HSRP_IN
   standby 1 track FastEthernet0/1 20
!
interface FastEthernet0/1
   ip address 172.16.200.2 255.255.255.0
   ip nat outside
   standby delay minimum 20 reload 20
   standby 2 ip 172.16.200.3
   standby 2 priority 95
   standby 2 preempt delay minimum 20 reload 20 sync 10
   standby 2 name HSRP_OUT
   standby 2 track FastEthernet0 20
```


Lab tests show the delay caused by failover is less than 30 seconds.

RELATED TECHNOLOGIES

NAT for MPLS VPNs

It may be more appropriate for the service provider to handle the NAT function; however, NAT can be deployed on the enterprise edge. An enterprise customer that purchases application services or outsources some portion of the processing workload would likely want to take advantage of NAT services, if possible.

Cisco NAT for MPLS VPNs extends NAT so that service providers can establish the translation function within an MPLS network. This is the subject of a separate paper.

CONCLUSION

Cisco continues to enhance core features to provide increased benefit in terms of productivity gained from deployment of a more resilient IP network. SNAT can provide higher availability to applications that use NAT services. Cisco will continue to develop more robust and automated features relative to NAT, which will lower administrative costs and increase the return on investment in network technology.

REFERENCES

Cisco High-Availability Initiatives

Cisco Globally Resilient IP: http://www.cisco.com/go/grip/

Cisco NAT: http://www.cisco.com/go/nat

HSRP Features and Functions: http://www.cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a0080094a91.shtml

Corporate 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: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
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

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco.com Website at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Printed in USA C11-363927-00 09/06