



Information About NAT

This chapter provides an overview of how Network Address Translation (NAT) works on the ASA. This chapter includes the following sections:

- [Why Use NAT?, page 3-1](#)
- [NAT Terminology, page 3-2](#)
- [NAT Types, page 3-3](#)
- [NAT in Routed and Transparent Mode, page 3-10](#)
- [NAT and IPv6, page 3-13](#)
- [How NAT is Implemented, page 3-13](#)
- [NAT Rule Order, page 3-18](#)
- [Routing NAT Packets, page 3-19](#)
- [NAT for VPN, page 3-22](#)
- [DNS and NAT, page 3-28](#)
- [Where to Go Next, page 3-33](#)



Note

To start configuring NAT, see [Chapter 4, “Configuring Network Object NAT,”](#) or [Chapter 5, “Configuring Twice NAT.”](#)

Why Use NAT?

Each computer and device within an IP network is assigned a unique IP address that identifies the host. Because of a shortage of public IPv4 addresses, most of these IP addresses are private, not routable anywhere outside of the private company network. RFC 1918 defines the private IP addresses you can use internally that should not be advertised:

- 10.0.0.0 through 10.255.255.255
- 172.16.0.0 through 172.31.255.255
- 192.168.0.0 through 192.168.255.255

One of the main functions of NAT is to enable private IP networks to connect to the Internet. NAT replaces a private IP address with a public IP address, translating the private addresses in the internal private network into legal, routable addresses that can be used on the public Internet. In this way, NAT conserves public addresses because it can be configured to advertise at a minimum only one public address for the entire network to the outside world.

Other functions of NAT include:

- Security—Keeping internal IP addresses hidden discourages direct attacks.
- IP routing solutions—Overlapping IP addresses are not a problem when you use NAT.
- Flexibility—You can change internal IP addressing schemes without affecting the public addresses available externally; for example, for a server accessible to the Internet, you can maintain a fixed IP address for Internet use, but internally, you can change the server address.
- Translating between IPv4 and IPv6 (Routed mode only) —If you want to connect an IPv6 network to an IPv4 network, NAT lets you translate between the two types of addresses.

**Note**

NAT is not required. If you do not configure NAT for a given set of traffic, that traffic will not be translated, but will have all of the security policies applied as normal.

NAT Terminology

This document uses the following terminology:

- Real address/host/network/interface—The real address is the address that is defined on the host, before it is translated. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the inside network would be the “real” network. Note that you can translate any network connected to the ASA, not just an inside network. Therefore if you configure NAT to translate outside addresses, “real” can refer to the outside network when it accesses the inside network.
- Mapped address/host/network/interface—The mapped address is the address that the real address is translated to. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the outside network would be the “mapped” network.
- Bidirectional initiation—Static NAT allows connections to be initiated *bidirectionally*, meaning both to the host and from the host.
- Source and destination NAT—For any given packet, both the source and destination IP addresses are compared to the NAT rules, and one or both can be translated/untranslated. For static NAT, the rule is bidirectional, so be aware that “source” and “destination” are used in commands and descriptions throughout this guide even though a given connection might originate at the “destination” address.

NAT Types

- [NAT Types Overview, page 3-3](#)
- [Static NAT, page 3-3](#)
- [Dynamic NAT, page 3-7](#)
- [Dynamic PAT, page 3-8](#)
- [Identity NAT, page 3-10](#)

NAT Types Overview

You can implement NAT using the following methods:

- **Static NAT**—A consistent mapping between a real and mapped IP address. Allows bidirectional traffic initiation. See the [“Static NAT” section on page 3-3](#).
- **Dynamic NAT**—A group of real IP addresses are mapped to a (usually smaller) group of mapped IP addresses, on a first come, first served basis. Only the real host can initiate traffic. See the [“Dynamic NAT” section on page 3-7](#).
- **Dynamic Port Address Translation (PAT)**—A group of real IP addresses are mapped to a single IP address using a unique source port of that IP address. See the [“Dynamic PAT” section on page 3-8](#).
- **Identity NAT**—A real address is statically translated to itself, essentially bypassing NAT. You might want to configure NAT this way when you want to translate a large group of addresses, but then want to exempt a smaller subset of addresses. See the [“Identity NAT” section on page 3-10](#).

Static NAT

This section describes static NAT and includes the following topics:

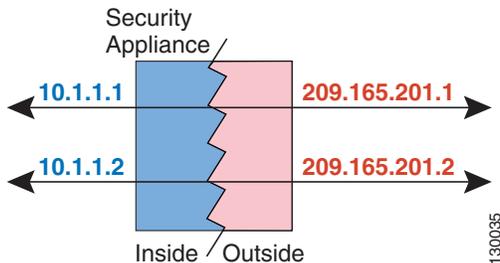
- [Information About Static NAT, page 3-3](#)
- [Information About Static NAT with Port Translation, page 3-4](#)
- [Information About One-to-Many Static NAT, page 3-5](#)
- [Information About Other Mapping Scenarios \(Not Recommended\), page 3-6](#)

Information About Static NAT

Static NAT creates a fixed translation of a real address to a mapped address. Because the mapped address is the same for each consecutive connection, static NAT allows bidirectional connection initiation, both to and from the host (if an access rule exists that allows it). With dynamic NAT and PAT, on the other hand, each host uses a different address or port for each subsequent translation, so bidirectional initiation is not supported.

Figure 3-1 shows a typical static NAT scenario. The translation is always active so both real and remote hosts can initiate connections.

Figure 3-1 Static NAT



Note

You can disable bidirectionality if desired.

Information About Static NAT with Port Translation

Static NAT with port translation lets you specify a real and mapped protocol (TCP or UDP) and port. This section includes the following topics:

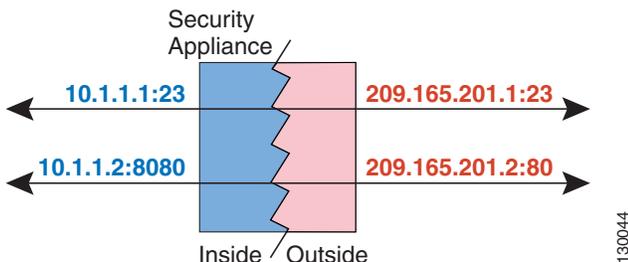
- [Information About Static NAT with Port Address Translation, page 3-4](#)
- [Static NAT with Identity Port Translation, page 3-5](#)
- [Static NAT with Port Translation for Non-Standard Ports, page 3-5](#)
- [Static Interface NAT with Port Translation, page 3-5](#)

Information About Static NAT with Port Address Translation

When you specify the port with static NAT, you can choose to map the port and/or the IP address to the same value or to a different value.

Figure 3-2 shows a typical static NAT with port translation scenario showing both a port that is mapped to itself and a port that is mapped to a different value; the IP address is mapped to a different value in both cases. The translation is always active so both translated and remote hosts can initiate connections.

Figure 3-2 Typical Static NAT with Port Translation Scenario



**Note**

For applications that require application inspection for secondary channels (for example, FTP and VoIP), the ASA automatically translates the secondary ports.

Static NAT with Identity Port Translation

The following static NAT with port translation example provides a single address for remote users to access FTP, HTTP, and SMTP. These servers are actually different devices on the real network, but for each server, you can specify static NAT with port translation rules that use the same mapped IP address, but different ports.

Static NAT with Port Translation for Non-Standard Ports

You can also use static NAT with port translation to translate a well-known port to a non-standard port or vice versa. For example, if inside web servers use port 8080, you can allow outside users to connect to port 80, and then undo translation to the original port 8080. Similarly, to provide extra security, you can tell web users to connect to non-standard port 6785, and then undo translation to port 80.

Static Interface NAT with Port Translation

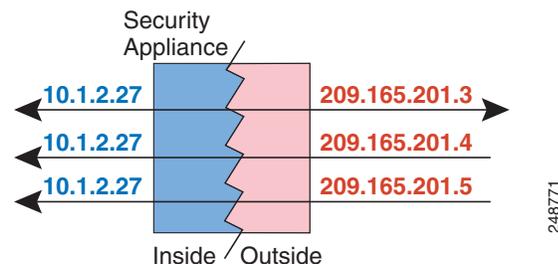
You can configure static NAT to map a real address to an interface address/port combination. For example, if you want to redirect Telnet access for the ASA outside interface to an inside host, then you can map the inside host IP address/port 23 to the ASA interface address/port 23. (Note that although Telnet to the ASA is not allowed to the lowest security interface, static NAT with interface port translation redirects the Telnet session instead of denying it).

Information About One-to-Many Static NAT

Typically, you configure static NAT with a one-to-one mapping. However, in some cases, you might want to configure a single real address to several mapped addresses (one-to-many). When you configure one-to-many static NAT, when the real host initiates traffic, it always uses the first mapped address. However, for traffic initiated to the host, you can initiate traffic to any of the mapped addresses, and they will be untranslated to the single real address.

Figure 3-3 shows a typical one-to-many static NAT scenario. Because initiation by the real host always uses the first mapped address, the translation of real host IP/1st mapped IP is technically the only bidirectional translation.

Figure 3-3 One-to-Many Static NAT



For example, you have a load balancer at 10.1.2.27. Depending on the URL requested, it redirects traffic to the correct web server.

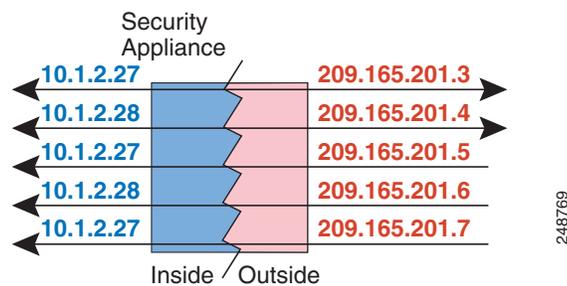
Information About Other Mapping Scenarios (Not Recommended)

The ASA has the flexibility to allow any kind of static mapping scenario: one-to-one, one-to-many, but also few-to-many, many-to-few, and many-to-one mappings. We recommend using only one-to-one or one-to-many mappings. These other mapping options might result in unintended consequences.

Functionally, few-to-many is the same as one-to-many; but because the configuration is more complicated and the actual mappings may not be obvious at a glance, we recommend creating a one-to-many configuration for each real address that requires it. For example, for a few-to-many scenario, the few real addresses are mapped to the many mapped addresses in order (A to 1, B to 2, C to 3). When all real addresses are mapped, the next mapped address is mapped to the first real address, and so on until all mapped addresses are mapped (A to 4, B to 5, C to 6). This results in multiple mapped addresses for each real address. Just like a one-to-many configuration, only the first mappings are bidirectional; subsequent mappings allow traffic to be initiated *to* the real host, but all traffic *from* the real host uses only the first mapped address for the source.

Figure 3-4 shows a typical few-to-many static NAT scenario.

Figure 3-4 Few-to-Many Static NAT



For a many-to-few or many-to-one configuration, where you have more real addresses than mapped addresses, you run out of mapped addresses before you run out of real addresses. Only the mappings between the lowest real IP addresses and the mapped pool result in bidirectional initiation. The remaining higher real addresses can initiate traffic, but traffic cannot be initiated to them (returning traffic for a connection is directed to the correct real address because of the unique 5-tuple (source IP, destination IP, source port, destination port, protocol) for the connection).

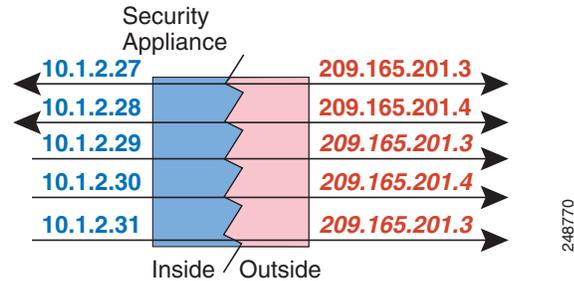


Note

Many-to-few or many-to-one NAT is not PAT. If two real hosts use the same source port number and go to the same outside server and the same TCP destination port, and both hosts are translated to the same IP address, then both connections will be reset because of an address conflict (the 5-tuple is not unique).

Figure 3-5 shows a typical many-to-few static NAT scenario.

Figure 3-5 Many-to-Few Static NAT



Instead of using a static rule this way, we suggest that you create a one-to-one rule for the traffic that needs bidirectional initiation, and then create a dynamic rule for the rest of your addresses.

Dynamic NAT

This section describes dynamic NAT and includes the following topics:

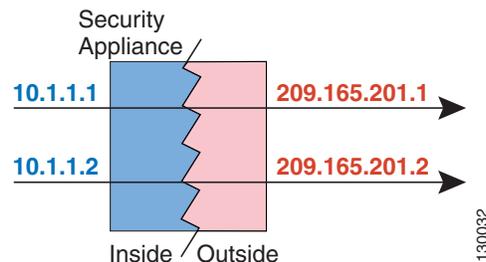
- [Information About Dynamic NAT, page 3-7](#)
- [Dynamic NAT Disadvantages and Advantages, page 3-8](#)

Information About Dynamic NAT

Dynamic NAT translates a group of real addresses to a pool of mapped addresses that are routable on the destination network. The mapped pool typically includes fewer addresses than the real group. When a host you want to translate accesses the destination network, the ASA assigns the host an IP address from the mapped pool. The translation is created only when the real host initiates the connection. The translation is in place only for the duration of the connection, and a given user does not keep the same IP address after the translation times out. Users on the destination network, therefore, cannot initiate a reliable connection to a host that uses dynamic NAT, even if the connection is allowed by an access rule.

Figure 3-6 shows a typical dynamic NAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back.

Figure 3-6 Dynamic NAT



**Note**

For the duration of the translation, a remote host can initiate a connection to the translated host if an access rule allows it. Because the address is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.

Dynamic NAT Disadvantages and Advantages

Dynamic NAT has these disadvantages:

- If the mapped pool has fewer addresses than the real group, you could run out of addresses if the amount of traffic is more than expected.

Use PAT or a PAT fallback method if this event occurs often because PAT provides over 64,000 translations using ports of a single address.

- You have to use a large number of routable addresses in the mapped pool, and routable addresses may not be available in large quantities.

The advantage of dynamic NAT is that some protocols cannot use PAT. PAT does not work with the following:

- IP protocols that do not have a port to overload, such as GRE version 0.
- Some multimedia applications that have a data stream on one port, the control path on another port, and are not open standard.

See the [“Default Settings” section on page 8-4](#) for more information about NAT and PAT support.

Dynamic PAT

This section describes dynamic PAT and includes the following topics:

- [Information About Dynamic PAT, page 3-8](#)
- [Per-Session PAT vs. Multi-Session PAT, page 3-9](#)
- [Dynamic PAT Disadvantages and Advantages, page 3-9](#)

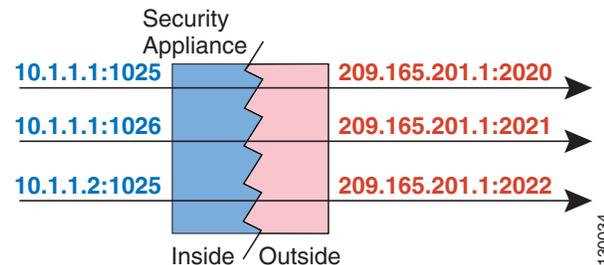
Information About Dynamic PAT

Dynamic PAT translates multiple real addresses to a single mapped IP address by translating the real address and source port to the mapped address and a unique port. If available, the real source port number is used for the mapped port. However, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 0 to 511, 512 to 1023, and 1024 to 65535. Therefore, ports below 1024 have only a small PAT pool that can be used. If you have a lot of traffic that uses the lower port ranges, you can specify a flat range of ports to be used instead of the three unequal-sized tiers.

Each connection requires a separate translation session because the source port differs for each connection. For example, 10.1.1.1:1025 requires a separate translation from 10.1.1.1:1026.

Figure 3-7 shows a typical dynamic PAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back. The mapped address is the same for each translation, but the port is dynamically assigned.

Figure 3-7 Dynamic PAT



After the connection expires, the port translation also expires. For multi-session PAT, the PAT timeout is used, 30 seconds by default. For per-session PAT, the xlate is immediately removed. Users on the destination network cannot reliably initiate a connection to a host that uses PAT (even if the connection is allowed by an access rule).



Note

For the duration of the translation, a remote host can initiate a connection to the translated host if an access rule allows it. Because the port address (both real and mapped) is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.

Per-Session PAT vs. Multi-Session PAT

The per-session PAT feature improves the scalability of PAT and, for clustering, allows each member unit to own PAT connections; multi-session PAT connections have to be forwarded to and owned by the master unit. At the end of a per-session PAT session, the ASA sends a reset and immediately removes the xlate. This reset causes the end node to immediately release the connection, avoiding the TIME_WAIT state. Multi-session PAT, on the other hand, uses the PAT timeout, by default 30 seconds. For “hit-and-run” traffic, such as HTTP or HTTPS, the per-session feature can dramatically increase the connection rate supported by one address. Without the per-session feature, the maximum connection rate for one address for an IP protocol is approximately 2000 per second. With the per-session feature, the connection rate for one address for an IP protocol is $65535/average-lifetime$.

By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate. For traffic that can benefit from multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT by creating a per-session deny rule. See the [“Configuring Per-Session PAT Rules”](#) section on page 4-16.

Dynamic PAT Disadvantages and Advantages

Dynamic PAT lets you use a single mapped address, thus conserving routable addresses. You can even use the ASA interface IP address as the PAT address.

Dynamic PAT does not work with some multimedia applications that have a data stream that is different from the control path. See the [“Default Settings”](#) section on page 8-4 for more information about NAT and PAT support.

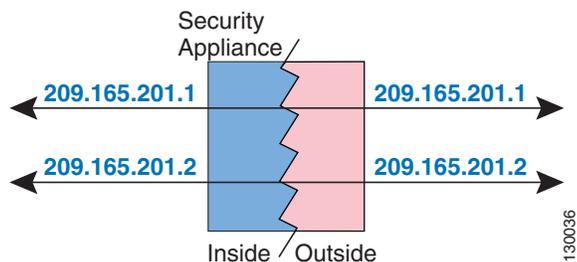
Dynamic PAT may also create a large number of connections appearing to come from a single IP address, and servers might interpret the traffic as a DoS attack. You can configure a PAT pool of addresses and use a round-robin assignment of PAT addresses to mitigate this situation.

Identity NAT

You might have a NAT configuration in which you need to translate an IP address to itself. For example, if you create a broad rule that applies NAT to every network, but want to exclude one network from NAT, you can create a static NAT rule to translate an address to itself. Identity NAT is necessary for remote access VPN, where you need to exempt the client traffic from NAT.

Figure 3-8 shows a typical identity NAT scenario.

Figure 3-8 Identity NAT



NAT in Routed and Transparent Mode

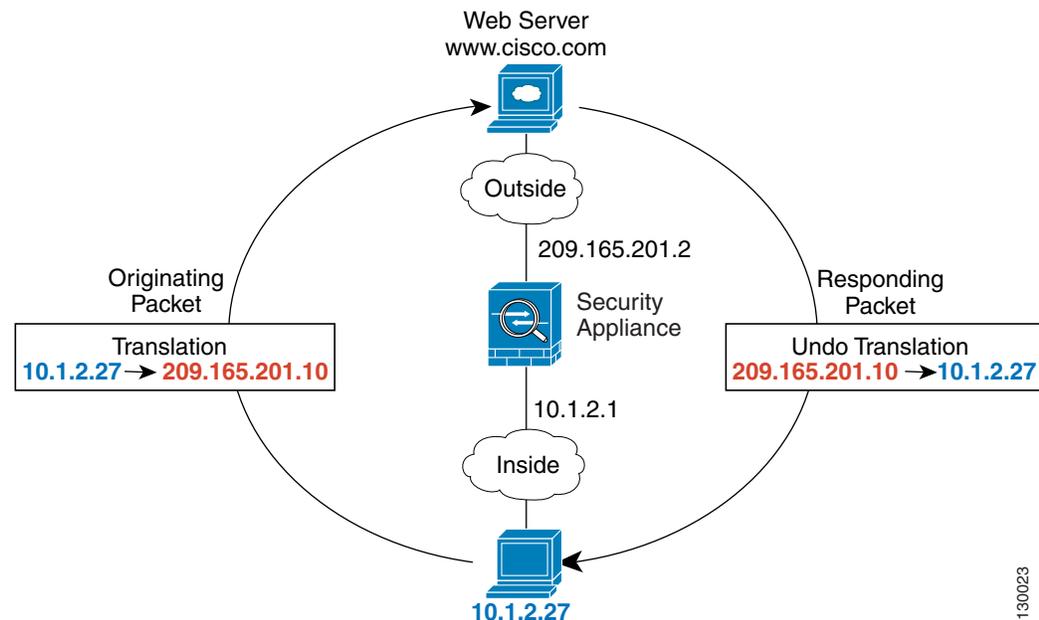
You can configure NAT in both routed and transparent firewall mode. This section describes typical usage for each firewall mode and includes the following topics:

- [NAT in Routed Mode, page 3-11](#)
- [NAT in Transparent Mode, page 3-11](#)

NAT in Routed Mode

Figure 3-9 shows a typical NAT example in routed mode, with a private network on the inside.

Figure 3-9 NAT Example: Routed Mode



1. When the inside host at 10.1.2.27 sends a packet to a web server, the real source address of the packet, 10.1.2.27, is changed to a mapped address, 209.165.201.10.
2. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the ASA receives the packet because the ASA performs proxy ARP to claim the packet.
3. The ASA then changes the translation of the mapped address, 209.165.201.10, back to the real address, 10.1.2.27, before sending it to the host.

NAT in Transparent Mode

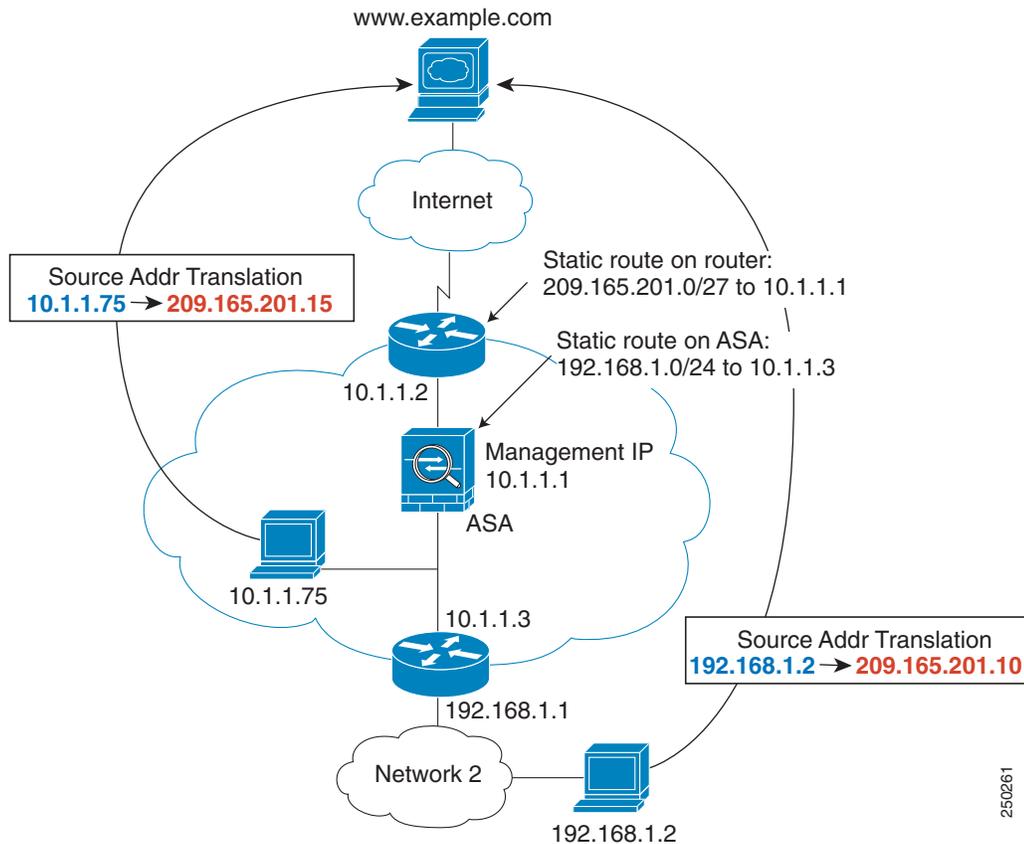
Using NAT in transparent mode eliminates the need for the upstream or downstream routers to perform NAT for their networks.

NAT in transparent mode has the following requirements and limitations:

- Because the transparent firewall does not have any interface IP addresses, you cannot use interface PAT.
- ARP inspection is not supported. Moreover, if for some reason a host on one side of the ASA sends an ARP request to a host on the other side of the ASA, and the initiating host real address is mapped to a different address on the same subnet, then the real address remains visible in the ARP request.
- Translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

Figure 3-10 shows a typical NAT scenario in transparent mode, with the same network on the inside and outside interfaces. The transparent firewall in this scenario is performing the NAT service so that the upstream router does not have to perform NAT.

Figure 3-10 NAT Example: Transparent Mode



1. When the inside host at 10.1.1.75 sends a packet to a web server, the real source address of the packet, 10.1.1.75, is changed to a mapped address, 209.165.201.15.
2. When the server responds, it sends the response to the mapped address, 209.165.201.15, and the ASA receives the packet because the upstream router includes this mapped network in a static route directed to the ASA management IP address. See the [“Mapped Addresses and Routing”](#) section on page 3-19 for more information about required routes.
3. The ASA then undoes the translation of the mapped address, 209.165.201.15, back to the real address, 10.1.1.75. Because the real address is directly-connected, the ASA sends it directly to the host.
4. For host 192.168.1.2, the same process occurs, except for returning traffic, the ASA looks up the route in its routing table and sends the packet to the downstream router at 10.1.1.3 based on the ASA static route for 192.168.1.0/24. See the [“Transparent Mode Routing Requirements for Remote Networks”](#) section on page 3-21 for more information about required routes.

NAT and IPv6

You can use NAT to translate between IPv6 networks, and also to translate between IPv4 and IPv6 networks (routed mode only). We recommend the following best practices:

- **NAT66 (IPv6-to-IPv6)**—We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (twice NAT only).
- **NAT46 (IPv4-to-IPv6)**—We recommend using static NAT. Because the IPv6 address space is so much larger than the IPv4 address space, you can easily accommodate a static translation. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (twice NAT only). When translating to an IPv6 subnet (/96 or lower), the resulting mapped address is by default an IPv4-embedded IPv6 address, where the 32-bits of the IPv4 address is embedded after the IPv6 prefix. For example, if the IPv6 prefix is a /96 prefix, then the IPv4 address is appended in the last 32-bits of the address. For example, if you map 192.168.1.0/24 to 201b::0/96, then 192.168.1.4 will be mapped to 201b::0.192.168.1.4 (shown with mixed notation). If the prefix is smaller, such as /64, then the IPv4 address is appended after the prefix, and a suffix of 0s is appended after the IPv4 address. You can also optionally translate the addresses net-tonet, where the first IPv4 address maps to the first IPv6 address, the second to the second, and so on.
- **NAT64 (IPv6-to-IPv4)**—You may not have enough IPv4 addresses to accommodate the number of IPv6 addresses. We recommend using a dynamic PAT pool to provide a large number of IPv4 translations.

For specific implementation guidelines and limitations, see the configuration chapters.

5.

How NAT is Implemented

The ASA can implement address translation in two ways: *network object NAT* and *twice NAT*. This section includes the following topics:

- [Main Differences Between Network Object NAT and Twice NAT, page 3-13](#)
- [Information About Network Object NAT, page 3-14](#)
- [Information About Twice NAT, page 3-14](#)

Main Differences Between Network Object NAT and Twice NAT

The main differences between these two NAT types are:

- How you define the real address.
 - **Network object NAT**—You define NAT as a parameter for a network object. A network object names an IP host, range, or subnet so you can then use the object in configuration instead of the actual IP addresses. The network object IP address serves as the real address. This method lets you easily add NAT to network objects that might already be used in other parts of your configuration.
 - **Twice NAT**—You identify a network object or network object group for both the real and mapped addresses. In this case, NAT is not a parameter of the network object; the network object or group is a parameter of the NAT configuration. The ability to use a network object *group* for the real address means that twice NAT is more scalable.

- How source and destination NAT is implemented.
 - Network object NAT— Each rule can apply to either the source or destination of a packet. So two rules might be used, one for the source IP address, and one for the destination IP address. These two rules cannot be tied together to enforce a specific translation for a source/destination combination.
 - Twice NAT—A single rule translates both the source and destination. A matching packet only matches the one rule, and further rules are not checked. Even if you do not configure the optional destination address for twice NAT, a matching packet still only matches one twice NAT rule. The source and destination are tied together, so you can enforce different translations depending on the source/destination combination. For example, sourceA/destinationA can have a different translation than sourceA/destinationB.
 - Order of NAT Rules.
 - Network object NAT—Automatically ordered in the NAT table.
 - Twice NAT—Manually ordered in the NAT table (before or after network object NAT rules).
- See the “[NAT Rule Order](#)” section on page 3-18 for more information.

We recommend using network object NAT unless you need the extra features that twice NAT provides. Network object NAT is easier to configure, and might be more reliable for applications such as Voice over IP (VoIP). (For VoIP, because twice NAT is applicable only between two objects, you might see a failure in the translation of indirect addresses that do not belong to either of the objects.)

Information About Network Object NAT

All NAT rules that are configured as a parameter of a network object are considered to be *network object NAT* rules. Network object NAT is a quick and easy way to configure NAT for a network object, which can be a single IP address, a range of addresses, or a subnet.

After you configure the network object, you can then identify the mapped address for that object, either as an inline address or as another network object or network object group.

When a packet enters the ASA, both the source and destination IP addresses are checked against the network object NAT rules. The source and destination address in the packet can be translated by separate rules if separate matches are made. These rules are not tied to each other; different combinations of rules can be used depending on the traffic.

Because the rules are never paired, you cannot specify that sourceA/destinationA should have a different translation than sourceA/destinationB. Use twice NAT for that kind of functionality (twice NAT lets you identify the source and destination address in a single rule).

To start configuring network object NAT, see [Chapter 4, “Configuring Network Object NAT.”](#)

Information About Twice NAT

Twice NAT lets you identify both the source and destination address in a single rule. Specifying both the source and destination addresses lets you specify that sourceA/destinationA can have a different translation than sourceA/destinationB.

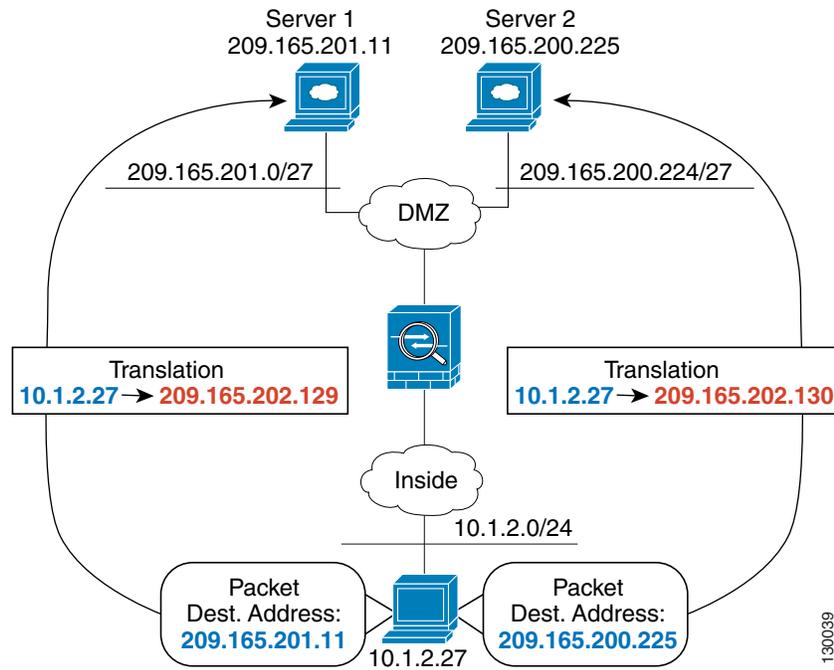
The destination address is optional. If you specify the destination address, you can either map it to itself (identity NAT), or you can map it to a different address. The destination mapping is always a static mapping.

Twice NAT also lets you use service objects for static NAT with port translation; network object NAT only accepts inline definition.

To start configuring twice NAT, see [Chapter 5, “Configuring Twice NAT.”](#)

[Figure 3-11](#) shows a host on the 10.1.2.0/24 network accessing two different servers. When the host accesses the server at 209.165.201.11, the real address is translated to 209.165.202.129. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130.

Figure 3-11 Twice NAT with Different Destination Addresses



130039

Figure 3-12 shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for web services, the real address is translated to 209.165.202.129. When the host accesses the same server for Telnet services, the real address is translated to 209.165.202.130.

Figure 3-12 Twice NAT with Different Destination Ports

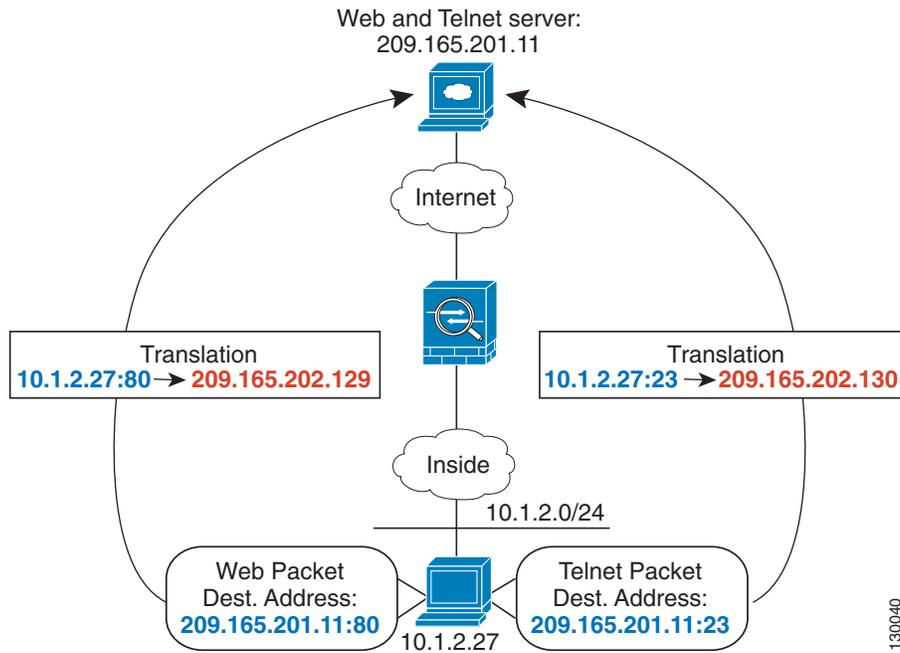
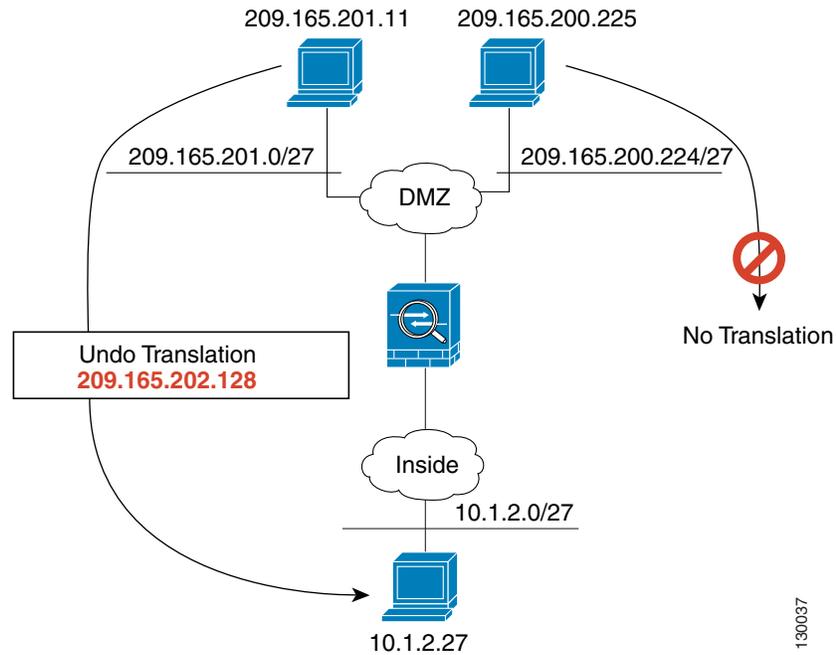


Figure 3-13 shows a remote host connecting to a mapped host. The mapped host has a twice static NAT translation that translates the real address only for traffic to and from the 209.165.201.0/27 network. A translation does not exist for the 209.165.200.224/27 network, so the translated host cannot connect to that network, nor can a host on that network connect to the translated host.

Figure 3-13 Twice Static NAT with Destination Address Translation



130037

NAT Rule Order

Network object NAT rules and twice NAT rules are stored in a single table that is divided into three sections. Section 1 rules are applied first, then section 2, and finally section 3. [Table 3-1](#) shows the order of rules within each section.

Table 3-1 NAT Rule Table

Table Section	Rule Type	Order of Rules within the Section
Section 1	Twice NAT	<p>Applied on a first match basis, in the order they appear in the configuration. By default, twice NAT rules are added to section 1.</p> <p>Note If you configure EasyVPN remote, the ASA dynamically adds invisible NAT rules to the end of this section. Be sure that you do not configure a twice NAT rule in this section that might match your VPN traffic, instead of matching the invisible rule. If VPN does not work due to NAT failure, consider adding twice NAT rules to section 3 instead.</p>
Section 2	Network object NAT	<p>Section 2 rules are applied in the following order, as automatically determined by the ASA:</p> <ol style="list-style-type: none"> 1. Static rules. 2. Dynamic rules. <p>Within each rule type, the following ordering guidelines are used:</p> <ol style="list-style-type: none"> a. Quantity of real IP addresses—From smallest to largest. For example, an object with one address will be assessed before an object with 10 addresses. b. For quantities that are the same, then the IP address number is used, from lowest to highest. For example, 10.1.1.0 is assessed before 11.1.1.0. c. If the same IP address is used, then the name of the network object is used, in alphabetical order. For example, abracadabra is assessed before catwoman.
Section 3	Twice NAT	<p>Section 3 rules are applied on a first match basis, in the order they appear in the configuration. You can specify whether to add a twice NAT rule to section 3 when you add the rule.</p>

For section 2 rules, for example, you have the following IP addresses defined within network objects:

```

192.168.1.0/24 (static)
192.168.1.0/24 (dynamic)
10.1.1.0/24 (static)
192.168.1.1/32 (static)
172.16.1.0/24 (dynamic) (object def)
172.16.1.0/24 (dynamic) (object abc)

```

The resultant ordering would be:

192.168.1.1/32 (static)
10.1.1.0/24 (static)
192.168.1.0/24 (static)
172.16.1.0/24 (dynamic) (object abc)
172.16.1.0/24 (dynamic) (object def)
192.168.1.0/24 (dynamic)

NAT Interfaces

You can configure a NAT rule to apply to any interface (in other words, all interfaces), or you can identify specific real and mapped interfaces. You can also specify any interface for the real address, and a specific interface for the mapped address, or vice versa.

For example, you might want to specify any interface for the real address and specify the outside interface for the mapped address if you use the same private addresses on multiple interfaces, and you want to translate them all to the same global pool when accessing the outside.

**Note**

For transparent mode, you must choose specific source and destination interfaces.

Routing NAT Packets

The ASA needs to be the destination for any packets sent to the mapped address. The ASA also needs to determine the egress interface for any packets it receives destined for mapped addresses. This section describes how the ASA handles accepting and delivering packets with NAT, and includes the following topics:

- [Mapped Addresses and Routing, page 3-19](#)
- [Transparent Mode Routing Requirements for Remote Networks, page 3-21](#)
- [Determining the Egress Interface, page 3-22](#)

Mapped Addresses and Routing

When you translate the real address to a mapped address, the mapped address you choose determines how to configure routing, if necessary, for the mapped address.

See additional guidelines about mapped IP addresses in [Chapter 4, “Configuring Network Object NAT,”](#) and [Chapter 5, “Configuring Twice NAT.”](#)

See the following mapped address types:

- Addresses on the same network as the mapped interface.

If you use addresses on the same network as the mapped interface, the ASA uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the ASA does not have to be the gateway for any additional networks. This solution is ideal if the outside network contains an adequate number of

free addresses, a consideration if you are using a 1:1 translation like dynamic NAT or static NAT. Dynamic PAT greatly extends the number of translations you can use with a small number of addresses, so even if the available addresses on the outside network is small, this method can be used. For PAT, you can even use the IP address of the mapped interface.



Note If you configure the mapped interface to be any interface, and you specify a mapped address on the same network as one of the mapped interfaces, then if an ARP request for that mapped address comes in on a *different* interface, then you need to manually configure an ARP entry for that network on the ingress interface, specifying its MAC address (see the **arp** command). Typically, if you specify any interface for the mapped interface, then you use a unique network for the mapped addresses, so this situation would not occur.

- Addresses on a unique network.

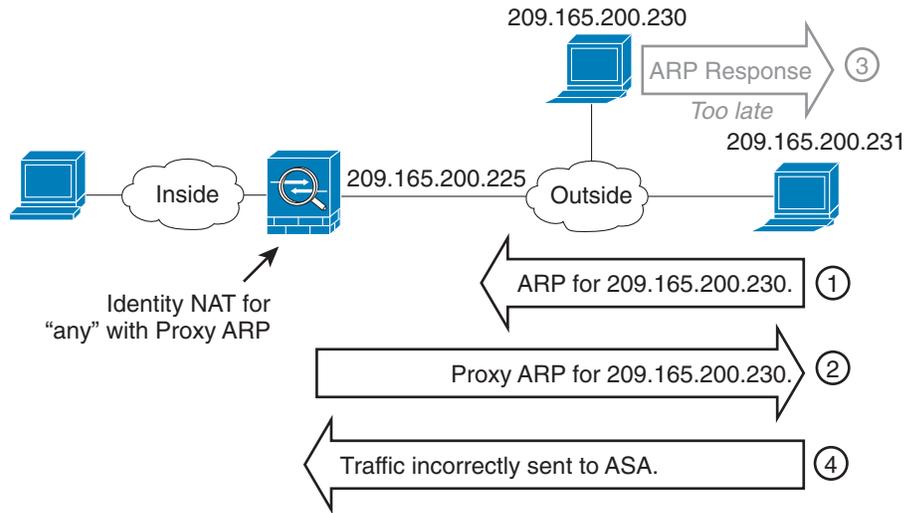
If you need more addresses than are available on the mapped interface network, you can identify addresses on a different subnet. The upstream router needs a static route for the mapped addresses that points to the ASA. Alternatively for routed mode, you can configure a static route on the ASA for the mapped addresses, and then redistribute the route using your routing protocol. For transparent mode, if the real host is directly-connected, configure the static route on the upstream router to point to the ASA: specify the bridge group IP address. For remote hosts in transparent mode, in the static route on the upstream router, you can alternatively specify the downstream router IP address.

- The same address as the real address (identity NAT).

The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. **Note:** You can also disable proxy ARP for regular static NAT if desired, in which case you need to be sure to have proper routes on the upstream router.

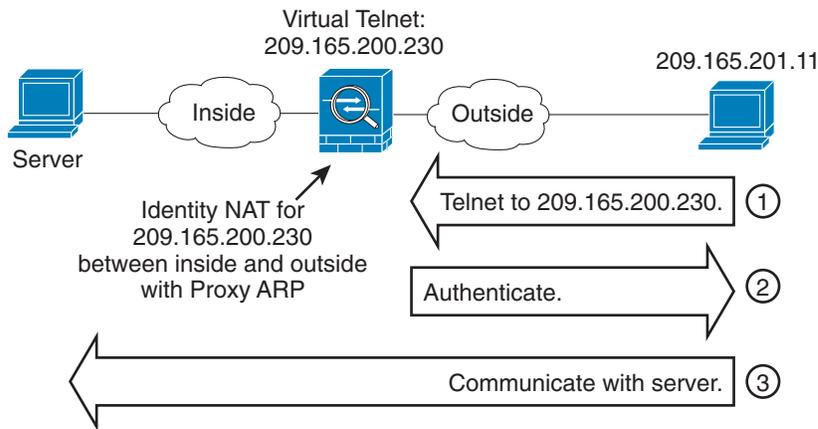
Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues. For example, if you configure a broad identity NAT rule for “any” IP address, then leaving proxy ARP enabled can cause problems for hosts on the network directly-connected to the mapped interface. In this case, when a host on the mapped network wants to communicate with another host on the same network, then the address in the ARP request matches the NAT rule (which matches “any” address). The ASA will then proxy ARP for the address, even though the packet is not actually destined for the ASA. (Note that this problem occurs even if you have a twice NAT rule; although the NAT rule must match both the source and destination addresses, the proxy ARP decision is made only on the “source” address). If the ASA ARP response is received before the actual host ARP response, then traffic will be mistakenly sent to the ASA (see [Figure 3-14](#)).

Figure 3-14 Proxy ARP Problems with Identity NAT



In rare cases, you need proxy ARP for identity NAT; for example for virtual Telnet. When using AAA for network access, a host needs to authenticate with the ASA using a service like Telnet before any other traffic can pass. You can configure a virtual Telnet server on the ASA to provide the necessary login. When accessing the virtual Telnet address from the outside, you must configure an identity NAT rule for the address specifically for the proxy ARP functionality. Due to internal processes for virtual Telnet, proxy ARP lets the ASA keep traffic destined for the virtual Telnet address rather than send the traffic out the source interface according to the NAT rule. (See Figure 3-15).

Figure 3-15 Proxy ARP and Virtual Telnet



Transparent Mode Routing Requirements for Remote Networks

When you use NAT in transparent mode, some types of traffic require static routes. See the “MAC Address vs. Route Lookups” section on page 4-5 for more information.

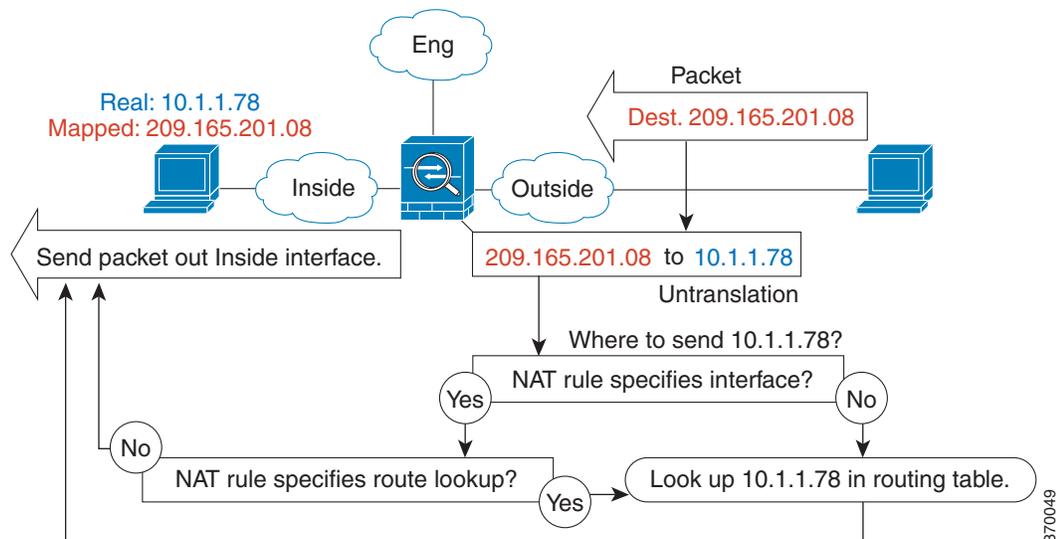
Determining the Egress Interface

When the ASA receives traffic for a mapped address, the ASA untranslates the destination address according to the NAT rule, and then it sends the packet on to the real address. The ASA determines the egress interface for the packet in the following ways:

- Transparent mode—The ASA determines the egress interface for the real address by using the NAT rule; you must specify the source and destination interfaces as part of the NAT rule.
- Routed mode—The ASA determines the egress interface in one of the following ways:
 - You configure the interface in the NAT rule—The ASA uses the NAT rule to determine the egress interface. However, you have the option to always use a route lookup instead. In certain scenarios, a route lookup override is required; for example, see the “[NAT and VPN Management Access](#)” section on page 3-26.
 - You do not configure the interface in the NAT rule—The ASA uses a route lookup to determine the egress interface.

Figure 3-16 shows the egress interface selection method in routed mode. In almost all cases, a route lookup is equivalent to the NAT rule interface, but in some configurations, the two methods might differ.

Figure 3-16 Routed Mode Egress Interface Selection



NAT for VPN

- [NAT and Remote Access VPN](#), page 3-23
- [NAT and Site-to-Site VPN](#), page 3-24
- [NAT and VPN Management Access](#), page 3-26
- [Troubleshooting NAT and VPN](#), page 3-28

NAT and Remote Access VPN

Figure 3-17 shows both an inside server (10.1.1.6) and a VPN client (209.165.201.10) accessing the Internet. Unless you configure split tunnelling for the VPN client (where only specified traffic goes through the VPN tunnel), then Internet-bound VPN traffic must also go through the ASA. When the VPN traffic enters the ASA, the ASA decrypts the packet; the resulting packet includes the VPN client local address (10.3.3.10) as the source. For both inside and VPN client local networks, you need a public IP address provided by NAT to access the Internet. The below example uses interface PAT rules. To allow the VPN traffic to exit the same interface it entered, you also need to enable intra-interface communication (AKA “hairpin” networking).

Figure 3-17 Interface PAT for Internet-Bound VPN Traffic (Intra-Interface)

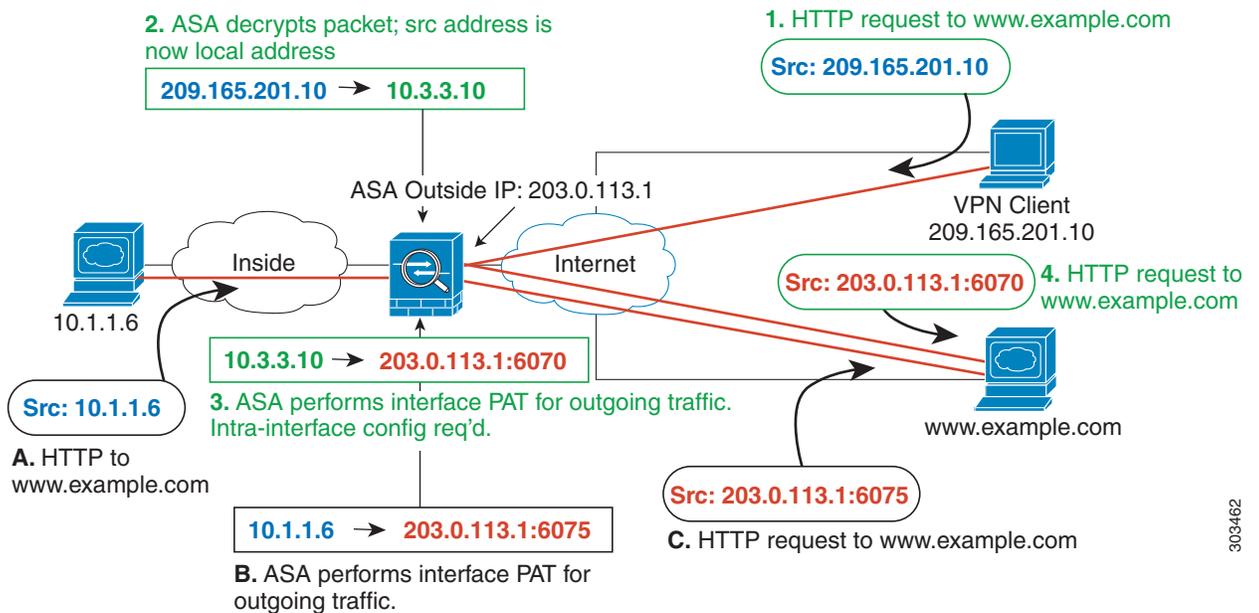
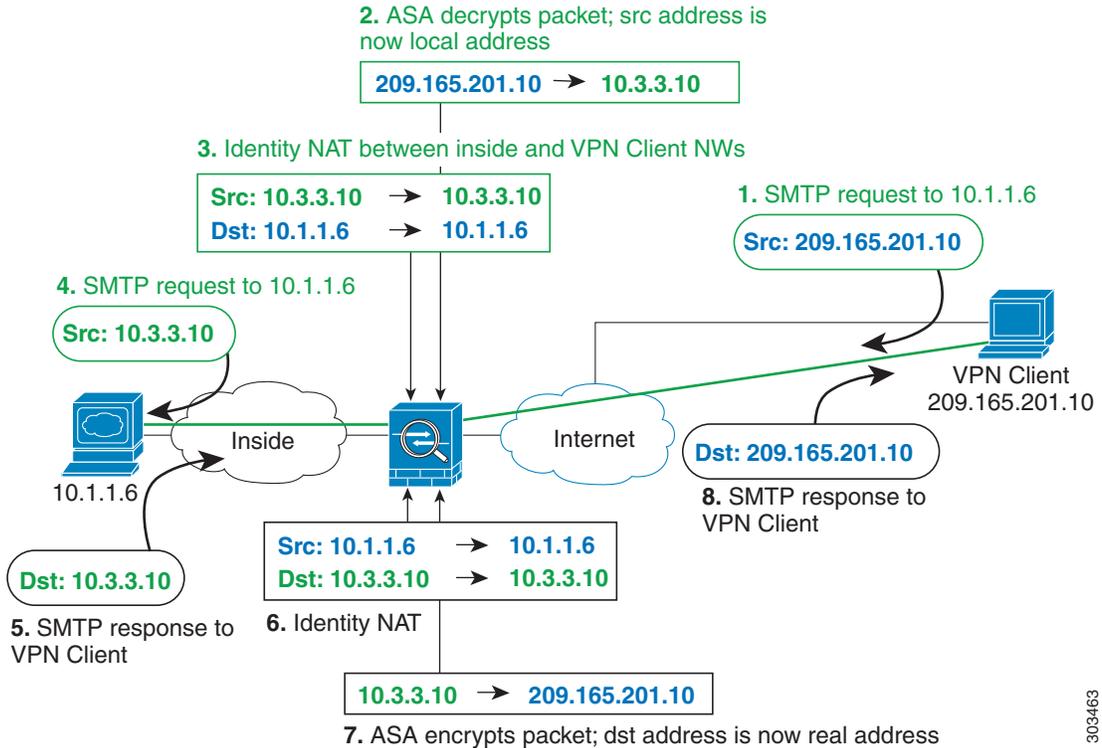


Figure 3-18 shows a VPN client that wants to access an inside mail server. Because the ASA expects traffic between the inside network and any outside network to match the interface PAT rule you set up for Internet access, traffic from the VPN client (10.3.3.10) to the SMTP server (10.1.1.6) will be dropped due to a reverse path failure: traffic from 10.3.3.10 to 10.1.1.6 does not match a NAT rule, but returning traffic from 10.1.1.6 to 10.3.3.10 *should* match the interface PAT rule for outgoing traffic. Because forward and reverse flows do not match, the ASA drops the packet when it is received. To avoid this failure, you need to exempt the inside-to-VPN client traffic from the interface PAT rule by using an identity NAT rule between those networks. Identity NAT simply translates an address to the same address.

Figure 3-18 Identity NAT for VPN Clients



303463

See the following sample NAT configuration for the above network:

```
! Enable hairpin for non-split-tunneled VPN client traffic:
same-security-traffic permit intra-interface

! Identify local VPN network, & perform object interface PAT when going to Internet:
object network vpn_local
  subnet 10.3.3.0 255.255.255.0
  nat (outside,outside) dynamic interface

! Identify inside network, & perform object interface PAT when going to Internet:
object network inside_nw
  subnet 10.1.1.0 255.255.255.0
  nat (inside,outside) dynamic interface

! Use twice NAT to pass traffic between the inside network and the VPN client without
! address translation (identity NAT):
nat (inside,outside) source static inside_nw inside_nw destination static vpn_local
vpn_local
```

NAT and Site-to-Site VPN

Figure 3-19 shows a site-to-site tunnel connecting the Boulder and San Jose offices. For traffic that you want to go to the Internet (for example from 10.1.1.6 in Boulder to www.example.com), you need a public IP address provided by NAT to access the Internet. The below example uses interface PAT rules. However, for traffic that you want to go over the VPN tunnel (for example from 10.1.1.6 in Boulder to 10.2.2.78 in San Jose), you do not want to perform NAT; you need to exempt that traffic by creating an identity NAT rule. Identity NAT simply translates an address to the same address.

Figure 3-19 Interface PAT and Identity NAT for Site-to-Site VPN

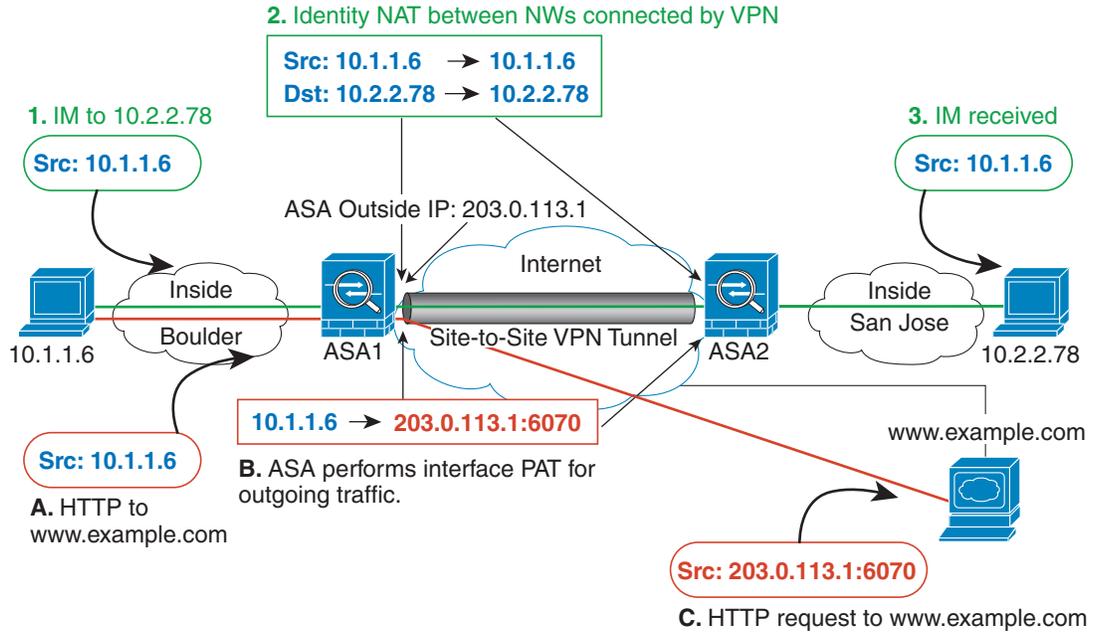
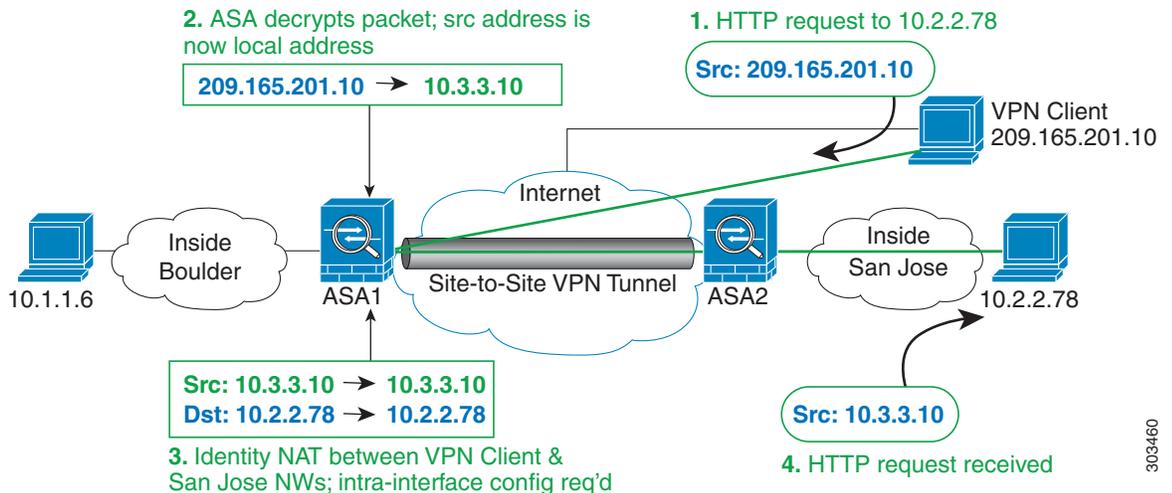


Figure 3-20 shows a VPN client connected to ASA1 (Boulder), with a Telnet request for a server (10.2.2.78) accessible over a site-to-site tunnel between ASA1 and ASA2 (San Jose). Because this is a hairpin connection, you need to enable intra-interface communication, which is also required for non-split-tunneled Internet-bound traffic from the VPN client. You also need to configure identity NAT between the VPN client and the Boulder & San Jose networks, just as you would between any networks connected by VPN to exempt this traffic from outbound NAT rules.

Figure 3-20 VPN Client Access to Site-to-Site VPN



See the following sample NAT configuration for ASA1 (Boulder):

```
! Enable hairpin for VPN client traffic:
same-security-traffic permit intra-interface
```

```
! Identify local VPN network, & perform object interface PAT when going to Internet:
```

```

object network vpn_local
  subnet 10.3.3.0 255.255.255.0
  nat (outside,outside) dynamic interface

! Identify inside Boulder network, & perform object interface PAT when going to Internet:
object network boulder_inside
  subnet 10.1.1.0 255.255.255.0
  nat (inside,outside) dynamic interface

! Identify inside San Jose network for use in twice NAT rule:
object network sanjose_inside
  subnet 10.2.2.0 255.255.255.0

! Use twice NAT to pass traffic between the Boulder network and the VPN client without
! address translation (identity NAT):
nat (inside,outside) source static boulder_inside boulder_inside destination static
vpn_local vpn_local

! Use twice NAT to pass traffic between the Boulder network and San Jose without
! address translation (identity NAT):
nat (inside,outside) source static boulder_inside boulder_inside destination static
sanjose_inside sanjose_inside

! Use twice NAT to pass traffic between the VPN client and San Jose without
! address translation (identity NAT):
nat (outside,outside) source static vpn_local vpn_local destination static sanjose_inside
sanjose_inside

```

See the following sample NAT configuration for ASA2 (San Jose):

```

! Identify inside San Jose network, & perform object interface PAT when going to Internet:
object network sanjose_inside
  subnet 10.2.2.0 255.255.255.0
  nat (inside,outside) dynamic interface

! Identify inside Boulder network for use in twice NAT rule:
object network boulder_inside
  subnet 10.1.1.0 255.255.255.0

! Identify local VPN network for use in twice NAT rule:
object network vpn_local
  subnet 10.3.3.0 255.255.255.0

! Use twice NAT to pass traffic between the San Jose network and Boulder without
! address translation (identity NAT):
nat (inside,outside) source static sanjose_inside sanjose_inside destination static
boulder_inside boulder_inside

! Use twice NAT to pass traffic between the San Jose network and the VPN client without
! address translation (identity NAT):
nat (inside,outside) source static sanjose_inside sanjose_inside destination static
vpn_local vpn_local

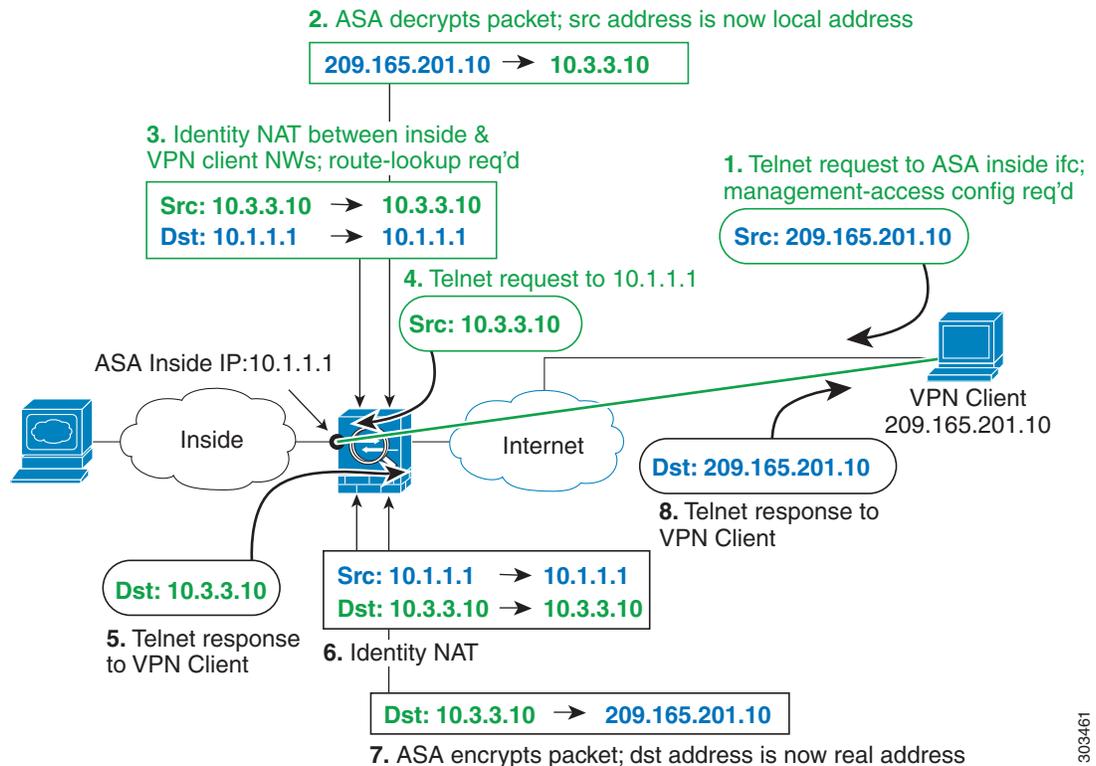
```

NAT and VPN Management Access

When using VPN, you can allow management access to an interface other than the one from which you entered the ASA (see the **management-access** command). For example, if you enter the ASA from the outside interface, the management-access feature lets you connect to the inside interface using ASDM, SSH, Telnet, or SNMP; or you can ping the inside interface.

Figure 3-21 shows a VPN client Telnetting to the ASA inside interface. When you use a management-access interface, and you configure identity NAT according to the “NAT and Remote Access VPN” or “NAT and Site-to-Site VPN” section, you must configure NAT with the route lookup option. Without route lookup, the ASA sends traffic out the interface specified in the NAT command, regardless of what the routing table says; in the below example, the egress interface is the inside interface. You do not want the ASA to send the management traffic out to the inside network; it will never return to the inside interface IP address. The route lookup option lets the ASA send the traffic directly to the inside interface IP address instead of to the inside network. For traffic from the VPN client to a host on the inside network, the route lookup option will still result in the correct egress interface (inside), so normal traffic flow is not affected. See the “Determining the Egress Interface” section on page 3-22 for more information about the route lookup option.

Figure 3-21 VPN Management Access



303461

See the following sample NAT configuration for the above network:

```
! Enable hairpin for non-split-tunneled VPN client traffic:
same-security-traffic permit intra-interface

! Enable management access on inside ifc:
management-access inside

! Identify local VPN network, & perform object interface PAT when going to Internet:
object network vpn_local
  subnet 10.3.3.0 255.255.255.0
  nat (outside,outside) dynamic interface

! Identify inside network, & perform object interface PAT when going to Internet:
object network inside_nw
  subnet 10.1.1.0 255.255.255.0
  nat (inside,outside) dynamic interface
```

```
! Use twice NAT to pass traffic between the inside network and the VPN client without
! address translation (identity NAT), w/route-lookup:
nat (outside,inside) source static vpn_local vpn_local destination static inside_nw
inside_nw route-lookup
```

Troubleshooting NAT and VPN

See the following monitoring tools for troubleshooting NAT issues with VPN:

- Packet tracer—When used correctly, a packet tracer shows which NAT rules a packet is hitting.
- **show nat detail**—Shows hit counts and untranslated traffic for a given NAT rule.
- **show conn all**—Lets you see active connections including to and from the box traffic.

To familiarize yourself with a non-working configuration vs. a working configuration, you can perform the following steps:

1. Configure VPN without identity NAT.
2. Enter **show nat detail** and **show conn all**.
3. Add the identity NAT configuration.
 - Repeat **show nat detail** and **show conn all**.

DNS and NAT

You might need to configure the ASA to modify DNS replies by replacing the address in the reply with an address that matches the NAT configuration. You can configure DNS modification when you configure each translation rule.

This feature rewrites the address in DNS queries and replies that match a NAT rule (for example, the A record for IPv4, the AAAA record for IPv6, or the PTR record for reverse DNS queries). For DNS replies traversing from a mapped interface to any other interface, the record is rewritten from the mapped value to the real value. Inversely, for DNS replies traversing from any interface to a mapped interface, the record is rewritten from the real value to the mapped value.



Note

DNS rewrite is not applicable for PAT because multiple PAT rules are applicable for each A-record, and the PAT rule to use is ambiguous.



Note

If you configure a twice NAT rule, you cannot configure DNS modification if you specify the source address as well as the destination address. These kinds of rules can potentially have a different translation for a single address when going to A vs. B. Therefore, the ASA cannot accurately match the IP address inside the DNS reply to the correct twice NAT rule; the DNS reply does not contain information about which source/destination address combination was in the packet that prompted the DNS request.



Note

This feature requires DNS application inspection to be enabled, which it is by default. See the [“DNS Inspection” section on page 9-1](#) for more information.

Figure 3-22 shows a DNS server that is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure the ASA to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network. In this case, you want to enable DNS reply modification on this static rule so that inside users who have access to ftp.cisco.com using the real address receive the real address from the DNS server, and not the mapped address. When an inside host sends a DNS request for the address of ftp.cisco.com, the DNS server replies with the mapped address (209.165.201.10). The ASA refers to the static rule for the inside server and translates the address inside the DNS reply to 10.1.3.14. If you do not enable DNS reply modification, then the inside host attempts to send traffic to 209.165.201.10 instead of accessing ftp.cisco.com directly.

Figure 3-22 DNS Reply Modification, DNS Server on Outside

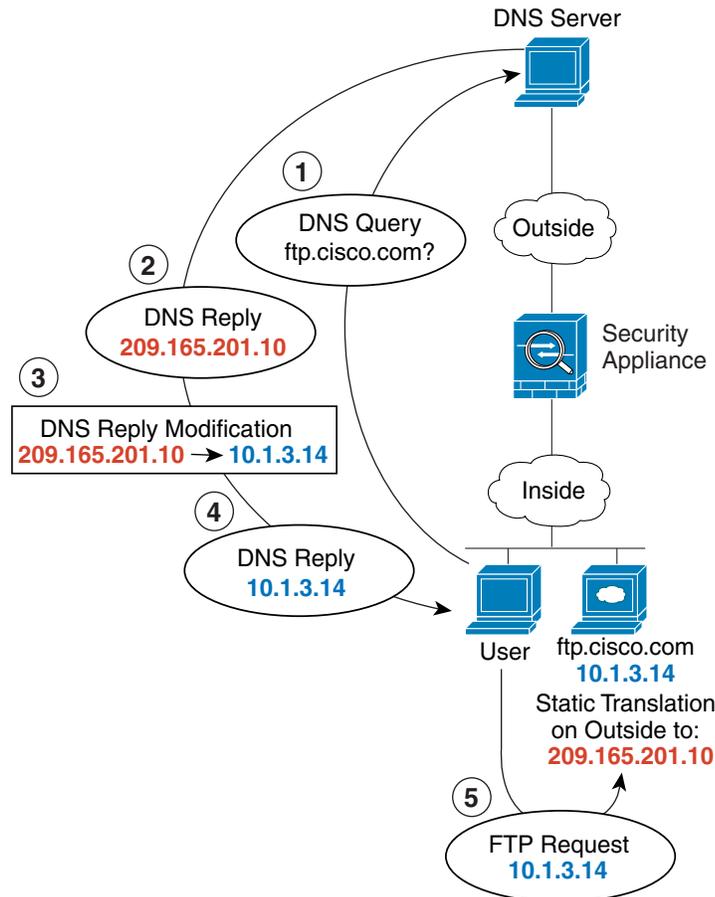


Figure 3-23 shows a user on the inside network requesting the IP address for ftp.cisco.com, which is on the DMZ network, from an outside DNS server. The DNS server replies with the mapped address (209.165.201.10) according to the static rule between outside and DMZ even though the user is not on the DMZ network. The ASA translates the address inside the DNS reply to 10.1.3.14. If the user needs to access ftp.cisco.com using the real address, then no further configuration is required. If there is also

a static rule between the inside and DMZ, then you also need to enable DNS reply modification on this rule. The DNS reply will then be modified two times. In this case, the ASA again translates the address inside the DNS reply to 192.168.1.10 according to the static rule between inside and DMZ.

Figure 3-23 DNS Reply Modification, DNS Server, Host, and Server on Separate Networks

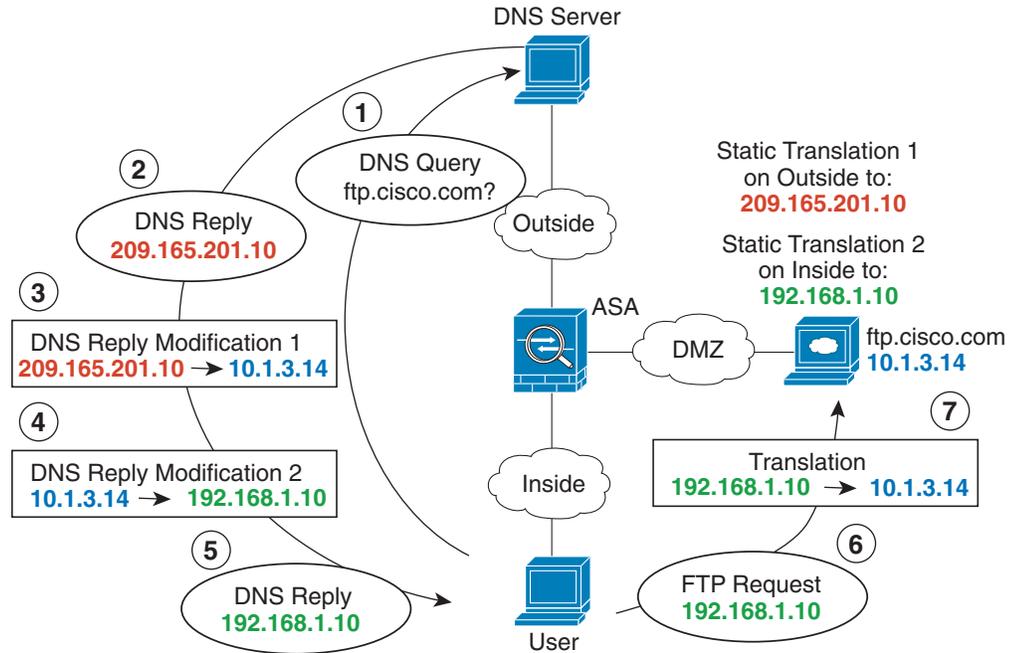


Figure 3-24 shows an FTP server and DNS server on the outside. The ASA has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.20.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.

Figure 3-24 DNS Reply Modification, DNS Server on Host Network

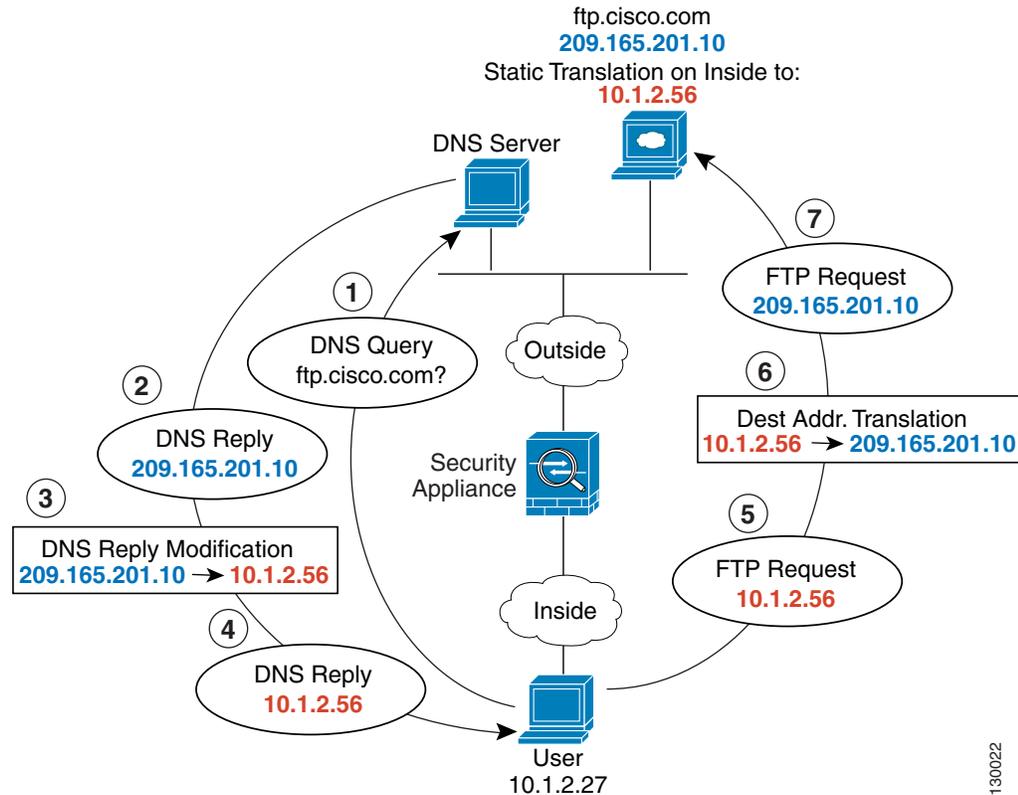
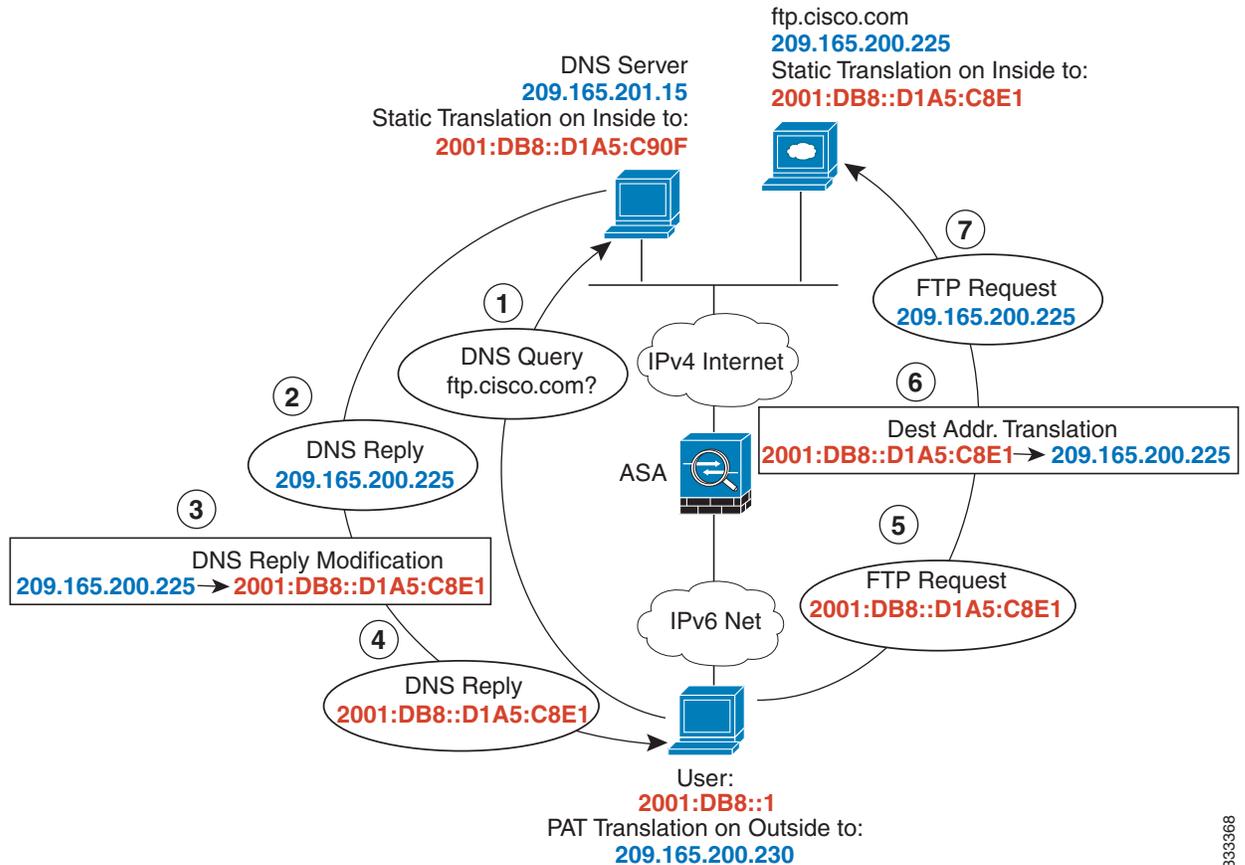


Figure 3-24 shows an FTP server and DNS server on the outside IPv4 network. The ASA has a static translation for the outside server. In this case, when an inside IPv6 user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.200.225.

Because you want inside users to use the mapped address for ftp.cisco.com (2001:DB8::D1A5:C8E1) you need to configure DNS reply modification for the static translation. This example also includes a static NAT translation for the DNS server, and a PAT rule for the inside IPv6 hosts.

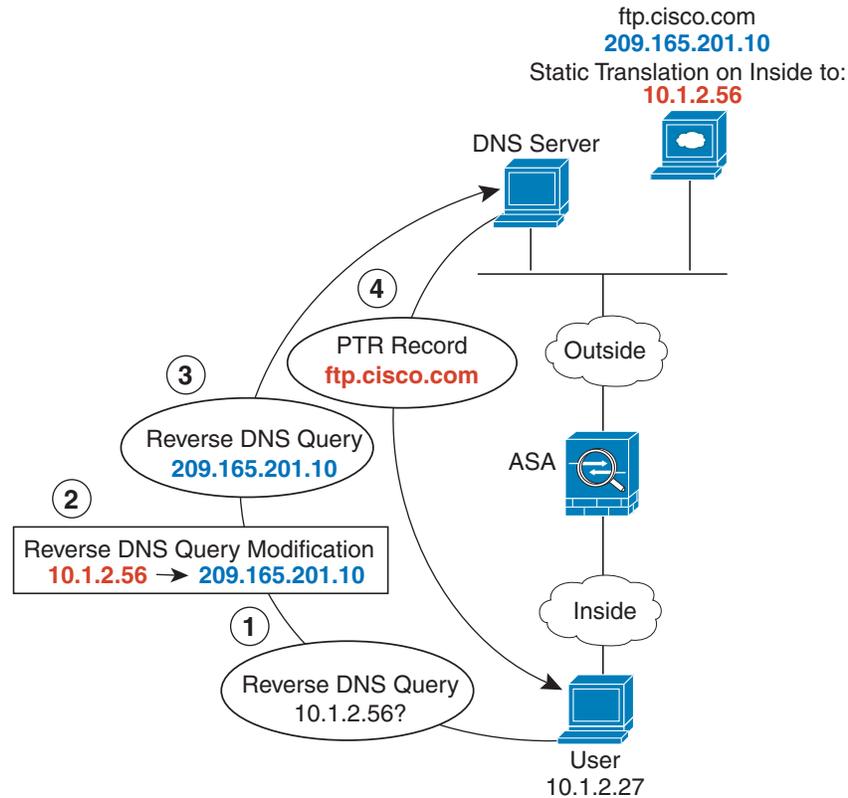
Figure 3-25 DNS64 Reply Modification Using Outside NAT



333368

Figure 3-26 shows an FTP server and DNS server on the outside. The ASA has a static translation for the outside server. In this case, when an inside user performs a reverse DNS lookup for 10.1.2.56, the ASA modifies the reverse DNS query with the real address, and the DNS server responds with the server name, ftp.cisco.com.

Figure 3-26 PTR Modification, DNS Server on Host Network



304002

Where to Go Next

To configure network object NAT, see [Chapter 4, “Configuring Network Object NAT.”](#)

To configure twice NAT, see [Chapter 5, “Configuring Twice NAT.”](#)

