Configuring NAT (ASA 8.2 and Earlier)

This chapter describes Network Address Translation, and includes the following sections:

NAT Overview

This section describes how NAT works on the ASA, and includes the following topics:

Introduction to NAT

Address translation substitutes the real address in a packet with a mapped address that is routable on the destination network. NAT is composed of two steps: the process by which a real address is translated into a mapped address, and the process to undo translation for returning traffic.

The ASA translates an address when a NAT rule matches the traffic. If no NAT rule matches, processing for the packet continues. The exception is when you enable NAT control. NAT control requires that packets traversing from a higher security interface (inside) to a lower security interface (outside) match a NAT rule, or processing for the packet stops. See the "Security Levels" section on page 13-1 in the

general operations configuration guide for more information about security levels. See the "NAT Control" section on page 6-4 for more information about NAT control.

Note In this document, all types of translation are referred to as NAT. When describing NAT, the terms *inside* and *outside* represent the security relationship between any two interfaces. The higher security level is inside and the lower security level is outside. For example, interface 1 is at 60 and interface 2 is at 50; therefore, interface 1 is "inside" and interface 2 is "outside."

Some of the benefits of NAT are as follows:

- You can use private addresses on your inside networks. Private addresses are not routable on the Internet.
- NAT hides the real addresses from other networks, so attackers cannot learn the real address of a host.
- You can resolve IP routing problems such as overlapping addresses.

See Table 10-1 on page 10-4 for information about protocols that do not support NAT.

NAT in Routed Mode

Figure 6-1 shows a typical NAT example in routed mode, with a private network on the inside. When the inside host at 10.1.1.27 sends a packet to a web server, the real source address, 10.1.1.27, of the packet is changed to a mapped address, 209.165.201.10. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the security appliance receives the packet. The security appliance then changes the translation of the mapped address, 209.165.201.10 back to the real address, 10.1.1.1.27 before sending it to the host.

Figure 6-1 NAT Example: Routed Mode

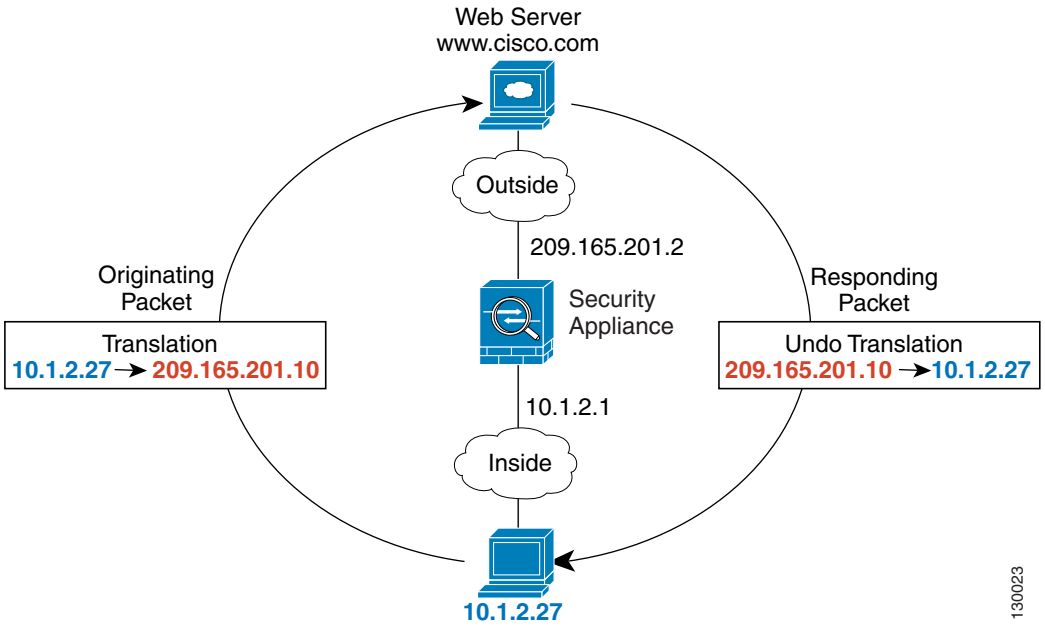

NAT in Transparent Mode

Using NAT in transparent mode eliminates the need for the upstream or downstream routers to perform NAT for their networks. For example, a transparent firewall ASA is useful between two VRFs so you can establish BGP neighbor relations between the VRFs and the global table. However, NAT per VRF might not be supported. In this case, using NAT in transparent mode is essential.

NAT in transparent mode has the following requirements and limitations:

- When the mapped addresses are not on the same network as the transparent firewall, then on the upstream router, you need to add a static route for the mapped addresses that points to the downstream router (through the ASA).
- When you have VoIP or DNS traffic with NAT and inspection enabled, to successfully translate the IP address inside VoIP and DNS packets, the ASA needs to perform a route lookup. Unless the host is on a directly-connected network, then you need to add a static route on the ASA for the real host address that is embedded in the packet.
- The **alias** command is not supported.
- 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 firewall sends an ARP request to a host on the other side of the firewall, 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.

Figure 6-2 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. When the inside host at 10.1.1.27 sends a packet to a web server, the real source address of the packet, 10.1.1.27, is changed to a mapped address, 209.165.201.10. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the ASA receives the packet because the upstream router includes this mapped network in a static route directed through the ASA. The ASA then undoes the translation of the mapped address, 209.165.201.10 back to the real address, 10.1.1.1.27. Because the real address is directly-connected, the ASA sends it directly to the host.

Figure 6-2 NAT Example: Transparent Mode

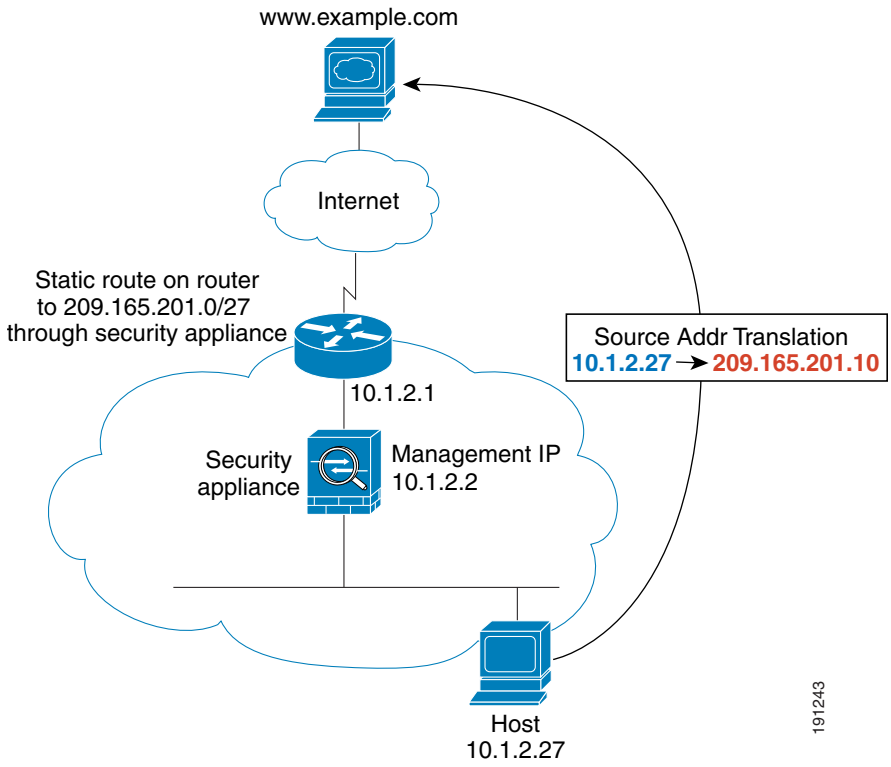

NAT Control

NAT control requires that packets traversing from an inside interface to an outside interface match a NAT rule; for any host on the inside network to access a host on the outside network, you must configure NAT to translate the inside host address, as shown in Figure 6-3.

Figure 6-3 NAT Control and Outbound Traffic

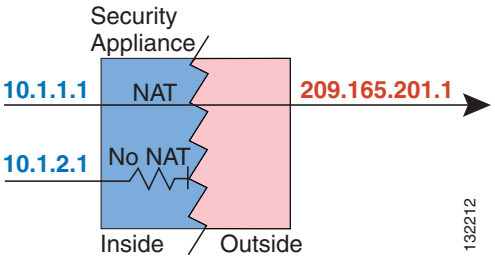

Interfaces at the same security level are not required to use NAT to communicate. However, if you configure dynamic NAT or PAT on a same security interface, then all traffic from the interface to a same security interface or an outside interface must match a NAT rule, as shown in Figure 6-4.

Figure 6-4 NAT Control and Same Security Traffic

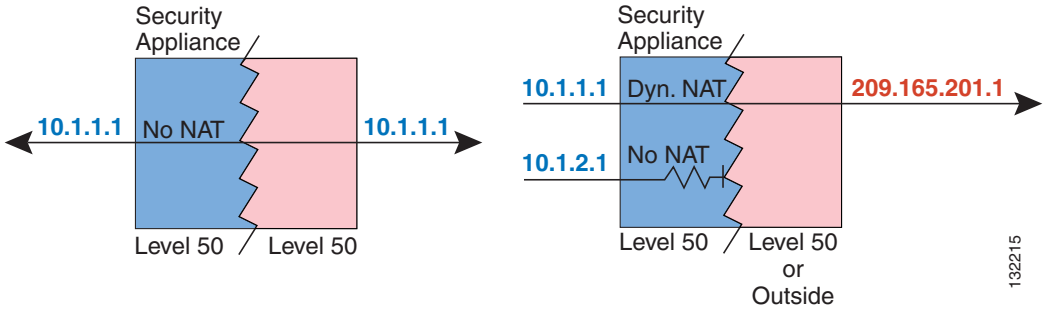

Similarly, if you enable outside dynamic NAT or PAT, then all outside traffic must match a NAT rule when it accesses an inside interface (see Figure 6-5).

Figure 6-5 NAT Control and Inbound Traffic

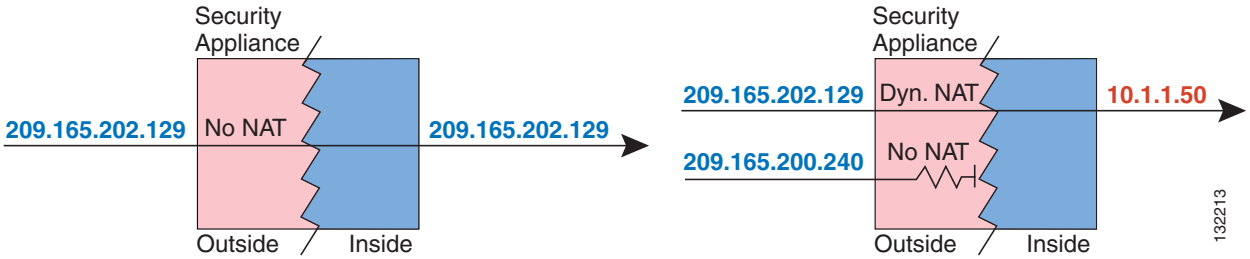

Static NAT does not cause these restrictions.

By default, NAT control is disabled; therefore, you do not need to perform NAT on any networks unless you want to do so. If you upgraded from an earlier version of software, however, NAT control might be enabled on your system. Even with NAT control disabled, you need to perform NAT on any addresses for which you configure dynamic NAT. See the "Dynamic NAT Implementation" section on page 6-17 for more information about how dynamic NAT is applied.

If you want the added security of NAT control but do not want to translate inside addresses in some cases, you can apply a NAT exemption or identity NAT rule on those addresses. (See the "Using NAT Exemption" section on page 6-33 for more information).

To configure NAT control, see the "Configuring NAT Control" section on page 6-16.

Note In multiple context mode, the packet classifier might rely on the NAT configuration to assign packets to contexts if you do not enable unique MAC addresses for shared interfaces. See the "How the ASA Classifies Packets" section on page 8-3 in the general operations configuration guide for more information about the relationship between the classifier and NAT.

NAT Types

This section describes the available NAT types, and includes the following topics:

You can implement address translation as dynamic NAT, Port Address Translation, static NAT, static PAT, or as a mix of these types. You can also configure rules to bypass NAT; for example, to enable NAT control when you do not want to perform NAT.

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 may include 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 added 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, although the connection is allowed by an ACL, and the ASA rejects any attempt to connect to a real host address directly. See the "Static NAT" or "Static PAT" section for information on how to obtain reliable access to hosts.

Note In some cases, a translation is added for a connection, although the session is denied by the ASA. This condition occurs with an outbound ACL, a management-only interface, or a backup interface in which the translation times out normally.

Figure 6-6 shows a remote host attempting to connect to the real address. The connection is denied, because the ASA only allows returning connections to the mapped address.

Figure 6-6 Remote Host Attempts to Connect to the Real Address

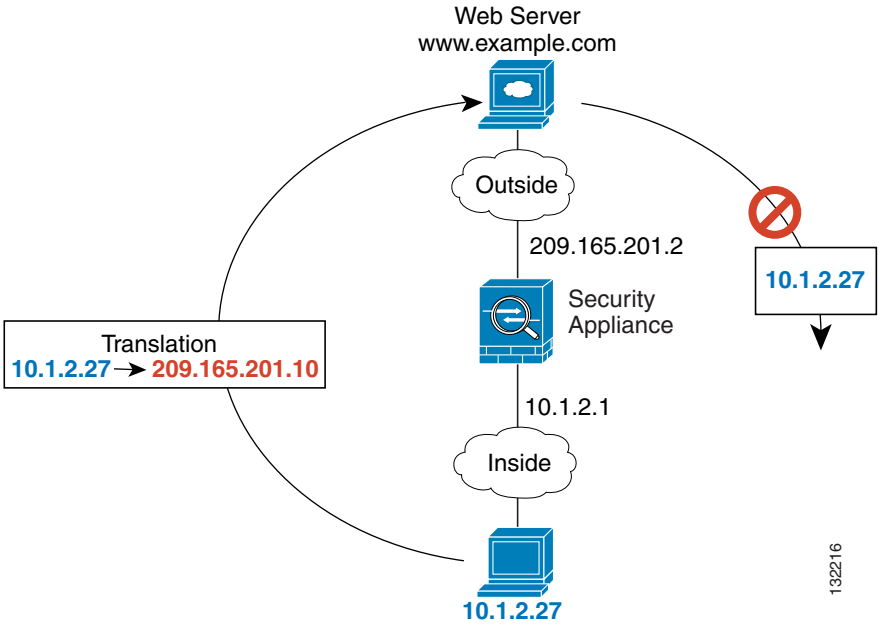

Figure 6-7 shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table; therefore, the ASA drops the packet.

Figure 6-7 Remote Host Attempts to Initiate a Connection to a Mapped Address

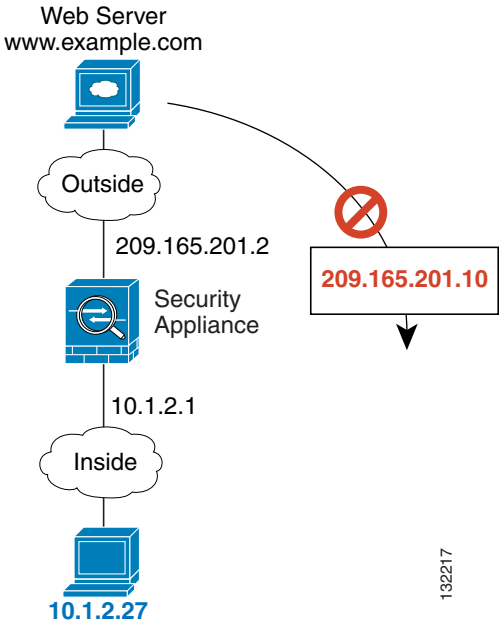

Note For the duration of the translation, a remote host can initiate a connection to the translated host if an ACL allows it. Because the address is unpredictable, a connection to the host is unlikely. Nevertheless, in this case, you can rely on the security of the ACL.

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 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; if the destination network requires registered addresses, such as the Internet, you might encounter a shortage of usable addresses.

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 "When to Use Application Protocol Inspection" section on page 10-2 for more information about NAT and PAT support.

PAT

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.

Each connection requires a separate translation, 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.

After the connection expires, the port translation also expires after 30 seconds of inactivity. The timeout is not configurable. Users on the destination network cannot reliably initiate a connection to a host that uses PAT (even if the connection is allowed by an ACL). Not only can you not predict the real or mapped port number of the host, but the ASA does not create a translation at all unless the translated host is the initiator. See the following "Static NAT" or "Static PAT" sections for reliable access to hosts.

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. PAT does not work with some multimedia applications that have a data stream that is different from the control path. See the "When to Use Application Protocol Inspection" section on page 10-2 for more information about NAT and PAT support.

Note For the duration of the translation, a remote host can initiate a connection to the translated host if an ACL 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 ACL. However, policy PAT does not support time-based ACLs.

Static NAT

Static NAT creates a fixed translation of real address(es) to mapped address(es).With dynamic NAT and PAT, each host uses a different address or port for each subsequent translation. Because the mapped address is the same for each consecutive connection with static NAT, and a persistent translation rule exists, static NAT allows hosts on the destination network to initiate traffic to a translated host (if an ACL exists that allows it).

The main difference between dynamic NAT and a range of addresses for static NAT is that static NAT allows a remote host to initiate a connection to a translated host (if an ACL exists that allows it), while dynamic NAT does not. You also need an equal number of mapped addresses as real addresses with static NAT.

Static PAT

Static PAT is the same as static NAT, except that it lets you specify the protocol (TCP or UDP) and port for the real and mapped addresses.

This feature lets you identify the same mapped address across many different static statements, provided the port is different for each statement. You cannot use the same mapped address for multiple static NAT statements.

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

For example, if you want to provide a single address for remote users to access FTP, HTTP, and SMTP, but these are all actually different servers on the real network, you can specify static PAT statements for each server that uses the same mapped IP address, but different ports (see Figure 6-8).

Figure 6-8 Static PAT

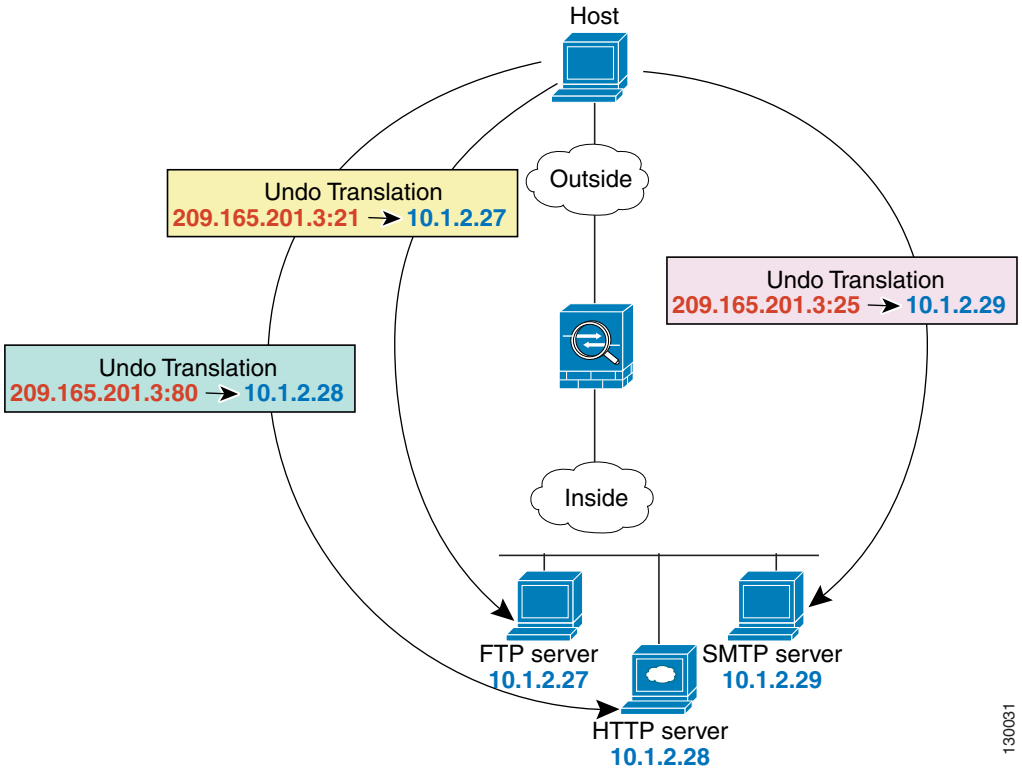

You can also use static PAT to translate a well-known port to a non-standard port or vice versa. For example, if inside web servers use port 8080, you can allow outside users to connect to port 80, and then undo translation to the original port 8080. Similarly, to provide extra security, you can tell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

Bypassing NAT When NAT Control is Enabled

If you enable NAT control, then inside hosts must match a NAT rule when accessing outside hosts. If you do not want to perform NAT for some hosts, then you can bypass NAT for those hosts or you can disable NAT control. You might want to bypass NAT, for example, if you are using an application that does not support NAT. See the "When to Use Application Protocol Inspection" section on page 10-2 for information about inspection engines that do not support NAT.

You can configure traffic to bypass NAT using one of three methods. All methods achieve compatibility with inspection engines. However, each method offers slightly different capabilities, as follows:

- Identity NAT—When you configure identity NAT (which is similar to dynamic NAT), you do not limit translation for a host on specific interfaces; you must use identity NAT for connections through all interfaces. Therefore, you cannot choose to perform normal translation on real addresses when you access interface A, but use identity NAT when accessing interface B. Regular dynamic NAT, on

the other hand, lets you specify a particular interface on which to translate the addresses. Make sure that the real addresses for which you use identity NAT are routable on all networks that are available according to your ACLs.

For identity NAT, even though the mapped address is the same as the real address, you cannot initiate a connection from the outside to the inside (even if the interface ACL allows it). Use static identity NAT or NAT exemption for this functionality.

- Static identity NAT—Static identity NAT lets you specify the interface on which you want to allow the real addresses to appear, so you can use identity NAT when you access interface A, and use regular translation when you access interface B. Static identity NAT also lets you use policy NAT, which identifies the real and destination addresses when determining the real addresses to translate (see the "Policy NAT" section on page 6-11 for more information about policy NAT). For example, you can use static identity NAT for an inside address when it accesses the outside interface and the destination is server A, but use a normal translation when accessing the outside server B.
- NAT exemption—NAT exemption allows both translated and remote hosts to initiate connections. Like identity NAT, you do not limit translation for a host on specific interfaces; you must use NAT exemption for connections through all interfaces. However, NAT exemption does let you specify the real and destination addresses when determining the real addresses to translate (similar to policy NAT), so you have greater control using NAT exemption. However unlike policy NAT, NAT exemption does not consider the ports in the ACL. NAT exemption also does not let you configure connection limits such as maximum TCP connections.

Policy NAT

Policy NAT lets you identify real addresses for address translation by specifying the source and destination addresses. You can also optionally specify the source and destination ports. Regular NAT can only consider the source addresses, and not the destination. For example, with policy NAT, you can translate the real address to mapped address A when it accesses server A, but translate the real address to mapped address B when it accesses server B.

For applications that require application inspection for secondary channels (for example, FTP and VoIP), the policy specified in the policy NAT rule should include the secondary ports. When the ports cannot be predicted, the policy should specify only the IP addresses for the secondary channel. With this configuration, the security appliance translates the secondary ports.

Figure 6-9 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. Consequently, the host appears to be on the same network as the servers, which can help with routing.

Figure 6-9 Policy NAT with Different Destination Addresses

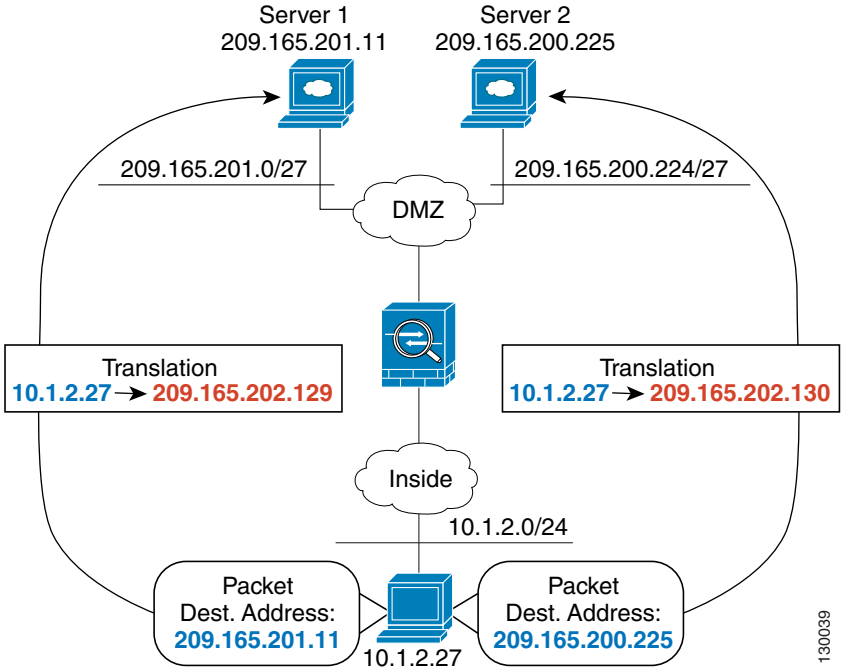

Figure 6-10 shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for web services, the real address is translated to 209.165.202.129. When the host accesses the same server for Telnet services, the real address is translated to 209.165.202.130.

Figure 6-10 Policy NAT with Different Destination Ports

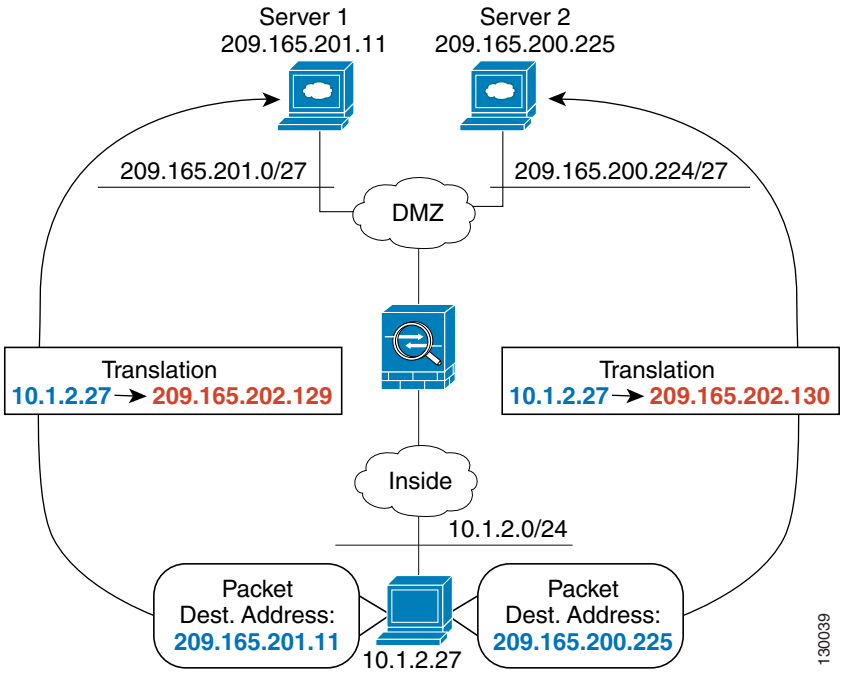

For policy static NAT, both translated and remote hosts can originate traffic. For traffic originated on the translated network, the NAT rule specifies the real addresses and the *destination* addresses, but for traffic originated on the remote network, the rule identifies the real addresses and the *source* addresses of remote hosts who are allowed to connect to the host using this translation.

Figure 6-11 shows a remote host connecting to a translated host. The translated host has a policy 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 6-11 Policy Static NAT with Destination Address Translation

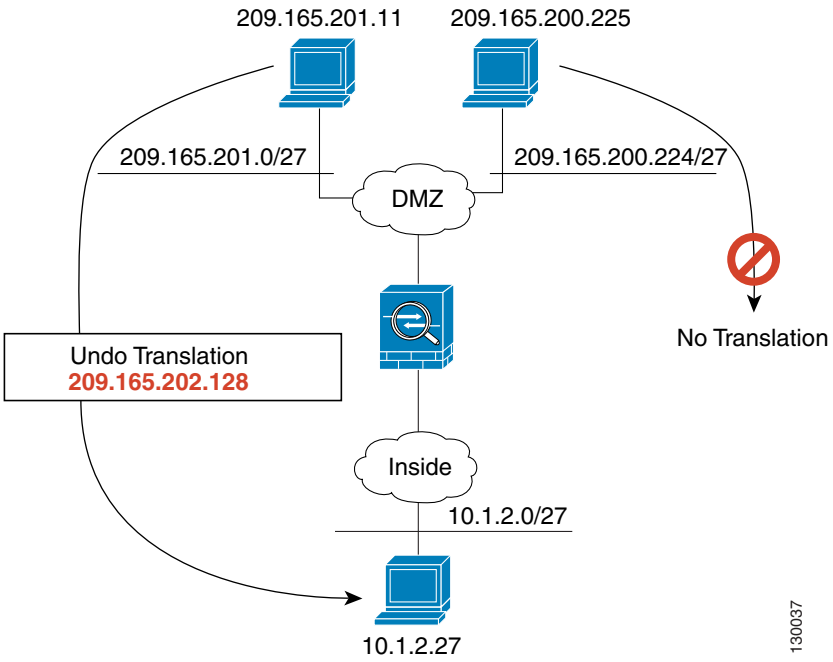

Note Policy NAT does not support SQL*Net, but it is supported by regular NAT. See the "When to Use Application Protocol Inspection" section on page 10-2 for information about NAT support for other protocols.

NAT and Same Security Level Interfaces

NAT is not required between same security level interfaces even if you enable NAT control. You can optionally configure NAT if desired. However, if you configure dynamic NAT when NAT control is enabled, then NAT is required. See the "NAT Control" section on page 6-4 for more information. Also, when you specify a group of IP address(es) for dynamic NAT or PAT on a same security interface, then you must perform NAT on that group of addresses when they access any lower or same security level interface (even when NAT control is not enabled). Traffic identified for static NAT is not affected.

Note The ASA does not support VoIP inspection engines when you configure NAT on same security interfaces. These inspection engines include Skinny, SIP, and H.323. See the "When to Use Application Protocol Inspection" section on page 10-2 for supported inspection engines.

Order of NAT Rules Used to Match Real Addresses

The ASA matches real addresses to NAT rules in the following order:

1. NAT exemption—In order, until the first match.
2. Static NAT and Static PAT (regular and policy)—In order, until the first match. Static identity NAT is included in this category.
3. Policy dynamic NAT—In order, until the first match. Overlapping addresses are allowed.
4. Regular dynamic NAT—Best match. Regular identity NAT is included in this category. The order of the NAT rules does not matter; the NAT rule that best matches the real address is used. For example, you can create a general rule to translate all addresses (0.0.0.0) on an interface. If you want to translate a subset of your network (10.1.1.1) to a different address, then you can create a rule to translate only 10.1.1.1. When 10.1.1.1 makes a connection, the specific rule for 10.1.1.1 is used because it matches the real address best. We do not recommend using overlapping rules; they use more memory and can slow the performance of the ASA.

Mapped Address Guidelines

When you translate the real address to a mapped address, you can use the following mapped addresses:

- Addresses on the same network as the mapped interface.

  If you use addresses on the same network as the mapped interface (through which traffic exits the ASA), the ASA uses proxy ARP to answer any requests for mapped addresses, and thus intercepts traffic destined for a real address. This solution simplifies routing, because the ASA does not have to be the gateway for any additional networks. However, this approach does put a limit on the number of available addresses used for translations.

  For PAT, you can even use the IP address of the mapped interface.

- 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 ASA uses proxy ARP to answer any requests for mapped addresses, and thus intercepts traffic destined for a real address. If you use OSPF, and you advertise routes on the mapped interface, then the ASA advertises the mapped addresses. If the mapped interface is passive (not advertising routes) or you are using static routing, then you need to add a static route on the upstream router that sends traffic destined for the mapped addresses to the ASA.

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.

For example, a DNS server 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 (see Figure 6-12). In this case, you want to enable DNS reply modification on this static statement 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 statement 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 6-12 DNS Reply Modification

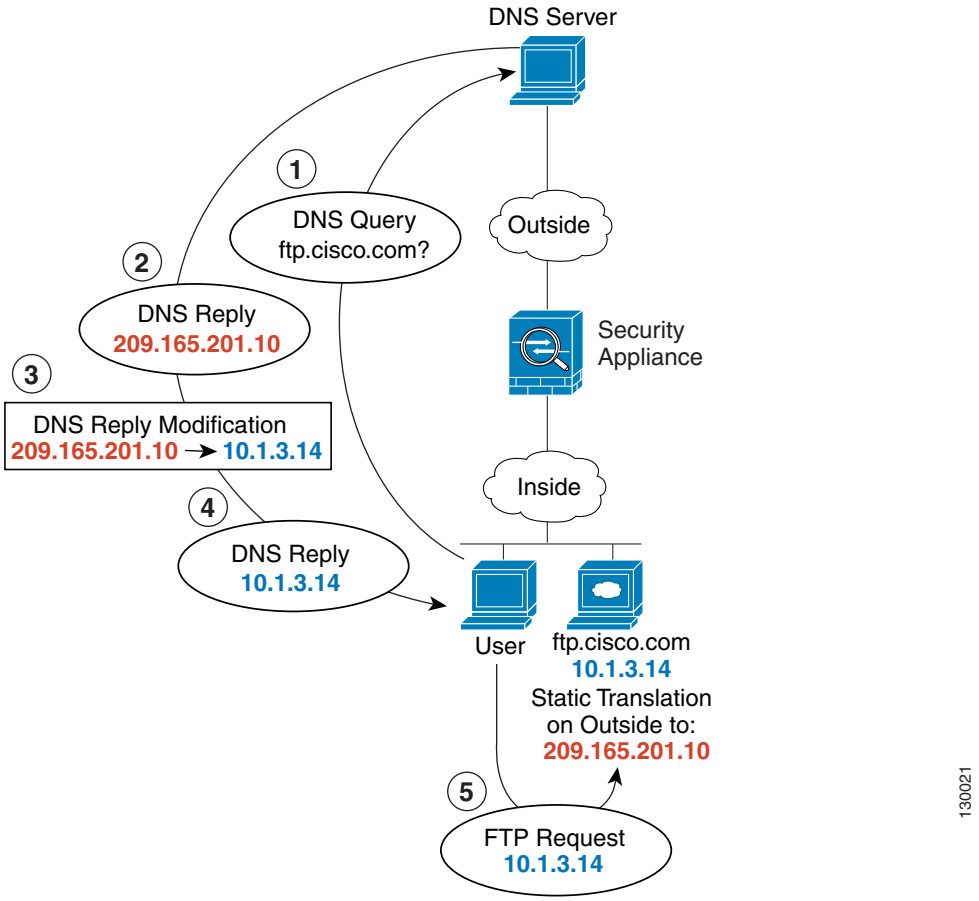

Note If a user on a different network (for example, DMZ) also requests the IP address for ftp.cisco.com from the outside DNS server, then the IP address in the DNS reply is also modified for this user, even though the user is not on the Inside interface referenced by the static rule.

Figure 6-13 shows a web 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 6-13 DNS Reply Modification Using Outside NAT

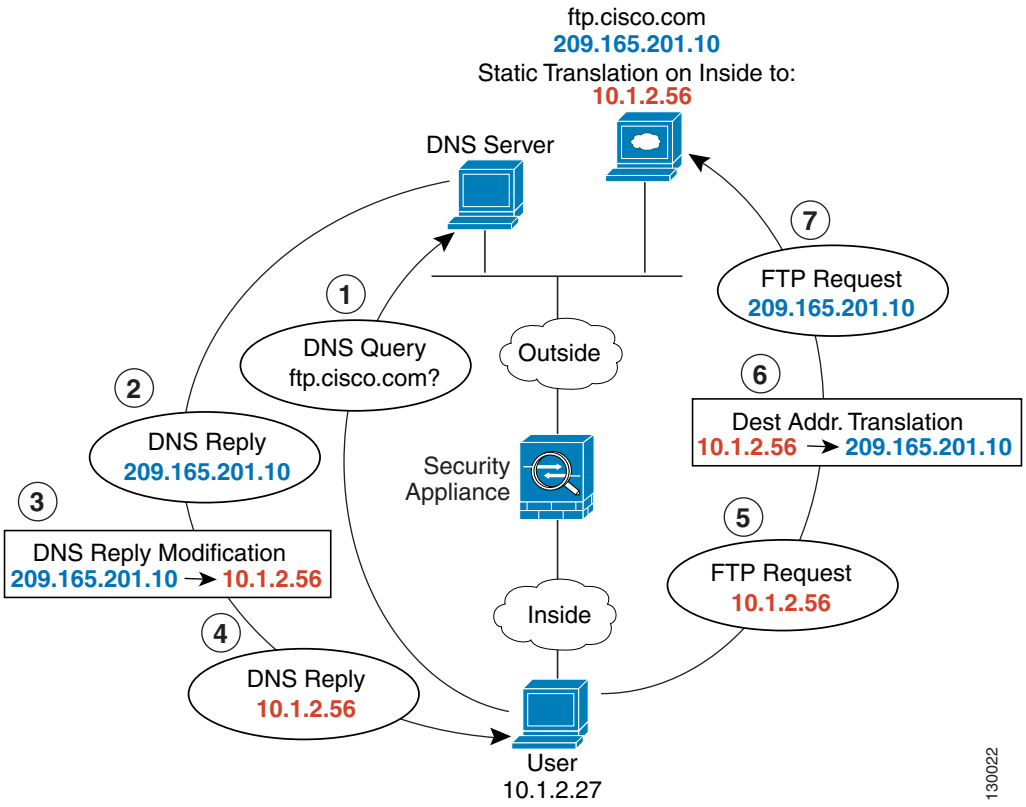

Configuring NAT Control

NAT control requires that packets traversing from an inside interface to an outside interface match a NAT rule. See the "NAT Control" section on page 6-4 for more information.

To enable NAT control, in the Configuration > Firewall > NAT Rules pane, check the **Enable traffic through the firewall without address translation** check box.

Using Dynamic NAT

This section describes how to configure dynamic NAT, including dynamic NAT and PAT, dynamic policy NAT and PAT, and identity NAT.

Policy NAT lets you identify real addresses for address translation by specifying the source and destination addresses. You can also optionally specify the source and destination ports. Regular NAT can only consider the source addresses, and not the destination. See the "Policy NAT" section on page 6-11 for more information.

This section includes the following topics:

Dynamic NAT Implementation

This section describes how dynamic NAT is implemented, and includes the following topics:

Real Addresses and Global Pools Paired Using a Pool ID

In a dynamic NAT rule, you specify real addresses and then pair them with a global pool of addresses to which the real addresses are mapped when they exit another interface (in the case of PAT, this is one address, and in the case of identity NAT, this is the same as the real address). Each global pool is assigned a pool ID.

NAT Rules on Different Interfaces with the Same Global Pools

You can create a NAT rule for each interface using the same global address pool. For example, you can configure NAT rules for Inside and DMZ interfaces, both using global pool 1 on the outside interface. Traffic from the Inside interface and the DMZ interface share a mapped pool or a PAT address when exiting the Outside interface (see Figure 6-14).

Figure 6-14 NAT Rules on Multiple Interfaces Using the Same Global Pool

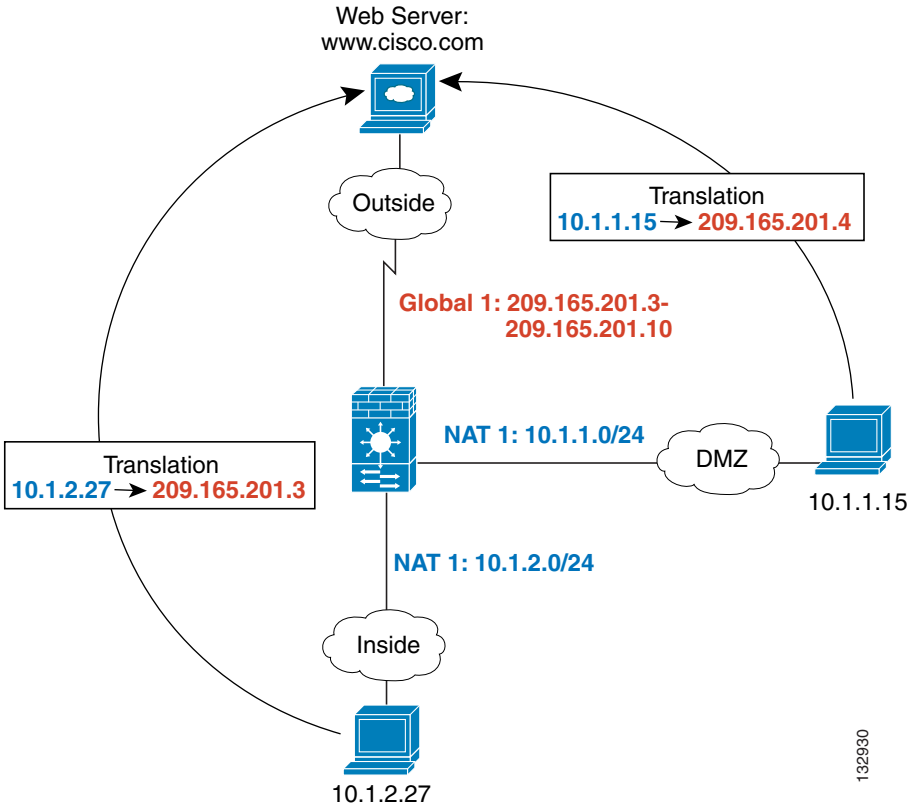

Global Pools on Different Interfaces with the Same Pool ID

You can create a global pool for each interface using the same pool ID. If you create a global pool for the Outside and DMZ interfaces on ID 1, then a single NAT rule associated with ID 1 identifies traffic to be translated when going to both the Outside and the DMZ interfaces. Similarly, if you create a NAT rule for the DMZ interface on ID 1, then all global pools on ID 1 are also used for DMZ traffic. (See

Figure 6-15).

Figure 6-15 NAT Rules and Global Pools using the Same ID on Multiple Interfaces

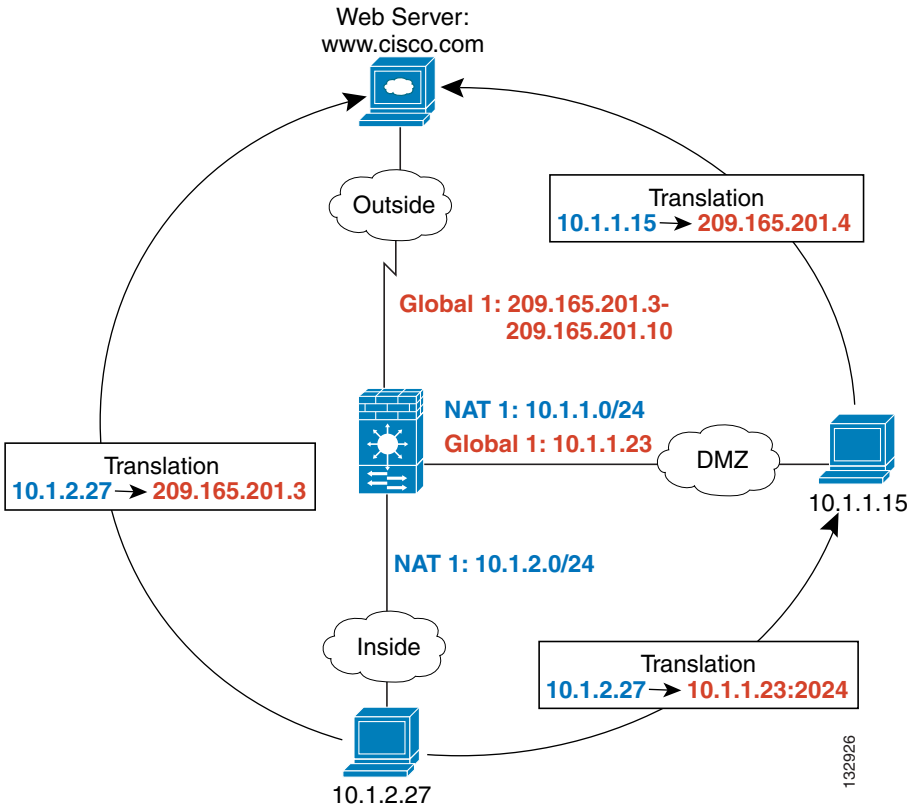

Multiple NAT Rules with Different Global Pools on the Same Interface

You can identify different sets of real addresses to have different mapped addresses. For example, on the Inside interface, you can have two NAT rules on two different pool IDs. On the Outside interface, you configure two global pools for these two IDs. Then, when traffic from Inside network A exits the Outside interface, the IP addresses are translated to pool 1 addresses; while traffic from Inside network B are translated to pool 2 addresses (see Figure 6-16). If you use policy NAT, you can specify the same real addresses for multiple NAT rules, as long as the destination addresses and ports are unique in each ACL.

Figure 6-16 Different NAT IDs

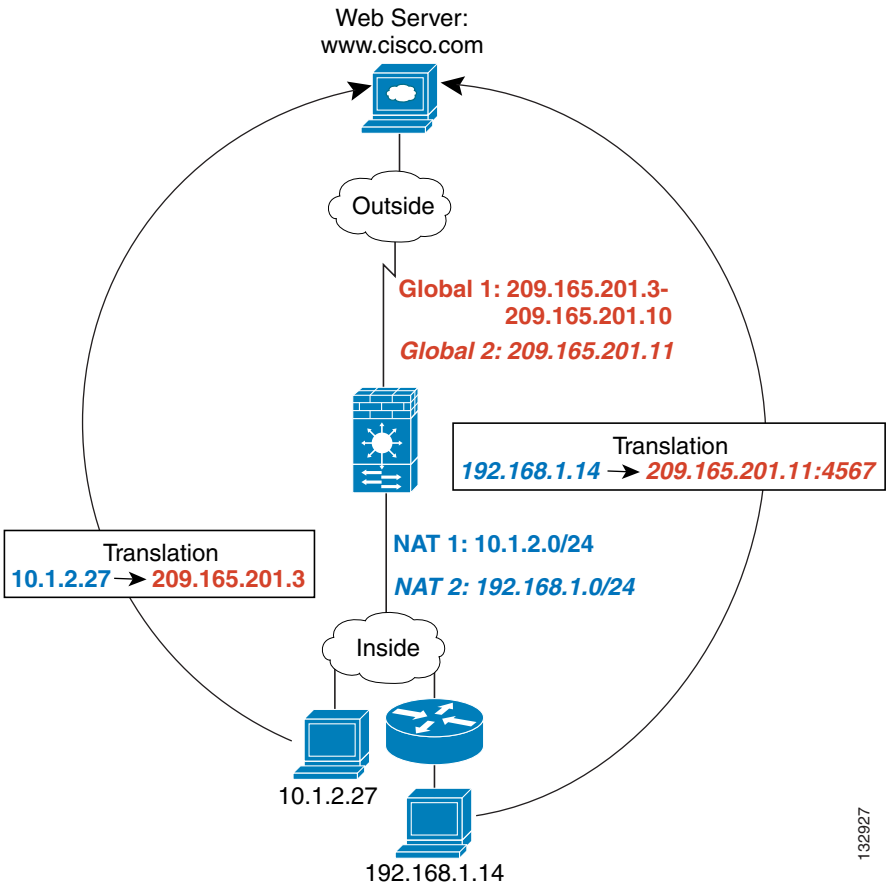

Multiple Addresses in the Same Global Pool

You can have multiple addresses in the same global pool; the ASA uses the dynamic NAT ranges of addresses first, in the order they are in the configuration, and then uses the PAT single addresses in order. You might want to add both a range of addresses and a PAT address if you need to use dynamic NAT for a particular application, but want to have a backup PAT rule in case all the dynamic NAT addresses are depleted. Similarly, you might want two PAT addresses in the pool if you need more than the approximately 64,000 PAT sessions that a single PAT mapped address supports (see Figure 6-17).

Figure 6-17 NAT and PAT Together

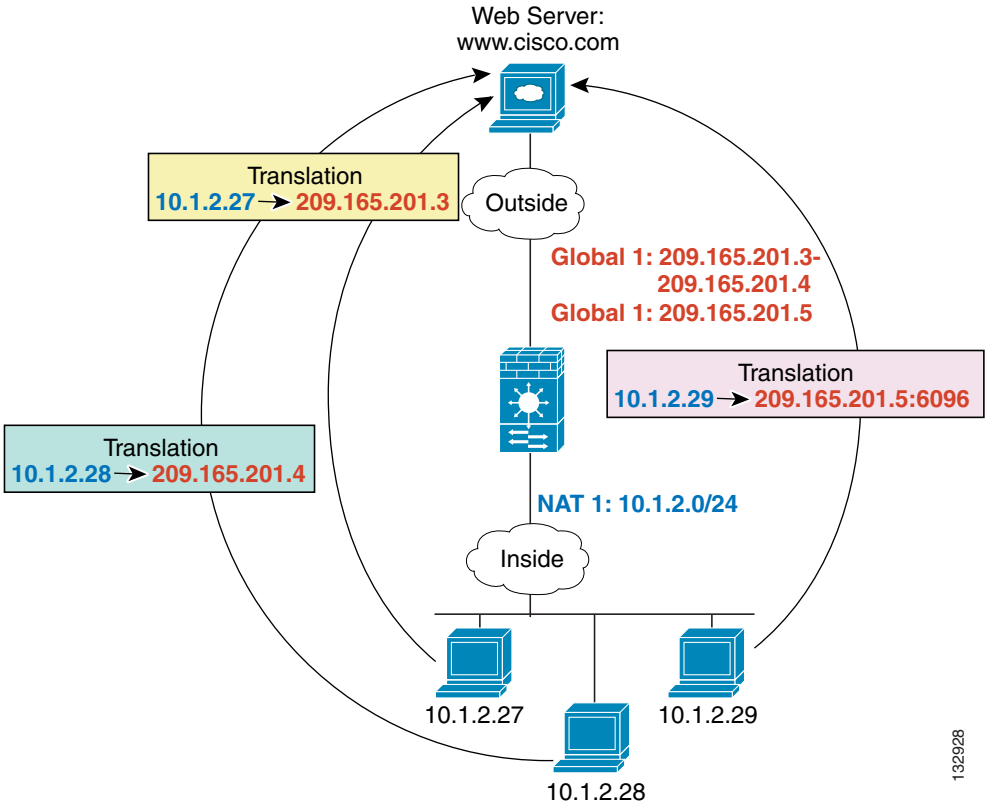

Outside NAT

If a NAT rule translates addresses from an outside interface to an inside interface, then the rule is an outside NAT rule, and you need to specify that it translates inbound traffic. If you also want to translate the same traffic when it accesses a lower security interface (for example, traffic on a DMZ is translated when accessing the Inside and the Outside interfaces), then you can create a second NAT rule using the same NAT ID (see Figure 6-18), but specifying outbound. Note that for outside NAT (DMZ interface to Inside interface), the inside host uses a static rule to allow outside access, so both the source and destination addresses are translated.

Figure 6-18 Outside NAT and Inside NAT Combined

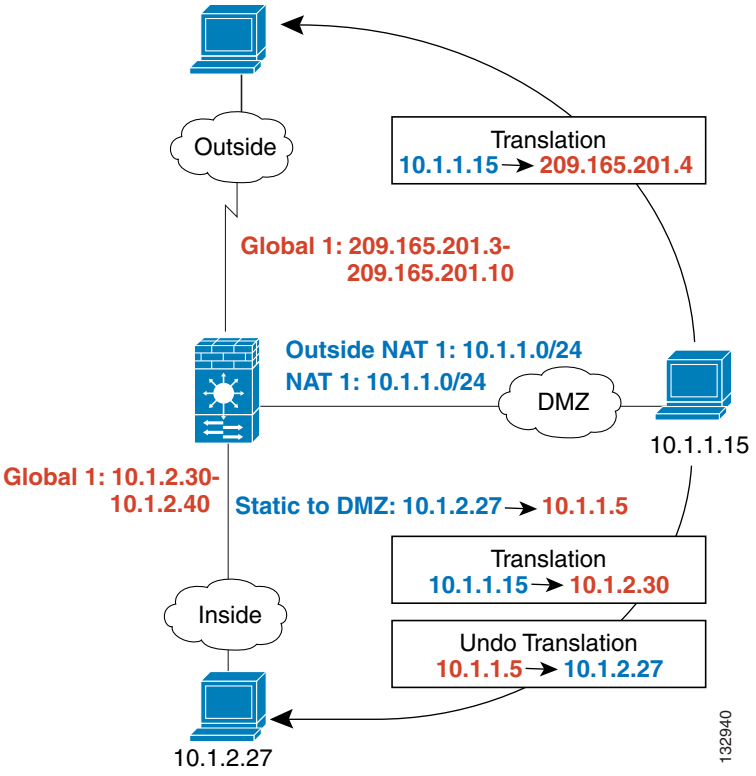

Real Addresses in a NAT Rule Must be Translated on All Lower or Same Security Interfaces

When you create a NAT rule for a group of IP addresses, then you must perform NAT on that group of addresses when they access any lower or same security level interface; you must create a global pool with the same pool ID on each interface, or use a static rule. NAT is not required for that group when it accesses a higher security interface. If you create an outside NAT rule, then the NAT requirements preceding come into effect for that group of addresses when they access all higher security interfaces. Traffic identified by a static rule is not affected.

Managing Global Pools

Dynamic NAT uses global pools for translation. For information about how global pools work, see the "Dynamic NAT Implementation" section on page 6-17.

To manage a global pool, perform the following steps:

Step 1 In the Configuration > Firewall > Objects > Global Pools pane, click **Add** to add a new pool, or select a pool, and click **Edit**.

You can also manage global pools from the Add/Edit Dynamic NAT Rule dialog box by clicking **Manage**.

The Add/Edit Global Address Pool dialog box appears.

Step 2 For a new pool, from the Interface drop-down list, choose the interface where you want to use the mapped IP addresses.

Step 3 For a new pool, in the Pool ID field, enter a number between 1 and 2147483647. Do not enter a pool ID that is already in use, or your configuration will be rejected.

Step 4 In the IP Addresses to Add area, click **Range**, **Port Address Translation (PAT)**, or **PAT Address Translation (PAT) Using IP Address of the interface**.

If you specify a range of addresses, the ASA performs dynamic NAT. If you specify a subnet mask in the Netmask field, the value specifies the subnet mask assigned to the mapped address when it is assigned to a host. If you do not specify a mask, then the default mask for the address class is used.

Step 5 Click **Add** to add the addresses to the Addresses Pool pane.

Step 6 (Optional) You can add multiple addresses to the global pool. If you want to add a PAT address after you configure a dynamic range, for example, then complete the value for PAT and click **Add** again. See the "Multiple Addresses in the Same Global Pool" section on page 6-20 for information about using multiple addresses on the same pool ID for an interface.

Step 7 Click **OK**.

Configuring Dynamic NAT, PAT, or Identity NAT

Figure 6-19 shows typical dynamic NAT, dynamic PAT, and identity NAT scenarios. Only real hosts can initiate connections.

Figure 6-19 Dynamic NAT Scenarios

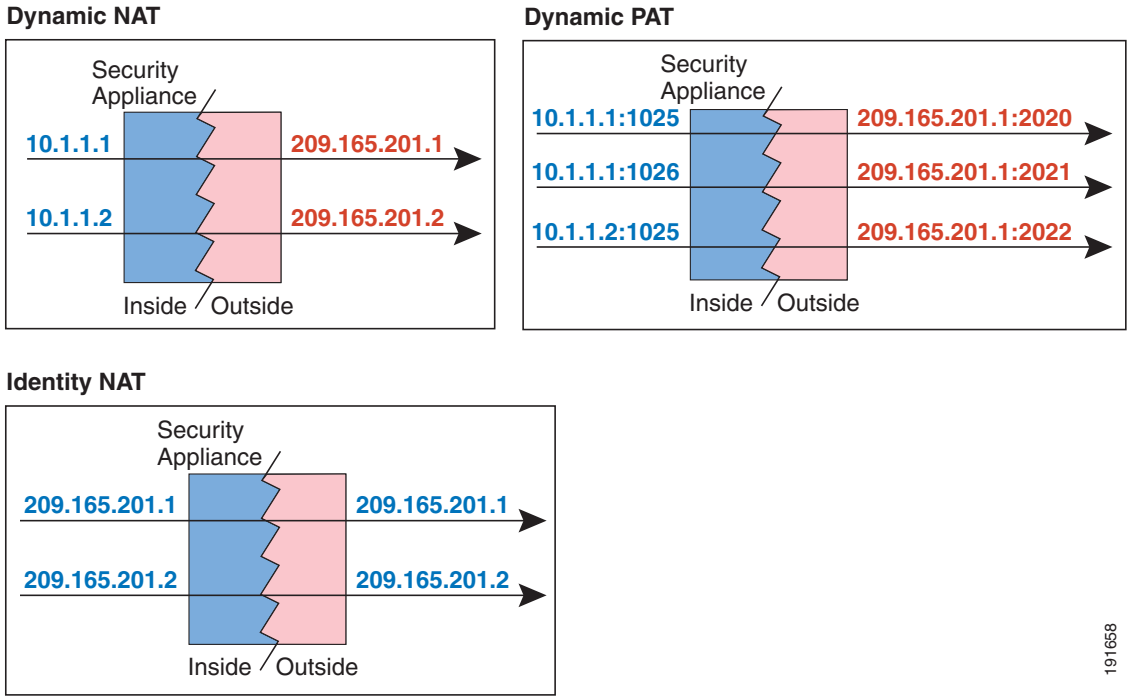

To configure a dynamic NAT, PAT, or identity NAT rule, perform the following steps.

Step 1 In the Configuration > Firewall > NAT Rules pane, choose **Add > Add Dynamic NAT Rule**.

The Add Dynamic NAT Rule dialog box appears.

Step 2 In the Original area, from the Interface drop-down list, choose the interface that is connected to the hosts with real addresses that you want to translate.

Step 3 Enter the real addresses in the Source field, or click the **...** button to select an IP address that you already defined in ASDM.

Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

Step 4 To choose a global pool, use one of the following options:

- Select an already-defined global pool.

  If the pool includes a range of addresses, then the ASA performs dynamic NAT. If the pool includes a single address, then the ASA performs dynamic PAT. If a pool includes both ranges and single addresses, then the ranges are used in order, and then the PAT addresses are used in order. See the "Multiple Addresses in the Same Global Pool" section on page 6-20 for more information.

  Pools are identified by a pool ID. If multiple global pools on different interfaces share the same pool ID, then they are grouped. If you choose a multi-interface pool ID, then traffic is translated as specified when it accesses any of the interfaces in the pool. For more information about pool IDs, see the "Dynamic NAT Implementation" section on page 6-17.

- Create a new global pool or edit an existing pool by clicking **Manage**. See the "Managing Global Pools" section on page 6-22.

- Choose identity NAT by selecting **global pool 0**.

Step 5 (Optional) To enable translation of addresses inside DNS replies, expand the **Connection Settings** area, and check the **Translate the DNS replies that match the translation rule** check box.

If your NAT rule includes the real address of a host that has an entry in a DNS server, and the DNS server is on a different interface from a client, then the client and the DNS server need different addresses for the host; one needs the mapped address and one needs the real address. This option rewrites the address in the DNS reply to the client. The mapped host needs to be on the same interface as either the client or the DNS server. Typically, hosts that need to allow access from other interfaces use a static translation, so this option is more likely to be used with a static rule. See the "DNS and NAT" section on page 6-14 for more information.

Step 6 (Optional) To enable connection settings, expand the **Connection Settings** area, and set one or more of the following options:

Note You can also set these values using a security policy rule (see Chapter 22, "Configuring Connection Settings"). If you set them in both places, then the ASA uses the lower limit. For TCP sequence randomization, if it is disabled using either method, then the ASA disables TCP sequence randomization.

- **Randomize sequence number**—With this check box checked (the default), the ASA randomizes the sequence number of TCP packets. Each TCP connection has two ISNs: one generated by the client and one generated by the server. The ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

TCP initial sequence number randomization can be disabled if required. For example:

- If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
- If you use eBGP multi-hop through the ASA, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
- You use a WAAS device that requires the ASA not to randomize the sequence numbers of connections.

- **Maximum TCP Connections**—Specifies the maximum number of TCP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum UDP Connections**—Specifies the maximum number of UDP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum Embryonic Connections**—Specifies the maximum number of embryonic connections per host up to 65,536. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. This limit enables the TCP Intercept feature. The default is **0**, which means the maximum embryonic connections. TCP Intercept protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. When the embryonic limit has been surpassed, the TCP intercept feature intercepts TCP SYN packets from clients to servers on a higher security level. SYN cookies are used during the validation process and help to minimize the amount of valid traffic being dropped. Thus, connection attempts from unreachable hosts will never reach the server.

Step 7 Click **OK**.

Configuring Dynamic Policy NAT or PAT

Figure 6-20 shows typical dynamic policy NAT and PAT scenarios. Only real hosts can initiate connections.

Figure 6-20 Dynamic Policy NAT Scenarios

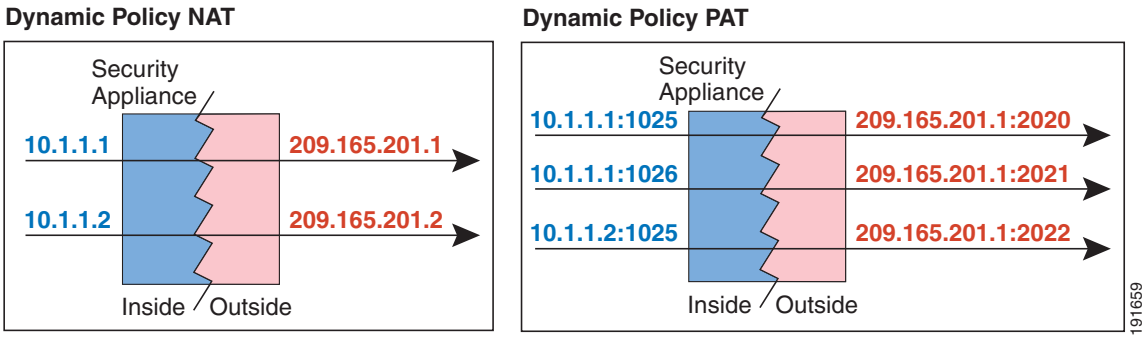

To configure dynamic policy NAT or PAT, perform the following steps:

Step 1 In the Configuration > Firewall > NAT Rules pane, choose **Add > Advanced > Add Dynamic Policy NAT Rule**.

The Add Dynamic Policy NAT Rule dialog box appears.

Step 2 In the Original area, from the Interface drop-down list, choose the interface that is connected to the hosts with real addresses that you want to translate.

Step 3 Enter the real addresses in the Source field, or click the **...** button to choose an IP address that you already defined in ASDM.

Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

Separate multiple real addresses by a comma.

Step 4 Enter the destination addresses in the Destination field, or click the **...** button to choose an IP address that you already defined in ASDM.

Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

Separate multiple destination addresses by a comma.

By default, the field shows **any**, which allows any destination address.

Step 5 To choose a global pool, use one of the following options:

- Choose an already-defined global pool.

  If the pool includes a range of addresses, then the ASA performs dynamic NAT. If the pool includes a single address, then the ASA performs dynamic PAT. If a pool includes both ranges and single addresses, then the ranges are used in order, and then the PAT addresses are used in order. See the "Multiple Addresses in the Same Global Pool" section on page 6-20 for more information.

  Pools are identified by a pool ID. If multiple global pools on different interfaces share the same pool ID, then they are grouped. If you choose a multi-interface pool ID, then traffic is translated as specified when it accesses any of the interfaces in the pool. For more information about pool IDs, see the "Dynamic NAT Implementation" section on page 6-17.

- Create a new global pool or edit an existing pool by clicking **Manage**. See the "Managing Global Pools" section on page 6-22.

- Choose identity NAT by choosing global pool 0.

Step 6 (Optional) Enter a description in the Description field.

Step 7 (Optional) To enable translation of addresses inside DNS replies, expand the **Connection Settings** area, and check the **Translate the DNS replies that match the translation rule** check box.

If your NAT rule includes the real address of a host that has an entry in a DNS server, and the DNS server is on a different interface from a client, then the client and the DNS server need different addresses for the host; one needs the mapped address and one needs the real address. This option rewrites the address in the DNS reply to the client. The mapped host needs to be on the same interface as either the client or the DNS server. Typically, hosts that need to allow access from other interfaces use a static translation, so this option is more likely to be used with a static rule. See the "DNS and NAT" section on page 6-14 for more information.

Step 8 (Optional) To enable connection settings, expand the **Connection Settings** area, and set one or more of the following options:

Note You can also set these values using a security policy rule. To set the number of rate intervals maintained for host statistics, on the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, choose **1**, **2**, or **3** from the User can specify the number of rate for Threat Detection Host drop-down list. Because host statistics use a lot of memory, reducing the number of rate intervals from the default of 3 reduces the memory usage. By default, the Firewall Dashboard Tab shows information for three rate intervals, for example, for the last 1 hour, 8 hours, and 24 hours. If you set this keyword to **1**, then only the shortest rate interval statistics are maintained. If you set the value to **2**, then the two shortest intervals are maintained. If you set them in both places, then the ASA uses the lower limit. For TCP sequence randomization, if it is disabled using either method, then the ASA disables TCP sequence randomization.

- **Randomize sequence number**—With this check box checked (the default), the ASA randomizes the sequence number of TCP packets. Each TCP connection has two ISNs: one generated by the client and one generated by the server. The ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

  TCP initial sequence number randomization can be disabled if required. For example:

  - If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
  - If you use eBGP multi-hop through the ASA, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
  - You use a WAAS device that requires the ASA not to randomize the sequence numbers of connections.
- **Maximum TCP Connections**—Specifies the maximum number of TCP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum UDP Connections**—Specifies the maximum number of UDP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum Embryonic Connections**—Specifies the maximum number of embryonic connections per host up to 65,536. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. This limit enables the TCP Intercept feature. The default is **0**, which means the maximum embryonic connections. TCP Intercept protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. When the embryonic limit has been surpassed, the TCP intercept feature intercepts TCP SYN packets from clients to servers on a higher security level. SYN cookies are used during the validation process and help to minimize the amount of valid traffic being dropped. Thus, connection attempts from unreachable hosts will never reach the server.

Step 9 Click **OK**.

Using Static NAT

This section describes how to configure a static translation, using regular or policy static NAT, PAT, or identity NAT.

For more information about static NAT, see the "Static NAT" section on page 6-9.

Policy NAT lets you identify real addresses for address translation by specifying the source and destination addresses. You can also optionally specify the source and destination ports. Regular NAT can only consider the source addresses, and not the destination. See the "Policy NAT" section on page 6-11 for more information.

Static PAT lets you translate the real IP address to a mapped IP address, as well as the real port to a mapped port. You can choose to translate the real port to the same port, which lets you translate only specific types of traffic, or you can take it further by translating to a different port. For applications that require application inspection for secondary channels (for example, FTP and VoIP), the ASA automatically translates the secondary ports. For more information about static PAT, see the "Static PAT" section on page 6-9.

You cannot use the same real or mapped address in multiple static rules between the same two interfaces unless you use static PAT. Do not use a mapped address in the static rule that is also defined in a global pool for the same mapped interface.

Static identity NAT translates the real IP address to the same IP address.

This section includes the following topics:

Configuring Static NAT, PAT, or Identity NAT

Figure 6-21 shows typical static NAT, static PAT, and static identity NAT scenarios. The translation is always active so both translated and remote hosts can originate connections.

Figure 6-21 Static NAT Scenarios

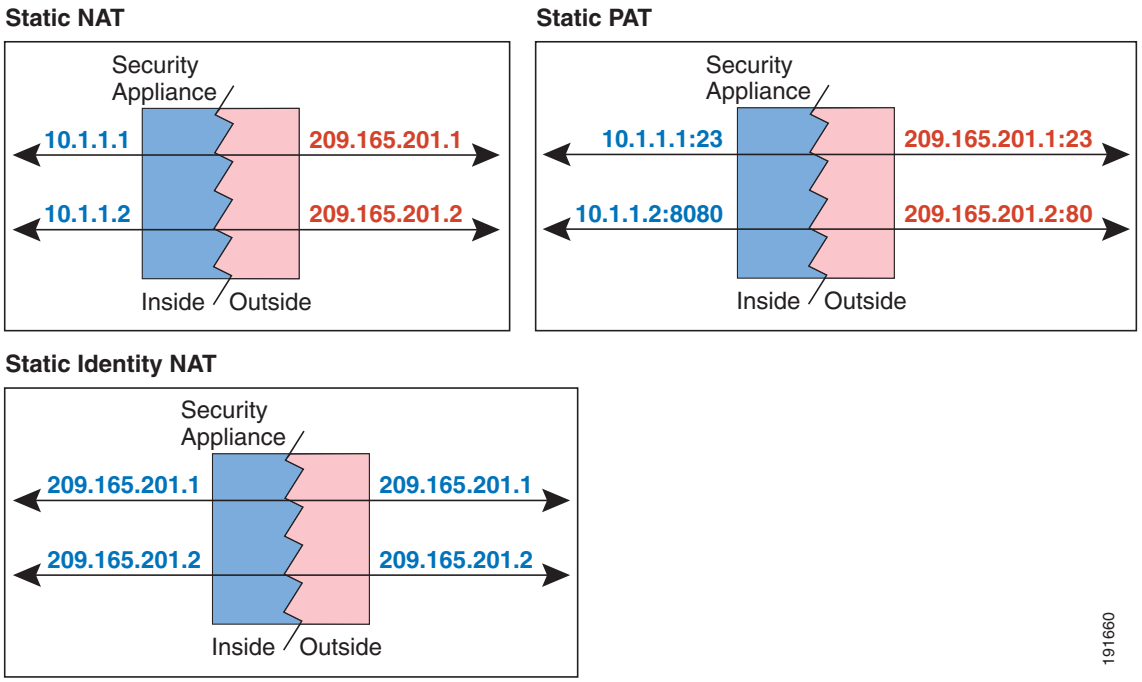

To configure static NAT, PAT, or identity NAT, perform the following steps:

Step 1 In the Configuration > Firewall > NAT Rules pane, choose **Add > Add Static NAT Rule**.

The Add Static NAT Rule dialog box appears.

Step 2 In the Original area, from the Interface drop-down list, choose the interface that is connected to the hosts with real addresses that you want to translate.

Step 3 Enter the real addresses in the Source field, or click the **...** button to choose an IP address that you already defined in ASDM.

Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

Step 4 In the Translated area, from the Interface drop-down list, choose the interface where you want to use the mapped addresses.

Step 5 Specify the mapped IP address by clicking one of the following:

- **Use IP Address**

  Enter the IP address or click the **...** button to choose an IP address that you already defined in ASDM.

  Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

- **Use Interface IP Address**

The real and mapped addresses must have the same subnet mask.

Note For identity NAT, enter the same IP address in the Original and Translated fields.

Step 6 (Optional) To use static PAT, check **Enable Port Address Translation (PAT)**.

a. For the Protocol, click **TCP** or **UDP**.

b. In the Original Port field, enter the real port number.

c. In the Translated Port field, enter the mapped port number.

Step 7 (Optional) To enable translation of addresses inside DNS replies, expand the **Connection Settings** area, and check the **Translate the DNS replies that match the translation rule** check box.

If your NAT rule includes the real address of a host that has an entry in a DNS server, and the DNS server is on a different interface from a client, then the client and the DNS server need different addresses for the host; one needs the mapped address and one needs the real address. This option rewrites the address in the DNS reply to the client. The mapped host needs to be on the same interface as either the client or the DNS server. See the "DNS and NAT" section on page 6-14 for more information.

Step 8 (Optional) To enable connection settings, expand the **Connection Settings** area, and set one or more of the following options:

Note You can also set these values using a security policy rule. To set the number of rate intervals maintained for host statistics, on the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, choose **1**, **2**, or **3** from the User can specify the number of rate for Threat Detection Host drop-down list. Because host statistics use a lot of memory, reducing the number of rate intervals from the default of 3 reduces the memory usage. By default, the Firewall Dashboard Tab shows information for three rate intervals, for example, for the last 1 hour, 8 hours, and 24 hours. If you set this keyword to **1**, then only the shortest rate interval statistics are maintained. If you set the value to **2**, then the two shortest intervals are maintained. If you set them in both places, then the ASA uses the lower limit. For TCP sequence randomization, if it is disabled using either method, then the ASA disables TCP sequence randomization.

- **Randomize sequence number**—With this check box checked (the default), the ASA randomizes the sequence number of TCP packets. Each TCP connection has two ISNs: one generated by the client and one generated by the server. The ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

  TCP initial sequence number randomization can be disabled if required. For example:

  - If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
  - If you use eBGP multi-hop through the ASA, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
  - You use a WAAS device that requires the ASA not to randomize the sequence numbers of connections.
- **Maximum TCP Connections**—Specifies the maximum number of TCP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum UDP Connections**—Specifies the maximum number of UDP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum Embryonic Connections**—Specifies the maximum number of embryonic connections per host up to 65,536. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. This limit enables the TCP Intercept feature. The default is **0**, which means the maximum embryonic connections. TCP Intercept protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. When the embryonic limit has been surpassed, the TCP intercept feature intercepts TCP SYN packets from clients to servers on a higher security level. SYN cookies are used during the validation process and help to minimize the amount of valid traffic being dropped. Thus, connection attempts from unreachable hosts will never reach the server.

Step 9 Click **OK**.

Configuring Static Policy NAT, PAT, or Identity NAT

Figure 6-22 shows typical static policy NAT, static policy PAT, and static policy identity NAT scenarios. The translation is always active so both translated and remote hosts can originate connections.

Figure 6-22 Static Policy NAT Scenarios

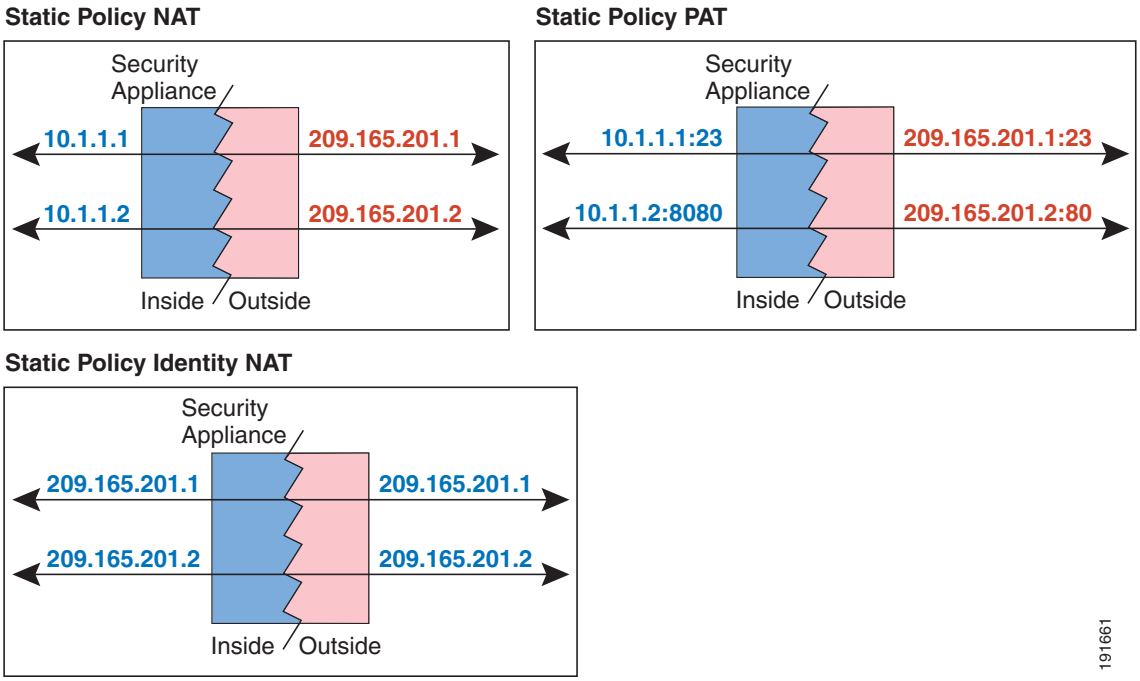

To configure static policy NAT, PAT, or identity NAT, perform the following steps:

Step 1 In the Configuration > Firewall > NAT Rules pane, choose **Add > Advanced > Add Static Policy NAT Rule**.

The Add Static Policy NAT Rule dialog box appears.

Step 2 In the Original area, from the Interface drop-down list, choose the interface that is connected to the hosts with real addresses that you want to translate.

Step 3 Enter the real addresses in the Source field, or click the **...** button to choose an IP address that you already defined in ASDM.

Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

Step 4 Enter the destination addresses in the Destination field, or click the **...** button to choose an IP address that you already defined in ASDM.

Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

Separate multiple destination addresses by a comma.

By default, the field shows **any**, which allows any destination address.

Step 5 In the Translated area, from the Interface drop-down list, choose the interface where you want to use the mapped addresses.

Step 6 Specify the mapped IP address by clicking one of the following:

- **Use IP Address**

  Enter the IP address or click the **...** button to choose an IP address that you already defined in ASDM.

  Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

- **Use Interface IP Address**

The real and mapped addresses must have the same subnet mask.

Step 7 (Optional) To use static PAT, check **Enable Port Address Translation (PAT)**.

a. For the Protocol, click **TCP** or **UDP**.

b. In the Original Port field, enter the real port number.

c. In the Translated Port field, enter the mapped port number.

Step 8 (Optional) Enter a description in the Description field.

Step 9 (Optional) To enable translation of addresses inside DNS replies, expand the **Connection Settings** area, and check the **Translate the DNS replies that match the translation rule** check box.

If your NAT rule includes the real address of a host that has an entry in a DNS server, and the DNS server is on a different interface from a client, then the client and the DNS server need different addresses for the host; one needs the mapped address and one needs the real address. This option rewrites the address in the DNS reply to the client. The mapped host needs to be on the same interface as either the client or the DNS server. See the "DNS and NAT" section on page 6-14 for more information.

Step 10 (Optional) To enable connection settings, expand the **Connection Settings** area, and set one or more of the following options:

Note You can also set these values using a security policy rule. To set the number of rate intervals maintained for host statistics, on the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, choose **1**, **2**, or **3** from the User can specify the number of rate for Threat Detection Host drop-down list. Because host statistics use a lot of memory, reducing the number of rate intervals from the default of 3 reduces the memory usage. By default, the the Firewall Dashboard Tab shows information for three rate intervals, for example, for the last 1 hour, 8 hours, and 24 hours. If you set this keyword to **1**, then only the shortest rate interval statistics are maintained. If you set the value to **2**, then the two shortest intervals are maintained. If you set them in both places, then the ASA uses the lower limit. For TCP sequence randomization, if it is disabled using either method, then the ASA disables TCP sequence randomization.

- **Randomize sequence number**—With this check box checked (the default), the ASA randomizes the sequence number of TCP packets. Each TCP connection has two ISNs: one generated by the client and one generated by the server. The ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

  TCP initial sequence number randomization can be disabled if required. For example:

  – If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.

  – If you use eBGP multi-hop through the ASA, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.

  - You use a WAAS device that requires the ASA not to randomize the sequence numbers of connections.
- **Maximum TCP Connections**—Specifies the maximum number of TCP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum UDP Connections**—Specifies the maximum number of UDP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum Embryonic Connections**—Specifies the maximum number of embryonic connections per host up to 65,536. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. This limit enables the TCP Intercept feature. The default is **0**, which means the maximum embryonic connections. TCP Intercept protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. When the embryonic limit has been surpassed, the TCP intercept feature intercepts TCP SYN packets from clients to servers on a higher security level. SYN cookies are used during the validation process and help to minimize the amount of valid traffic being dropped. Thus, connection attempts from unreachable hosts will never reach the server.

Step 11 Click **OK**.

Using NAT Exemption

NAT exemption exempts addresses from translation and allows both real and remote hosts to originate connections. NAT exemption lets you specify the real and destination addresses when determining the real traffic to exempt (similar to policy NAT), so you have greater control using NAT exemption than dynamic identity NAT. However unlike policy NAT, NAT exemption does not consider the ports. Use static policy identity NAT to consider ports.

For more information about NAT exemption, see the "Bypassing NAT When NAT Control is Enabled" section on page 6-10.

Figure 6-23 shows a typical NAT exemption scenario.

Figure 6-23 NAT Exemption

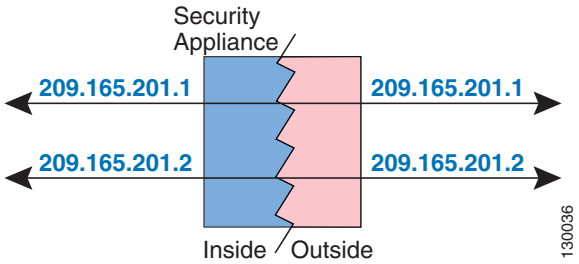

To configure NAT exemption, perform the following steps:

Step 1 In the Configuration > Firewall > NAT Rules pane, choose **Add > Add NAT Exempt Rule**.

The Add NAT Exempt Rule dialog box appears.

Step 2 Click **Action: Exempt**.

Step 3 In the Original area, from the Interface drop-down list, choose the interface that is connected to the hosts with real addresses that you want to exempt.

Step 4 Enter the real addresses in the Source field, or click the **...** button to choose an IP address that you already defined in ASDM.

Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

Note You can later specify addresses that you do not want to exempt. For example, you can specify a subnet to exempt such as 10.1.1.0/24, but if you want to translate 10.1.1.50, then you can create a separate rule for that address that removes the exemption.

Separate multiple real addresses by a comma.

Step 5 Enter the destination addresses in the Destination field, or click the **...** button to choose an IP address that you already defined in ASDM.

Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

Separate multiple destination addresses by a comma.

By default, the field shows **any**, which allows any destination address.

Step 6 In the NAT Exempt Direction area, choose whether you want to exempt traffic going to lower security interfaces (the default) or to higher security interfaces by clicking the appropriate radio button.

Step 7 (Optional) Enter a description in the Description field.

Step 8 Click **OK**.

Step 9 (Optional) If you do not want to exempt some addresses that were included in your NAT exempt rule, then create another rule to remove the exemption. Right-click the existing NAT Exempt rule, and choose **Insert**.

The Add NAT Exempt Rule dialog box appears.

a. Click **Action: Do not exempt**.

b. Complete Steps 3 through 8 to complete the rule.

The No Exempt rule is added before the Exempt rule. The order of Exempt and No Exempt rules is important. When the ASA decides whether to exempt a packet, the ASA tests the packet against each NAT exempt and No Exempt rule in the order in which the rules are listed. After a match is found, no more rules are checked.