

Network Address Translation (NAT)

The following topics explain Network Address Translation (NAT) and how to configure it.

- Why use NAT?, on page 1
- NAT basics, on page 2
- Guidelines for NAT, on page 7
- Dynamic NAT, on page 14
- Dynamic PAT, on page 23
- Static NAT, on page 42
- Identity NAT, on page 54
- Monitoring NAT, on page 61
- History for NAT, on page 61

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 basics

The following topics explain some of the basics of NAT.

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 device, 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.



Note During address translation, IP addresses configured for the device interfaces are not translated.

- 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

You can implement NAT using the following methods:

- 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 Dynamic NAT, on page 14.
- 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 Dynamic PAT, on page 23.
- Static NAT—A consistent mapping between a real and mapped IP address. Allows bidirectional traffic initiation. See Static NAT, on page 42.
- 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 Identity NAT, on page 54.

Network object NAT and Twice NAT

You can implement address translation in two ways: *network object NAT* and *twice NAT*.

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

Network object NAT

All NAT rules that are configured as a parameter of a network object are considered to be *network object NAT* rules. This is a quick and easy way to configure NAT for a network object. You cannot create these rules for a group object, however.

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 an interface, 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, where you can identify the source and destination address in a single rule.

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.

Note 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. For example, if you configure static NAT with port address translation, and specify the source address as a Telnet server, and you want all traffic going to that Telnet server to have the port translated from 2323 to 23, then you must specify the *source* ports to be translated (real: 23, mapped: 2323). You specify the source ports because you specified the Telnet server address as the source address.

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.

Comparing 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 the NAT 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 packet matches one rule only, and further rules are not checked. Even if you do not configure the optional destination address, a matching packet still matches one twice NAT rule only. 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).

NAT rule order

Network object NAT 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, until a match is found. For example, if

a match is found in section 1, sections 2 and 3 are not evaluated. The following table shows the order of rules within each section.

Note There is also a Section 0, which contains any NAT rules that the system creates for its own use. These rules have priority over all others. The system automatically creates these rules and clears xlates as needed. You cannot add, edit, or modify rules in Section 0.

Table 1: NAT Rule Table

Table Section	Rule Type	Order of Rules within the Section
Section 1	Twice NAT	Applied on a first match basis, in the order they appear in the configuration. Because the first match is applied, you must ensure that specific rules come before more general rules, or the specific rules might not be applied as desired. By default, twice NAT rules are added to section 1. By "specific rules first," we mean: • Static rules should come before dynamic rules. • Rules that include destination translation should come before rules with source translation only. If you cannot eliminate overlapping rules, where more than one rule might apply based on the source or destination address, be especially careful to follow these recommendations. If you do not follow the above recommendation, ensure that accurate routes are configured for the mapped addresses to match the correct NAT rule.
Section 2	Network object NAT	If a match in section 1 is not found, section 2 rules are applied in the following order: **1.** Static rules. **2.** Dynamic rules. **Within each rule type, the following ordering guidelines are used:** **1.** 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. **2.** 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. **3.** 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.

Table Section	Rule Type	Order of Rules within the Section
Section 3	Twice NAT	If a match is still not found, section 3 rules are applied on a first match basis, in the order they appear in the configuration. This section should contain your most general rules. You must also ensure that any specific rules in this section come before general rules that would otherwise apply.

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:

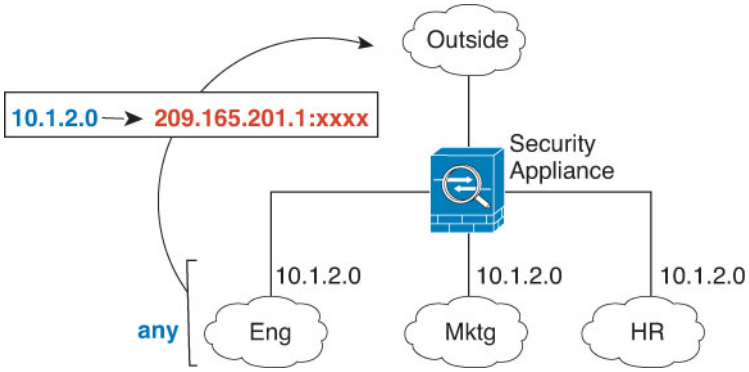
- 192.168.1.1/32 (static)
- 10.1.1.0/24 (static)
- 192.168.1.0/24 (static)
- 172.16.1.0/24 (dynamic) (object abc)
- 172.16.1.0/24 (dynamic) (object def)
- 192.168.1.0/24 (dynamic)

NAT interfaces

Except for bridge group member 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.

Figure 1: Specifying Any Interface



However, the concept of "any" interface does not apply to bridge group member interfaces. When you specify "any" interface, all bridge group member interfaces are excluded. Thus, to apply NAT to bridge group members, you must specify the member interface. This could result in many similar rules where only one interface is different. You cannot configure NAT for the Bridge Virtual Interface (BVI) itself, you can configure NAT for member interfaces only.

Guidelines for NAT

The following topics provide detailed guidelines for implementing NAT.

Firewall mode guidelines for NAT

NAT is supported in routed and transparent firewall mode.

However, configuring NAT on bridge group member interfaces (interfaces that are part of a Bridge Group Virtual Interface, or BVI) has the following restrictions:

- When configuring NAT for the members of a bridge group, you specify the member interface. You cannot configure NAT for the bridge group interface (BVI) itself.
- When doing NAT between bridge group member interfaces, you must specify the real and mapped addresses. You cannot specify "any" as the interface.
- You cannot configure interface PAT when the mapped address is a bridge group member interface, because there is no IP address attached to the interface.
- You cannot translate between IPv4 and IPv6 networks (NAT64/46) when the source and destination interfaces are members of the same bridge group. Static NAT/PAT 44/66, dynamic NAT44/66, and dynamic PAT44 are the only allowed methods; dynamic PAT66 is not supported. However, you can do NAT64/46 between members of different bridge groups, or between a bridge group member (source) and standard routed interface (destination).

IPv6 NAT guidelines

NAT supports IPv6 with the following guidelines and restrictions.

- For standard routed mode interfaces, you can also translate between IPv4 and IPv6.

- You cannot translate between IPv4 and IPv6 for interfaces that are members of the same bridge group. You can translate between two IPv6 or two IPv4 networks only. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.
- You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.
- For static NAT, you can specify an IPv6 subnet up to /64. Larger subnets are not supported.
- When using FTP with NAT46, when an IPv4 FTP client connects to an IPv6 FTP server, the client must use either the extended passive mode (EPSV) or extended port mode (EPRT); PASV and PORT commands are not supported with IPv6.

IPv6 NAT best practices

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-to-net, 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.

Additional guidelines for NAT

- NAT rules apply to through the device traffic only, they do not apply to traffic initiated by the device, such as a RADIUS authentication.
- For interfaces that are members of a bridge group, you write NAT rules for the member interfaces. You cannot write NAT rules for the Bridge Virtual Interface (BVI) itself.
- You cannot write NAT rules for a Virtual Tunnel Interface (VTI), which are used in site-to-site VPN. Writing rules for the VTI's source interface will not apply NAT to the VPN tunnel. To write NAT rules that will apply to VPN traffic tunneled on a VTI, you must use "any" as the interface; you cannot explicitly specify interface names.

- (Network object NAT only.) You can only define a single NAT rule for a given object; if you want to configure multiple NAT rules for an object, you need to create multiple objects with different names that specify the same IP address.
- If a VPN is defined on an interface, inbound ESP traffic on the interface is not subject to the NAT rules. The system allows the ESP traffic for established VPN tunnels only, dropping traffic not associated with an existing tunnel. This restriction applies to ESP and UDP ports 500 and 4500.
- If you define a site-to-site VPN on a device that is behind a device that is applying dynamic PAT, so that UDP ports 500 and 4500 are not the ones actually used, you must initiate the connection from the device that is behind the PAT device. The responder cannot initiate the security association (SA) because it does not know the correct port numbers.
- If you change the NAT configuration, and you do not want to wait for existing translations to time out before the new NAT configuration is used, you can clear the translation table using the **clear xlate** command in the device CLI. However, clearing the translation table disconnects all current connections that use translations.

  If you create a new NAT rule that should apply to an existing connection (such as a VPN tunnel), you need to use **clear conn** to end the connection. Then, the attempt to re-establish the connection should hit the NAT rule and the connection should be NAT'ed correctly.



Note If you remove a dynamic NAT or PAT rule, and then add a new rule with mapped addresses that overlap the addresses in the removed rule, then the new rule will not be used until all connections associated with the removed rule time out or are cleared using the **clear xlate** or **clear conn** commands. This safeguard ensures that the same address is not assigned to multiple hosts.

- When translating SCTP traffic, use static network object NAT only. Dynamic NAT/PAT is not allowed. Although you can configure static twice NAT, this is not recommended because the topology of the destination part of the SCTP association is unknown.
- Objects and object groups used in NAT cannot be undefined; they must include IP addresses.
- You cannot use an object group with both IPv4 and IPv6 addresses; the object group must include only one type of address.
- (Twice NAT only.) When using **any** as the source address in a NAT rule, the definition of “any” traffic (IPv4 vs. IPv6) depends on the rule. Before the ASA performs NAT on a packet, the packet must be IPv6-to-IPv6 or IPv4-to-IPv4; with this prerequisite, the ASA can determine the value of **any** in a NAT rule. For example, if you configure a rule from “any” to an IPv6 server, and that server was mapped from an IPv4 address, then **any** means “any IPv6 traffic.” If you configure a rule from “any” to “any,” and you map the source to the interface IPv4 address, then **any** means “any IPv4 traffic” because the mapped interface address implies that the destination is also IPv4.
- You can use the same mapped object or group in multiple NAT rules.
- The mapped IP address pool cannot include:
  - The mapped interface IP address. If you specify “any” interface for the rule, then all interface IP addresses are disallowed. For interface PAT (routed mode only), specify the interface name instead of the interface address.

  - The failover interface IP address.
  - (Transparent mode.) The management IP address.
  - (Dynamic NAT.) The standby interface IP address when VPN is enabled.
  - Existing VPN pool addresses.

- Avoid using overlapping addresses in static and dynamic NAT policies. For example, with overlapping addresses, a PPTP connection can fail to get established if the secondary connection for PPTP hits the static instead of dynamic xlate.
- You cannot use overlapping addresses in the source address of a NAT rule and a remote access VPN address pool.
- For application inspection limitations with NAT or PAT, see Default Inspections and NAT Limitations.
- The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. See Routing NAT Packets for more information.
- If you enable the **arp permit-nonconnected** command, the system does not respond to ARP requests if the mapped address is not part of any connected subnet and you also do not specify the mapped interface in the NAT rule (that is, you specify "any" interface). To resolve this problem, specify the mapped interface.
- If you specify a destination interface in a rule, then that interface is used as the egress interface rather than looking up the route in the routing table. However, for identity NAT, you have the option to use a route lookup instead.
- If you use PAT on Sun RPC traffic, which is used to connect to NFS servers, be aware that the NFS server might reject connections if the PAT'ed port is above 1024. The default configuration of NFS servers is to reject connections from ports higher than 1024. The error is typically "Permission Denied." Mapping ports above 1024 might happen if you use the "flat range" option to use the higher port numbers if a port in the lower range is not available, especially if you do not select the option to include the lower range in the flat range.Mapping ports above 1024 happens if you do not select the option to include the reserved ports (1-1023) in the port range of a PAT pool. You can avoid this problem by changing the NFS server configuration to allow all port numbers.
- NAT applies to through traffic only. Traffic generated by the system is not subject to NAT.
- You can improve system performance and reliability by using the transactional commit model for NAT. See the basic settings chapter in the general operations configuration guide for more information. The option is under **Configurations > Device Management > Advanced > Rule Engine**.
- Do not name a network object or group pat-pool, using any combination of upper- or lower-case letters.
- The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.
- You cannot use NAT on the internal payload of Protocol Independent Multicast (PIM) registers.
- (Twice NAT) When writing NAT rules for a dual ISP interface setup (primary and backup interfaces using service level agreements in the routing configuration), do not specify destination criteria in the rule. Ensure the rule for the primary interface comes before the rule for the backup interface. This allows the device to choose the correct NAT destination interface based on the current routing state when the

primary ISP is unavailable. If you specify destination objects, the NAT rule will always select the primary interface for the otherwise duplicate rules.

- If you get the ASP drop reason nat-no-xlate-to-pat-pool for traffic that should not match the NAT rules defined for the interface, configure identity NAT rules for the affected traffic so the traffic can pass untranslated.
- If you configure NAT for GRE tunnel endpoints, you must disable keepalives on the endpoints or the tunnel cannot be established. The endpoints send keepalives to the original addresses.
- DHCP and BOOTP share ports UDP/67-68. Because BOOTP is obsolete, writing NAT rules for the bootps port can cause port allocation problems when also running DHCP. Consider using DHCP relay instead for transmitting DHCP requests between network segments.
- In rare cases, return traffic (server to client) with an existing translation (xlate) might be logged as a new flow in Connection Events. This can occur when the client has already terminated the connection and the server sends another packet that reaches the device during the brief interval between connection closure and xlate removal, often due to application behavior or TCP stack cleanup. Because the device removes the xlate only after deleting the connection, a server packet can arrive while the xlate still exists. If no valid connection entry is found, the device logs a separate connection event based on the matched Access Control Policy rule.

Network Object NAT Guidelines for Mapped Address Objects

For dynamic NAT, you must use an object or group for the mapped addresses. For the other NAT types, you can use an object or group, or you have the option of using inline addresses. Network object groups are particularly useful for creating a mapped address pool with discontinuous IP address ranges or multiple hosts or subnets.

Consider the following guidelines when creating objects for mapped addresses.

- A network object group can contain objects or inline addresses of either IPv4 or IPv6 addresses. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.
- See Additional guidelines for NAT, on page 8 for information about disallowed mapped IP addresses.
- Do not name a network object or group pat-pool, using any combination of upper- or lower-case letters.
- Dynamic NAT:
  - You cannot use an inline address; you must configure a network object or group.
  - The object or group cannot contain a subnet; the object must define a range; the group can include hosts and ranges.
  - If a mapped network object contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.
- Dynamic PAT (Hide):
  - Instead of using an object, you can optionally configure an inline host address or specify the interface address.
  - If you use an object, the object or group cannot contain a subnet. The object must define a host, or for a PAT pool, a range. The group (for a PAT pool) can include hosts and ranges.

- Static NAT or Static NAT with port translation:
  - Instead of using an object, you can configure an inline address or specify the interface address (for static NAT-with-port-translation).
  - If you use an object, the object or group can contain a host, range, or subnet.
- Identity NAT
  - Instead of using an object, you can configure an inline address.
  - If you use an object, the object must match the real addresses you want to translate.

Twice NAT Guidelines for Real and Mapped Address Objects

For each NAT rule, configure up to four network objects or groups for:

- Source real address
- Source mapped address
- Destination real address
- Destination mapped address

Objects are required unless you specify the **any** keyword inline to represent all traffic, or for some types of NAT, the **interface** keyword to represent the interface address. Network object groups are particularly useful for creating a mapped address pool with discontinuous IP address ranges or multiple hosts or subnets.

Consider the following guidelines when creating objects for twice NAT.

- A network object group can contain objects or inline addresses of either IPv4 or IPv6 addresses. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.
- See Additional guidelines for NAT, on page 8 for information about disallowed mapped IP addresses.
- Do not name a network object or group pat-pool, using any combination of upper- or lower-case letters.
- Source Dynamic NAT:
  - You typically configure a larger group of real addresses to be mapped to a smaller group.
  - The mapped object or group cannot contain a subnet; the object must define a range; the group can include hosts and ranges.
  - If a mapped network object contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.
- Source Dynamic PAT (Hide):
  - If you use an object, the object or group cannot contain a subnet. The object must define a host, or for a PAT pool, a range. The group (for a PAT pool) can include hosts and ranges.
- Source Static NAT or Static NAT with port translation:
  - The mapped object or group can contain a host, range, or subnet.

    - The static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired.

- Source Identity NAT
    - The real and mapped objects must match. You can use the same object for both, or you can create separate objects that contain the same IP addresses.

- Destination Static NAT or Static NAT with port translation (the destination translation is always static):
    - Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 4.
    - For identity NAT, the real and mapped objects must match. You can use the same object for both, or you can create separate objects that contain the same IP addresses.
    - The static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired.
    - For static interface NAT with port translation (routed mode only), you can specify the **interface** keyword instead of a network object/group for the mapped address.
    - You can use a fully-qualified domain name, such as www.example.com, as the translated (mapped) destination. For details, see FQDN destination guidelines, on page 13.

FQDN destination guidelines

You can specify the translated (mapped) destination in a twice NAT rule using a fully-qualified domain name (FQDN) network object instead of an IP address. For example, you can create a rule based on traffic that is destined for the www.example.com web server.

When using an FQDN, the system obtains the DNS resolution and writes the NAT rule based on the returned address. If you are using multiple DNS server groups, the filter domains are honored and the address is requested from the appropriate group based on the filters. If more than one address is obtained from the DNS server, the address used is based on the following:

- If there is an address on the same subnet as the specified interface, that address is used. If there isn't one on the same subnet, the first address returned is used.
- The IP type for the translated source and translated destination must match. For example, if the translated source address is IPv6, the FQDN object must specify IPv6 as the address type. If the translated source is IPv4, the FQDN object must specify IPv4 as the address type.

You cannot include an FQDN object in a network group that is used for manual NAT destination. In NAT, an FQDN object must be used alone, as only a single destination host makes sense for this type of NAT rule.

If the FQDN cannot be resolved to an IP address, the rule is not functional until a DNS resolution is obtained.

Twice NAT Guidelines for Service Objects for Real and Mapped Ports

You can optionally configure service objects for:

- **Source real port (Static only)** *or* **Destination real port**
- **Source mapped port (Static only)** *or* **Destination mapped port**

Consider the following guidelines when creating objects for twice NAT.

- NAT supports TCP, UDP, and SCTP only. When translating a port, be sure the protocols in the real and mapped service objects are identical (for example, both TCP). Although you can configure static twice NAT rules with SCTP port specifications, this is not recommended, because the topology of the destination part of the SCTP association is unknown. Use static object NAT instead for SCTP.
- The "not equal" (**neq**) operator is not supported.
- For identity port translation, you can use the same service object for both the real and mapped ports.
- Source Dynamic NAT—Source Dynamic NAT does not support port translation.
- Source Dynamic PAT (Hide)—Source Dynamic PAT does not support port translation.
- Source Static NAT, Static NAT with port translation, or Identity NAT—A service object can contain both a source and destination port; however, you should specify *either* the source *or* the destination port for both service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. For example, if you want to translate the port for the source host, then configure the source service.
- Destination Static NAT or Static NAT with port translation (the destination translation is always static)—For non-static source NAT, you can only perform port translation on the destination. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored.

Dynamic NAT

The following topics explain dynamic NAT and how to configure it.

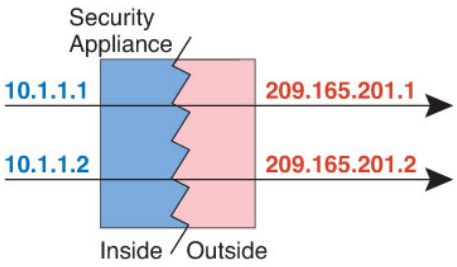
About dynamic NAT

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, NAT 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.

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. A successful connection from a remote host can reset the idle timer for the connection.

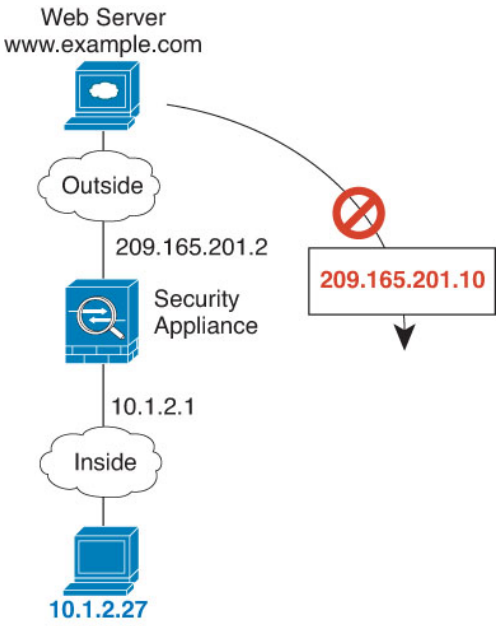
The following figure shows a typical dynamic NAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back.

Figure 2: Dynamic NAT



The following figure shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table; therefore, the packet is dropped.

Figure 3: Remote Host Attempts to Initiate a Connection to a Mapped Address



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 fall-back 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.

Configure Dynamic Network Object NAT

This section describes how to configure network object NAT for dynamic NAT.

Procedure

Step 1 Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration** > **Firewall** > **NAT Rules**, then click **Add > Add Network Object NAT Rule**.
- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then edit a network object.

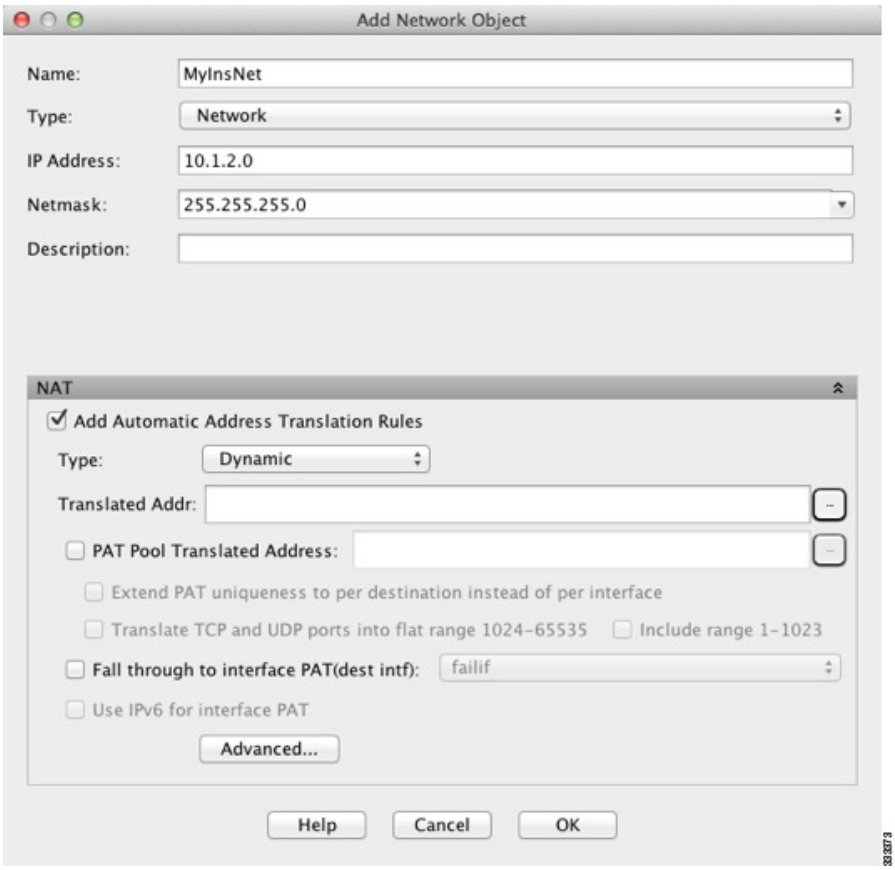
Step 2 For a new object, enter values for the following fields:

a) **Name**—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.

b) **Type**—Host, Network, or Range.

c) **IP Addresses**—IPv4 or IPv6 addresses, a single address for a host, a starting and ending address for a range, and for subnet, either an IPv4 network address and mask (for example, 10.100.10.0 255.255.255.0) or IPv6 address and prefix length (for example, 2001:DB8:0:CD30::/60).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.



Step 4 Check the **Add Automatic Translation Rules** check box.

Step 5 From the Type drop-down list, choose **Dynamic**.

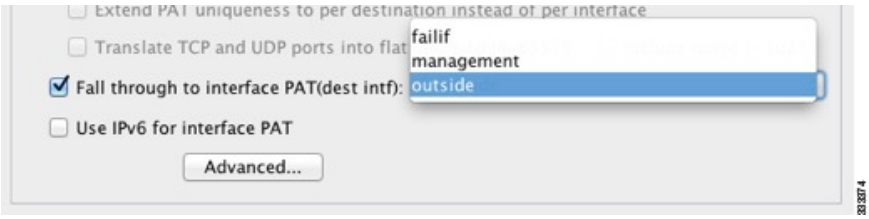
Step 6 To the right of the Translated Addr field, click the browse button and choose the network object or network object group that contains the mapped addresses.

You can create a new object if necessary.

The object or group cannot contain a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

If a mapped network object contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.

Step 7 (Optional, when mapped interface is a non-bridge group member only.) To use the interface IP address as a backup method when the other mapped addresses are already allocated, check the **Fall through to interface PAT (dest intf)** check box, and choose the interface from the drop-down list. To use the IPv6 address of the interface, also check the **Use IPv6 for interface PAT** check box.



Step 8 (Optional) Click **Advanced**, configure the following options in the Advanced NAT Settings dialog box, and click **OK**.

- **Translate DNS replies for rule**—Translates the IP address in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). See Rewriting DNS queries and responses using NAT for more information.
- (Required for bridge group member interfaces.) **Interface**—Specifies the real interface (**Source**) and the mapped interface (**Destination**) where this NAT rule applies. By default, the rule applies to all interfaces except for bridge group members.

Step 9 Click **OK**, and then **Apply**.

Configure Dynamic Twice NAT

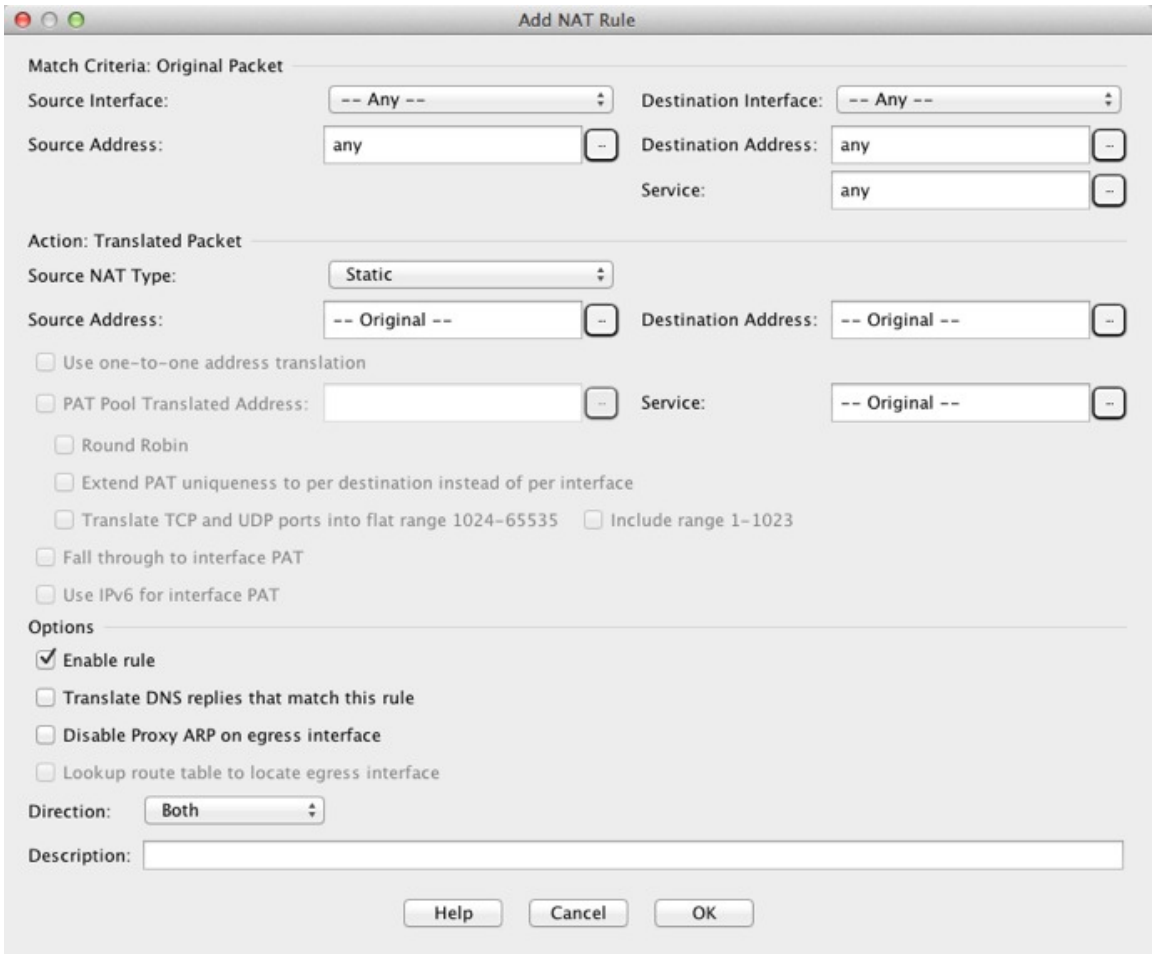
This section describes how to configure twice NAT for dynamic NAT.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**, and then do one of the following:

- Click **Add**, or **Add > Add NAT Rule Before Network Object NAT Rules**.
- Click **Add > Add NAT Rule After Network Object NAT Rules**.
- Select a twice NAT rule and click **Edit**.

The Add NAT Rule dialog box appears.

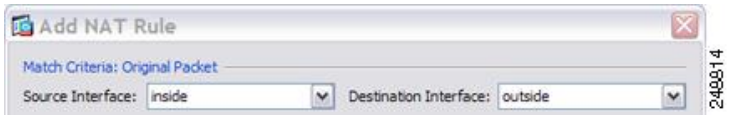


Step 2 (Required for bridge group member interfaces.) Set the source and destination interfaces.

In routed mode, the default is any interface for both source and destination. You can select specific interfaces for one or both options. However, you must select the interface when writing a rule for bridge group member interfaces; "any" does not include these interfaces.

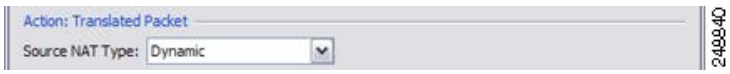
a) From the **Match Criteria: Original Packet > Source Interface** drop-down list, choose the source interface.

b) From the **Match Criteria: Original Packet > Destination Interface** drop-down list, choose the destination interface.

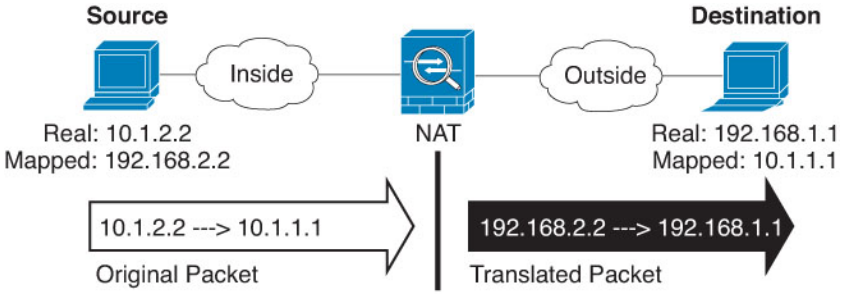


Step 3 Choose **Dynamic** from the **Action: Translated Packet** > **Source NAT Type** drop-down list.

This setting only applies to the source address; the destination translation is always static.



Step 4 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.



a) For the **Match Criteria: Original Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**.

b) (Optional.) For the **Match Criteria: Original Packet > Destination Address**, click the browse button and choose an existing network object, group, or interface, or create a new object or group from the Browse Original Destination Address dialog box. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 4.

For static interface NAT with port translation only, choose an interface from the Browse dialog box. Be sure to also configure a service translation. For this option, you must configure a specific interface for the Source Interface. See Static NAT with port translation, on page 43 for more information.

Step 5 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired.

a) For **Action: Translated Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated Source Address dialog box.

For dynamic NAT, you typically configure a larger group of source addresses to be mapped to a smaller group.

If a mapped network object contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.

Note
The object or group cannot contain a subnet.

b) For **Action: Translated Packet > Destination Address**, click the browse button and choose an existing network object or group, or create a new object or group from the Browse Translated Destination Address dialog box. You can also use an FQDN network object for the destination mapped address.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

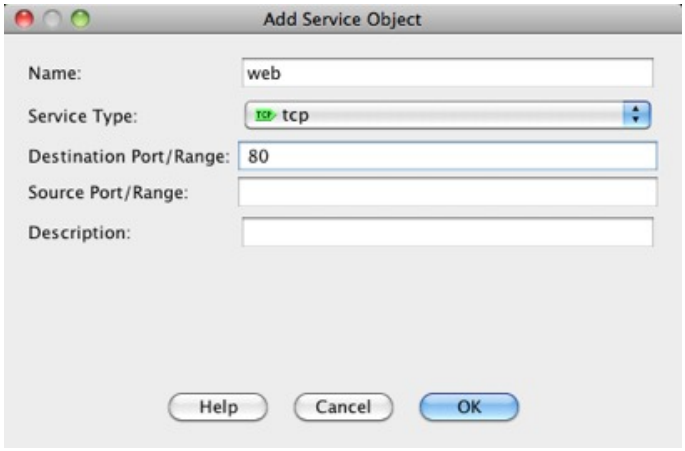
If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see Static NAT, on page 42. See Additional guidelines for NAT, on page 8 for information about disallowed mapped IP addresses.

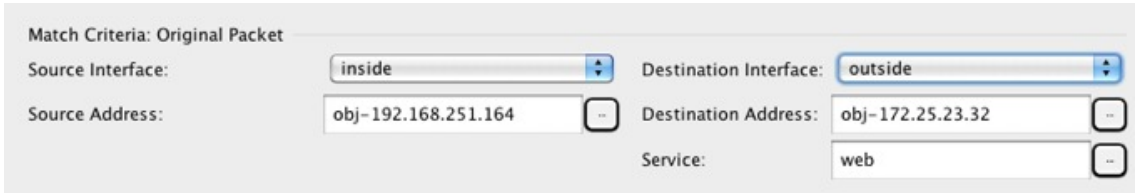
Step 6 (Optional.) Identify the destination service ports for service translation.

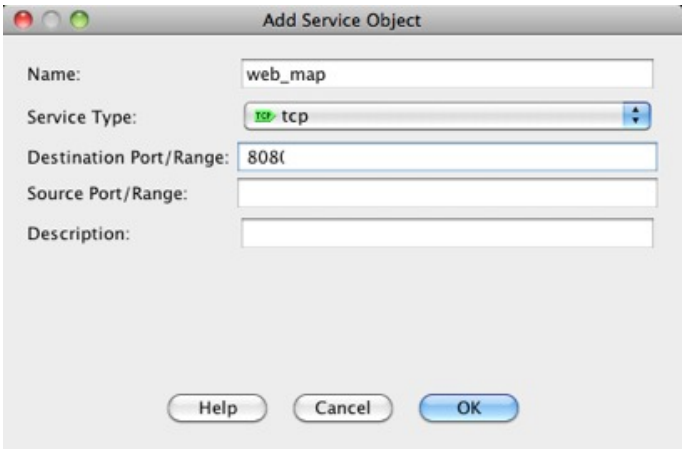
- Identify the original packet port (the *mapped destination port*). For **Match Criteria: Original Packet > Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Original Service dialog box.
- Identify the translated packet port (the *real destination port*). For **Action: Translated Packet > Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Translated Service dialog box.

Dynamic NAT does not support port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

For example:



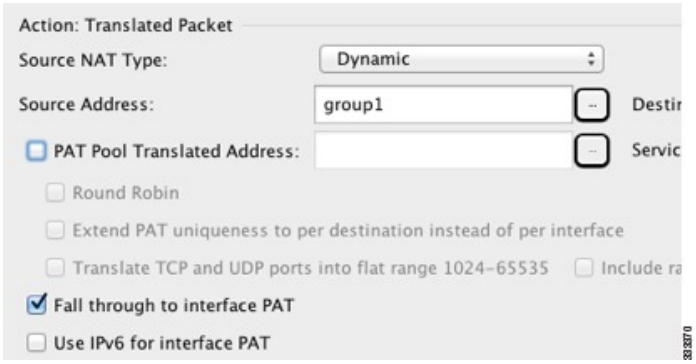




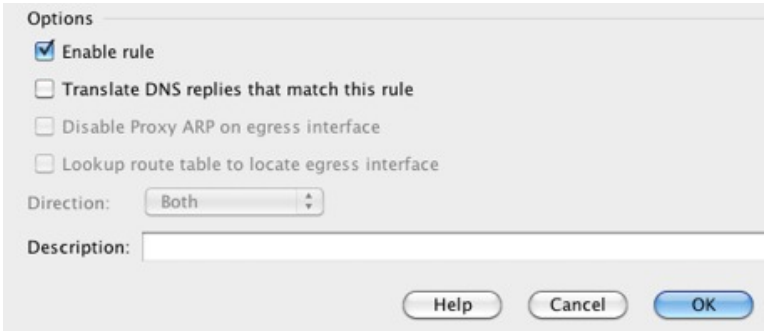


Step 7 (Optional, when mapped interface is a non-bridge group member only.) To use the interface IP address as a backup method if the other mapped source addresses are already allocated, check the **Fall through to interface PAT** check box. To use the IPv6 interface address, also check the **Use IPv6 for interface PAT** check box.

The destination interface IP address is used. This option is only available if you configure a specific Destination Interface.



Step 8 (Optional.) Configure NAT options in the Options area.



- **Enable rule** —Enables this NAT rule. The rule is enabled by default.
- (For a source-only rule) **Translate DNS replies that match this rule**—Rewrites the DNS A record in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). You cannot configure DNS modification if you configure a destination address. See Rewriting DNS queries and responses using NAT for more information.
- **Description**—Adds a description about the rule up to 200 characters in length.

Step 9 Click **OK**, then click **Apply**.

Dynamic PAT

The following topics describe dynamic PAT.

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.

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.

The following figure 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 4: Dynamic PAT



For the duration of the translation, a remote host on the destination network 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.

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.

Note We recommend that you use different PAT pools for each interface. If you use the same pool for multiple interfaces, especially if you use it for "any" interface, the pool can be quickly exhausted, with no ports available for new translations.

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.

You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.

Dynamic PAT does not work with some multimedia applications that have a data stream that is different from the control path. See Default Inspections and NAT Limitations for more information about NAT and PAT support.

Dynamic PAT might 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.

PAT pool object guidelines

When creating network objects for a PAT pool, follow these guidelines.

For a PAT pool

- Ports are mapped to an available port in the 1024 to 65535 range. You can optionally include the reserved ports, those below 1024, to make the entire port range available for translations.

  When operating in a cluster, blocks of 512 ports per address are allocated to the members of the cluster, and mappings are made within these port blocks. If you also enable block allocation, the ports are distributed according to the block allocation size, whose default is also 512. If you change the cluster unit limit (the size of the cluster), ensure that you clear xlates or reboot the devices so that PAT pools can be appropriately reallocated to the cluster units.

- If you enable block allocation for a PAT pool, port blocks are allocated in the 1024-65535 range only. Thus, if an application requires a low port number (1-1023), it might not work. For example, an application requesting port 22 (SSH) will get a mapped port within the range of 1024-65535 and within the block allocated to the host.

- If you use the same PAT pool object in two separate rules, then be sure to specify the same options for each rule. For example, if one rule specifies extended PAT, then the other rule must also specify extended PAT.

- If a host has an existing connection, then subsequent connections from that host use the same PAT IP address. If no ports are available, this can prevent the connection. Use the round robin option to avoid this problem.

- For best performance, limit the number of IP addresses within a PAT pool to 10,000.

For extended PAT for a PAT pool

- Many application inspections do not support extended PAT. See Default Inspections and NAT Limitations for a complete list of unsupported inspections.

- If you enable extended PAT for a dynamic PAT rule, then you cannot also use an address in the PAT pool as the PAT address in a separate static NAT with port translation rule. For example, if the PAT pool includes 10.1.1.1, then you cannot create a static NAT-with-port-translation rule using 10.1.1.1 as the PAT address.

- If you use a PAT pool and specify an interface for fallback, you cannot specify extended PAT.
- For VoIP deployments that use ICE or TURN, do not use extended PAT. ICE and TURN rely on the PAT binding to be the same for all destinations.
- You cannot use extended PAT on units in a cluster.
- Extended PAT increases memory usage on the device.

For round robin for a PAT pool

- If a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. However, this "stickiness" does not survive a failover. If the device fails over, then subsequent connections from a host might not use the initial IP address.
- IP address "stickiness" is also impacted if you mix PAT pool/round robin rules with interface PAT rules on the same interface. For any given interface, choose either a PAT pool or interface PAT; do not create competing PAT rules.
- Round robin, especially when combined with extended PAT, can consume a large amount of memory. Because NAT pools are created for every mapped protocol/IP address/port range, round robin results in a large number of concurrent NAT pools, which use memory. Extended PAT results in an even larger number of concurrent NAT pools.

Configure Dynamic Network Object PAT (Hide)

This section describes how to configure network object NAT for dynamic PAT (hide), which uses a single address for translation instead of a PAT pool.

Procedure

Step 1 Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration** > **Firewall** > **NAT Rules**, then click **Add > Add Network Object NAT Rule**.
- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then edit a network object.

Step 2 For a new object, enter values for the following fields:

a) **Name**—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.

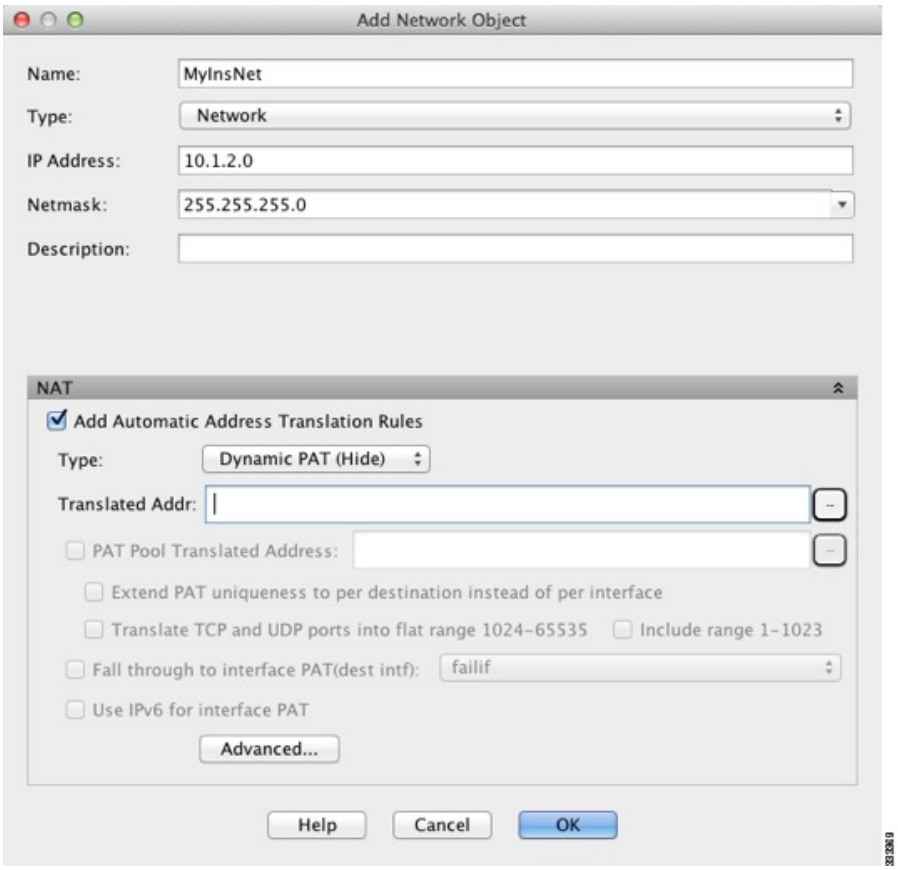
b) **Type**—Host, Network, or Range.

c) **IP Addresses**—IPv4 or IPv6 addresses, a single address for a host, a starting and ending address for a range, and for subnet, either an IPv4 network address and mask (for example, 10.100.10.0 255.255.255.0) or IPv6 address and prefix length (for example, 2001:DB8:0:CD30::/60).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

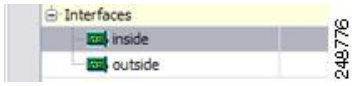
Step 4 Check the **Add Automatic Translation Rules** check box.

Step 5 From the Type drop-down list, choose **Dynamic PAT (Hide)**.



Step 6 Specify a single mapped address. In the Translated Addr. field, specify the mapped IP address by doing one of the following:

- Type a host IP address.
- Click the browse button and select a host network object (or create a new one).
- (Non-bridge group member interfaces only.) Type an interface name or click the browse button, and choose an interface from the Browse Translated Addr dialog box.



If you specify an interface name, then you enable *interface PAT*, where the specified interface IP address is used as the mapped address. To use the IPv6 interface address, you must also check the **Use IPv6 for interface PAT** check box. With interface PAT, the NAT rule only applies to the specified mapped interface, which cannot be a member of a bridge group. (If you do not use interface PAT, then the rule applies to all interfaces by default.) You cannot specify an interface in transparent mode.

Step 7 (Optional.) Click **Advanced**, configure the following options in the Advanced NAT Settings dialog box, and click **OK**.

- (Required for bridge group member interfaces.) **Interface**—Specifies the real interface (**Source**) and the mapped interface (**Destination**) where this NAT rule applies. By default, the rule applies to all interfaces except for bridge group members.

Step 8 Click **OK**, and then **Apply**.

Configure Dynamic Network Object PAT Using a PAT Pool

This section describes how to configure network object NAT for dynamic PAT using a PAT pool.

Procedure

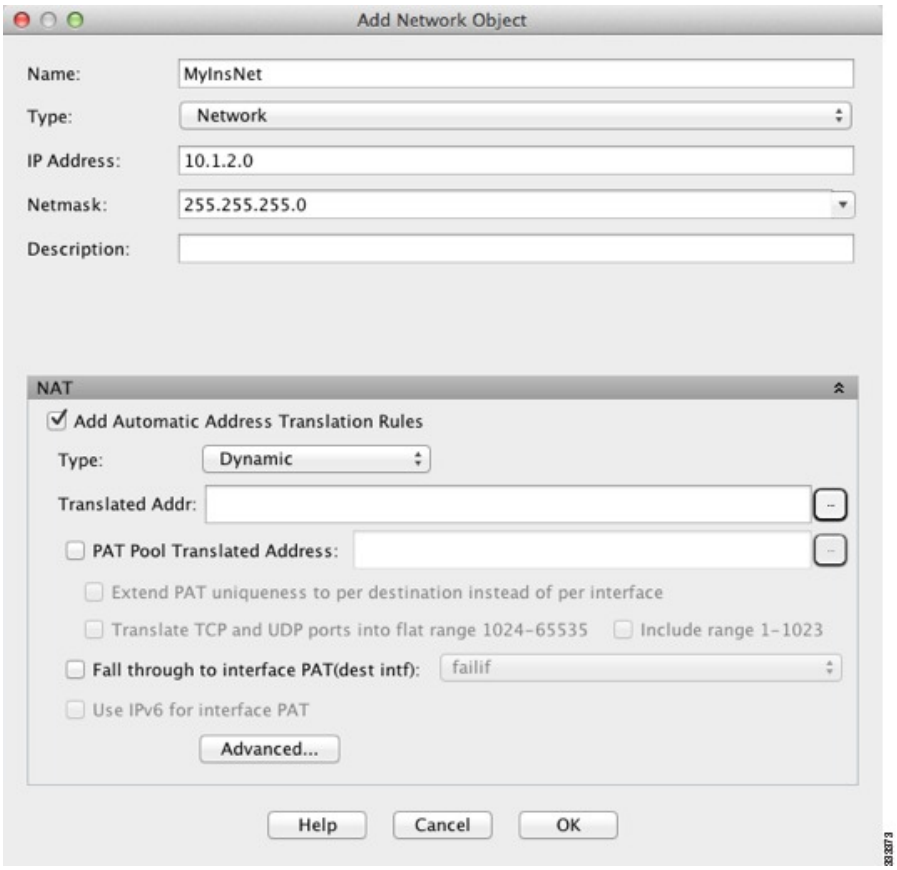
Step 1 Add NAT to a new or existing network object:

- To add a new network object NAT rule, choose **Configuration** > **Firewall** > **NAT Rules**, then click **Add** > **Add Network Object NAT Rule**.
- To add NAT to an existing network object, choose **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups**, and then edit a network object.

Step 2 For a new object, enter values for the following fields:

a) **Name**—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.
b) **Type**—Host, Network, or Range.
c) **IP Addresses**—IPv4 or IPv6 addresses, a single address for a host, a starting and ending address for a range, and for subnet, either an IPv4 network address and mask (for example, 10.100.10.0 255.255.255.0) or IPv6 address and prefix length (for example, 2001:DB8:0:CD30::/60).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.



Step 4 Check the **Add Automatic Translation Rules** check box.

Step 5 From the Type drop-down list, choose **Dynamic** even though you are configuring dynamic PAT with a PAT pool.

Step 6 To configure the PAT pool:

a) Do not enter a value for the Translated Addr. field; leave it blank.

b) Check the **PAT Pool Translated Address** check box, then click the browse button and choose the network object or group that contains the PAT pool addresses. Or create a new object from the Browse Translated PAT Pool Address dialog box.

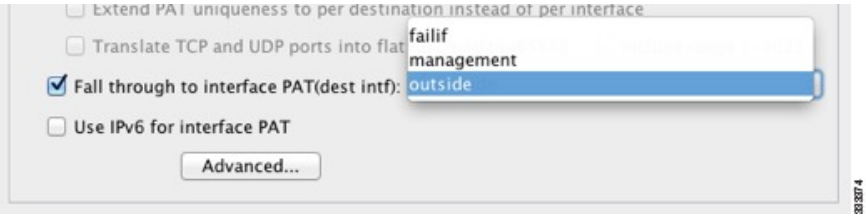
Note
The PAT pool object or group cannot contain a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.



c) (Optional) Select the following options as needed:

- **Round Robin**—To assign addresses and ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.
- **Extend PAT uniqueness to per destination instead of per interface** (8.4(3) and later, not including 8.5(1) or 8.6(1))—To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80.
- **Include Reserved Ports (1 to 1023)**—Includes the reserved ports, 1-1023, in the range of ports that are available for address translation. If you do not specify this option, addresses are translated to ports in the 1024-65535 range only.
- **Enable Block Allocation** (9.5.1 and later)—Enables port block allocation. For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time. If you allocate a block of ports, subsequent connections from the host use new randomly selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Port blocks are allocated in the 1024-65535 range only. Port block allocation is compatible with round robin, but you cannot use it with extended PAT. You also cannot use interface PAT fallback.

Step 7 (Optional, when mapped interface is a non-bridge group member only.) To use the interface IP address as a backup method when the other mapped addresses are already allocated, check the **Fall through to interface PAT** check box, and choose the interface from the drop-down list. To use the IPv6 address of the interface, also check the **Use IPv6 for interface PAT** check box.



Step 8 (Optional) Click **Advanced**, configure the following options in the Advanced NAT Settings dialog box, and click **OK**.

- (Required for bridge group member interfaces.) **Interface**—Specifies the real interface (**Source**) and the mapped interface (**Destination**) where this NAT rule applies. By default, the rule applies to all interfaces except for bridge group members.

Step 9 Click **OK**, and then **Apply**.

Configure Dynamic Twice PAT (Hide)

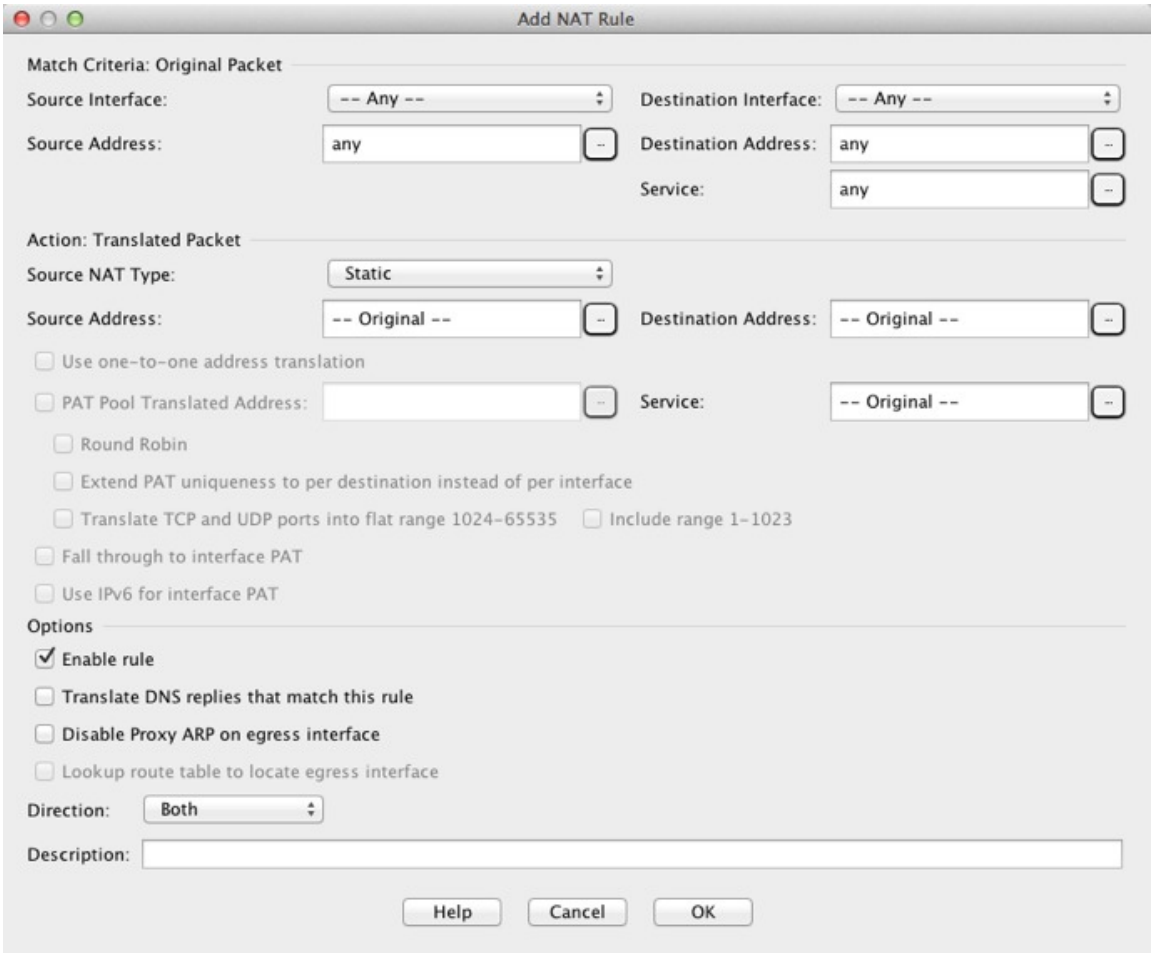
This section describes how to configure twice NAT for dynamic PAT (hide), which uses a single address for translation instead of a PAT pool.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**, and then do one of the following:

- Click **Add**, or **Add > Add NAT Rule Before Network Object NAT Rules**.
- Click **Add > Add NAT Rule After Network Object NAT Rules**.
- Select a twice NAT rule and click **Edit**.

The Add NAT Rule dialog box appears.

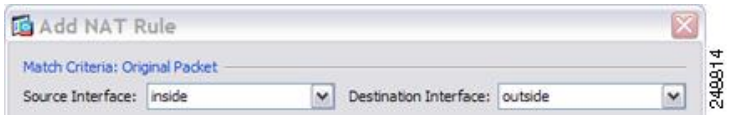


Step 2 (Required for bridge group member interfaces.) Set the source and destination interfaces.

In routed mode, the default is any interface for both source and destination. You can select specific interfaces for one or both options. However, you must select the interface when writing a rule for bridge group member interfaces; “any” does not include these interfaces.

a) From the **Match Criteria: Original Packet > Source Interface** drop-down list, choose the source interface.

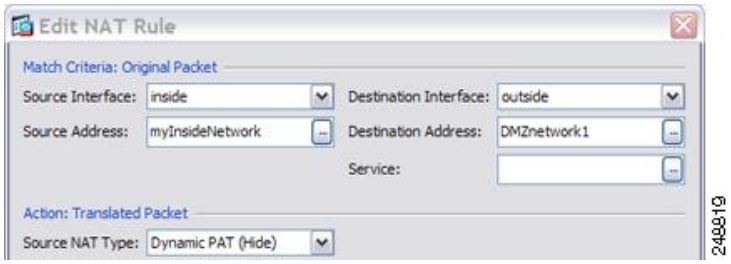
b) From the **Match Criteria: Original Packet > Destination Interface** drop-down list, choose the destination interface.



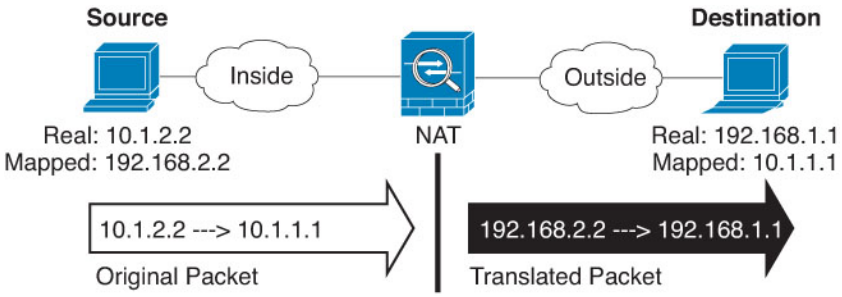
Step 3 Choose **Dynamic PAT (Hide)** from the **Action: Translated Packet** > **Source NAT Type** drop-down list.

This setting only applies to the source address; the destination translation is always static.

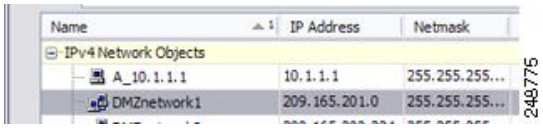
Note
To configure dynamic PAT using a PAT pool, choose **Dynamic** instead of Dynamic PAT (Hide), see Configure Dynamic Twice PAT Using a PAT Pool, on page 34.



Step 4 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.



a) For **Match Criteria: Original Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**.



b) (Optional) For **Match Criteria: Original Packet > Destination Address**, click the browse button and choose an existing network object, group, or interface (non-bridge group member interfaces only), or create a new object or group from the Browse Original Destination Address dialog box. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

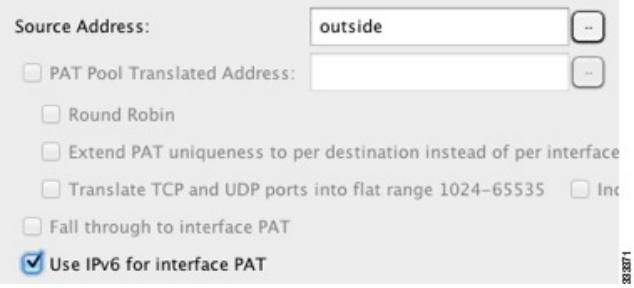
Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 4.

For static interface NAT with port translation only, choose an interface from the Browse dialog box. Be sure to also configure a service translation. For this option, you must configure a specific interface for the Source Interface. See Static NAT with port translation, on page 43 for more information.

Step 5 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired.

a) For **Action: Translated Packet > Source Address**, click the browse button and choose an existing network object that defines a host address, or an interface, or create a new object from the Browse Translated Source Address dialog box. The interface cannot be a bridge group member.

If you want to use the IPv6 address of the interface, check the **Use IPv6 for interface PAT** check box.



b) (Optional.) For **Action: Translated Packet > Destination Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated Destination Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. You can also use an FQDN network object for the destination mapped address.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

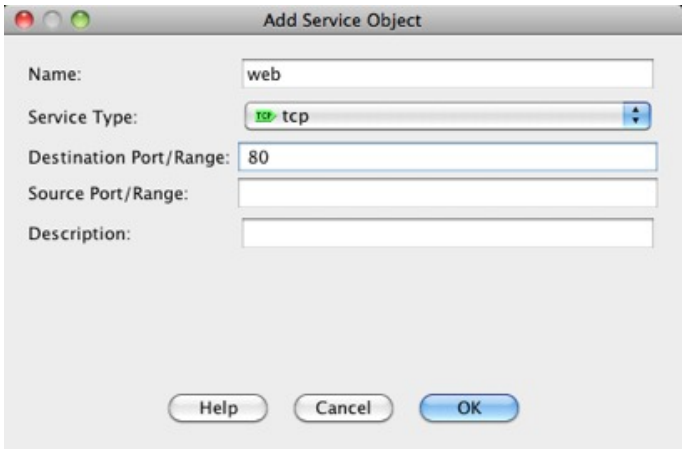
If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see Static NAT, on page 42. See Guidelines for NAT, on page 7 for information about disallowed mapped IP addresses.

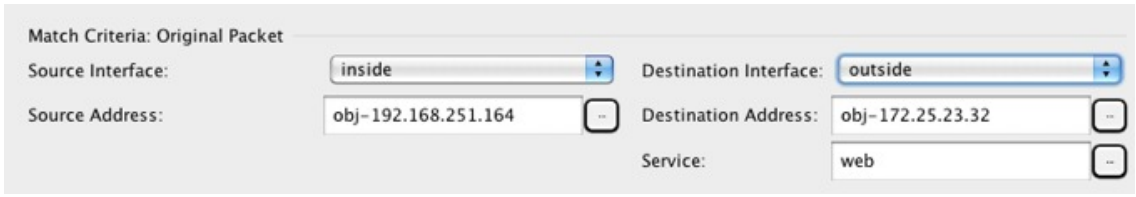
Step 6 (Optional.) Identify the destination service ports for service translation.

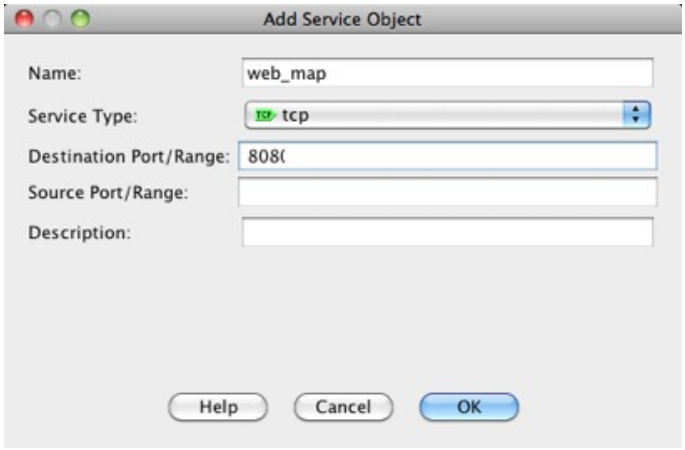
- Identify the original packet port (the *mapped destination port*). For **Match Criteria: Original Packet > Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Original Service dialog box.
- Identify the translated packet port (the *real destination port*). For **Action: Translated Packet > Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Translated Service dialog box.

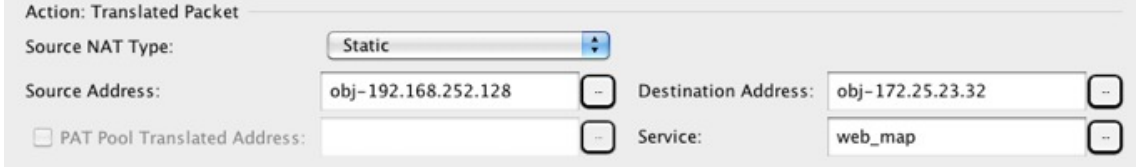
Dynamic NAT does not support port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

For example:

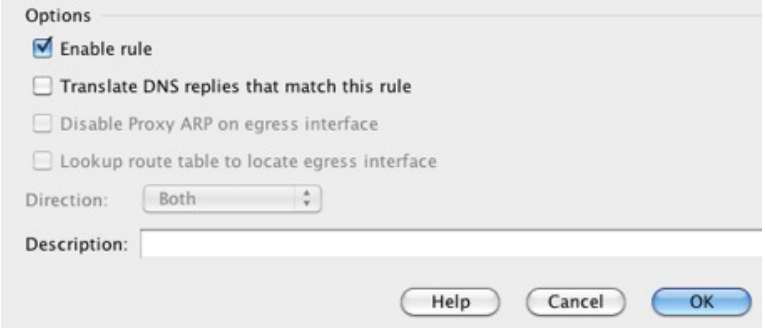








Step 7 (Optional) Configure NAT options in the Options area.



- **Enable rule**—Enables this NAT rule. The rule is enabled by default.
- **Description**—Adds a description about the rule up to 200 characters in length.

Step 8 Click **OK**, then click **Apply**.

Configure Dynamic Twice PAT Using a PAT Pool

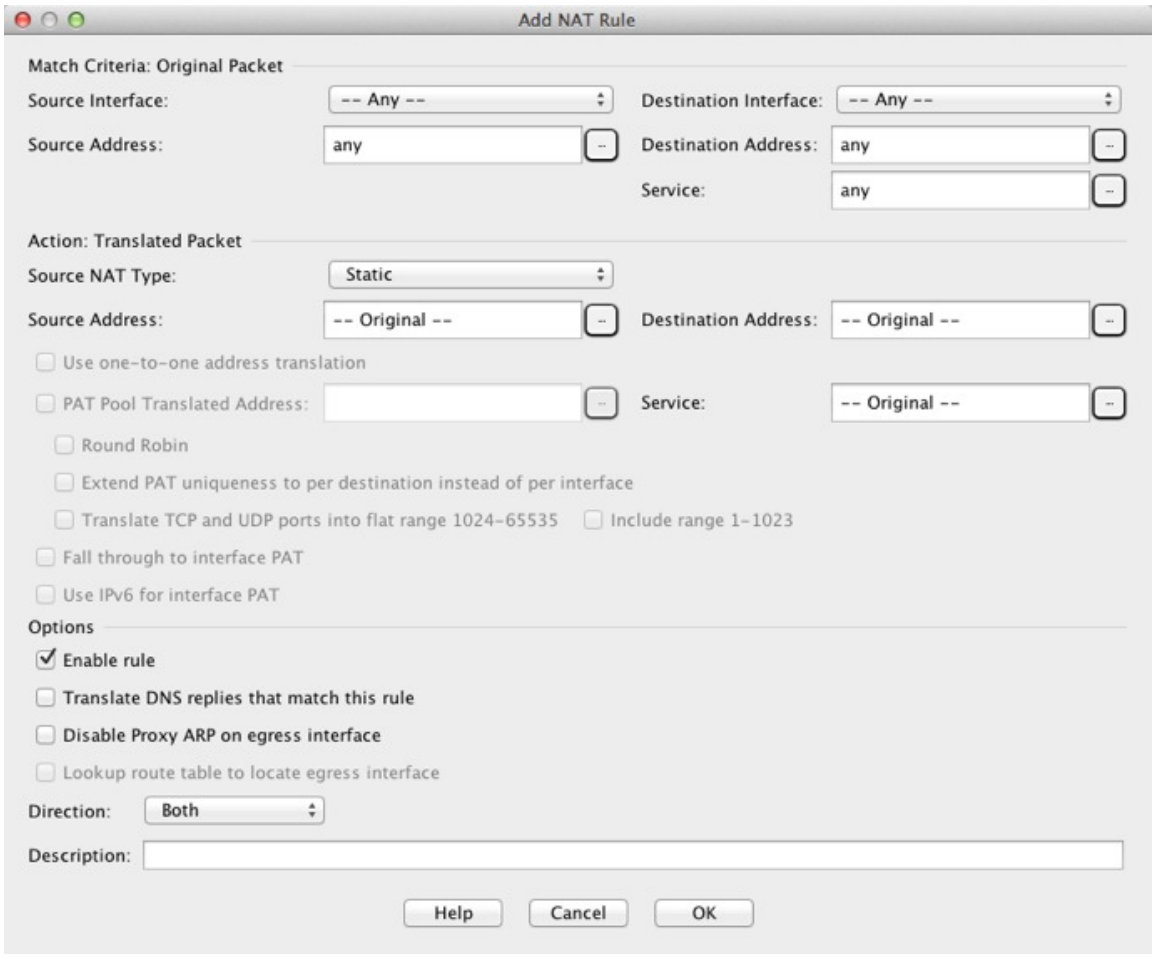
This section describes how to configure twice NAT for dynamic PAT using a PAT pool.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**, and then do one of the following:

- Click **Add**, or **Add** > **Add NAT Rule Before Network Object NAT Rules**.
- Click **Add** > **Add NAT Rule After Network Object NAT Rules**.
- Select a twice NAT rule and click **Edit**.

The Add NAT Rule dialog box appears.

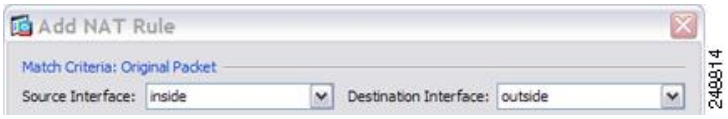


Step 2 (Required for bridge group member interfaces.) Set the source and destination interfaces.

In routed mode, the default is any interface for both source and destination. You can select specific interfaces for one or both options. However, you must select the interface when writing a rule for bridge group member interfaces; "any" does not include these interfaces.

a) From the **Match Criteria: Original Packet** > **Source Interface** drop-down list, choose the source interface.

b) From the **Match Criteria: Original Packet** > **Destination Interface** drop-down list, choose the destination interface.

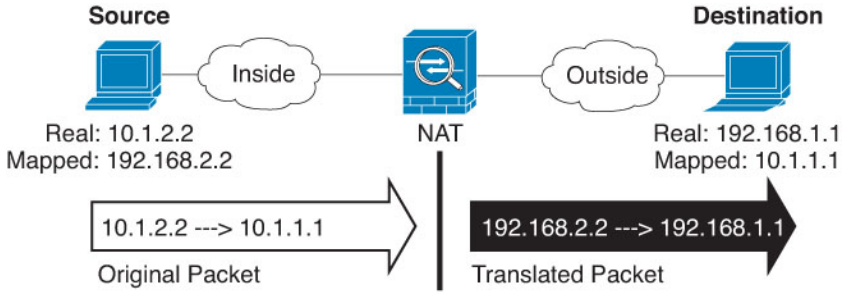


Step 3 Choose **Dynamic** from the **Action: Translated Packet** > **Source NAT Type** drop-down list.

This setting only applies to the source address; the destination translation is always static.



Step 4 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.



a) For the **Match Criteria: Original Packet** > **Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**.

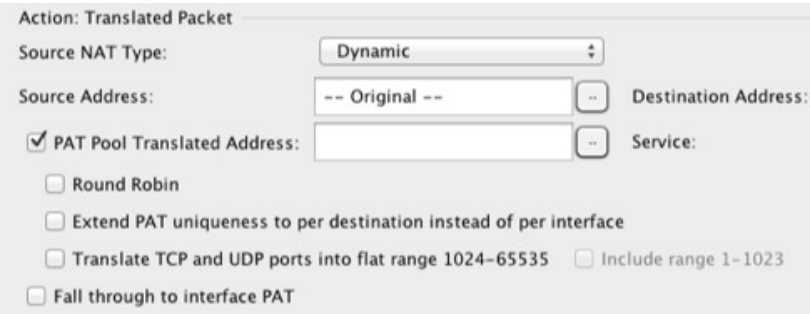
b) (Optional.) For the **Match Criteria: Original Packet** > **Destination Address**, click the browse button and choose an existing network object, group, or interface (non-bridge group member interfaces only), or create a new object or group from the Browse Original Destination Address dialog box. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 4.

For static interface NAT with port translation only, choose an interface from the Browse dialog box. Be sure to also configure a service translation. For this option, you must configure a specific interface for the Source Interface. See Static NAT with port translation, on page 43 for more information.

Step 5 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired.

a) Check the **PAT Pool Translated Address** check box, then click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated PAT Pool Address dialog box. **Note**: Leave the Source Address field empty.



Note

The object or group cannot contain a subnet.

b) (Optional.) For **Action: Translated Packet** > **Destination Address**, click the browse button and choose an existing network object or group, or create a new object or group from the Browse Translated Destination Address dialog box.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

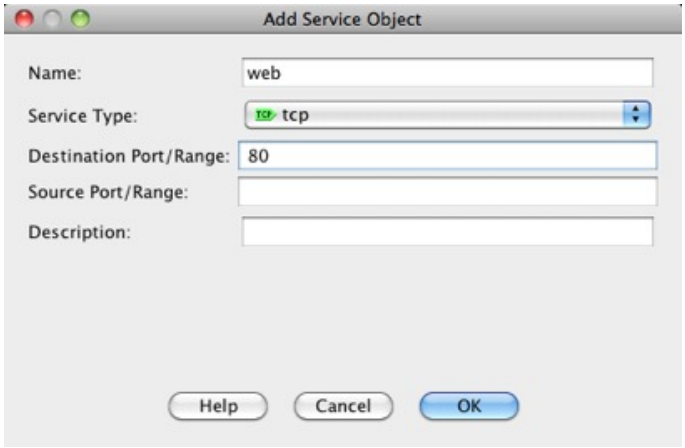
If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see Static NAT, on page 42. See Guidelines for NAT, on page 7 for information about disallowed mapped IP addresses.

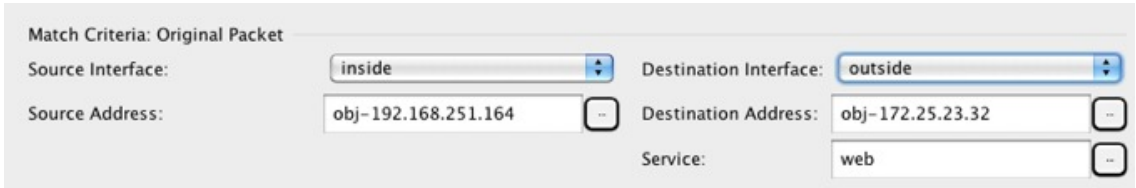
Step 6 (Optional.) Identify the destination service ports for service translation.

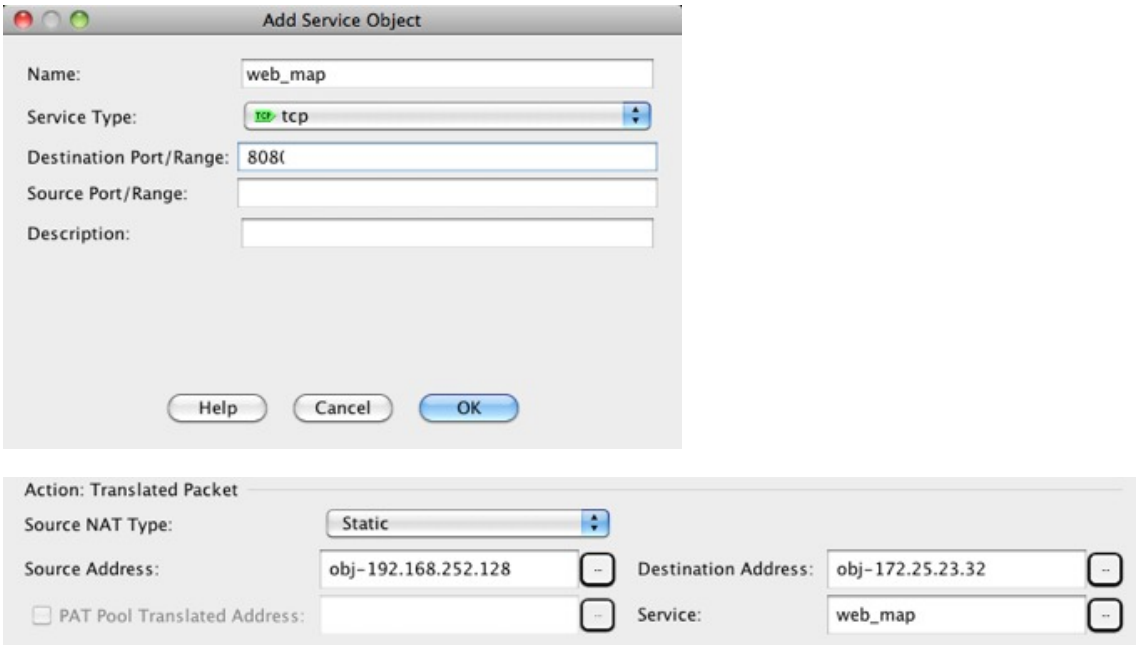
- Identify the original packet port (the *mapped destination port*). For **Match Criteria: Original Packet** > **Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Original Service dialog box.
- Identify the translated packet port (the *real destination port*). For **Action: Translated Packet** > **Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Translated Service dialog box.

Dynamic NAT does not support port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

For example:







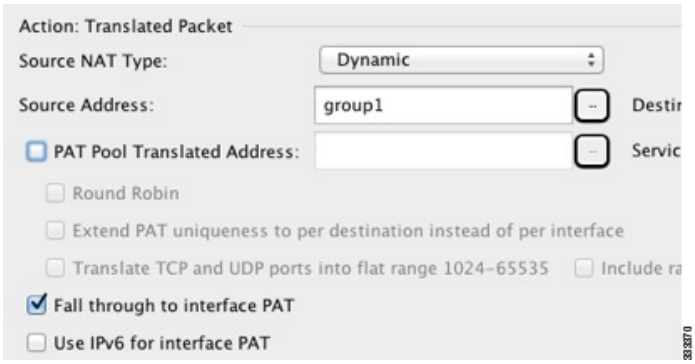
Step 7 (Optional.) For a PAT pool, configure the following options as needed:

- **Round Robin** —To assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.
- **Extend PAT uniqueness to per destination instead of per interface** (8.4(3) and later, not including 8.5(1) or 8.6(1).)—To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80.
- **Translate TCP or UDP ports into flat range (1024-65535)** (8.4(3) and later, not including 8.5(1) or 8.6(1).)—To use the 1024 to 65535 port range as a single flat range when allocating ports. When choosing the mapped port number for a translation, the ASA uses the real source port number if it is available. However, without this option, 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: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include range 1 to 1023** check box.
- **Include Reserved Ports (1 to 1023)**—Includes the reserved ports, 1-1023, in the range of ports that are available for address translation. If you do not specify this option, addresses are translated to ports in the 1024-65535 range only.
- **Enable Block Allocation** (9.5.1 and later)—Enables port block allocation. For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time. If you allocate a block of ports, subsequent connections from the host use new randomly selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Port blocks are allocated in the 1024-65535 range only. Port block allocation

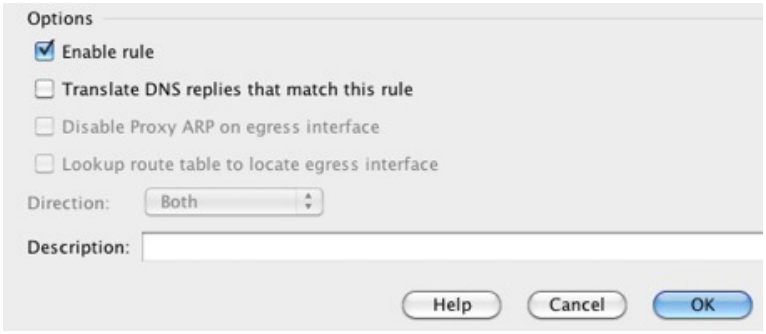
is compatible with round robin, but you cannot use it with extended PAT. You also cannot use interface PAT fallback.

Step 8 (Optional, when mapped interface is a non-bridge group member only.) To use the interface IP address as a backup method if the other mapped source addresses are already allocated, check the **Fall through to interface PAT** check box. To use the IPv6 interface address, also check the **Use IPv6 for interface PAT** check box.

The destination interface IP address is used. This option is only available if you configure a specific Destination Interface.



Step 9 (Optional.) Configure NAT options in the Options area.



- **Enable rule**—Enables this NAT rule. The rule is enabled by default.
- **Description**—Adds a description about the rule up to 200 characters in length.

Step 10 Click **OK**, then click **Apply**.

Configure PAT with Port Block Allocation

For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time (see RFC 6888). If you allocate a block of ports, subsequent connections from the host use new randomly-selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Blocks are freed when the last xlate that uses a port in the block is removed.

The main reason for allocating port blocks is reduced logging. The port block allocation is logged, connections are logged, but xlates created within the port block are not logged. On the other hand, this makes log analysis more difficult.

Port blocks are allocated in the 1024-65535 range only. There are separate blocks for TCP, UDP, and ICMP connections, and these blocks can overlap. Thus, if an application requires a low port number (1-1023), it might not work. For example, an application requesting port 22 (SSH) will get a mapped port within the range of 1024-65535 and within the block allocated to the host. You can create a separate NAT rule that does not use block allocation for applications that use low port numbers; for twice NAT, ensure the rule comes before the block allocation rule.

Before you begin

Usage notes for NAT rules:

- You can include the **Round Robin** option, but you cannot include the options for extending PAT uniqueness, or falling through to interface PAT. Other source/destination address and port information is also allowed.
- As with all NAT changes, if you replace an existing rule, you must clear xlates related to the replaced rule to have the new rule take effect. You can clear them explicitly or simply wait for them to time out. However, when operating in a cluster, you must reload the cluster so that the units can distribute blocks correctly.



Note If you are switching between a regular PAT and block allocation PAT rule, for object NAT, you must first delete the rule, then clear xlates. You can then create the new object NAT rule. Otherwise, you will see pat-port-block-state-mismatch drops in the **show asp drop** output.

- For a given PAT pool, you must specify (or not specify) block allocation for all rules that use the pool. You cannot allocate blocks in one rule and not in another. PAT pools that overlap also cannot mix block allocation settings. You also cannot overlap static NAT with port translation rules with the pool.

Procedure

Step 1 Select **Configuration** > **Firewall** > **Advanced** > **PAT Port Block Allocation** and configure the following settings:

- **Size of the block**—The number of ports in each block. The range is 32-4096. The default is 512.

  If you do not use the default, ensure that the size you choose divides evenly into 64,512 (the number of ports in the 1024-65535 range). Otherwise, there will be ports that cannot be used. For example, if you specify 100, there will be 12 unused ports.

- **Maximum block allocation per host**—The maximum number of blocks that can be allocated per host. The limit is per protocol, so a limit of 4 means at most 4 UDP blocks, 4 TCP blocks, and 4 ICMP blocks per host. The range is 1-8, the default is 4.
- **PBA Interim Logging**—If you enter a value, the system enables interim logging. By default, the system generates syslog messages during port block creation and deletion. If you enable interim logging, the

system generates the following message at the interval you specify. The messages report all active port blocks allocated at that time, including the protocol (ICMP, TCP, UDP) and source and destination interface and IP address, and the port block. You can specify an interval from 21600-604800 seconds (6 hours to 7 days).

%ASA-6-305017: Pba-interim-logging: Active *protocol* block of ports for translation from *real_interface*:*real_host_ip* to *mapped_interface*:*mapped_ip_address*/*start_port_num-end_port_num*

Step 2 Add NAT rules that use PAT pool block allocation.

a) Select **Configuration** > **Firewall** > **NAT Rules**.
b) Add or edit an object NAT or twice NAT rule.
c) Configure at least the following options:

- (Twice NAT.) Select the object that defines the source address in **Original Packet** > **Source Address**.
- **Type** = **Dynamic**.
- **Pat Pool Translated Address**. Select a network object that defines the pat pool network.
- **Enable Block Allocation**.

d) Click **OK**.

Configure Per-Session PAT or Multi-Session PAT (Version 9.0(1) and Higher)

By default, all TCP PAT traffic and all UDP DNS traffic uses per-session PAT. To use multi-session PAT for traffic, you can configure per-session PAT rules: a permit rule uses per-session PAT, and a deny rule uses multi-session PAT.

Per-session PAT 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 control 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, per-session PAT can dramatically increase the connection rate supported by one address. Without per-session PAT, the maximum connection rate for one address for an IP protocol is approximately 2000 per second. With per-session PAT, the connection rate for one address for an IP protocol is 65535/*average-lifetime*.

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. However, if you also want to use per-session PAT for the UDP ports used by these protocols, you must create the permit rules for them.

Before you begin

By default, the following rules are installed:

- Permit TCP from any (IPv4 and IPv6) to any (IPv4 and IPv6).
- Permit UDP from any (IPv4 and IPv6) to the domain port.

These rules do not show up in the table.

You cannot remove these rules, and they always exist after any manually-created rules. Because rules are evaluated in order, you can override the default rules. For example, to completely negate these rules, you could add the following:

- Deny TCP from any (IPv4 and IPv6) to any (IPv4 and IPv6).
- Deny UDP from any (IPv4 and IPv6) to the domain port.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Advanced** > **Per-Session NAT Rules**.

Step 2 Do one of the following:

- Choose **Add > Add Per-Session NAT Rule**.
- Select a rule and click **Edit**.

Step 3 Configure the rule:

- **Action**—Click **Permit** or **Deny**. A permit rule uses per-session PAT; a deny rule uses multi-session PAT.
- Source—Specify the Source Address either by typing an address or clicking the **...** button to choose an object. For the service, select UDP or TCP. You can optionally specify a source port, although normally you only specify the destination port. Either type in UDP/*port* or TCP/*port*, or click the **...** button to select a common value or object.
- Destination—Specify the Destination Address either by typing an address or clicking the **...** button to choose an object. For the service, select UDP or TCP; this must match the source service. You can optionally specify a destination port. Either type in UDP/*port* or TCP/*port*, or click the **...** button to select a common value or object. You can use the operators != (not equal to), > (greater than), < (less than), or specify a range using a hyphen, for example, 100-200.

Step 4 Click **OK**, then click **Apply**.

Static NAT

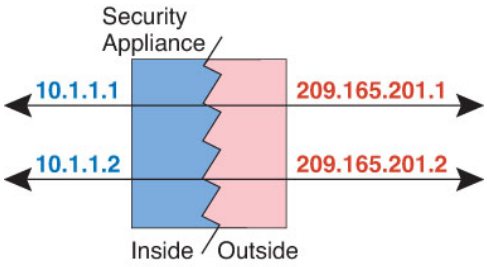
The following topics explain static NAT and how to implement it.

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.

The following figure shows a typical static NAT scenario. The translation is always active so both real and remote hosts can initiate connections.

Figure 5: Static NAT



Note You can disable bidirectionality if desired.

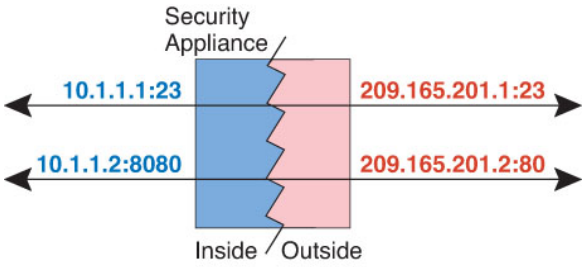
Static NAT with port translation

Static NAT with port translation lets you specify a real and mapped protocol and port.

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.

The following figure 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 6: Typical Static NAT with Port Translation Scenario



Static NAT-with-port-translation rules limit access to the destination IP address for the specified port only. If you try to access the destination IP address on a different port not covered by a NAT rule, then the connection is blocked. In addition, for twice NAT, traffic that does not match the source IP address of the NAT rule will be dropped if it matches the destination IP address, regardless of the destination port. Therefore, you must add additional rules for all other traffic allowed to the destination IP address. For example, you can configure a static NAT rule for the IP address, without port specification, and place it after the port translation rule.

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

Following are some other uses of static NAT with port translation.

Static NAT with Identity Port Translation

You can simplify external access to internal resources. For example, if you have three separate servers that provide services on different ports (such as FTP, HTTP, and SMTP), you can give external users a single IP address to access those services. You can then configure static NAT with identity port translation to map the single external IP address to the correct IP addresses of the real servers based on the port they are trying to access. You do not need to change the port, because the servers are using the standard ones (21, 80, and 25 respectively). For details on how to configure this example, see Single Address for FTP, HTTP, and SMTP (Static NAT-with-Port-Translation).

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

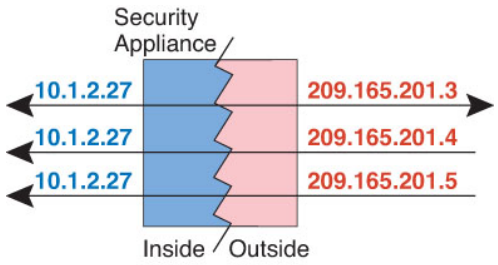
You can configure static NAT to map a real address to an interface address/port combination. For example, if you want to redirect Telnet access for the device's outside interface to an inside host, then you can map the inside host IP address/port 23 to the outside interface address/port 23.

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.

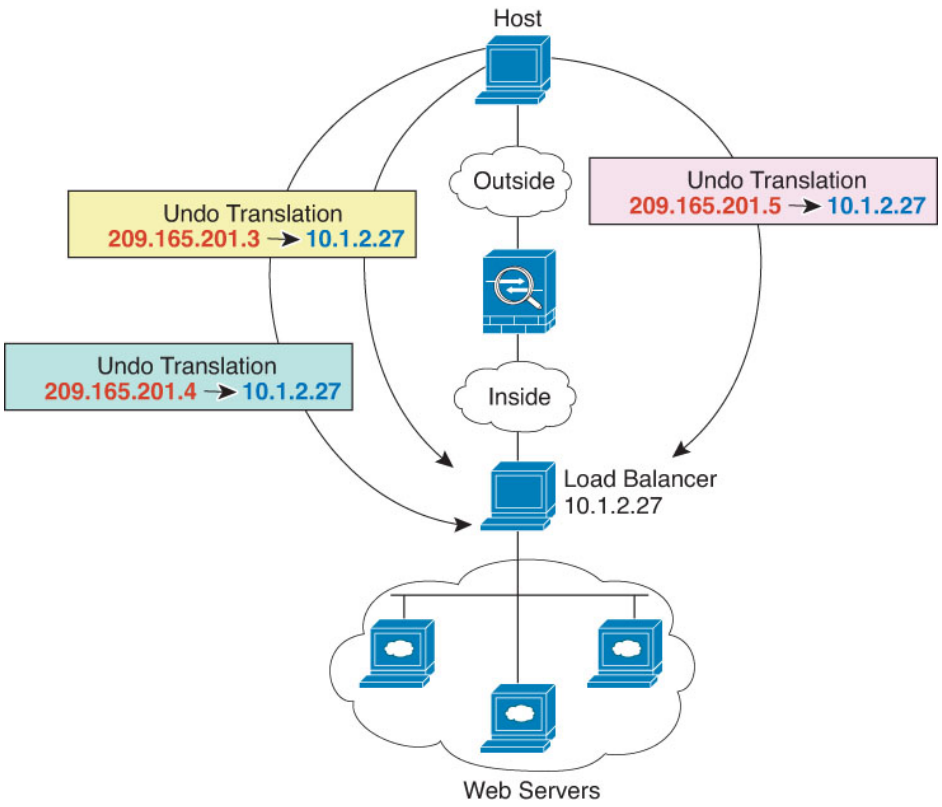
The following figure shows a typical one-to-many static NAT scenario. Because initiation by the real host always uses the first mapped address, the translation of real host IP/first mapped IP is technically the only bidirectional translation.

Figure 7: 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. For details on how to configure this example, see Inside Load Balancer with Multiple Mapped Addresses (Static NAT, One-to-Many).

Figure 8: One-to-Many Static NAT Example



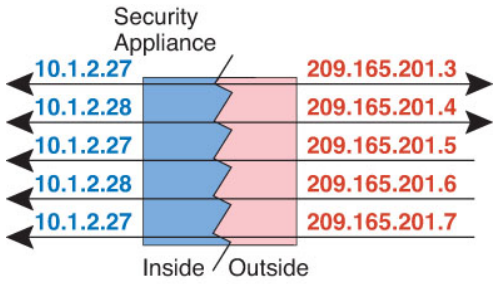
Other mapping scenarios (not recommended)

NAT 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.

The following figure shows a typical few-to-many static NAT scenario.

Figure 9: Few-to-Many Static NAT

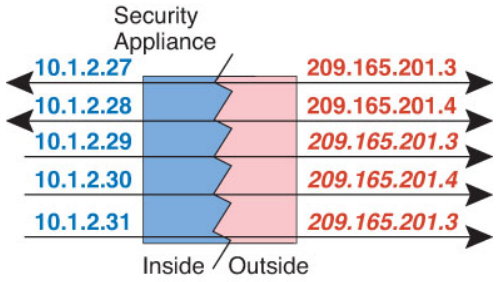


For a many-to-few or many-to-one configuration, where you have more real addresses than mapped addresses, you run out of mapped addresses before you run out of real addresses. Only the mappings between the lowest real IP addresses and the mapped pool result in bidirectional initiation. The remaining 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).

The following figure shows a typical many-to-few static NAT scenario.

Figure 10: 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.

Configure Static Network Object NAT or Static NAT-with-Port-Translation

This section describes how to configure a static NAT rule using network object NAT.

Procedure

Step 1 Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration** > **Firewall** > **NAT Rules**, then click **Add > Add Network Object NAT Rule**.
- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then edit a network object.

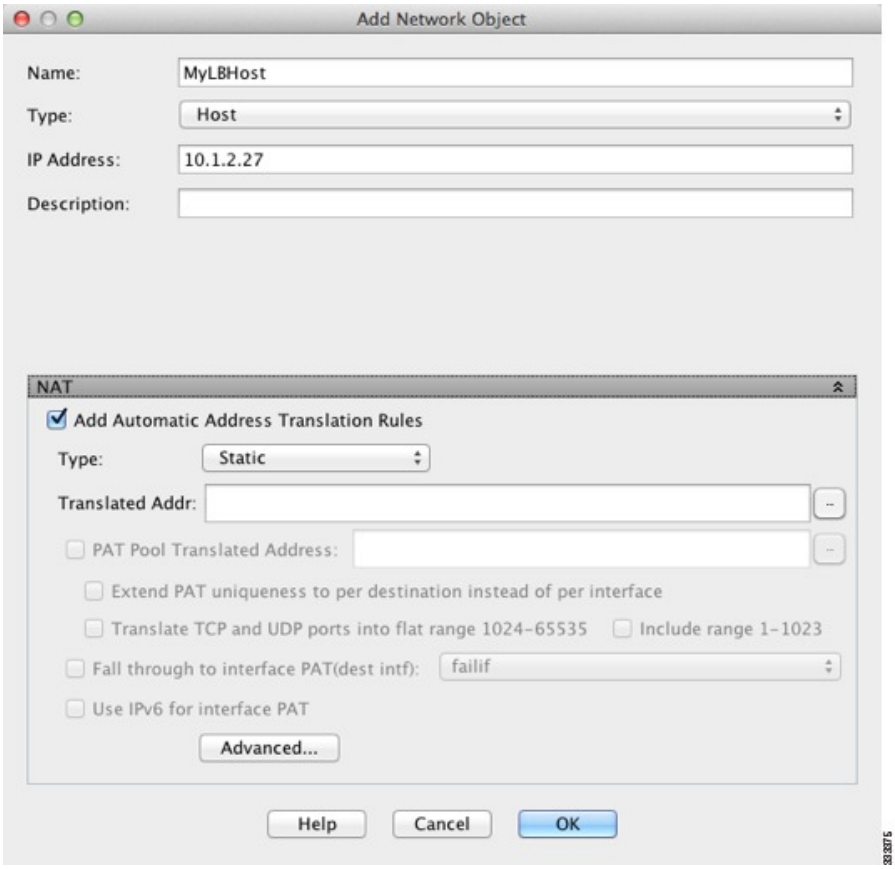
Step 2 For a new object, enter values for the following fields:

- **Name**—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.
- **Type**—Host, Network, or Range.
- **IP Addresses**—IPv4 or IPv6 addresses, a single address for a host, a starting and ending address for a range, and for subnet, either an IPv4 network address and mask (for example, 10.100.10.0 255.255.255.0) or IPv6 address and prefix length (for example, 2001:DB8:0:CD30::/60).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

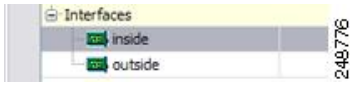
Step 4 Check the **Add Automatic Translation Rules** check box.

Step 5 From the Type drop-down list, choose **Static**.



Step 6 In the Translated Addr. field, specify the mapped IP address as one of the following. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses. For more information, see Static NAT, on page 42.

- Type a host IP address. This provides a one-to-one mapping for host objects. For subnet objects, the same netmask is used for the inline host address, and you get one-to-one translations for addresses in the mapped inline host's subnet. For range objects, the mapped address includes the same number of hosts that are in the range object, starting with the mapped host address. For example, if the real address is defined as a range from 10.1.1.1 through 10.1.1.6, and you specify 172.20.1.1 as the mapped address, then the mapped range will include 172.20.1.1 through 172.20.1.6. For NAT46 or NAT66 translations, this can be an IPv6 network address.
- Click the browse button and select a network object (or create a new one). To do a one-to-one mapping for a range of IP addresses, select an object that contains a range with the same number of addresses.
- (For static NAT-with-port-translation only.) Type an interface name or click the browse button, and choose an interface from the Browse Translated Addr dialog box. You cannot select a bridge group member interface.



To use the IPv6 interface address, you must also check the **Use IPv6 for interface PAT** check box. Be sure to also click **Advanced** and configure a service port translation. (You cannot specify an interface in transparent mode.)

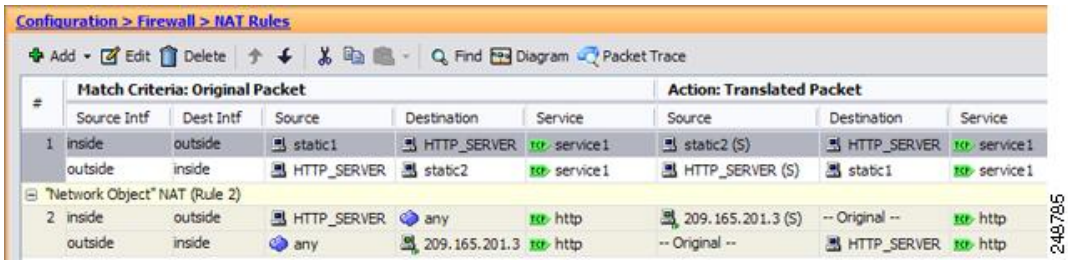
Step 7 (Optional.) For NAT46, check **Use one-to-one address translation**. For NAT 46, specify one-to-one to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this keyword.

Step 8 (Optional) Click **Advanced**, configure the following options in the Advanced NAT Settings dialog box, and click **OK**.

- **Translate DNS replies for rule**—Translates the IP address in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). See Rewriting DNS queries and responses using NAT for more information.
- **Disable Proxy ARP on egress interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. For information on the conditions which might require the disabling of proxy ARP, see Mapped Addresses and Routing.
- (Required for bridge group member interfaces.) **Interface**—Specifies the real interface (**Source**) and the mapped interface (**Destination**) where this NAT rule applies. By default, the rule applies to all interfaces except for bridge group members.
- **Service**—Configures static NAT-with-port-translation. Choose the protocol, then enter the real port and the mapped port. You can use port numbers or a well-known port name such as http.

Step 9 Click **OK**, and then **Apply**.

Because static rules are bidirectional (allowing initiation to and from the real host), the NAT Rules table shows two rows for each static rule, one for each direction.



Configure Static Twice NAT or Static NAT-with-Port-Translation

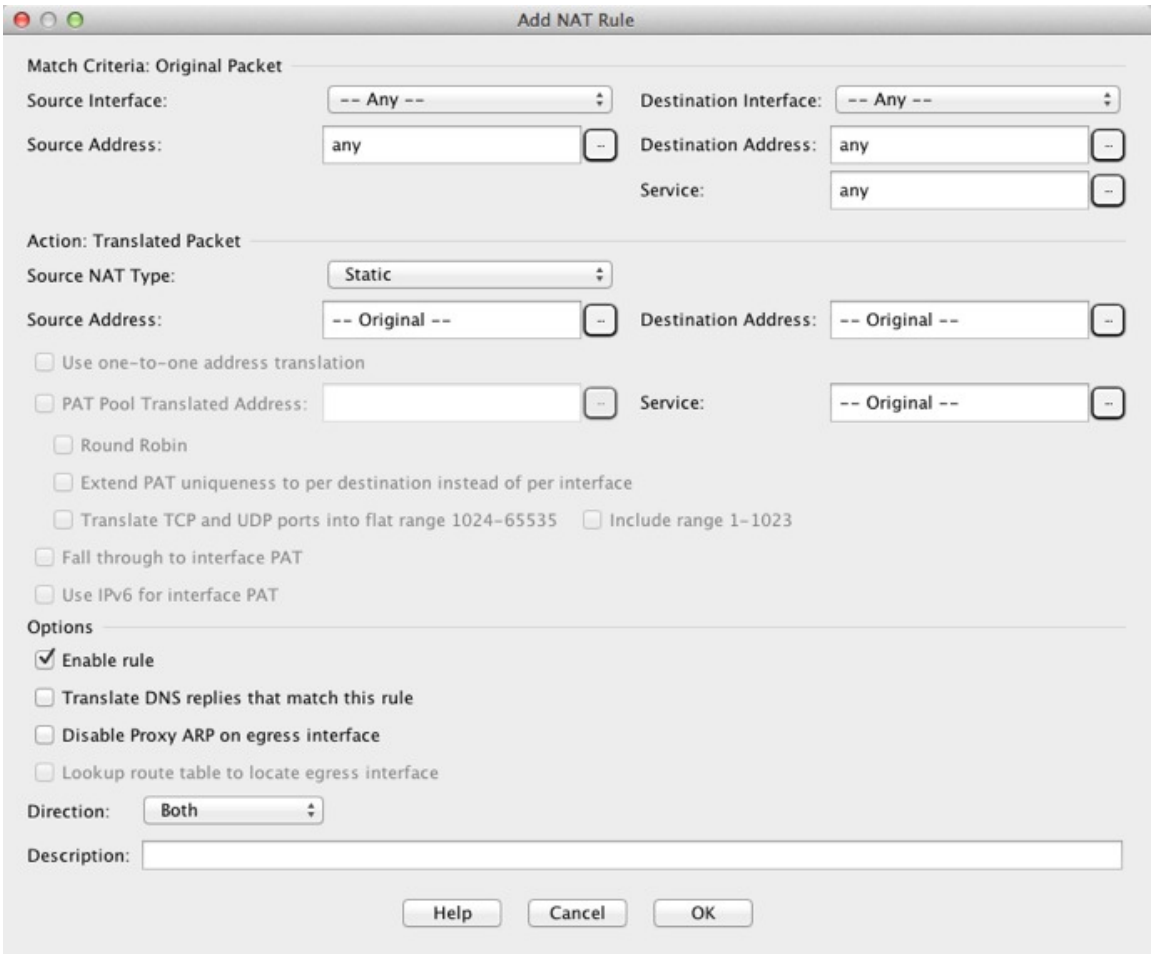
This section describes how to configure a static NAT rule using twice NAT.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**, and then do one of the following:

- Click **Add**, or **Add > Add NAT Rule Before Network Object NAT Rules**.
- Click **Add > Add NAT Rule After Network Object NAT Rules**.
- Select a twice NAT rule and click **Edit**.

The Add NAT Rule dialog box appears.

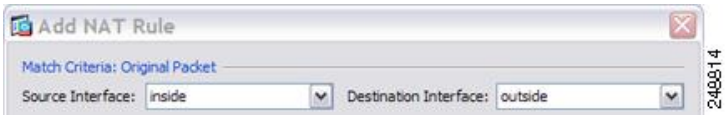


Step 2 (Required for bridge group member interfaces.) Set the source and destination interfaces.

In routed mode, the default is any interface for both source and destination. You can select specific interfaces for one or both options. However, you must select the interface when writing a rule for bridge group member interfaces; "any" does not include these interfaces.

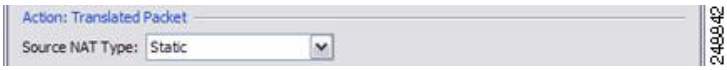
a) From the **Match Criteria: Original Packet > Source Interface** drop-down list, choose the source interface.

b) From the **Match Criteria: Original Packet > Destination Interface** drop-down list, choose the destination interface.

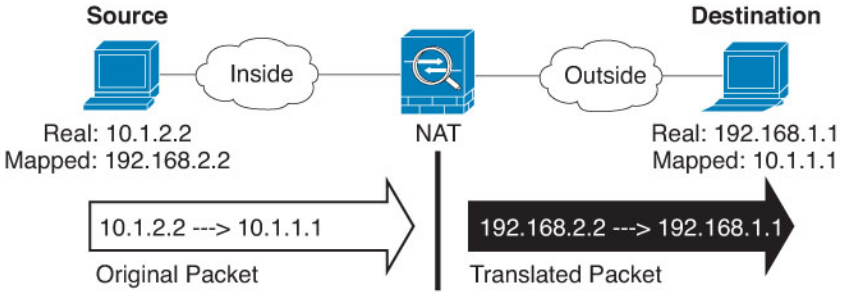


Step 3 Choose **Static** from the **Action: Translated Packet** > **Source NAT Type** drop-down list. Static is the default setting.

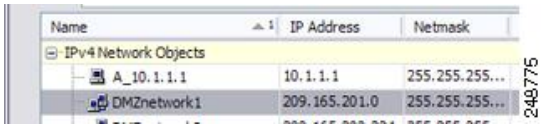
This setting only applies to the source address; the destination translation is always static.



Step 4 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.



a) For **Match Criteria: Original Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**, but do not use this option except for identity NAT.



b) (Optional) For **Match Criteria: Original Packet > Destination Address**, click the browse button and choose an existing network object, group, or interface, or create a new object or group from the Browse Original Destination Address dialog box.

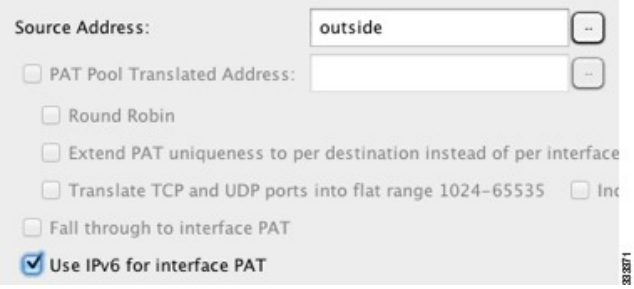
Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 4.

Step 5 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired.

a) For **Action: Translated Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated Source Address dialog box.

For static NAT, the mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired.

For static interface NAT with port translation, you can specify the interface instead of a network object/group for the mapped address. If you want to use the IPv6 address of the interface, check the **Use IPv6 for interface PAT** check box. You cannot select a bridge group member interface.



For more information, see Static NAT with port translation, on page 43. See Guidelines for NAT, on page 7 for information about disallowed mapped IP addresses.

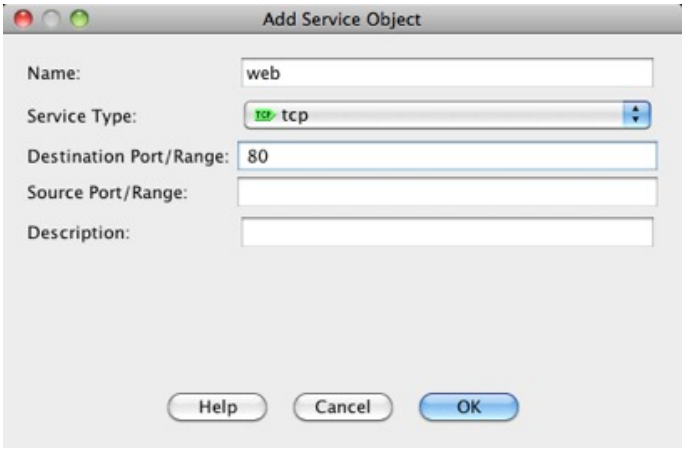
b) (Optional.) For **Action: Translated Packet > Destination Address**, click the browse button and choose an existing network object or group, or create a new object or group from the Browse Translated Destination Address dialog box. You can also use an FQDN network object for the destination mapped address.

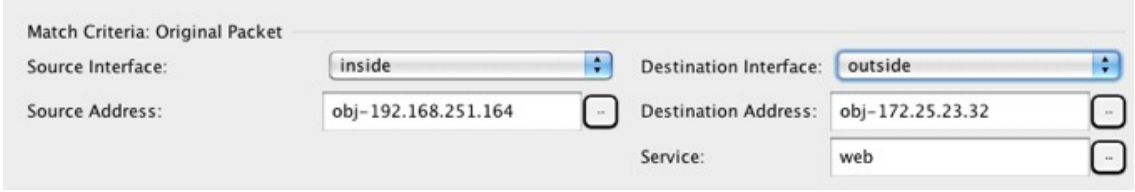
Step 6 (Optional.) Identify the source or destination service ports for service translation.

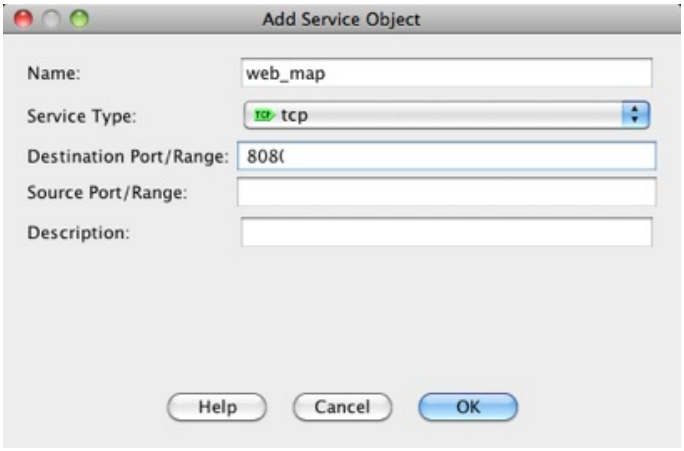
- Identify the original packet source or destination port (the *real source port* **or** the *mapped destination port*). For **Match Criteria: Original Packet > Service**, click the browse button and choose an existing service object that specifies ports, or create a new object from the Browse Original Service dialog box.
- Identify the translated packet source or destination port (the *mapped source port* **or** the *real destination port*). For **Action: Translated Packet > Service**, click the browse button and choose an existing service object that specifies ports, or create a new object from the Browse Translated Service dialog box.

A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both the real and mapped service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real destination port. When translating a port, be sure the protocols in the real and mapped service objects are identical (for example, both TCP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

For example:



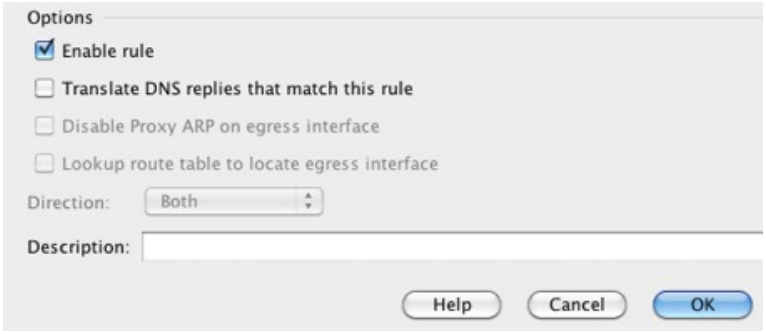






Step 7 (Optional.) For NAT46, check the **Use one-to-one address translation** check box. For NAT46, specify one-to-one to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this keyword.

Step 8 (Optional.) Configure NAT options in the Options area.



- **Enable rule** —Enables this NAT rule. The rule is enabled by default.
- (For a source-only rule.) **Translate DNS replies that match this rule**—Rewrites the DNS A record in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). You cannot configure DNS modification if you configure a destination address. See Rewriting DNS queries and responses using NAT for more information.

- **Disable Proxy ARP on egress interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. See Mapped Addresses and Routing for more information.
- **Direction**—To make the rule unidirectional, choose **Unidirectional**. The default is Both. Making the rule unidirectional prevents destination addresses from initiating connections to the real addresses.
- **Description**—Adds a description about the rule up to 200 characters in length.

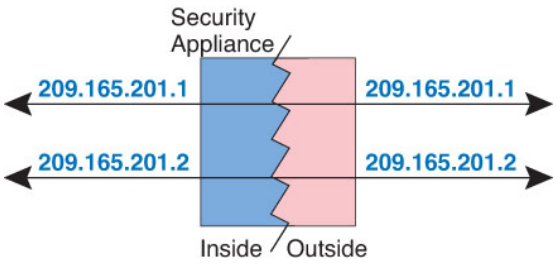
Step 9 Click **OK**, then click **Apply**.

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 that network to itself. Identity NAT is necessary for remote access VPN, where you need to exempt the client traffic from NAT.

The following figure shows a typical identity NAT scenario.

Figure 11: Identity NAT



The following topics explain how to configure identity NAT.

Configure Identity Network Object NAT

This section describes how to configure an identity NAT rule using network object NAT.

Procedure

Step 1 Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration** > **Firewall** > **NAT Rules**, then click **Add** > **Add Network Object NAT Rule**.
- To add NAT to an existing network object, choose **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups**, and then edit a network object.

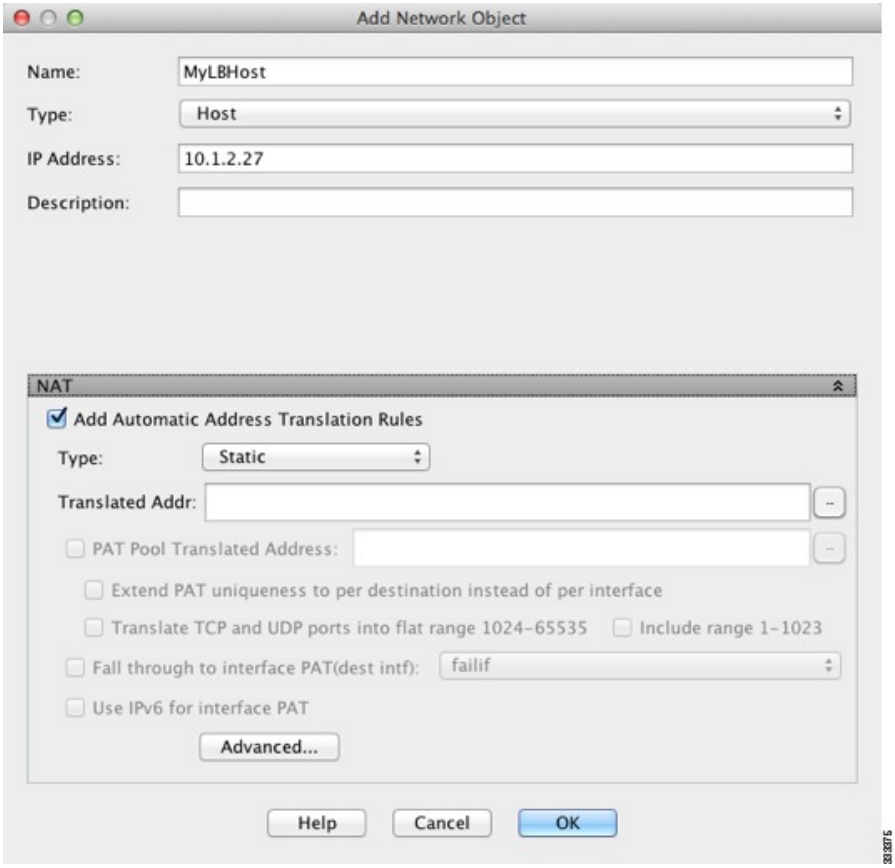
Step 2 For a new object, enter values for the following fields:

- **Name**—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.
- **Type**—Host, Network, or Range.
- **IP Addresses**—IPv4 or IPv6 addresses, a single address for a host, a starting and ending address for a range, and for subnet, either an IPv4 network address and mask (for example, 10.100.10.0 255.255.255.0) or IPv6 address and prefix length (for example, 2001:DB8:0:CD30::/60).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

Step 4 Check the **Add Automatic Translation Rules** check box.

Step 5 From the Type drop-down list, choose **Static**.



Step 6 In the Translated Addr. field, do one of the following:

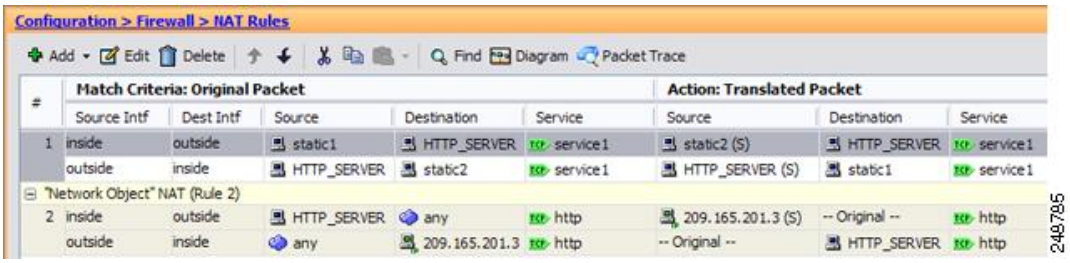
- For host objects, enter the same address. For range objects, enter the first address in the real range (the same number of addresses in the range will be used). For subnet objects, enter any address within the real subnet (all addresses in the subnet will be used).
- Click the browse button and select a network object (or create a new one). Use this option when configuring identity NAT for a range of addresses.

Step 7 (Optional) Click **Advanced**, configure the following options in the Advanced NAT Settings dialog box, and click **OK**.

- **Translate DNS replies for rule**—Do not configure this option for identity NAT.
- **Disable Proxy ARP on egress interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. For information on the conditions which might require the disabling of proxy ARP, see Mapped Addresses and Routing.
- (Routed mode; interfaces specified.) **Lookup route table to locate egress interface**—Determines the egress interface using a route lookup instead of using the interface specified in the NAT command. See Determining the Egress Interface for more information.
- (Required for bridge group member interfaces.) **Interface**—Specifies the real interface (**Source**) and the mapped interface (**Destination**) where this NAT rule applies. By default, the rule applies to all interfaces except for bridge group members.
- **Service**—Do not configure this option for identity NAT.

Step 8 Click **OK**, and then **Apply**.

Because static rules are bidirectional (allowing initiation to and from the real host), the NAT Rules table shows two rows for each static rule, one for each direction, unless you select the route lookup option.



Configure Identity Twice NAT

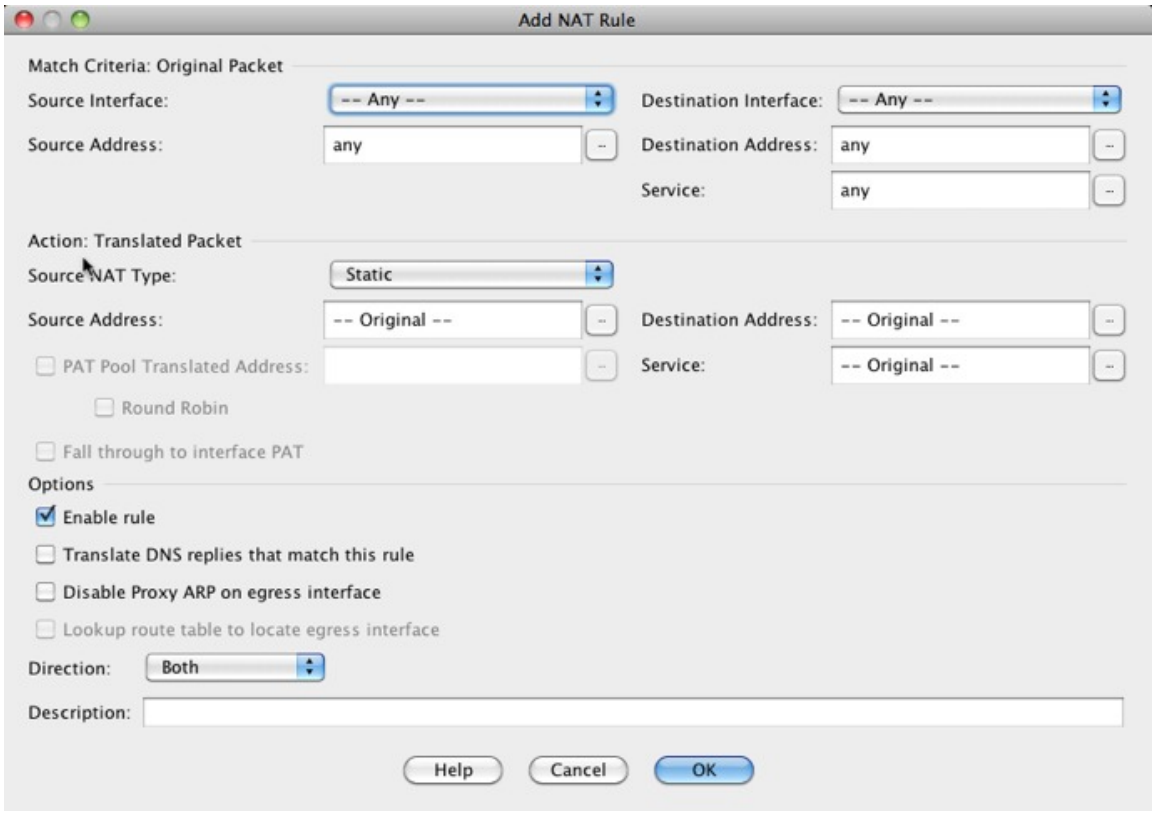
This section describes how to configure an identity NAT rule using twice NAT.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**, and then do one of the following:

- Click **Add**, or **Add > Add NAT Rule Before Network Object NAT Rules**.
- Click **Add > Add NAT Rule After Network Object NAT Rules**.
- Select a twice NAT rule and click **Edit**.

The Add NAT Rule dialog box appears.

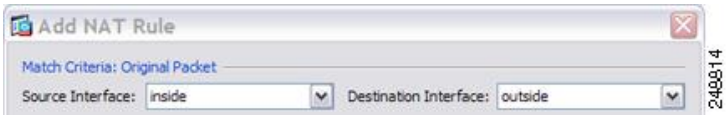


Step 2 (Required for bridge group member interfaces.) Set the source and destination interfaces.

In routed mode, the default is any interface for both source and destination. You can select specific interfaces for one or both options. However, you must select the interface when writing a rule for bridge group member interfaces; "any" does not include these interfaces.

a) From the **Match Criteria: Original Packet > Source Interface** drop-down list, choose the source interface.

b) From the **Match Criteria: Original Packet > Destination Interface** drop-down list, choose the destination interface.



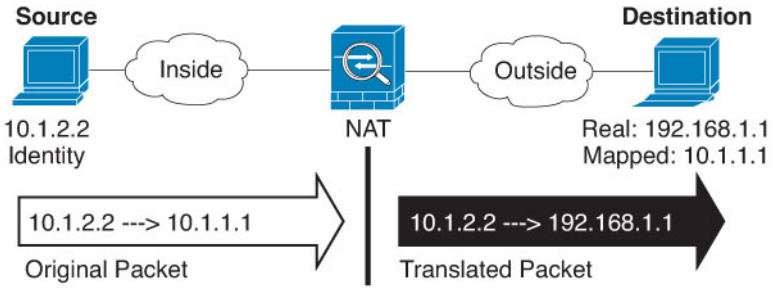
Step 3 Choose **Static** from the **Action: Translated Packet** > **Source NAT Type** drop-down list. Static is the default setting.

This setting only applies to the source address; the destination translation is always static.

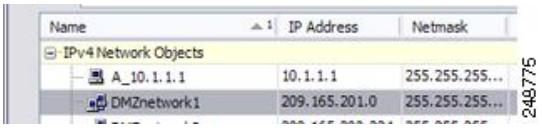


Step 4 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following

figure for an example of the original packet vs. the translated packet where you perform identity NAT on the inside host but translate the outside host.



a) For **Match Criteria: Original Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**; only use this option when also setting the mapped address to **any**.



b) (Optional.) For **Match Criteria: Original Packet > Destination Address**, click the browse button and choose an existing network object, group, or interface, or create a new object or group from the Browse Original Destination Address dialog box.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 4.

For static interface NAT with port translation only, choose an interface. If you specify an interface, be sure to also configure a service translation. For more information, see Static NAT with port translation, on page 43.

Step 5 Identify the translated packet addresses; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*).

a) For **Action: Translated Packet > Source Address**, click the browse button and choose the same network object or group from the Browse Translated Source Address dialog box that you chose for the real source address. Use **any** if you specified **any** for the real address.

b) For **Match Criteria: Translated Packet > Destination Address**, click the browse button and choose an existing network object or group, or create a new object or group from the Browse Translated Destination Address dialog box. You can also use an FQDN network object for the destination mapped address.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

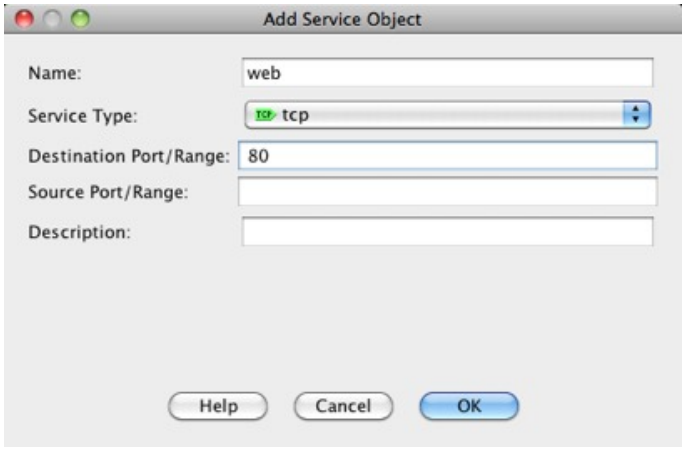
If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see Static NAT, on page 42. See Guidelines for NAT, on page 7 for information about disallowed mapped IP addresses.

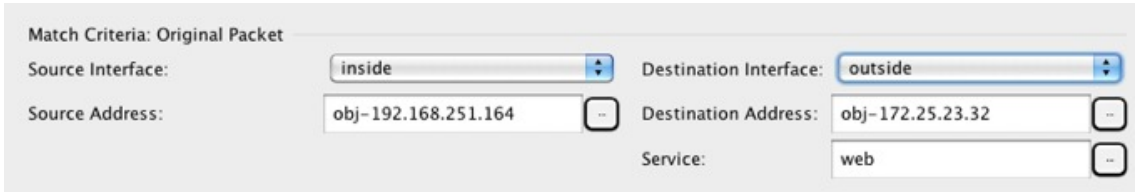
Step 6 (Optional.) Identify the source or destination service ports for service translation.

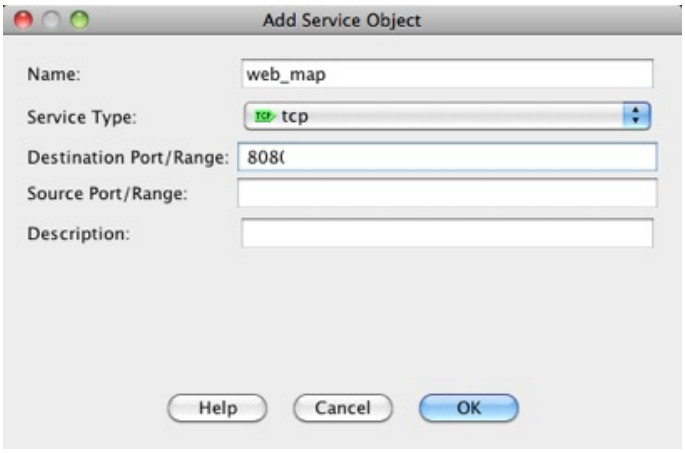
- Identify the original packet source or destination port (the *real source port* **or** the *mapped destination port*). For **Match Criteria: Original Packet > Service**, click the browse button and choose an existing service object that specifies ports, or create a new object from the Browse Original Service dialog box.
- Identify the translated packet source or destination port (the *mapped source port* **or** the *real destination port*). For **Action: Translated Packet > Service**, click the browse button and choose an existing service object that specifies ports, or create a new object from the Browse Translated Service dialog box.

A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both the real and mapped service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real destination port. When translating a port, be sure the protocols in the real and mapped service objects are identical (for example, both TCP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

For example:

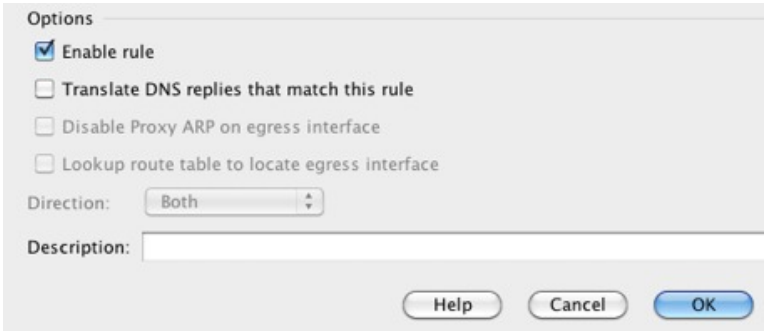








Step 7 (Optional) Configure NAT options in the Options area.



- **Enable rule** —Enables this NAT rule. The rule is enabled by default.
- (For a source-only rule.) **Translate DNS replies that match this rule**—Although this option is available if you do not configure a destination address, it is not applicable to identity NAT because you are translating the address to itself, so the DNS reply does not need modification.
- **Disable Proxy ARP on egress interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. See Mapped Addresses and Routing for more information.
- (Routed mode; interfaces specified.) **Lookup route table to locate egress interface**—Determines the egress interface using a route lookup instead of using the interface specified in the NAT command. See Determining the Egress Interface for more information.
- **Direction**—To make the rule unidirectional, choose **Unidirectional**. The default is Both. Making the rule unidirectional prevents traffic from initiating connections to the real addresses. You might want to use this setting for testing purposes.
- **Description**—Adds a description about the rule up to 200 characters in length.

Step 8 Click **OK**, then click **Apply**.

Monitoring NAT

You can view NAT related graphs from the following pages:

- **Monitoring > Properties > Connection Graphs > Xlates**—Select the Xlate Utilization graph to view the in-use and most-used xlates. This is equivalent to the **show xlate** command.
- **Monitoring > Properties > Connection Graphs > Perfmon**—Select the Xlate Perfmon graph to see NAT performance information. This is equivalent to the xlate information from the **show perfmon** command.

History for NAT

Feature Name	Platform Releases	Description
Network Object NAT	8.3(1)	Configures NAT for a network object IP address(es). We introduced or modified the following screens: Configuration > Firewall > NAT Rules Configuration > Firewall > Objects > Network Objects/Groups
Twice NAT	8.3(1)	Twice NAT lets you identify both the source and destination address in a single rule. We modified the following screen: Configuration > Firewall > NAT Rules.

Feature Name	Platform Releases	Description
Identity NAT configurable proxy ARP and route lookup	8.4(2)/8.5(1)	In earlier releases for identity NAT, proxy ARP was disabled, and a route lookup was always used to determine the egress interface. You could not configure these settings. In 8.4(2) and later, the default behavior for identity NAT was changed to match the behavior of other static NAT configurations: proxy ARP is enabled, and the NAT configuration determines the egress interface (if specified) by default. You can leave these settings as is, or you can enable or disable them discretely. Note that you can now also disable proxy ARP for regular static NAT. For pre-8.3 configurations, the migration of NAT exempt rules (the **nat 0 access-list** command) to 8.4(2) and later now includes the following keywords to disable proxy ARP and to use a route lookup: **no-proxy-arp** and **route-lookup**. The **unidirectional** keyword that was used for migrating to 8.3(2) and 8.4(1) is no longer used for migration. When upgrading to 8.4(2) from 8.3(1), 8.3(2), and 8.4(1), all identity NAT configurations will now include the **no-proxy-arp** and **route-lookup** keywords, to maintain existing functionality. The **unidirectional** keyword is removed. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit Network Object > Advanced NAT Settings; Configuration > Firewall > NAT Rules > Add/Edit NAT Rule.
PAT pool and round robin address assignment	8.4(2)/8.5(1)	You can now specify a pool of PAT addresses instead of a single address. You can also optionally enable round-robin assignment of PAT addresses instead of first using all ports on a PAT address before using the next address in the pool. These features help prevent a large number of connections from a single PAT address from appearing to be part of a DoS attack and makes configuration of large numbers of PAT addresses easy. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit Network Object; Configuration > Firewall > NAT Rules > Add/Edit NAT Rule.
Round robin PAT pool allocation uses the same IP address for existing hosts	8.4(3)	When using a PAT pool with round robin allocation, if a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. We did not modify any screens. *This feature is not available in 8.5(1) or 8.6(1).*

Feature Name	Platform Releases	Description
Flat range of PAT ports for a PAT pool	8.4(3)	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. If you have a lot of traffic that uses the lower port ranges, when using a PAT pool, you can now specify a flat range of ports to be used instead of the three unequal-sized tiers: either 1024 to 65535, or 1 to 65535. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit Network Object; Configuration > Firewall > NAT Rules > Add/Edit NAT Rule. *This feature is not available in 8.5(1) or 8.6(1).*
Extended PAT for a PAT pool	8.4(3)	Each PAT IP address allows up to 65535 ports. If 65535 ports do not provide enough translations, you can now enable extended PAT for a PAT pool. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit Network Object; Configuration > Firewall > NAT Rules > Add/Edit NAT Rule. *This feature is not available in 8.5(1) or 8.6(1).*

Feature Name	Platform Releases	Description
Automatic NAT rules to translate a VPN peer's local IP address back to the peer's real IP address	8.4(3)	In rare situations, you might want to use a VPN peer's real IP address on the inside network instead of an assigned local IP address. Normally with VPN, the peer is given an assigned local IP address to access the inside network. However, you might want to translate the local IP address back to the peer's real public IP address if, for example, your inside servers and network security is based on the peer's real IP address. You can enable this feature on one interface per tunnel group. Object NAT rules are dynamically added and deleted when the VPN session is established or disconnected. You can view the rules using the **show nat** command. Because of routing issues, we do not recommend using this feature unless you know you need it; contact Cisco TAC to confirm feature compatibility with your network. See the following limitations: • Only supports Cisco IPsec and Secure Client. • Return traffic to the public IP addresses must be routed back to the ASA so the NAT policy and VPN policy can be applied. • Does not support load-balancing (because of routing issues). • Does not support roaming (public IP changing). ASDM does not support this command; enter the command using the Command Line Tool.
NAT support for IPv6	9.0(1)	NAT now supports IPv6 traffic, as well as translating between IPv4 and IPv6. Translating between IPv4 and IPv6 is not supported in transparent mode. We modified the following screen: Configuration > Firewall > Objects > Network Objects/Group; Configuration > Firewall > NAT Rules.
NAT support for reverse DNS lookups	9.0(1)	NAT now supports translation of the DNS PTR record for reverse DNS lookups when using IPv4 NAT, IPv6 NAT, and NAT64 with DNS inspection enabled for the NAT rule.

Feature Name	Platform Releases	Description
Per-session PAT	9.0(1)	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 control 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 requires multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT by creating a per-session deny rule. We introduced the following screen: Configuration > Firewall > Advanced > Per-Session NAT Rules.
Transactional Commit Model on NAT Rule Engine	9.3(1)	When enabled, a NAT rule update is applied after the rule compilation is completed; without affecting the rule matching performance. We added NAT to the following screen: Configuration > Device Management > Advanced > Rule Engine.
Carrier Grade NAT enhancements	9.5(1)	For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time (see RFC 6888). We added the following command: **Configuration** > **Firewall** > **Advanced** > **PAT Port Block Allocation**. We added **Enable Block Allocation** the object NAT and twice NAT dialog boxes.
NAT support for SCTP	9.5(2)	You can now specify SCTP ports in static network object NAT rules. Using SCTP in static twice NAT is not recommended. Dynamic NAT/PAT does not support SCTP. We modified the following screen: **Configuration** > **Firewall** > **NAT** add/edit static network object NAT rule, **Advanced NAT Settings** dialog box.

Feature Name	Platform Releases	Description
Interim logging for NAT port block allocation.	9.12(1)	When you enable port block allocation for NAT, the system generates syslog messages during port block creation and deletion. If you enable interim logging, the system generates message 305017 at the interval you specify. The messages report all active port blocks allocated at that time, including the protocol (ICMP, TCP, UDP) and source and destination interface and IP address, and the port block. We modified the following screen: **Configuration** > **Firewall** > **Advanced** > **PAT Port Block Allocation**.
Changes to PAT address allocation in clustering. The PAT pool **flat** option is now enabled by default and it is not configurable.	9.15(1)	The way PAT addresses are distributed to the members of a cluster is changed. Previously, addresses were distributed to the members of the cluster, so your PAT pool would need a minimum of one address per cluster member. Now, the control unit instead divides each PAT pool address into equal-sized port blocks and distributes them across cluster members. Each member has port blocks for the same PAT addresses. Thus, you can reduce the size of the PAT pool, even to as few as one IP address, depending on the amount of connections you typically need to PAT. Port blocks are allocated in 512-port blocks from the 1024-65535 range. You can optionally included the reserved ports, 1-1023, in this block allocation when you configure PAT pool rules. For example, in a 4-node cluster, each node gets 32 blocks with which it will be able to handle 16384 connections per PAT pool IP address compared to a single node handling all 65535 connections per PAT pool IP address. As part of this change, PAT pools for all systems, whether standalone or operating in a cluster, now use a flat port range of 1023 - 65535. Previously, you could optionally use a flat range by including the **flat** keyword in a PAT pool rule. The **flat** keyword is no longer supported: the PAT pool is now always flat. The **include-reserve** keyword, which was previously a sub-keyword to **flat**, is now an independent keyword within the PAT pool configuration. With this option, you can include the 1 - 1023 port range within the PAT pool. Note that if you configure port block allocation (the **block-allocation** PAT pool option), your block allocation size is used rather than the default 512-port block. In addition, you cannot configure extended PAT for a PAT pool for systems in a cluster. New/Modified screens: NAT PAT Pool configuration.

Feature Name	Platform Releases	Description
New Section 0 for system-defined NAT rules.	9.16(1)	A new Section 0 has been added to the NAT rule table. This section is exclusively for the use of the system. Any NAT rules that the system needs for normal functioning are added to this section, and these rules take priority over any rules you create. Previously, system-defined rules were added to Section 1, and user-defined rules could interfere with proper system functioning. You cannot add, edit, or delete Section 0 rules, but you will see them in **show nat detail** command output.
Twice NAT support for fully-qualified domain name (FQDN) objects as the translated (mapped) destination.	9.17(1)	You can use an FQDN network object, such as one specifying www.example.com, as the translated (mapped) destination address in twice NAT rules. The system configures the rule based on the IP address returned from the DNS server.