

Network Address Translation

The following topics explain Network Address Translation (NAT) and how to configure it on Firewall Threat Defense devices.

- NAT functions, on page 1
- NAT basics, on page 2
- Requirements and prerequisites for NAT policies, on page 11
- Guidelines for NAT, on page 11
- Manage NAT policies, on page 17
- Configure NAT for threat defense, on page 19
- Translating IPv6 networks, on page 59
- Monitoring NAT, on page 72
- Examples for NAT, on page 73
- History for Firewall Threat Defense NAT, on page 120

NAT functions

NAT enables private IP networks to connect to the Internet by replacing private IP addresses with public IP addresses. It conserves public addresses by advertising at minimum only one public address for the entire network to the outside world.

Private IP address ranges

Most devices within an IP network are assigned private IP addresses, which are not routable outside the company network due to the shortage of public IPv4 addresses. RFC 1918 defines the private IP address ranges that should not be advertised externally:

- 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 primary functions of NAT is to enable private IP networks to access the Internet by translating private IP addresses into legal, routable public addresses. This approach conserves public IP addresses since NAT can be configured to advertise only a single public address for the entire network.

Other functions of NAT include:

- Security: Hides internal IP addresses, discouraging direct attacks.
- IP routing solutions: Handles overlapping IP addresses without conflict.
- Flexibility: Allows internal IP addressing schemes to change without affecting externally visible public addresses.
- Translating between IPv4 and IPv6 (Routed mode only): Supports translating between IPv4 and IPv6, enabling interoperability between different network types.



Note NAT is optional. If NAT is not configured for specific traffic, that traffic will not be translated but will still be subject to normal security policies.

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 24.
- 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 30.
- Static NAT—A consistent mapping between a real and mapped IP address. Allows bidirectional traffic initiation. See Static NAT, on page 40.
- 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 49.

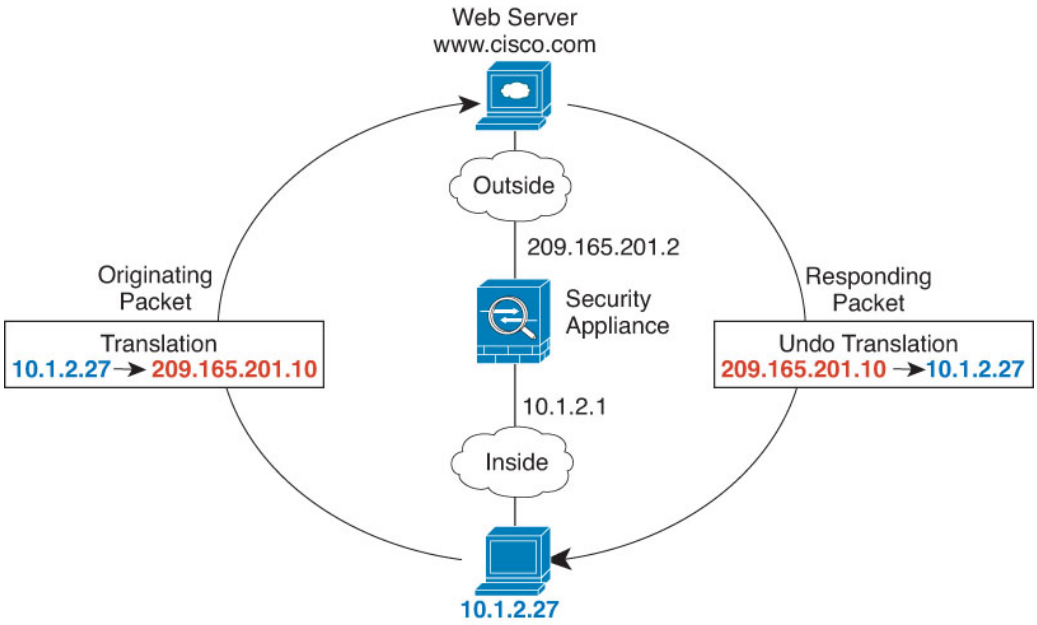
NAT in routed and transparent mode

You can configure NAT in both routed and transparent firewall mode. You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes. The following sections describe typical usage for each firewall mode.

NAT in routed mode

The following figure shows a typical NAT example in routed mode, with a private network on the inside.

Figure 1: NAT Example: Routed Mode



1. When the inside host at 10.1.2.27 sends a packet to a web server, the real source address of the packet, 10.1.2.27, is translated to a mapped address, 209.165.201.10.
2. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the Firewall Threat Defense receives the packet because the Firewall Threat Defense performs proxy ARP to claim the packet.

3. The Firewall Threat Defense then changes the translation of the mapped address, 209.165.201.10, back to the real address, 10.1.2.27, before sending it to the host.

NAT in transparent mode or within a bridge group

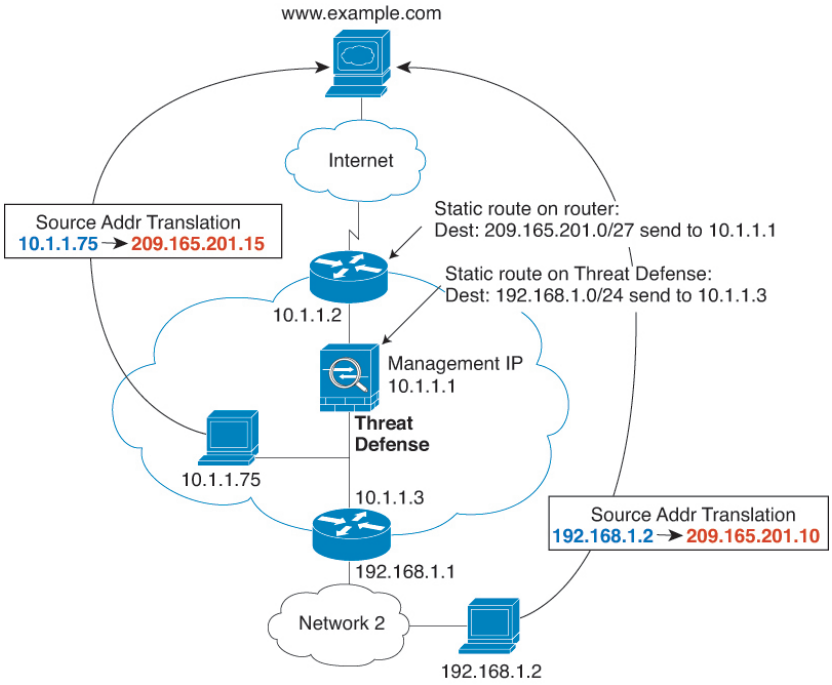
Using NAT in transparent mode eliminates the need for the upstream or downstream routers to perform NAT for their networks. It can perform a similar function within a bridge group in routed mode.

NAT in transparent mode, or in routed mode between members of the same bridge group, has the following requirements and limitations:

- 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.
- ARP inspection is not supported. Moreover, if for some reason a host on one side of the Firewall Threat Defense sends an ARP request to a host on the other side of the Firewall Threat Defense, and the initiating host real address is mapped to a different address on the same subnet, then the real address remains visible in the ARP request.
- Translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

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

Figure 2: NAT Example: Transparent Mode



1. When the inside host at 10.1.1.75 sends a packet to a web server, the real source address of the packet, 10.1.1.75, is changed to a mapped address, 209.165.201.15.

2. When the server responds, it sends the response to the mapped address, 209.165.201.15, and the Firewall Threat Defense receives the packet because the upstream router includes this mapped network in a static route directed to the Firewall Threat Defense management IP address.
3. The Firewall Threat Defense then undoes the translation of the mapped address, 209.165.201.15, back to the real address, 10.1.1.1.75. Because the real address is directly-connected, the Firewall Threat Defense sends it directly to the host.
4. For host 192.168.1.2, the same process occurs, except for returning traffic, the Firewall Threat Defense looks up the route in its routing table and sends the packet to the downstream router at 10.1.1.3 based on the Firewall Threat Defense static route for 192.168.1.0/24.

Auto NAT and Manual NAT

You can implement address translation in two ways: *auto NAT* and *manual NAT*.

We recommend using auto NAT unless you need the extra features that manual NAT provides. It is easier to configure auto 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.)

Auto NAT

All NAT rules that are configured as a parameter of a network object are considered to be *auto 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.

Although these rules are configured as part of the object itself, you cannot see the NAT configuration in the object definition through the object manager.

When a packet enters an interface, both the source and destination IP addresses are checked against the auto 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 manual NAT for that kind of functionality, where you can identify the source and destination address in a single rule.

Manual NAT

Manual 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 Auto NAT and Manual NAT

The main differences between these two NAT types are:

- How you define the real address.
  - Auto NAT—The NAT rule becomes a parameter for a network object. The network object IP address serves as the original (real) address.
  - Manual 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 manual NAT is more scalable.
- How source and destination NAT is implemented.
  - Auto 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.
  - Manual 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 manual 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.
  - Auto NAT—Automatically ordered in the NAT table.
  - Manual NAT—Manually ordered in the NAT table (before or after auto NAT rules).

NAT rule order

Auto NAT and manual 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	Manual 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, manual 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	Auto 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.
Section 3	Manual 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:

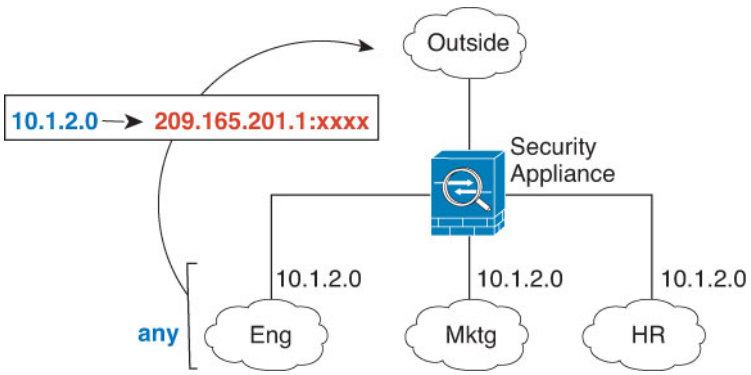
- 192.168.1.1/32 (static)
- 10.1.1.0/24 (static)
- 192.168.1.0/24 (static)
- 172.16.1.0/24 (dynamic) (object abc)
- 172.16.1.0/24 (dynamic) (object def)
- 192.168.1.0/24 (dynamic)

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 3: 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.

Note You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes. When specifying interfaces, you do so indirectly by selecting the interface object that contains the interface.

NAT exemption

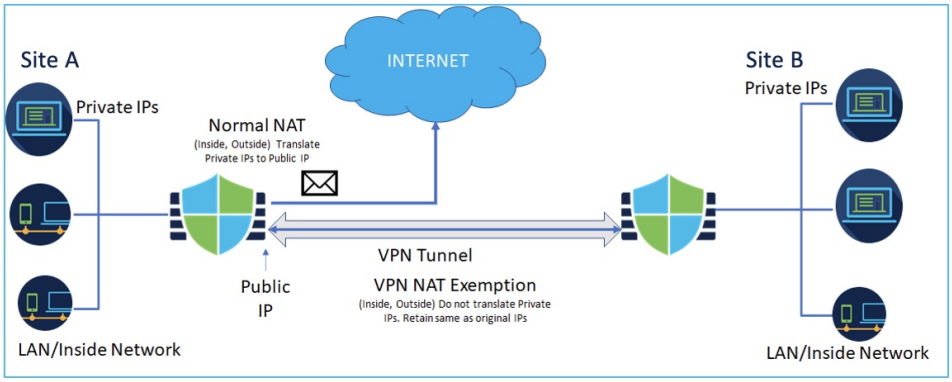
When internet edge devices have a site-to-site VPN configured on an interface and also have NAT rules for that interface, you must exempt the VPN traffic from the NAT rules. If you do not exempt the VPN traffic from NAT translation, the traffic gets dropped or is not routed through the VPN tunnel to the remote peer.

NAT exemption allows you to exclude traffic from being translated by NAT rules. When you create a policy-based site-to-site VPN using the Firewall Management Center VPN wizard, you can select the **NAT Exempt** option to create the rules automatically (**Devices** > **VPN** > **Site to Site**). By default, this option is enabled. You can view the NAT exemptions for a device in the NAT policy page (**Devices** > **NAT**, and then click **NAT Exemptions**).

The Firewall Management Center supports NAT exemption for all policy-based site-to-site VPN topology types. For more information, see Configure a policy-based site-to-site VPN.

Consider the following example, which shows a site-to-site VPN tunnel connecting Site A and Site B. For traffic that must go to the Internet, NAT translates the private IPs to a public IP address to access the Internet. For traffic that must go over the VPN tunnel, you must configure NAT exemption for the device in the VPN wizard.

Figure 4: Site-to-site VPN Topology with NAT Exemption



Note that when you upgrade from versions earlier than 7.4, if you have enabled NAT exemption, Firewall Management Center will disable this option. You must enable this option in the policy-based site-to-site VPN wizard.

Configuring routing for NAT

The Firewall Threat Defense device needs to be the destination for any packets sent to the translated (mapped) address.

When sending packets, the device uses the destination interface if you specify one, or a routing table lookup if you do not, to determine the egress interface. For identity NAT, you have the option to use a route lookup even if you specify a destination interface.

The type of routing configuration needed depends on the type of mapped address, as explained in the following topics.

Addresses on the same network as the mapped interface

If you use addresses on the same network as the destination (mapped) interface, the Firewall Threat Defense uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the Firewall Threat Defense does not have to be the gateway for any additional networks. This solution is ideal if the outside network contains an adequate number of free addresses, a consideration if you are using a 1:1 translation like dynamic NAT or static NAT. Dynamic PAT greatly extends the number of translations you can use with a small number of addresses, so even if the available addresses on the outside network is small, this method can be used. For PAT, you can even use the IP address of the mapped interface.



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

Addresses on a unique network

If you need more addresses than are available on the destination (mapped) interface network, you can identify addresses on a different subnet. The upstream router needs a static route for the mapped addresses that points to the Firewall Threat Defense.

Alternatively for routed mode, you can configure a static route on the Firewall Threat Defense for the mapped addresses using any IP address on the destination network as the gateway, and then redistribute the route using your routing protocol. For example, if you use NAT for the inside network (10.1.1.0/24) and use the mapped IP address 209.165.201.5, then you can configure a static route for 209.165.201.5 255.255.255.255 (host address) to the 10.1.1.99 gateway that can be redistributed.

For transparent mode, if the real host is directly-connected, configure the static route on the upstream router to point to the Firewall Threat Defense: specify the bridge group IP address. For remote hosts in transparent mode, in the static route on the upstream router, you can alternatively specify the downstream router IP address.

The same address as the real address (identity NAT)

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

Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues. For example, if you configure a broad identity NAT rule for “any” IP address, then leaving proxy ARP enabled can cause problems for hosts on the network directly connected to the mapped interface. In this case, when a host on the mapped network wants to communicate with another host on the same network, then the address in the ARP request matches the NAT rule (which matches “any” address). The Firewall Threat Defense will then proxy ARP for the address, even though the packet is not actually destined for the Firewall Threat Defense. (Note that this problem occurs even if you have a manual NAT rule; although the NAT rule must match both the source and destination addresses, the proxy ARP decision is made only on the “source” address). If the

Firewall Threat Defense ARP response is received before the actual host ARP response, then traffic will be mistakenly sent to the Firewall Threat Defense.

Requirements and prerequisites for NAT policies

Ensure you meet these requirements for configuring NAT policies:

Supported domains

Any

User roles

You must have one of these user privileges to configure and use NAT policies:

- Admin
- Access Admin
- Network Admin

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).

Note You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes.

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 (manual 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 (manual 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.

NAT support for inspected protocols

Some application layer protocols that open secondary connections, or that embedded IP addresses in packets, are inspected to provide the following services:

- Pinhole creation—Some application protocols open secondary TCP or UDP connections either on standard or negotiated ports. Inspection opens pinholes for these secondary ports so that you do not need to create access control rules to allow them.

- NAT rewrite— Protocols such as FTP embed IP addresses and ports for the secondary connections in packet data as part of the protocol. If there is NAT translation involved for either of the endpoints, the inspection engines rewrite the packet data to reflect the NAT translation of the embedded addresses and ports. The secondary connections would not work without NAT rewrite.
- Protocol enforcement—Some inspections enforce some degree of conformance to the RFCs for the inspected protocol.

The following table lists the inspected protocols that apply NAT rewrite and their NAT limitations. Keep these limitations in mind when writing NAT rules that include these protocols. Inspected protocols not listed here do not apply NAT rewrite. These inspections include GTP, HTTP, IMAP, POP, SMTP, SSH, and SSL.



Note NAT rewrite is supported on the listed ports only. For some of these protocols, you can extend inspection to other ports using Network Analysis Policies, but NAT rewrite is not extended to those ports. This includes DCERPC, DNS, FTP, and Sun RPC inspection. If you use these protocols on non-standard ports, do not use NAT on the connections.

Table 2: NAT Supported Application Inspection

Application	Inspected Protocol, Port	NAT Limitations	Pinholes Created
DCERPC	TCP/135	No NAT64.	Yes
DNS over UDP	UDP/53	No NAT support is available for name resolution through WINS.	No
ESMTP	TCP/25	No NAT64.	No
FTP	TCP/21	(Clustering) No static PAT.	Yes
H.323 H.225 (Call signaling) H.323 RAS	TCP/1720 UDP/1718 For RAS, UDP/1718-1719	(Clustering) No static PAT. No extended PAT. No NAT64.	Yes
ICMP ICMP Error	ICMP (ICMP traffic directed to a device interface is never inspected.)	No limitations.	No
IP Options	RSVP	No NAT64.	No
NetBIOS Name Server over IP	UDP/137, 138 (Source ports)	No extended PAT. No NAT64.	No
RSH	TCP/514	No PAT. No NAT64. (Clustering) No static PAT.	Yes

Application	Inspected Protocol, Port	NAT Limitations	Pinholes Created
RTSP	TCP/554 (No handling for HTTP cloaking.)	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes
SIP	TCP/5060 UDP/5060	No extended PAT. No NAT64 or NAT46. (Clustering) No static PAT.	Yes
Skinny (SCCP)	TCP/2000	No extended PAT. No NAT64, NAT46, or NAT66. (Clustering) No static PAT.	Yes
SQL*Net (versions 1, 2)	TCP/1521	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes
Sun RPC	TCP/111 UDP/111	No extended PAT. No NAT64.	Yes
TFTP	UDP/69	No NAT64. (Clustering) No static PAT. Payload IP addresses are not translated.	Yes
XDMCP	UDP/177	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes

FQDN destination guidelines

You can specify the translated (mapped) destination in a manual 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 can specify IPv4 or both IPv4 and IPv6. In this case, an IPv4 address is selected.

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.

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.
- (Auto 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.

- You cannot use an object group with both IPv4 and IPv6 addresses; the object group must include only one type of address.
- A network object used in NAT cannot include more than 131,838 IP addresses, either explicitly or implied in a range of addresses or a subnet. Break up the address space into smaller ranges and write separate rules for the smaller objects.

- (Manual 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 Firewall Threat Defense performs NAT on a packet, the packet must be IPv6-to-IPv6 or IPv4-to-IPv4; with this prerequisite, the Firewall Threat Defense 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.
- 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.
- 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 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.
- 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.
- (Manual 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.

Manage NAT policies

NAT converts the IP address of an incoming packet to a different address in the outgoing packet. 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 routable addresses that can be used on the public Internet. NAT keeps track of the translations, also known as xlates, to ensure that return traffic is directed to the correct untranslated host address.

Procedure

Step 1 Choose **Devices** > **NAT**.

Step 2 Manage your NAT policies:

- Create—Click **New Policy**. Refer to Create NAT policies, on page 18.
- Copy—Click **Copy** ( ) next to the policy you want to copy. You are prompted to give the copy a new, unique name. The copy includes all policy rules and configurations, but does not include device assignments.
- Report—Click **Report** ( ) for the policy. You are prompted to save the PDF report, which includes policy attributes, device assignments, rules, and object usage information.
- Edit—Click **Edit** ( ) next to the policy you want to edit. Refer to Configure NAT for threat defense, on page 19.
- Delete—Click **Delete** ( ) next to the policy you want to delete, then click **OK**. When prompted whether to continue, you are also informed if another user has unsaved changes in the policy.

Caution After you have deployed a NAT policy to a managed device, you cannot delete the policy from the device. Instead, you must deploy a NAT policy with no rules to remove the NAT rules already present on the managed device. You also cannot delete a policy that is the last deployed policy on any of its target devices, even if it is out of date. Before you can delete the policy completely, you must deploy a different policy to those targets.

Create NAT policies

Create a NAT policy to define how network addresses are translated between different network segments.

When you create a new NAT policy you must, at minimum, give it a unique name. Although you are not required to identify policy targets at policy creation time, you must perform this step before you can deploy the policy. If you apply a NAT policy with no rules to a device, the system removes all NAT rules from that device.

Procedure

Step 1 Choose **Devices** > **NAT**.

Step 2 Click **New Policy** and from the drop-down list, choose **Threat Defense NAT** for Firewall Threat Defense devices.

Step 3 Enter a unique **Name**.

Step 4 (Optional) Enter a **Description**.

Step 5 Choose the devices where you want to deploy the policy:

- Choose a device in the **Available Devices and Templates** list, and then click **Add to Policy**.

  Alternatively, click and drag a device from the **Available Devices and Templates** list to the **Selected Devices and Templates** list.

- Remove a device from the **Selected Devices and Templates** list by clicking **Delete** ( ) next to the device.

Step 6 Click **Save**.

Configure NAT policy targets

Identify the managed devices that you want to target with your policy while creating or editing a policy and assign NAT policies to identified devices to control NAT behavior. You can search a list of available devices and high-availability pairs, and add them to a list of selected devices.

Procedure

Step 1 Choose **Devices** > **NAT**.

Step 2 Click **Edit** (✎) next to the NAT policy you want to modify.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Policy Assignments**.

Step 4 Do any of the following:

- To assign a device, high-availability pair, or device group to the policy, select it in the **Available Devices and Templates** list and click **Add to Policy**. Alternatively, you can also drag and drop selected devices.
- To remove a device assignment, click **Delete** (🗑) next to a device, high-availability pair, or device group in the **Selected Devices and Templates** list.

Step 5 Click **OK**.

Configure NAT for threat defense

The NAT policy is a shared policy. Assign it to devices that need similar NAT rules.

Each rule applies to a device if its interface objects (security zones or interface groups) include an interface for that device. When the interface objects include one or more interface for the device, the rule is deployed to the device. You can target subsets of devices within a shared policy by designing interface objects accordingly. Rules using the "any" interface object apply to all devices.

If you change an interface type to one not valid for NAT on an assigned device, the policy marks the interface as deleted. Click **Save** in the NAT policy to automatically remove the interface from the policy.

You can configure multiple NAT policies if groups of your devices require significantly different rules.

Procedure

Step 1 Select **Devices** > **NAT**.

- Click **New Policy** to create a new policy. Give the policy a name, optionally assign devices to it, and click **Save**.

  You can change device assignments later by editing the policy and clicking **Policy Assignments**.

- Click **Edit** (✎) to edit an existing threat defense NAT policy.

  If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Decide what kinds of rules you need.

You can create dynamic NAT, dynamic PAT, static NAT, and identity NAT rules. For an overview, refer to NAT types, on page 2.

Step 3 Decide which rules should be implemented as manual or auto NAT.

For a comparison of these two implementation options, refer to Auto NAT and Manual NAT, on page 5.

Step 4 Decide which rules should be custom per device.

Because you can assign a NAT policy to multiple devices, you can configure a single rule on many devices. However, you might have rules that should be interpreted differently by each device, or some rules that should apply to a subset of devices only.

Use interface objects to control on which devices a rule is configured. Then, use object overrides on network objects to customize the addresses used per device.

For detailed information, refer to Customize NAT rules for multiple devices, on page 21.

Step 5 Create the rules as explained in these sections.

- Dynamic NAT, on page 24
- Dynamic PAT, on page 30
- Static NAT, on page 40
- Identity NAT, on page 49

Step 6 Do these to manage your NAT policy:

- To edit the policy name or description, click in those fields, type in your changes, and click outside the fields.
- To view only those rules that apply to a specific device, click **Filter by Device** and select the desired device. A rule applies to a device if it uses an interface object that includes an interface on the device.
- To view any warnings or errors in the policy, click **Show Warnings**, then choose a **Device**. Warnings and errors mark configurations that could adversely affect traffic flow or prevent the policy from deploying.
- To change the devices to which the policy is assigned, click the **Policy Assignments** link and modify the selected devices list as desired.
- To change the number of rules displayed on the page, use the **Rows Per Page** drop-down list.

Step 7 Do these to manage your NAT rules:

- To add a rule, click the **Add Rule** button.
- To edit a rule, click the **Edit** (✎) icon next to the rule.
- To delete a rule, click the **Delete** (🗑) icon next to the rule.
- To enable or disable a rule, right-click the rule and choose **Enable** or **Disable** from the **State** menu. You can temporarily disable a rule without deleting it using these controls.
- To select more than one rule to enable, disable, or delete, click the checkbox for the rules, or the checkbox in the header, then perform the action.

Step 8 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Customize NAT rules for multiple devices

Because the NAT policy is shared, you can assign a given policy to more than one device. However, you can configure at most one auto NAT rule for a given object. Thus, if you want to configure different translations for an object based on the specific device doing the translation, you need to carefully configure the interface objects (security zones or interface groups) and define network object overrides for the translated address. The interface objects determine on which devices a rule gets configured. The network object overrides determine what IP addresses are used by a given device for that object.

Consider the following scenario:

- FTD-A and FTD-B have inside networks 192.168.1.0/24 attached to the interface named “inside.”
- On FTD-A, you want to translate all 192.168.1.0/24 addresses to a NAT pool in the 10.100.10.10 - 10.100.10.200 range when going to the “outside” interface.
- On FTD-B, you want to translate all 192.168.1.0/24 addresses to a NAT pool in the 10.200.10.10 - 10.200.10.200 range when going to the “outside” interface.

Do these steps to accomplish the above scenario. Although this example rule is for dynamic auto NAT, you can generalize the technique for any type of NAT rule.

Procedure

Step 1 Create the security zones for the inside and outside interfaces.

a) Choose **Objects** > **Object Management**.
b) Select **Interface** from the table of contents and click **Add** > **Security Zone**.

You can use interface groups instead of zones.

c) Configure the inside zone properties.

- **Name**—Enter a name, for example, **inside-zone**.
- **Type**—Select **Routed** for routed-mode devices, **Switched** for transparent mode.
- **Selected Interfaces**—Add the FTD-A/inside and FTD-B/inside interfaces to the selected list.

d) Click **Save**.
e) Click **Add** > **Security Zone** and define the outside zone properties.

- **Name**—Enter a name, for example, **outside-zone**.
- **Interface Type**—Select **Routed** for routed-mode devices, **Switched** for transparent mode.
- **Selected Interfaces**—Add the FTD-A/outside and FTD-B/outside interfaces to the selected list.

f) Click **Save**.

Step 2 Create the network object for the original inside network on the Object Management page.

a) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

b) Configure the inside network properties.

- **Name**—Enter a name, for example, **inside-network**.
- **Network**—Enter the network address, for example, **192.168.1.0/24**.

c) Click **Save**.

Step 3 Create the network object for the translated NAT pool and define overrides.

a) Click **Add Network** > **Add Object**.

b) Configure the NAT pool properties for FTD-A.

- **Name**—Enter a name, for example, **NAT-pool**.
- **Network**—Enter the range of addresses to include in the pool for FTD-A, for example, **10.100.10.10-10.100.10.200**.

c) Select **Allow Overrides**.

d) Click the **Overrides** heading to open the list of object overrides.

e) Click **Add** to open the Add Object Override dialog box.

f) Select FTD-B and **Add** it to the **Selected Devices** list.

g) Click **Override** and change **Network** to **10.200.10.10-10.200.10.200**.

h) Click **Add** to add the override to the device.

By defining an override for FTD-B, whenever the system configures this object on FTD-B, it will use the override value instead of the value defined in the original object.

i) Click **Save**.

Step 4 Configure the NAT rule.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside-zone.
- **Destination Interface Objects** = outside-zone.



Note The interface objects control on which devices the rule is configured. Because in this example the zones contain interfaces for FTD-A and FTD-B only, even if the NAT policy were assigned to additional devices, the rule would be deployed to those 2 devices only.

e) On **Translation**, configure the following:

- **Original Source** = inside-network object.
- **Translated Source** > **Address**= NAT-pool object.

f) Click **Save**.

You now have a single rule that will be interpreted differently for FTD-A and FTD-B, providing unique translations for the inside networks protected by each firewall.

Search and filter the NAT rule table

You can search and filter the NAT rule table to help you find rules that you need to modify or view. When you filter the table, only matching rules are shown. Note that although the rule numbers change to be sequentially 1, 2, and so forth, filtering does not change the actual rule number or the rule's location in the table relative to hidden rules. Filtering simply changes what you can see to help you locate rules that interest you.

When editing the NAT policy, you can use the fields above the table to do these types of search or filter:

- **Filter by Device**—Click **Filter by Device**, then select the devices whose rules you want to see and click **OK**. Whether a rule applies to a device is determined by the rule's interface constraints. If you specify a security zone or interface group for either the source or destination interface, the rule applies to a device if at least one interface for the device is in the zone or group. If a NAT rule applies to any source and any destination interface, then it applies to all devices.

  If you also do a text or multiple-attribute search, the results are constrained to the selected devices.

  To remove this filter, click **Filter by Device** and deselect the devices, or select **All**, and click **OK**.

- **Simple Text Search**—In the **Filter** box, type a string and press Enter. The string is compared to all values in the rules. For example, if you enter "network-object-1," which is the name of a network object, you would get rules that use the object in source, destination, and PAT pool attributes.

  For network and port objects, the string is also compared to the contents of the objects used in the rule. For example, if a PAT pool object includes the range 10.100.10.3-10.100.10.100, searching on either 10.100.10.3 or 10.100.10.100 (or a partial 10.100.10) will include rules that use that PAT pool object. However, the match must be exact: searching on 10.100.10.5 will not match this PAT pool object, even though the IP address is within the object's IP address range.

  To remove the filter, click the **x** on the right side of the Filter box.

- **Multiple-Attribute Search**—If a simple text search gives you too many hits, you can configure multiple values for the search. Click in the **Filter** box to open the list of attributes, then select or enter strings for the attributes you intend to search and click the **Filter** button. These attributes are the same as the ones you would configure within a NAT rule. The attributes are and'ed, so filtered results include only those rules that match all attributes you configured.

  - For binary attributes, such as the rule state (enabled/disabled), whether a PAT pool is configured (enabled/disabled), the direction of the rule (uni/bi), or rule type (static/dynamic), simply check or uncheck the boxes as appropriate. Select both boxes if you do not care about the attribute value. If you deselect both boxes, no rules will match the filter.

- For string attributes, type a full or partial string relevant to that attribute. These will be object names, either for security zones/interface groups, network objects, or port objects. It can also be the network or port object contents, which are matched the same way they are for simple text searches.

To remove the filter, click the **x** on the right side of the Filter box, or click in the Filter box to open the drop-down list, and click the Clear button.

Enable, disable, or delete multiple rules

You can enable or disable manual NAT rules, or delete any NAT rule, one by one. You can also select multiple rules and apply changes to all of them at once. Because enable/disable applies to manual NAT only, if you select a mix of rule types, you can delete them only. Note that when you enable or disable rules, it does not matter if you select some rules that were already enabled or disabled. For example, enabling an already enabled rule simply leaves the rule enabled.

Procedure

Step 1 Select **Devices** > **NAT**, and edit a threat defense NAT policy.

Step 2 (Optional) Filter the NAT rules to locate the ones you want to change.

Filtering is especially useful if you have a large NAT policy. For example, you could search on disabled rules to find ones that need to be enabled.

Step 3 Select the rules you want to change.

- Click the checkbox in the left column of the rule to select (or deselect) individual rules.
- Click the checkbox in the table header to select all the rules on the currently displayed page.

Your selection is preserved as you go from page to page. However, in practice, it makes most sense to perform your actions on the rules selected on a page before going to the next page.

Step 4 Perform the desired action. When selecting multiple rules, you are asked to confirm the action.

Note that these actions are also available on the right-click menu.

- To enable all rules, click **Select Bulk Action** > **Enable**.
- To disable all rules, click **Select Bulk Action** > **Disable**.
- To delete all rules, click **Select Bulk Action** > **Delete**.

Dynamic NAT

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

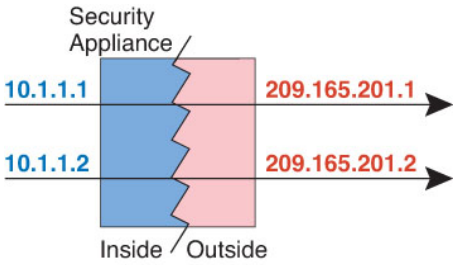
About dynamic NAT

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.

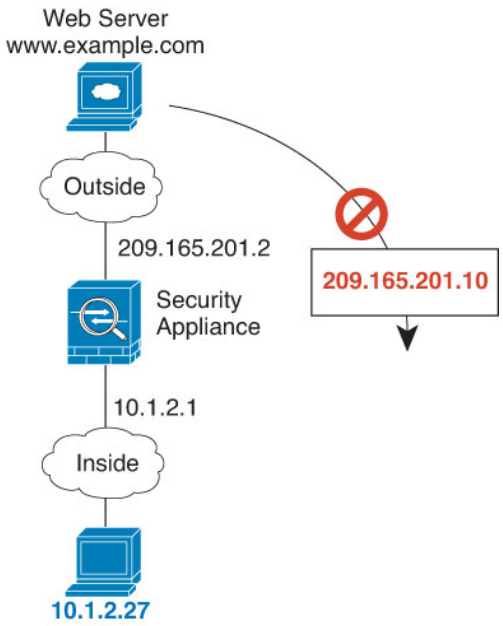
The following figure shows a typical dynamic NAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back.

Figure 5: 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 6: 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 auto NAT

Use dynamic auto NAT rules to translate addresses to different IP addresses that are routable on the destination network.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.

- **Translated Source**—This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group 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.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Dynamic**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—The network object or group that contains the mapped addresses.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see DNS query and response rewriting using NAT, on page 105.
- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure dynamic manual NAT

Use dynamic manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Dynamic NAT translates addresses to different IP addresses that are routable on the destination network.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—This can be a network object or group, but it cannot include a subnet. If a group 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.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule.

For dynamic NAT, you can also perform port translation on the destination. In the Object Manager, ensure that there are port objects you can use for the **Original Destination Port** and **Translated Destination Port**. If you specify the source port, it will be ignored.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

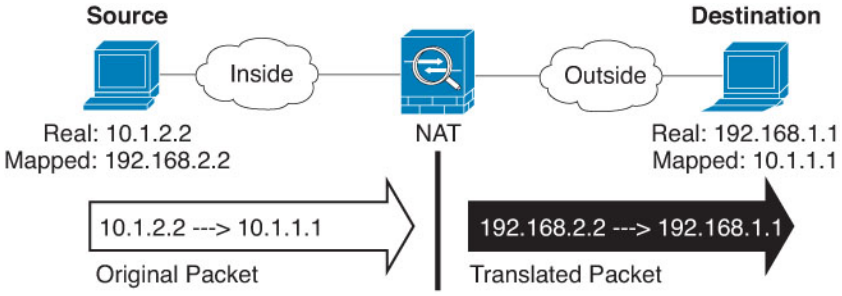
- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.



- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—The network object or group that contains the mapped addresses.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the destination service ports for service translation: **Original Destination Port**, **Translated Destination Port**.

Dynamic NAT does not support port translation, so leave the **Original Source Port** and **Translated Source Port** fields empty. However, because the destination translation is always static, you can perform port translation for the destination port.

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.

Step 8 (Optional.) On **Advanced**, select the desired options:

- (For source translation only.) **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see DNS query and response rewriting using NAT, on page 105.
- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 9 Click **Save** to add the rule.

Step 10 Click **Save** on the NAT page to save your changes.

Dynamic PAT

The following topics describe dynamic PAT.

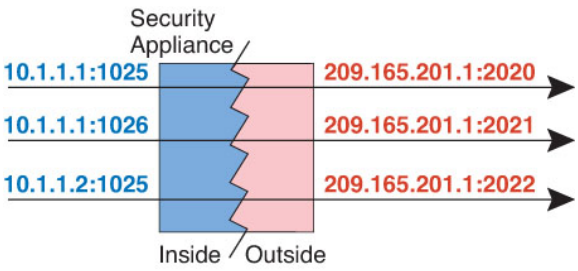
About dynamic PAT

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 7: 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.

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 Firewall Threat Defense 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. For more information, see NAT support for inspected protocols, on page 12.

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.
- 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 auto PAT

Use dynamic auto PAT rules to translate addresses to unique IP address/port combinations, rather than to multiple IP addresses only. You can translate to a single address (either the destination interface's address or another address), or use a PAT pool of addresses to provide a larger number of possible translations.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Source**—You have the following options to specify the PAT address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object.
  - **Single PAT address**—Create a network object containing a single host.
  - **PAT pool**—Create a network object that includes a range, or create a network object group that contains hosts, ranges, or both. You cannot include subnets. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Dynamic**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Skip the step for configuring a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Skip the step for configuring a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Do not select the destination interface IP or a single address as the **Translated Source**.

Step 6 If you are using a PAT pool, select the **PAT Pool** page and do the following:

a) Select **Enable PAT pool**.

b) Select the network object group that contains the addresses for the pool in the **PAT** > **Address** field.

You can alternatively select **Destination Interface IP**, which is another way to implement interface PAT.

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

- **Use Round Robin Allocation**—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.

- **Extended PAT Table**—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. You cannot use this option with interface PAT or interface PAT fallback.
- **Flat Port Range**, **Include Reserved Ports**—To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. (Pre-6.7) When choosing the mapped port number for a translation, PAT 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 Reserved Ports** option. For the Firewall Threat Defense devices running version 6.7 or higher, the flat port range is always configured, whether you select the option or not. You can still select the **Include Reserved Ports** option for these systems, and that setting is honored.
- **Block Allocation**—To enable 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 the extended PAT or flat port range options. You also cannot use interface PAT fallback.

Step 7 (Optional.) On **Advanced**, select the desired options:

- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option. You cannot select this option if you already configured interface PAT as the translated address or PAT pool.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 8 Click **Save** to add the rule.

Step 9 Click **Save** on the NAT page to save your changes.

Configure dynamic manual PAT

Use dynamic manual PAT rules when auto PAT does not meet your needs. For example, if you want to do different translations based on the destination. Dynamic PAT translates addresses to unique IP address/port combinations, rather than to multiple IP addresses only. You can translate to a single address (either the destination interface's address or another address), or use a PAT pool of addresses to provide a larger number of possible translations.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—You have the following options to specify the PAT address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object.
  - **Single PAT address**—Create a network object containing a single host.
  - **PAT pool**—Create a network object that includes a range, or create a network object group that contains hosts, ranges, or both. You cannot include subnets.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule.

For dynamic NAT, you can also perform port translation on the destination. In the Object Manager, ensure that there are port objects you can use for the **Original Destination Port** and **Translated Destination Port**. If you specify the source port, it will be ignored.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

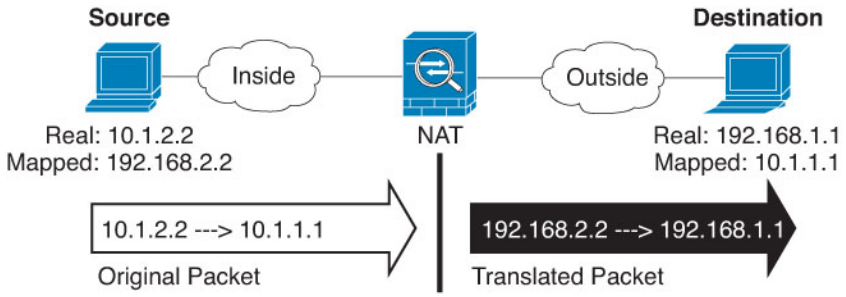
- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.



- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Skip the step for configuring a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Skip the step for configuring a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Do not select the destination interface IP or a single address as the **Translated Source**.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the destination service ports for service translation: **Original Destination Port**, **Translated Destination Port**.

Dynamic NAT does not support port translation, so leave the **Original Source Port** and **Translated Source Port** fields empty. However, because the destination translation is always static, you can perform port translation for the destination port.

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.

Step 8 If you are using a PAT pool, select the **PAT Pool** page and do the following:

a) Select **Enable PAT pool**.

b) Select the network object group that contains the addresses for the pool in the **PAT** > **Address** field.

You can alternatively select **Destination Interface IP**, which is another way to implement interface PAT.

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

- **Use Round Robin Allocation**—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.
- **Extended PAT Table**—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. You cannot use this option with interface PAT or interface PAT fallback.
- **Flat Port Range**, **Include Reserved Ports**—To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. (Pre-6.7) When choosing the mapped port number for a translation, PAT 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 Reserved Ports** option. For the Firewall Threat Defense devices running version 6.7 or higher, the flat port range is always configured, whether you select the option or not. You can still select the **Include Reserved Ports** option for these systems, and that setting is honored.
- **Block Allocation**—To enable 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 the extended PAT or flat port range options. You also cannot use interface PAT fallback.

Step 9 (Optional.) On **Advanced**, select the desired options:

- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 10 Click **Save** to add the rule.

Step 11 Click **Save** on the NAT page to save your changes.

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 **Use Round Robin Allocation** option, but you cannot include the options for extending PAT uniqueness, using a flat range, including the reserved ports, 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.

- When you form a cluster while the device that will become the control node has an existing NAT policy that uses PAT pools, PAT port block distribution might not initialize correctly. The control node does not own or distribute the port blocks, and the **show nat pool cluster summary** command might show all port blocks as *RESERVED* on every cluster node.
  - Before forming a cluster, remove any existing NAT policy that uses PAT pools from the device that will become the control node. After cluster formation is complete, reconfigure and deploy the NAT policy.
  - If the cluster is already formed and all port blocks appear as *RESERVED*, delete the affected NAT policy and deploy the change. Then re-create and deploy the policy. Use the **show nat pool cluster summary** command to verify that the port blocks are distributed correctly across the cluster nodes.
- 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 (Optional.) Configure global PAT port block allocation settings.

There are a few global settings that control port block allocation. If you want to change the defaults for these options, you must configure a FlexConfig object and add it to your FlexConfig policy.

a) Select **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object** and create a new object.

b) Configure the block allocation size, which is the number of ports in each block.

xlate block-allocation size *value*

The range is 32-4096. The default is 512. Use the "no" form to return to the default.

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.

Using the default block-allocation size, a single PAT IP can serve 126 unique, simultaneous users (number of ports 64,512 divided by 512).

c) Configure the maximum blocks that can be allocated per host.

xlate block-allocation maximum-per-host *number*

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. Use the "no" form to return to the default.

d) (Optional.) Enable interim syslog generation.

xlate block-allocation pba-interim-logging *seconds*

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*

Example:

The following example sets the block allocation size to 64, the per-host maximum to 8, and enables interim logging every 6 hours.

```
xlate block-allocation size 64
xlate block-allocation maximum-per-host 8
xlate block-allocation pba-interim-logging 21600
```

e) Select the following options in the FlexConfig object:

- **Deployment** = **Everytime**
- **Type** = **Append**

f) Click **Save** to create the FlexConfig object.

g) Select **Devices** > **FlexConfig**, and create or edit the FlexConfig policy that is assigned to the devices that need to have these settings adjusted.
h) Select your object in the available objects list and click > to move it to the selected objects list.
i) Click **Save**.

You can click **Preview Config**, select one of the target devices, and verify that the xlate commands appear correctly.

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

a) Select **Devices** > **NAT** and add or edit the threat defense NAT policy.
b) Add or edit a NAT rule and configure at least the following options.

- **Type** = **Dynamic**
- In **Translation** > **Original Source**, select the object that defines the source address.
- On **PAT Pool**, configure the following options:
  - Select **Enable PAT Pool**
  - In **PAT** > **Address**, select a network object or group that defines the pat pool.
  - Select the **Block Allocation** option.

c) Save your changes to the rule and to the NAT policy.

Static NAT

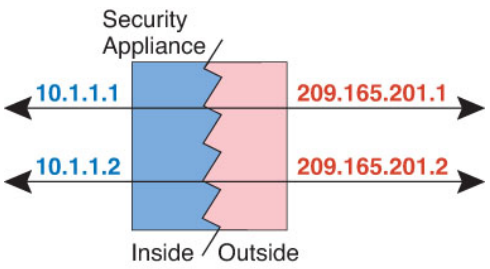
The following topics explain static NAT and how to implement it.

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 8: Static NAT



Note You can disable bidirectionality if desired.

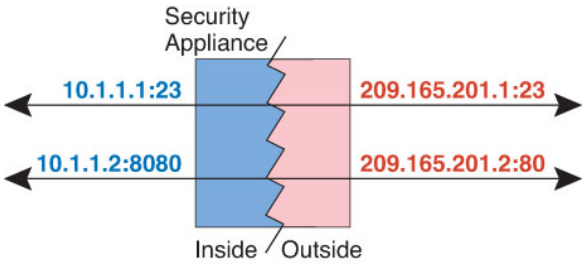
Static NAT with port translation

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 9: 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 manual 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).

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

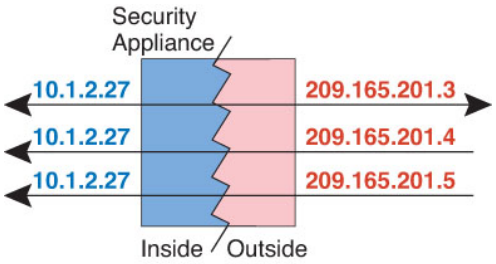
You can configure static NAT to map a real address to an interface address/port combination. For example, if you want to redirect Telnet access for the device's outside interface to an inside host, then you can map the inside host IP address/port 23 to the outside interface address/port 23.

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.

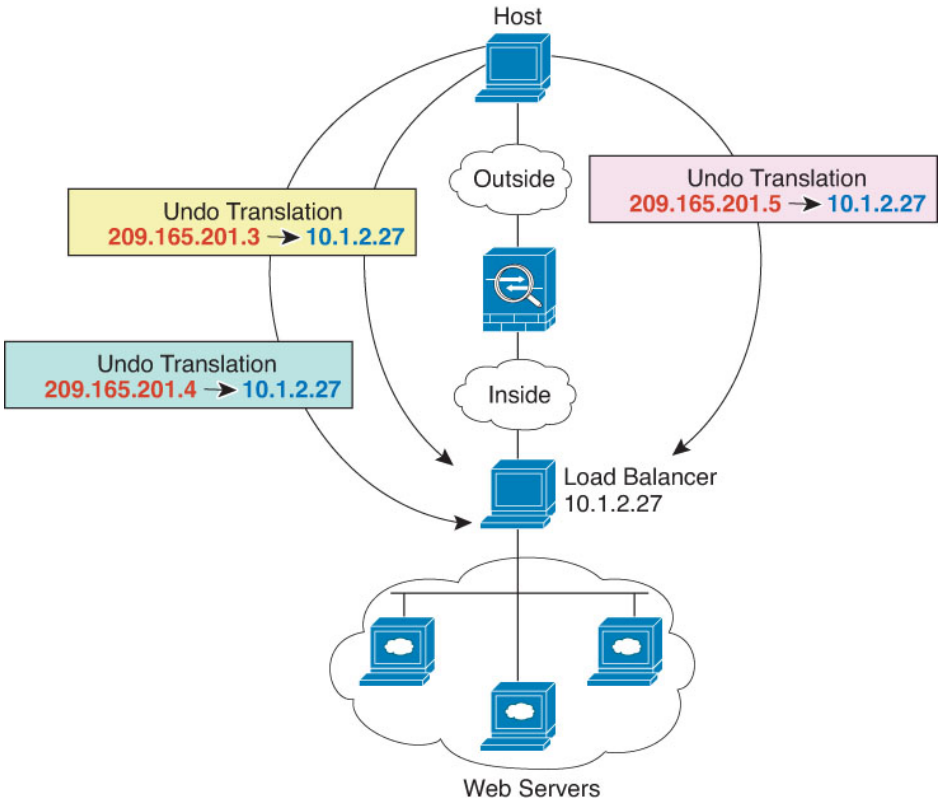
The following figure shows a typical one-to-many static NAT scenario. Because initiation by the real host always uses the first mapped address, the translation of real host IP/first mapped IP is technically the only bidirectional translation.

Figure 10: 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.

Figure 11: One-to-Many Static NAT Example



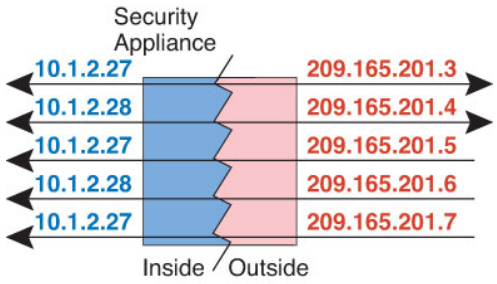
Other mapping scenarios (not recommended)

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 12: Few-to-Many Static NAT

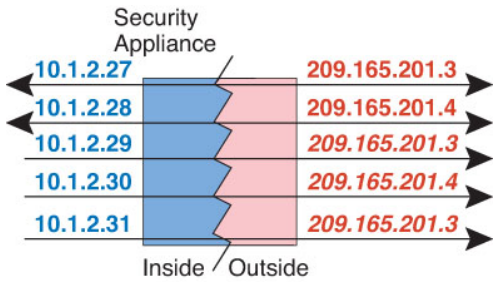


For a many-to-few or many-to-one configuration, where you have more real addresses than mapped addresses, you run out of mapped addresses before you run out of real addresses. Only the mappings between the lowest real IP addresses and the mapped pool result in bidirectional initiation. The remaining 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 13: 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 auto NAT

Use static auto NAT rules to translate addresses to different IP addresses that are routable on the destination network. You can also do port translation with the static NAT rule.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.

- **Translated Source**—You have the following options to specify the translated address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
  - **Address**—Create a network object or group containing hosts, ranges, or subnets. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. 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.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Static**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. 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.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- (Optional.) **Original Port**, **Translated Port**—If you need to translate a TCP or UDP port, select the protocol in **Original Port**, and type the original and translated port numbers. For example, you can translate TCP/80 to 8080 if necessary.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see DNS query and response rewriting using NAT, on page 105. This option is not available if you are doing port translation.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.
- **Net to Net Mapping**—For NAT 46, select this option 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 option.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure static manual NAT

Use static manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Static NAT translates addresses to different IP addresses that are routable on the destination network. You can also do port translation with the static NAT rule.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—You have the following options to specify the translated address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
  - **Address**—Create a network object or group containing hosts, range, or subnets. 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.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule. If you want to configure destination

static interface NAT with port translation only, you can skip adding an object for the destination mapped addresses and specify the interface in the rule.

You can also perform port translation on the source, destination, or both. In the Object Manager, ensure that there are port objects you can use for the original and translated ports.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

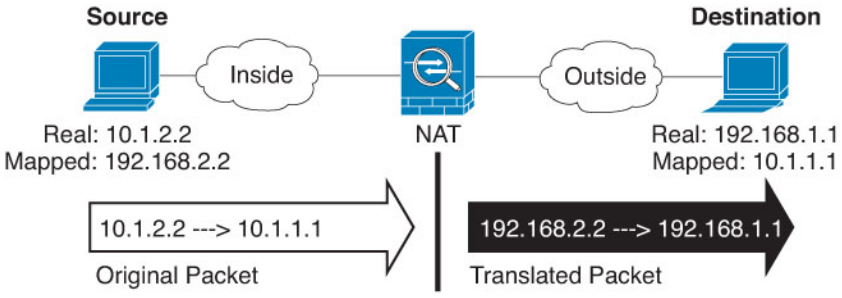
- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.



- **Original Source**—The network object or group that contains the addresses you are translating.

- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. 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.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the source or destination service ports for service translation.

If you are configuring static NAT with port translation, you can translate ports for the source, destination, or both. For example, you can translate between TCP/80 and TCP/8080.

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.

- **Original Source Port**, **Translated Source Port**—Defines a port translation for the source address.
- **Original Destination Port**, **Translated Destination Port**—Defines a port translation for the destination address.

Step 8 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see DNS query and response rewriting using NAT, on page 105. This option is not available if you are doing port translation.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

- **Net to Net Mapping**—For NAT 46, select this option 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 option.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Unidirectional**—Select this option to prevent the destination addresses from initiating traffic to the source addresses. 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.

Step 9 Click **Save** to add the rule.

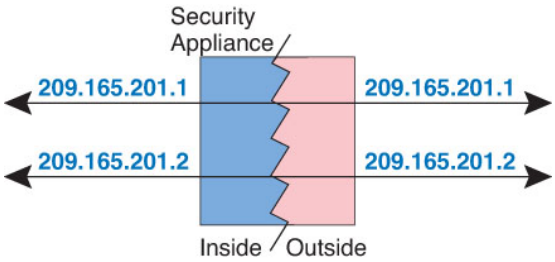
Step 10 Click **Save** on the NAT page to save your changes.

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.

The following figure shows a typical identity NAT scenario.

Figure 14: Identity NAT



The following topics explain how to configure identity NAT.

Configure identity auto NAT

Use static identity auto NAT rules to prevent the translation of an address. That is, to translate the address to itself.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.

- **Translated Source**—A network object or group with the exact same contents as the original source object. You can use the same object.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Static**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Do not configure the **Original Port** and **Translated Port** options for identity NAT.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Do not configure this option for identity NAT.
- **IPv6**—Do not configure this option for identity NAT.
- **Net to Net Mapping**—Do not configure this option for identity NAT.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Perform Route Lookup for Destination Interface**— If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure identity manual NAT

Use static identity manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Use static identity NAT rules to prevent the translation of an address. That is, to translate the address to itself.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—The same object or group as the original source. Optionally, you can select a different object that has the exact same contents.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule. If you want to configure destination static interface NAT with port translation only, you can skip adding an object for the destination mapped addresses and specify the interface in the rule.

You can also perform port translation on the source, destination, or both. In the Object Manager, ensure that there are port objects you can use for the original and translated ports. You can use the same object for identity NAT.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

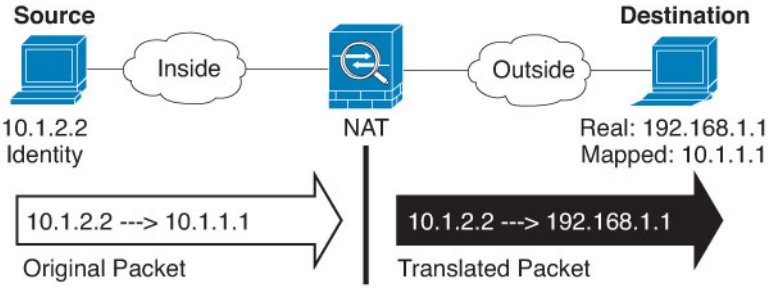
- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

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.



- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Interface Object** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—The same object or group as the original source. Optionally, you can select a different object that has the exact same contents.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the source or destination service ports for service translation.

If you are configuring static NAT with port translation, you can translate ports for the source, destination, or both. For example, you can translate between TCP/80 and TCP/8080.

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.

- **Original Source Port**, **Translated Source Port**—Defines a port translation for the source address.

- **Original Destination Port**, **Translated Destination Port**—Defines a port translation for the destination address.

Step 8 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Do not configure this option for identity NAT.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Perform Route Lookup for Destination Interface**— If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.
- **Unidirectional**—Select this option to prevent the destination addresses from initiating traffic to the source addresses. 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.

Step 9 Click **Save** to add the rule.

Step 10 Click **Save** on the NAT page to save your changes.

NAT rule properties for Firewall Threat Defense

Use network address translation (NAT) rules to translate IP addresses to other IP addresses. You would typically use NAT rules to convert private addresses to publicly routable addresses. The translation can be from one address to another, or you can use port address translation (PAT) to translate many addresses to one or a few addresses, using port numbers to distinguish among the source addresses.

NAT rules include the following basic properties. The properties are the same for auto NAT and manual NAT rules except where indicated.

NAT Type

Whether you want to configure a **Manual NAT Rule** or an **Auto NAT Rule**. Auto NAT translates the source address only, and you cannot make different translations based on the destination address. Because auto NAT is more simple to configure, use it unless you need the added features of manual NAT. For more information on the differences, see Auto NAT and Manual NAT, on page 5.

Type

Whether the translation rule is **Dynamic** or **Static**. Dynamic translation automatically chooses the mapped address from a pool of addresses, or an address/port combination when implementing PAT. Use static translation if you want to precisely define the mapped address/port.

Enable (Manual NAT only.)

Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page. You cannot disable auto NAT rules.

Insert (Manual NAT only.)

Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Description (Optional. Manual NAT only.)

A description of the purpose of the rule.

The following topics describe the tabs for the NAT rules properties.

Interface objects NAT properties

Interface objects (security zones or interface groups) define the interfaces to which a NAT rule applies. In routed mode, you can use the default "any" for both source and destination to apply to all interfaces of all assigned devices. However, you typically want to select specific source and destination interfaces.

Notes

- 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. You cannot configure NAT for the bridge virtual interface (BVI) itself, you can configure NAT for member interfaces only.

  If you select interface objects, a NAT rule will be configured on an assigned device only if the device has interfaces included in all selected objects. For example, if you select both source and destination security zones, both zones must contain one or more interface for a given device.

- If more than one interface in an interface object exists on a given device, identical rules are created for each interface. This can become an issue for static NAT rules that include destination translation. Because NAT rules are applied based on first hit rule, only the rule created for the first interface configured for the object matches traffic. When configuring static NAT with destination translation, use interface objects that include at most one interface per device assigned to the NAT policy to ensure you are getting the desired results.

Source Interface Objects, Destination Interface Objects

(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Translation properties for auto NAT

Use the **Translation** options to define the source addresses and the mapped translated addresses. The following properties apply to auto NAT only.

Original Source (Always required.)

The network object that contains the addresses you are translating. This must be a network object (not a group), and it can be a host, range, or subnet.

You cannot create auto NAT rules for the system-defined any-ipv4 or any-ipv6 objects.

Translated Source (Usually required.)

The mapped addresses, the ones to which you are translating. What you select here depends on the type of translation rule you are defining.

- **Dynamic NAT**—The network object or group that contains the mapped addresses. This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group 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**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Do not configure a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Do not configure a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Select the PAT pool object on **PAT Pool**.
- **Static NAT**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. The object or group can contain hosts, ranges, or subnets. 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.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on the **Advanced** tab. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- **Identity NAT**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Original Port, Translated Port (Static NAT only.)

If you need to translate a TCP or UDP port, select the protocol in **Original Port**, and type the original and translated port numbers. For example, you can translate TCP/80 to 8080 if necessary. Do not configure these options for identity NAT.

Translation properties for manual NAT

Use the **Translation** options to define the source addresses and the mapped translated addresses. The following properties apply to manual NAT only. All are optional except as indicated.

Original Source (Always required.)

The network object or group that contains the addresses you are translating. This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can specify **Any** in the rule.

Translated Source (Usually required.)

The mapped addresses, the ones to which you are translating. What you select here depends on the type of translation rule you are defining.

- **Dynamic NAT**—The network object or group that contains the mapped addresses. This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group 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**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Do not configure a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Do not configure a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Select the PAT pool object on **PAT Pool**.
- **Static NAT**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. The object or group can contain hosts, ranges, or subnets. 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.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on the **Advanced** tab. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- **Identity NAT**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Original Destination

The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Translated Destination

The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

You can use a network object that specifies a fully-qualified domain name as the translated destination; for more information, see FQDN destination guidelines, on page 14.

Original Source Port, Translated Source Port, Original Destination Port, Translated Destination Port

The port objects that define the source and destination services for the original and translated packets. You can translate the ports, or select the same object to make the rule sensitive to the service without translating the ports. Keep the following rules in mind when configuring services:

- (Dynamic NAT or PAT.) You cannot do translation on the **Original Source Port** and **Translated Source Port**. You can do translation on the destination port only.
- 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.

PAT pool NAT properties

When you configure dynamic NAT, you can define a pool of addresses to use for port address translation using the properties on the **PAT Pool** tab.

Enable PAT Pool

Select this option to configure a pool of addresses for PAT.

PAT

The addresses to use for the PAT pool, one of the following:

- **Address**—The object that defines the PAT pool addresses, either a network object that includes a range, or a network object group that contains hosts, ranges, or both. You cannot include subnets. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.
- **Destination Interface IP**—Indicates that you want to use the destination interface as the PAT address. For this option, you must select a specific **Destination Interface Object**; you cannot use **Any** as the destination interface. This is another way to implement interface PAT.

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.

Extended PAT Table

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. You cannot use this option with interface PAT or interface PAT fallback.

Flat Port Range; Include Reserved Ports

To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. (Pre-6.7) When choosing the mapped port number for a translation, PAT 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 Reserved Ports** option. For the Firewall Threat Defense devices running

version 6.7 or higher, the flat port range is always configured, whether you select the option or not. You can still select the **Include Reserved Ports** option for these systems, and that setting is honored.

Block Allocation

To enable 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 the extended PAT or flat port range options. You also cannot use interface PAT fallback.

Advanced NAT properties

When you configure NAT, you can configure properties that provide specialized services in the **Advanced** options. All of these properties are optional: configure them only if you need the service.

Translate DNS replies that match this rule

Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see DNS query and response rewriting using NAT, on page 105. This option is not available if you are doing port translation in a static NAT rule.

Fallthrough to Interface PAT (Destination Interface) (Dynamic NAT only.)

Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option. You cannot select this option if you already configured interface PAT as the translated address. You also cannot select the option if you configure a PAT pool.

IPv6

Whether to use the IPv6 address of the destination interface for interface PAT.

Net to Net Mapping (Static NAT only.)

For NAT 46, select this option 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 option.

Do not proxy ARP on Destination Interface (Static NAT only.)

Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.

Perform Route Lookup for Destination Interface (Static Identity NAT only. Routed mode only.)

If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.

Unidirectional (Manual NAT only, static NAT only.)

Select this option to prevent the destination addresses from initiating traffic to the source addresses. 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.

Translating IPv6 networks

In cases where you need to pass traffic between IPv6-only and IPv4-only networks, you need to use NAT to convert between the address types. Even with two IPv6 networks, you might want to hide internal addresses from the outside network.

You can use the following translation types with IPv6 networks:

- NAT64, NAT46—Translates IPv6 packets into IPv4 and vice versa. You need to define two policies, one for the IPv6 to IPv4 translation, and one for the IPv4 to IPv6 translation. Although you can accomplish this with a single manual NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a manual NAT rule when you specify a destination, creating two auto NAT rules is the better solution.



Note NAT46 supports static mappings only.

- NAT66—Translates IPv6 packets to a different IPv6 address. 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.



Note NAT64 and NAT 46 are possible on standard routed interfaces only. NAT66 is possible on both routed and bridge group member interfaces.

NAT64/46: translating IPv6 addresses to IPv4

When traffic goes from an IPv6 network to an IPv4-only network, you need to convert the IPv6 address to IPv4, and return traffic from IPv4 to IPv6. You need to define two address pools, an IPv4 address pool to bind IPv6 addresses in the IPv4 network, and an IPv6 address pool to bind IPv4 addresses in the IPv6 network.

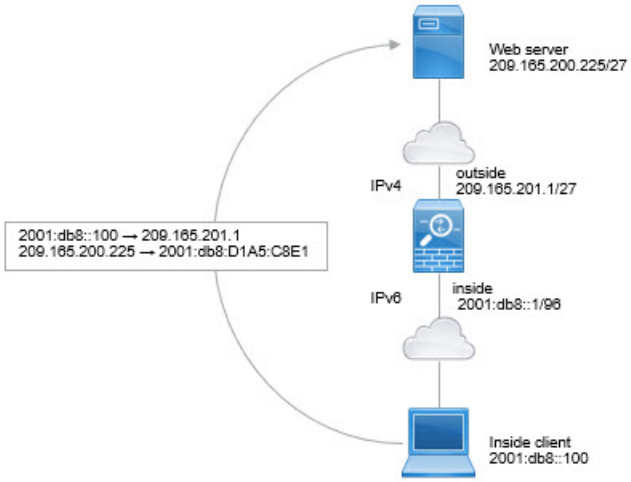
- The IPv4 address pool for the NAT64 rule is normally small and typically might not have enough addresses to map one-to-one with the IPv6 client addresses. Dynamic PAT might more easily meet the possible large number of IPv6 client addresses compared to dynamic or static NAT.

- The IPv6 address pool for the NAT46 rule can be equal to or larger than the number of IPv4 addresses to be mapped. This allows each IPv4 address to be mapped to a different IPv6 address. NAT46 supports static mappings only, so you cannot use dynamic PAT.

You need to define two policies, one for the source IPv6 network, and one for the destination IPv4 network. Although you can accomplish this with a single manual NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a manual NAT rule when you specify a destination, creating two auto NAT rules is the better solution.

NAT64/46 example: inside IPv6 network with outside IPv4 internet

Following is a straight-forward example where you have an inside IPv6-only network, and you want to convert to IPv4 for traffic sent to the Internet. This example assumes you do not need DNS translation, so you can perform both the NAT64 and NAT46 translations in a single manual NAT rule.



In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network.

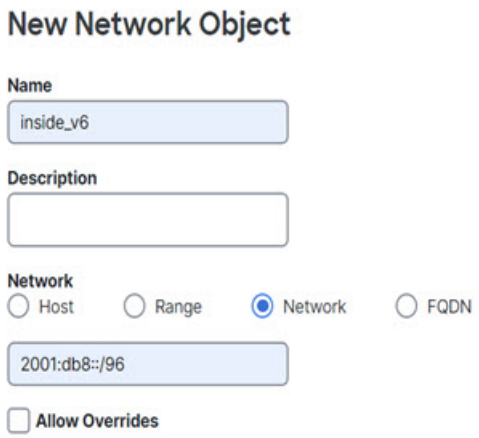
Procedure

Step 1 Create the network object that defines the inside IPv6 network.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:db8::/96.



c) Click **Save**.

Step 2 Create the manual NAT rule to translate the IPv6 network to IPv4 and back again.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.
b) Click **Add Rule**.
c) Configure the following properties:

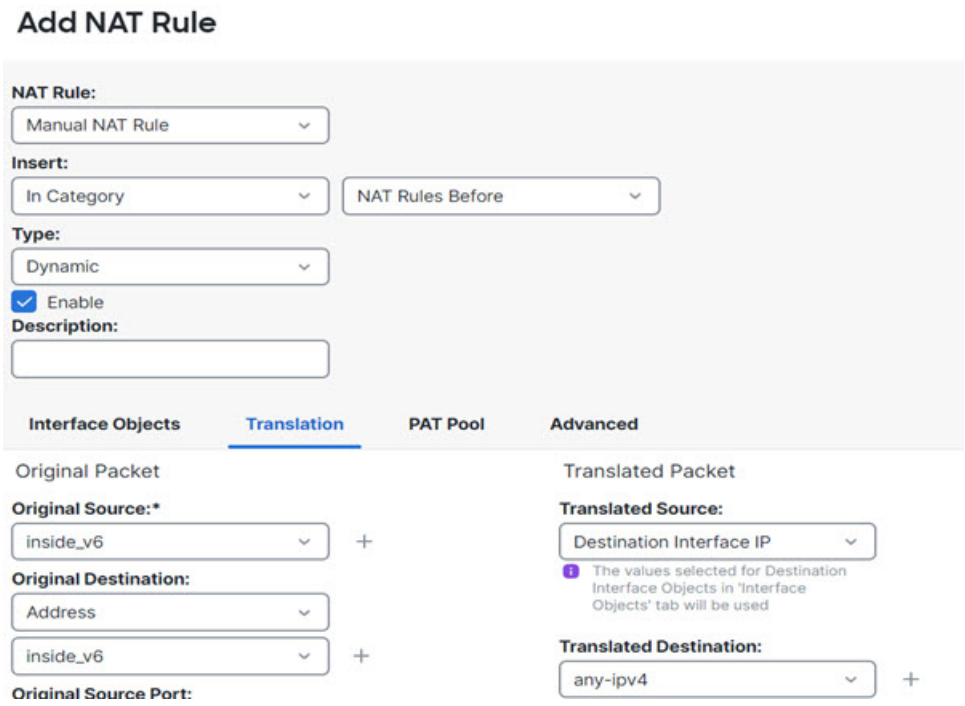
- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** = **Destination Interface IP**.
- **Original Destination** = inside_v6 network object.
- **Translated Destination** = any-ipv4 network object.



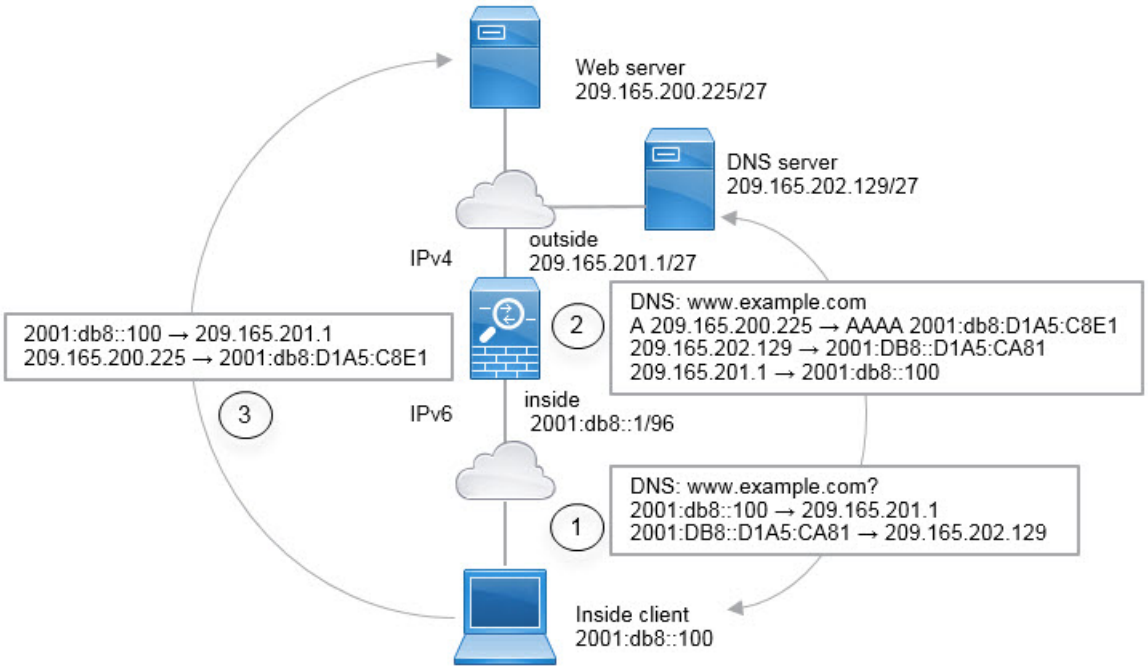
f) Click **OK**.

With this rule, any traffic from the 2001:db8::/96 subnet on the inside interface going to the outside interface gets a NAT64 PAT translation using the IPv4 address of the outside interface. Conversely, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method.

g) Click **Save** on the NAT rules page.

NAT64/46 example: inside IPv6 network with outside IPv4 internet and DNS translation

Following is a typical example where you have an inside IPv6-only network, but there are some IPv4-only services on the outside Internet that internal users need.



In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network. You enable DNS rewrite on the NAT46 rule, so that replies from the external DNS server can be converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.

Following is a typical sequence for a web request where a client at 2001:DB8::100 on the internal IPv6 network tries to open www.example.com.

1. The client's computer sends a DNS request to the DNS server at 2001:DB8::D1A5:CA81. The NAT rules make the following translations to the source and destination in the DNS request:
   - 2001:DB8::100 to a unique port on 209.165.201.1 (The NAT64 interface PAT rule.)
   - 2001:DB8::D1A5:CA81 to 209.165.202.129 (The NAT46 rule. D1A5:CA81 is the IPv6 equivalent of 209.165.202.129.)
2. The DNS server responds with an A record indicating that www.example.com is at 209.165.200.225. The NAT46 rule, with DNS rewrite enabled, converts the A record to the IPv6-equivalent AAAA record, and translates 209.165.200.225 to 2001:db8:D1A5:C8E1in the AAAA record. In addition, the source and destination addresses in the DNS response are untranslated:
   - 209.165.202.129 to 2001:DB8::D1A5:CA81
   - 209.165.201.1 to 2001:db8::100
3. The IPv6 client now has the IP address of the web server, and makes an HTTP request to www.example.com at 2001:db8:D1A5:C8E1. (D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225.) The source and destination of the HTTP request are translated:
   - 2001:DB8::100 to a unique port on 209.156.101.54 (The NAT64 interface PAT rule.)
   - 2001:db8:D1A5:C8E1 to 209.165.200.225 (The NAT46 rule.)

The following procedure explains how to configure this example.

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

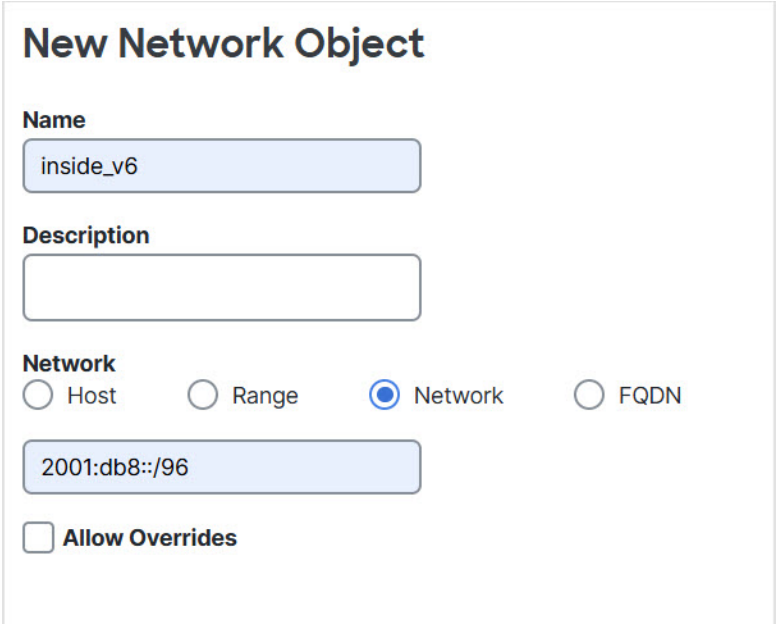
Procedure

Step 1 Create the network objects that define the inside IPv6 and outside IPv4 networks.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the inside IPv6 network.

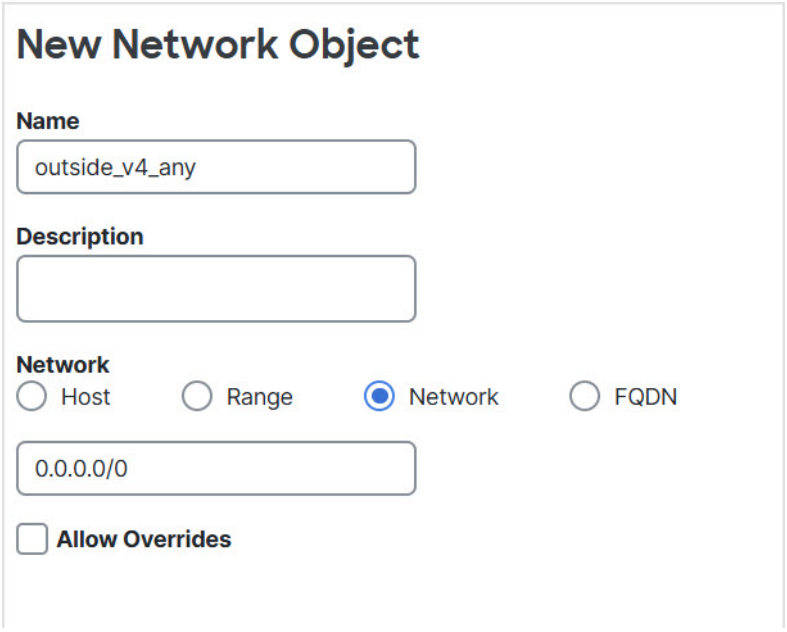
Name the network object (for example, inside_v6) and enter the network address, 2001:db8::/96.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the outside IPv4 network.

Name the network object (for example, outside_v4_any) and enter the network address 0.0.0.0/0.



e) Click **Save**.

Step 2 Configure the NAT64 dynamic PAT rule for the inside IPv6 network.

Step 3 Configure the static NAT46 rule for the outside IPv4 network.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

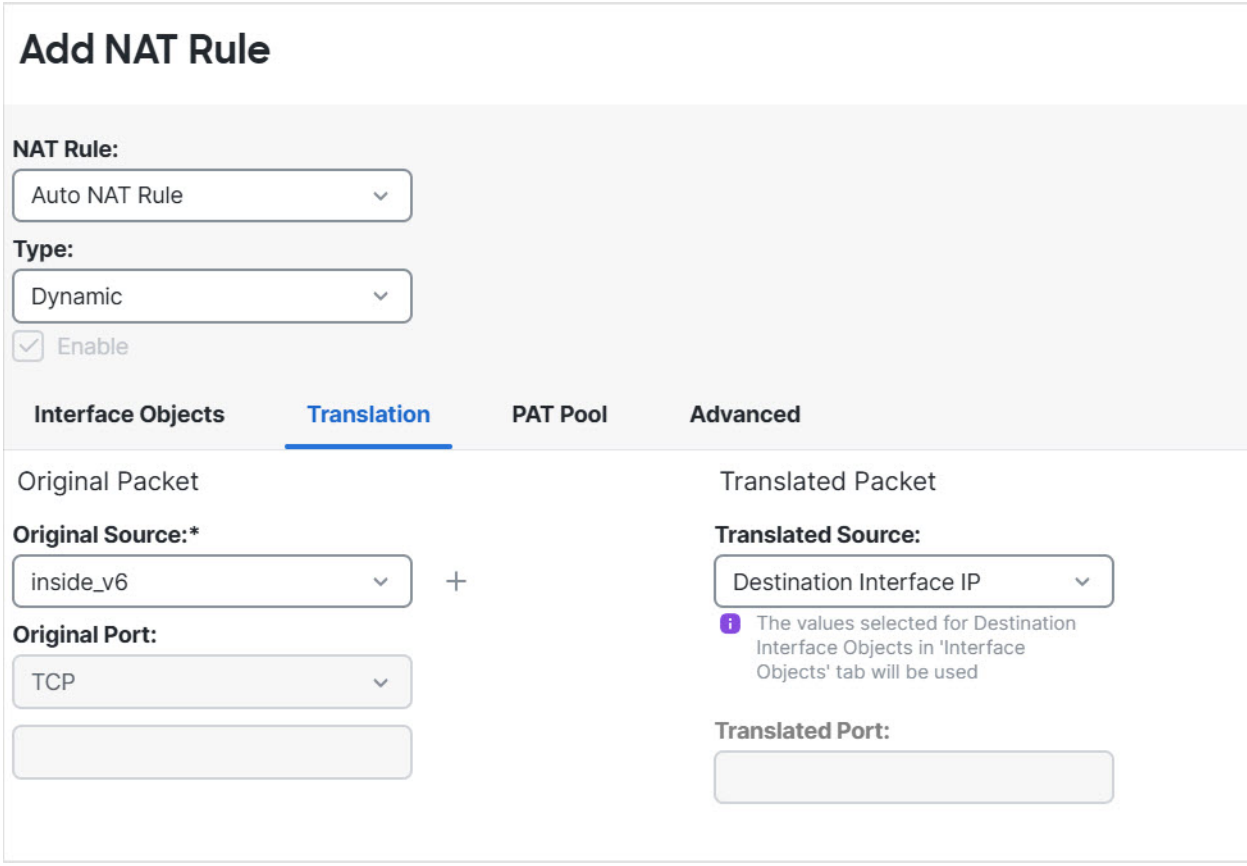
c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

d) On **Translation**, configure the following:

- **Original Source** = outside_v4_any network object.
- **Translated Source** > **Address** = inside_v6 network object.

e) On **Advanced**, select **Translate DNS replies that match this rule**.



f) Click **OK**.

With this rule, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method. In addition, DNS responses are converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.

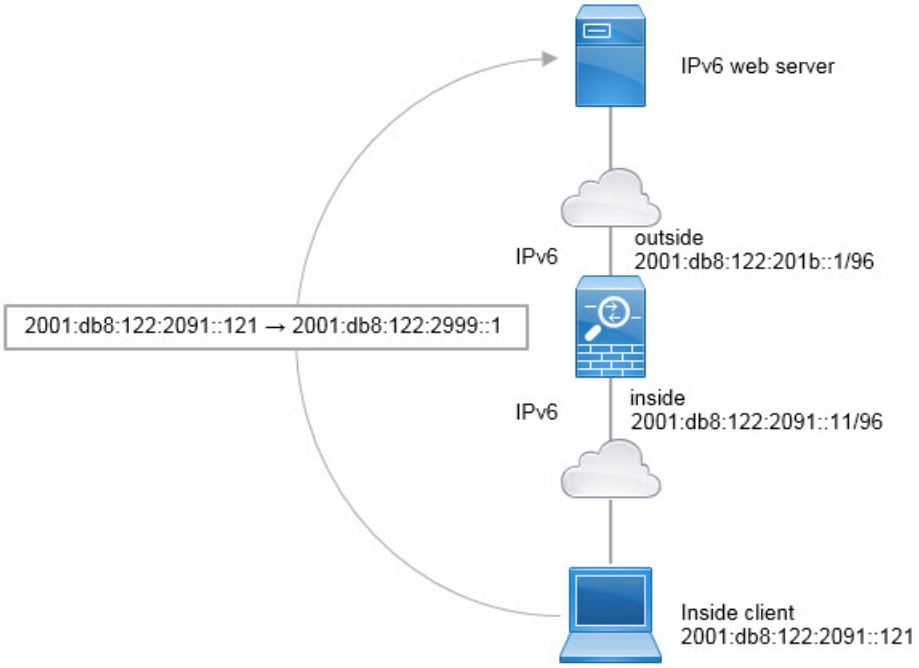
NAT66: translating IPv6 addresses to different IPv6 addresses

When going from an IPv6 network to another IPv6 network, you can translate the addresses to different IPv6 addresses on the outside network. 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.

Because you are not translating between different address types, you need a single rule for NAT66 translations. You can easily model these rules using auto NAT. However, if you do not want to allow returning traffic, you can make the static NAT rule unidirectional using manual NAT only.

NAT66 example, static translation between networks

You can configure a static translation between IPv6 address pools using auto NAT. The following example explains how to convert inside addresses on the 2001:db8:122:2091::/96 network to outside addresses on the 2001:db8:122:2999::/96 network.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

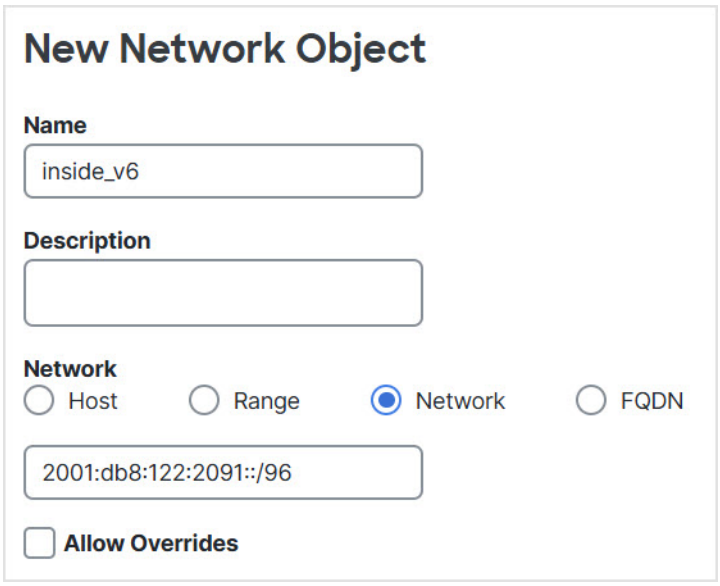
Procedure

Step 1 Create the network objects that define the inside IPv6 and outside IPv6 NAT networks.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the inside IPv6 network.

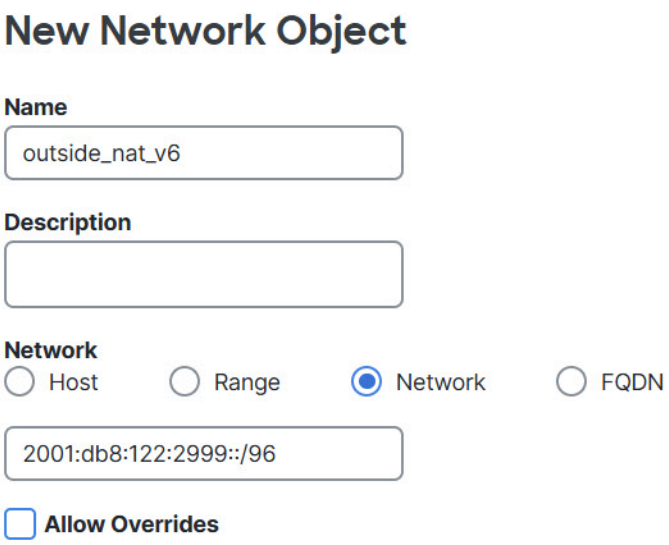
Name the network object (for example, inside_v6) and enter the network address, 2001:db8:122:2091::/96.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the outside IPv6 NAT network.

Name the network object (for example, outside_nat_v6) and enter the network address 2001:db8:122:2999::/96.



e) Click **Save**.

Step 2 Configure the static NAT rule for the inside IPv6 network.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

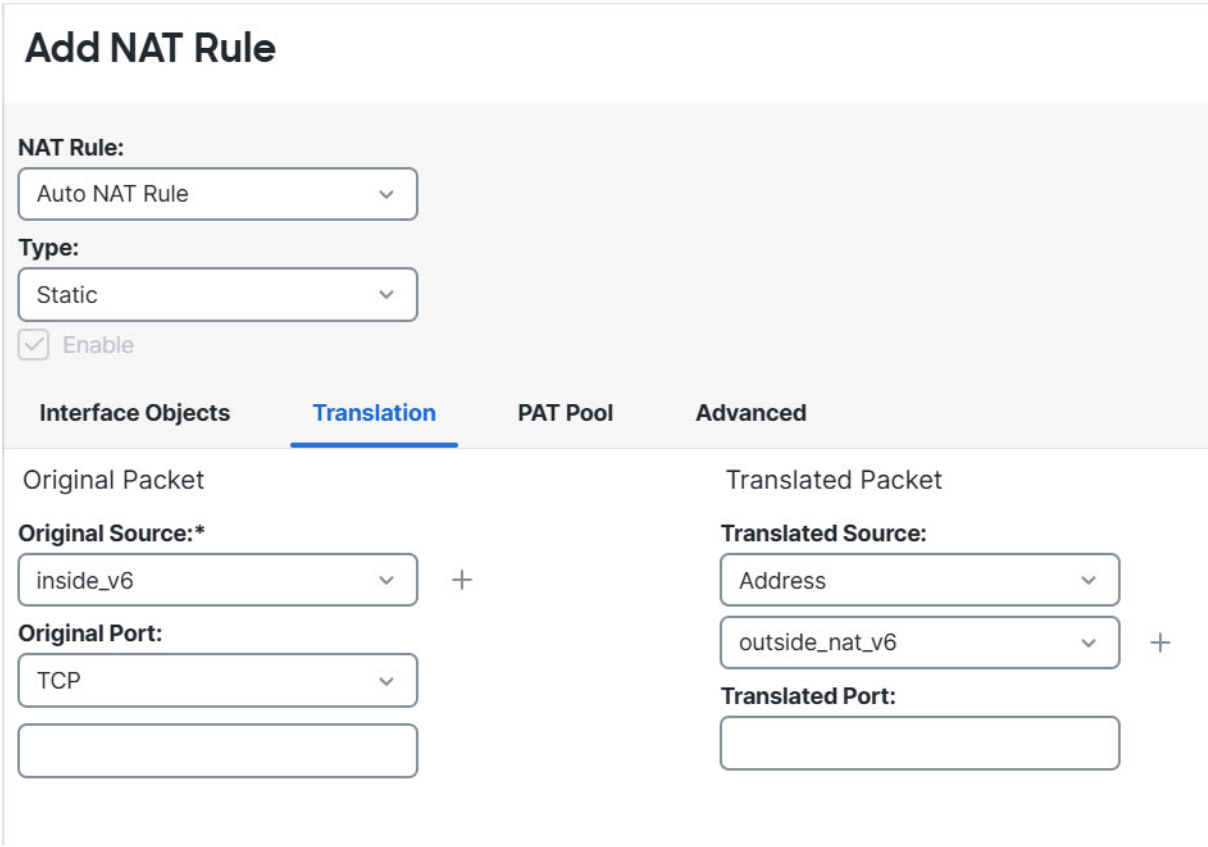
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** > **Address** = outside_nat_v6 network object.



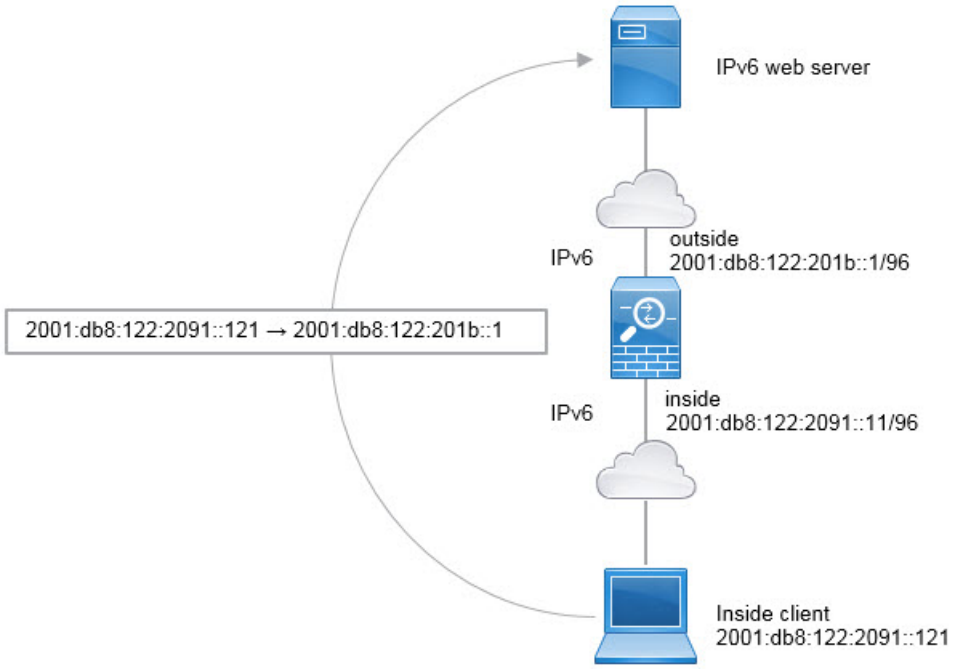
f) Click **OK**.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a static NAT66 translation to an address on the 2001:db8:122:2999::/96 network.

NAT66 example, simple IPv6 interface PAT

A simple approach for implementing NAT66 is to dynamically assign internal addresses to different ports on the outside interface IPv6 address.

When you configure an interface PAT rule for NAT66, all the global addresses that are configured on that interface are used for PAT mapping. Link-local or site-local addresses for the interface are not used for PAT.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

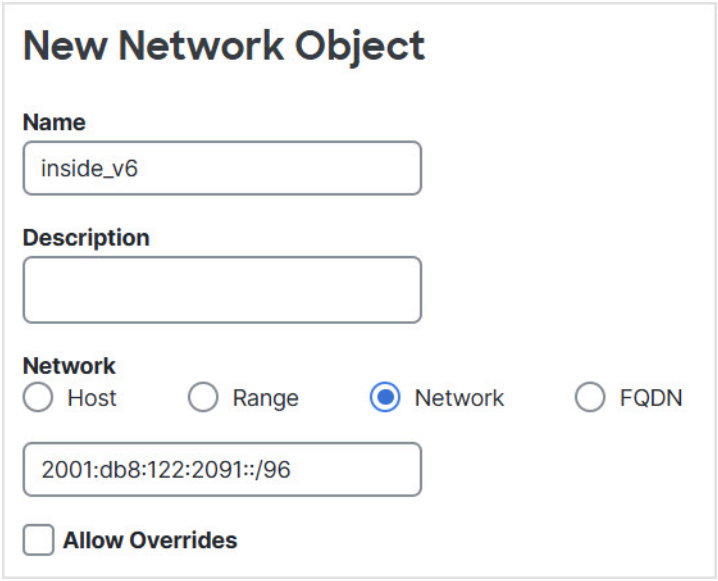
Procedure

Step 1 Create the network object that defines the inside IPv6 network.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the inside IPv6 network.

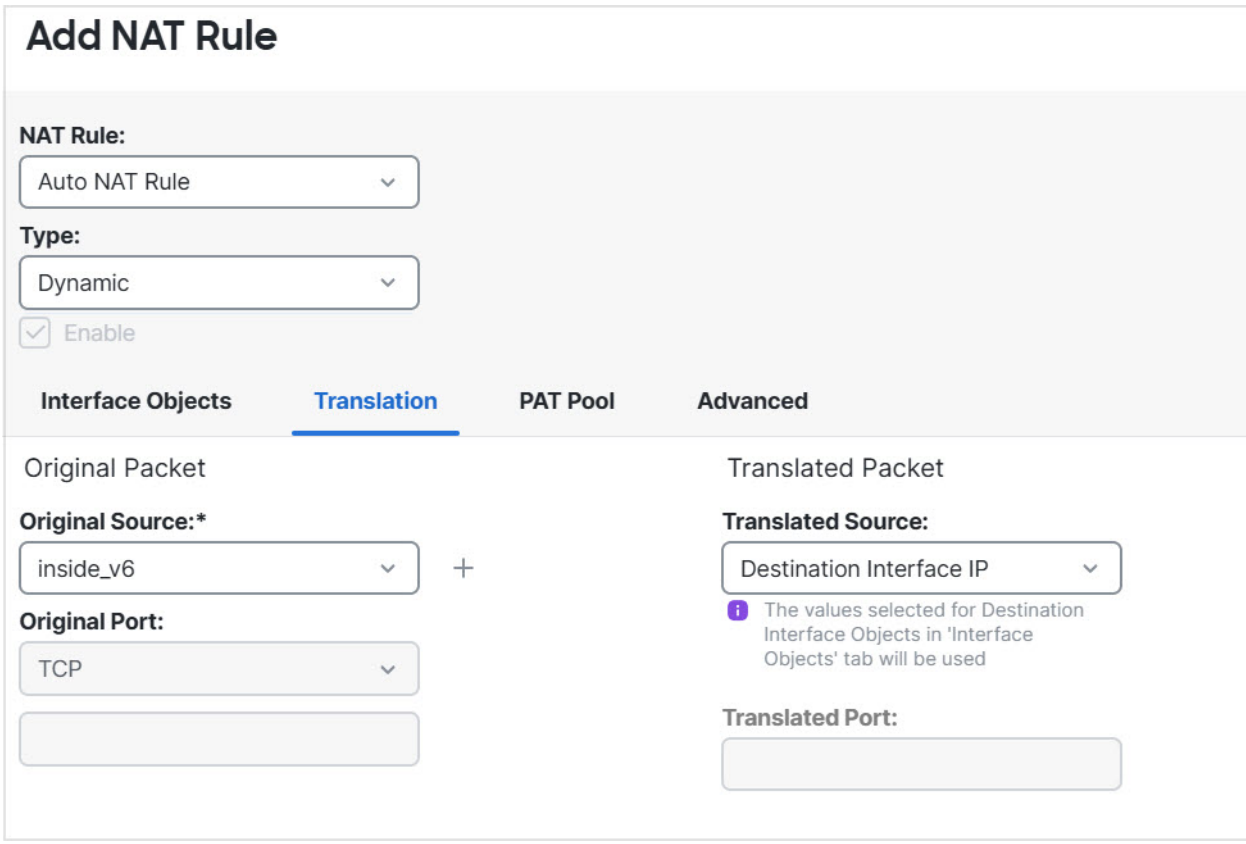
Name the network object (for example, inside_v6) and enter the network address, 2001:db8:122:2091::/96.



c) Click **Save**.

Step 2 Configure the dynamic PAT rule for the inside IPv6 network.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.
b) Click **Add Rule**.
c) Configure the following properties:
   - **NAT Rule** = Auto NAT Rule.
   - **Type** = Dynamic.
d) On **Interface Objects**, configure the following:
   - **Source Interface Objects** = inside.
   - **Destination Interface Objects** = outside.
e) On **Translation**, configure the following:
   - **Original Source** = inside_v6 network object.
   - **Translated Source** = **Destination Interface IP**.
f) On **Advanced**, select **IPv6**, which indicates that the IPv6 address of the destination interface should be used.



g) Click **OK**.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a NAT66 PAT translation to one of the IPv6 global addresses configured for the outside interface.

Monitoring NAT

To monitor and troubleshoot NAT connections, log into the device CLI and use the following commands.

- **show nat** displays the NAT rules and per-rule hit counts. There are additional keywords to show other aspects of NAT.
- **show xlate** displays the actual NAT translations that are currently active.
- **clear xlate** lets you remove an active NAT translation. You might need to remove active translations if you alter NAT rules, because existing connections continue to use the old translation slot until the connection ends. Clearing a translation allows the system to build a new translation for a client on the client's next connection attempt based on your new rules.

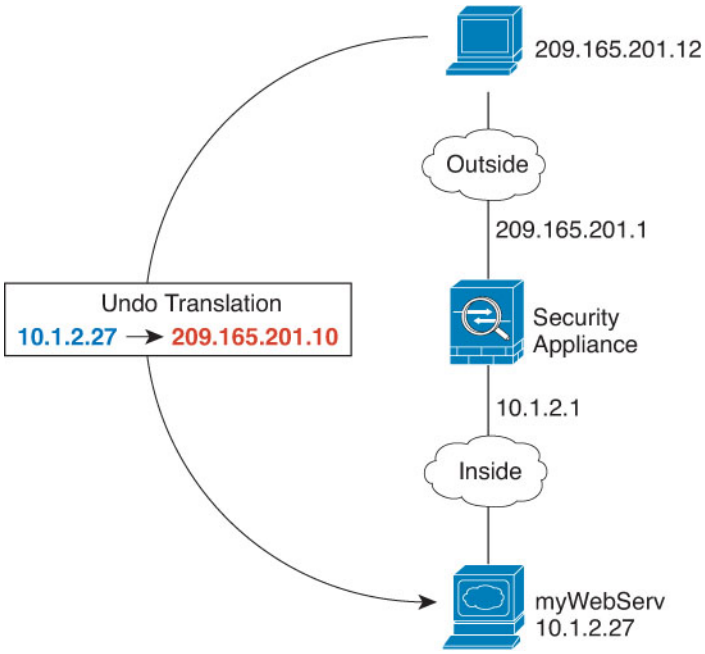
Examples for NAT

The following topics provide examples for configuring NAT.

Providing access to an inside web server (static auto NAT)

This example performs static NAT for an inside web server. The real address is on a private network, so a public address is required. Static NAT is necessary so hosts can initiate traffic to the web server at a fixed address.

Figure 15: Static NAT for an inside web server



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

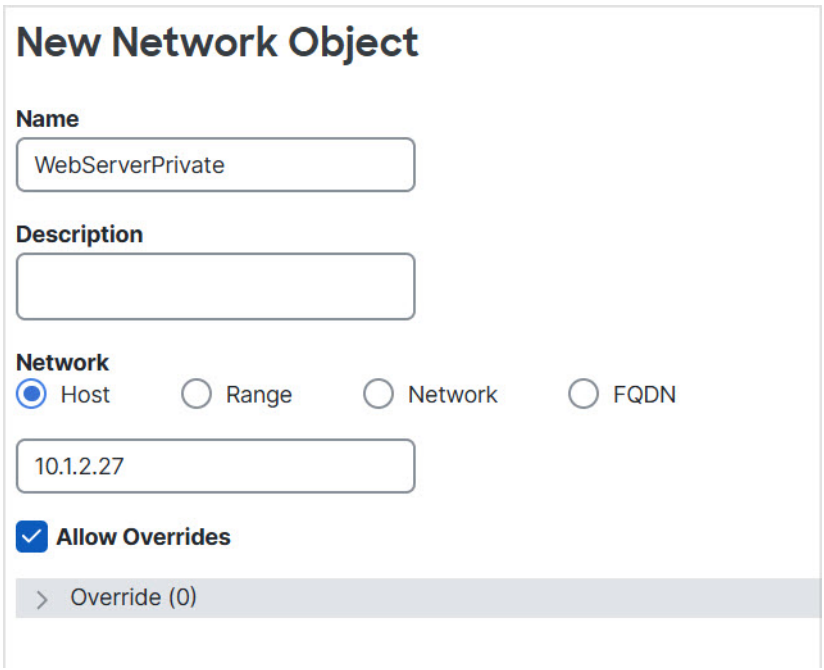
Procedure

Step 1 Create the network objects that define the server's private and public host addresses.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the web server's private address.

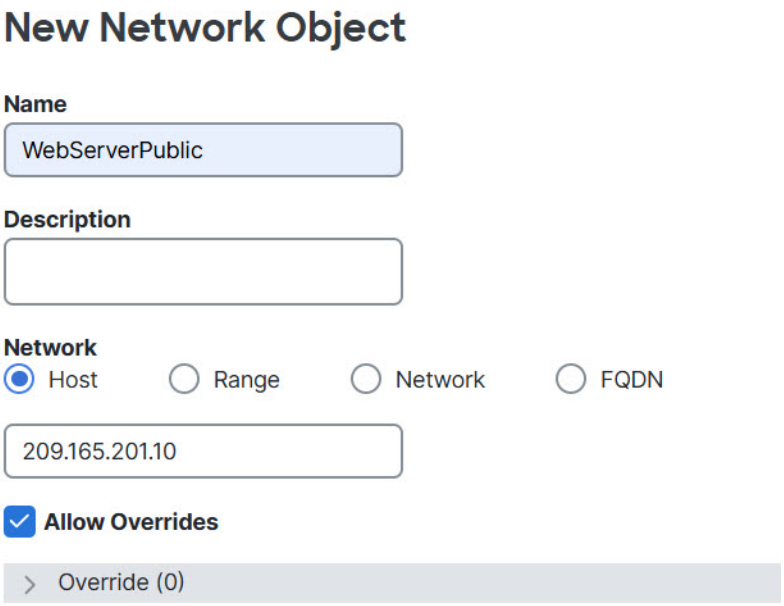
Name the network object (for example, WebServerPrivate) and enter the real host IP address, 10.1.2.27.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the public address.

Name the network object (for example, WebServerPublic) and enter the host address 209.165.201.10.

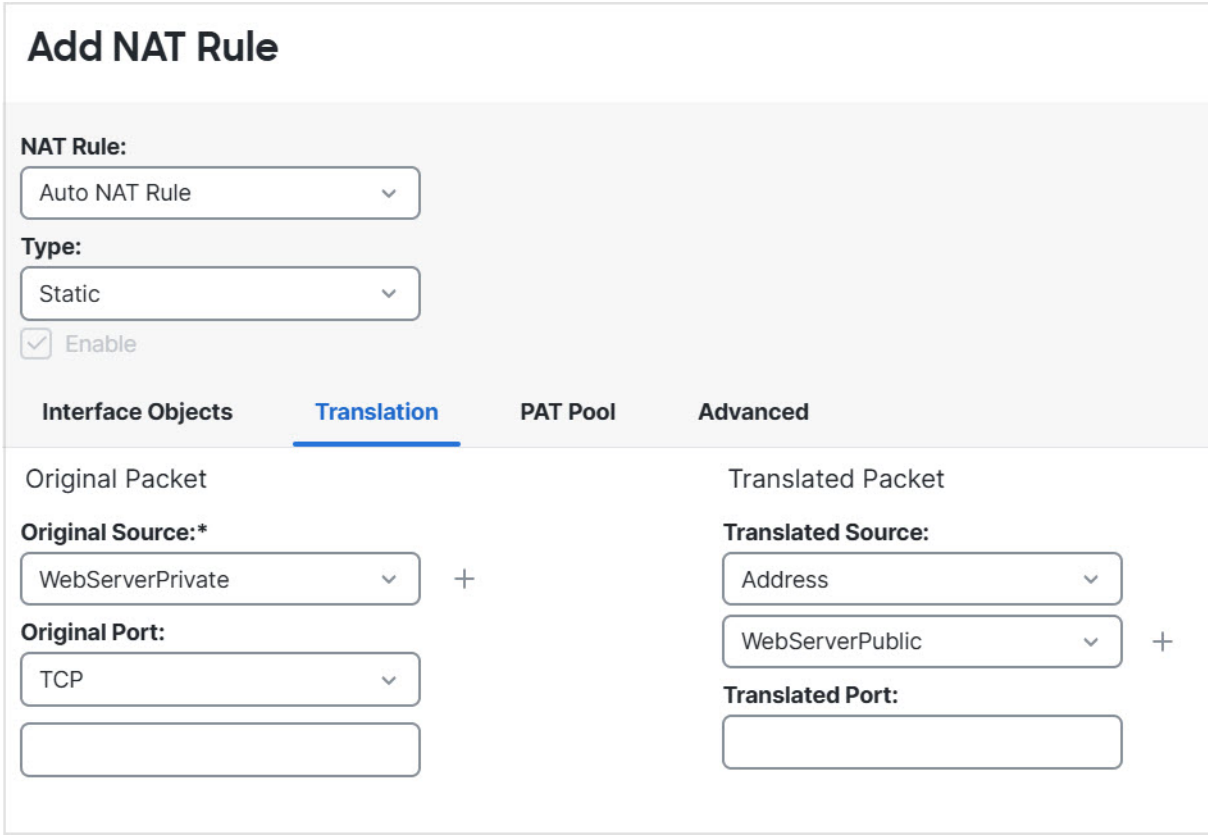


e) Click **Save**.

Step 2 Configure static NAT for the object.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.
c) Configure the following properties:
   - **NAT Rule** = Auto NAT Rule.
   - **Type** = Static.
d) On **Interface Objects**, configure the following:
   - **Source Interface Objects** = inside.
   - **Destination Interface Objects** = outside.
e) On **Translation**, configure the following:
   - **Original Source** = WebServerPrivate network object.
   - **Translated Source** > **Address** = WebServerPublic network object.



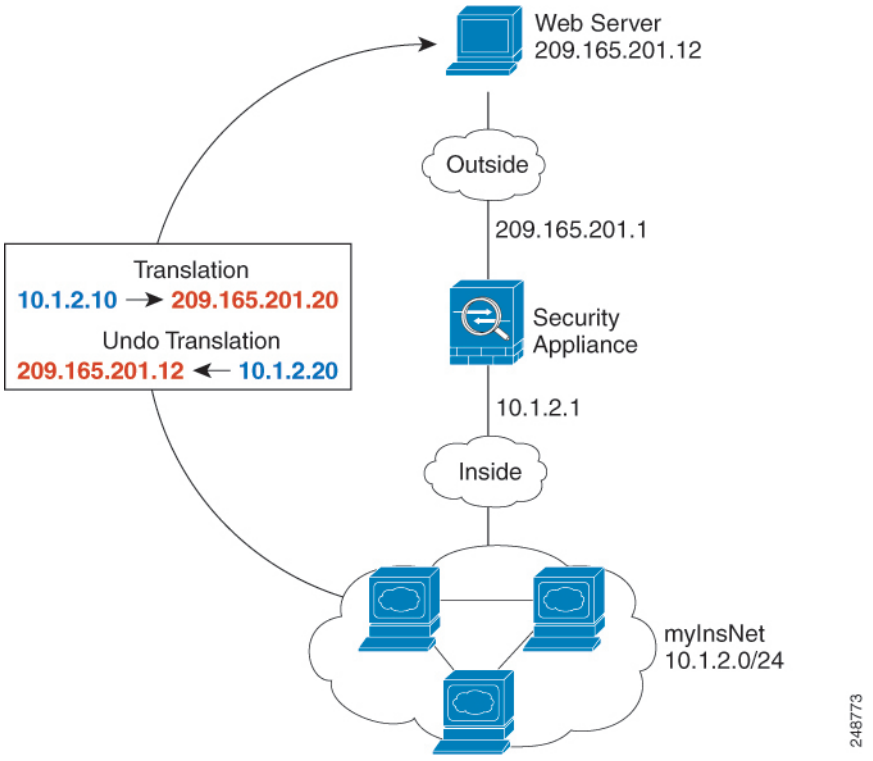
f) Click **Save**.

Step 3 Click **Save** on the NAT rule page.

Dynamic auto NAT for inside hosts and static NAT for an outside web server

The following example configures dynamic NAT for inside users on a private network when they access the outside. Also, when inside users connect to an outside web server, that web server address is translated to an address that appears to be on the inside network.

Figure 16: Dynamic NAT for Inside, Static NAT for Outside Web Server



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select and click **Add Network** > **Add ObjectObjects** > **Object Management** > **Interface**.

Procedure

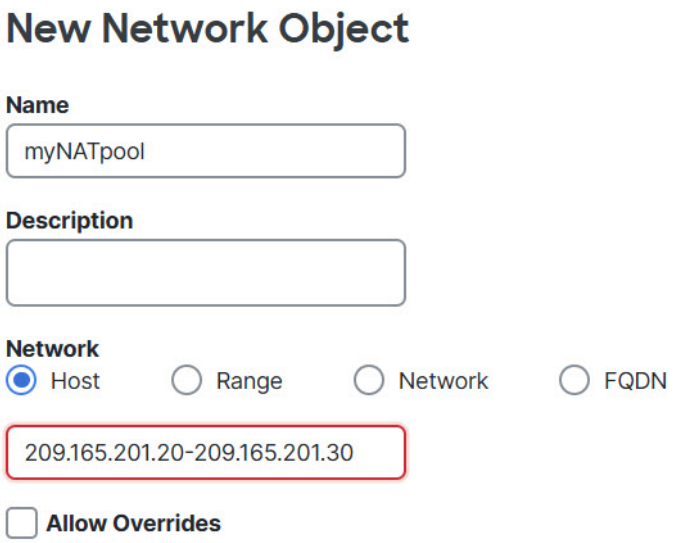
Step 1 Create a network object for the dynamic NAT pool to which you want to translate the inside addresses.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the dynamic NAT pool.

Name the network object (for example, myNATpool) and enter the network range 209.165.201.20-209.165.201.30.

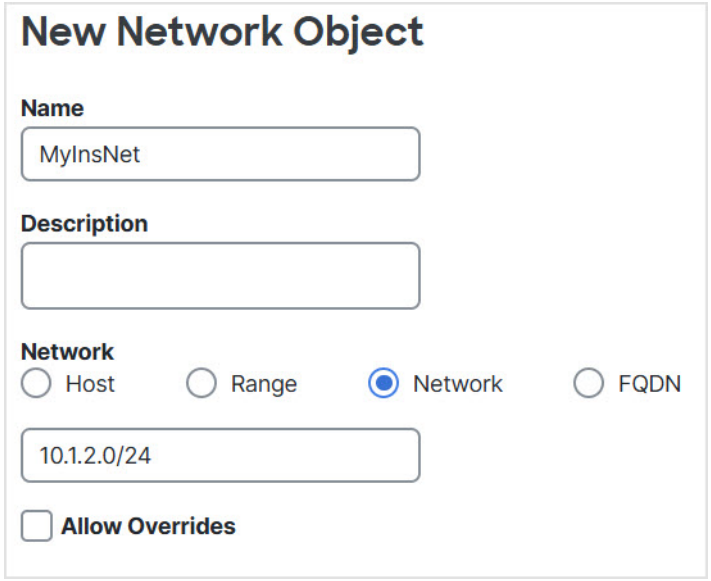


d) Click **Save**.

Step 2 Create a network object for the inside network.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, MyInsNet) and enter the network address 10.1.2.0/24.

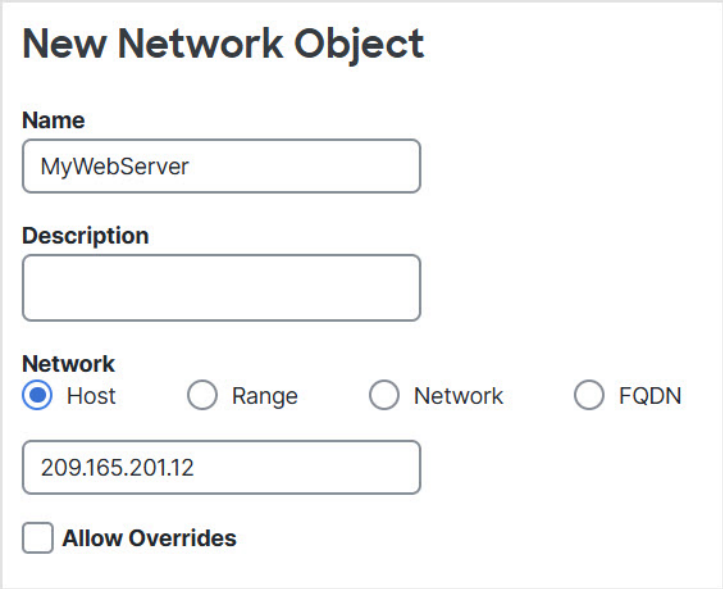


c) Click **Save**.

Step 3 Create a network object for the outside web server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, MyWebServer) and enter the host address 209.165.201.12.

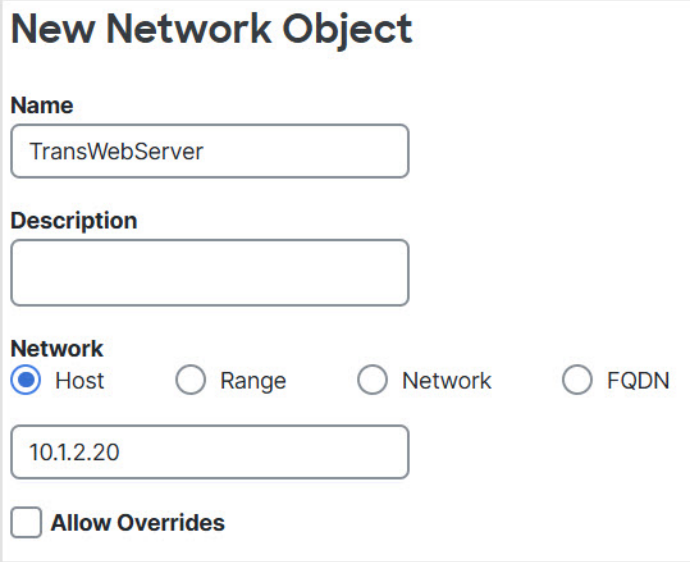


c) Click **Save**.

Step 4 Create a network object for the translated web server address.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, TransWebServer) and enter the host address 10.1.2.20.



c) Click **Save**.

Step 5 Configure dynamic NAT for the inside network using the dynamic NAT pool object.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

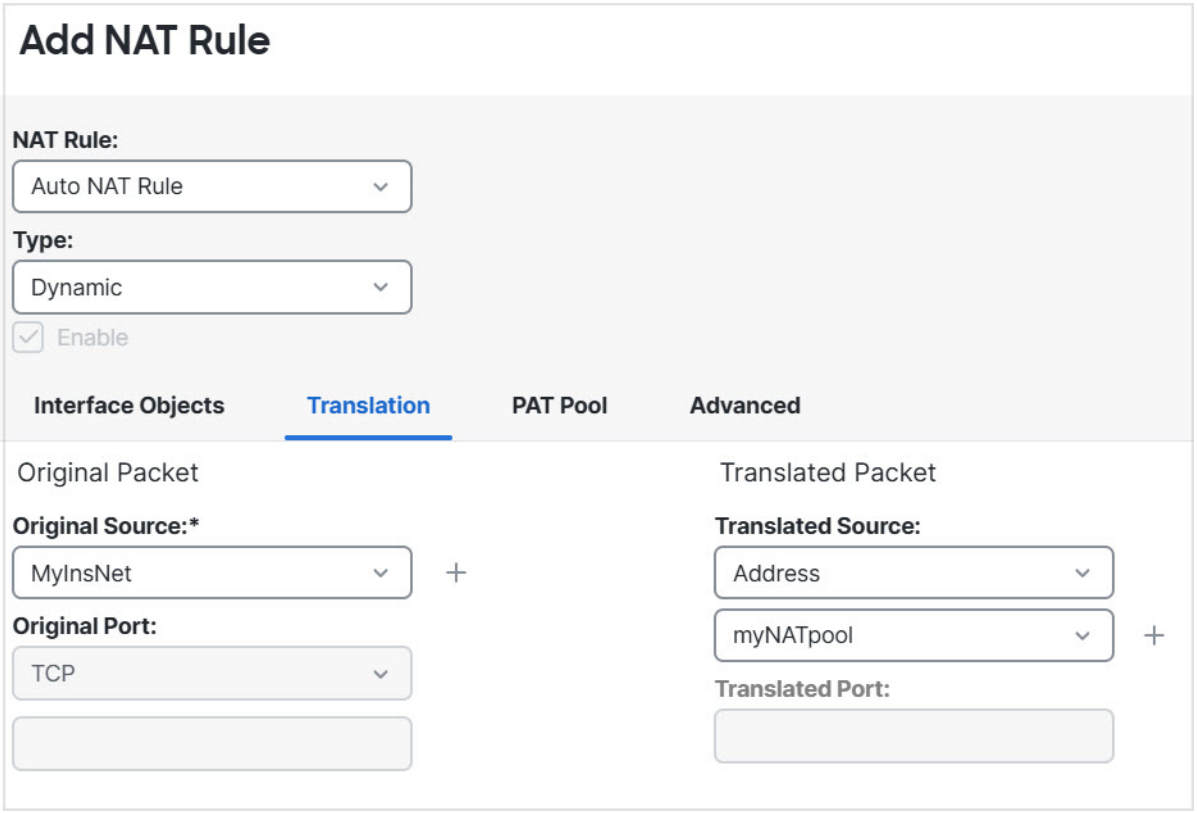
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = myInsNet network object.
- **Translated Source** > **Address**= myNATpool network group.



f) Click **Save**.

Step 6 Configure static NAT for the web server.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

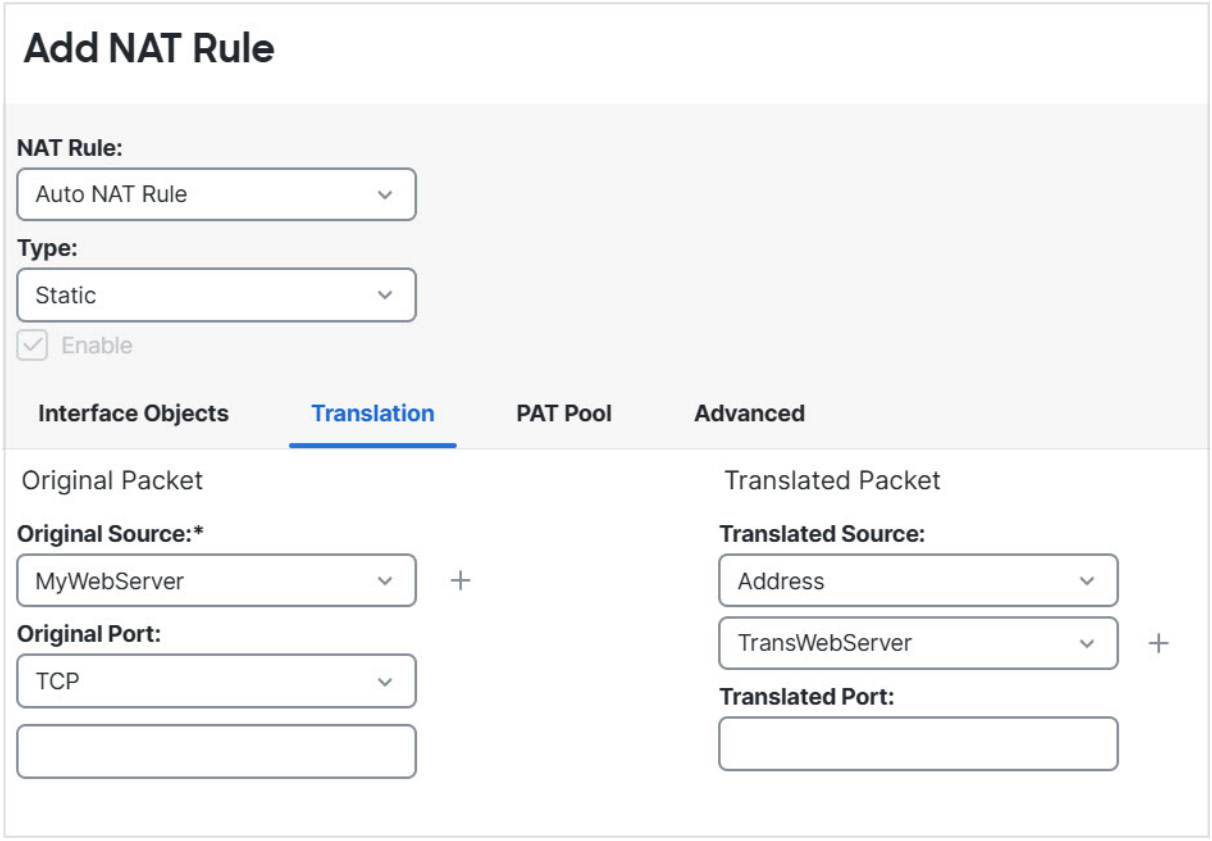
c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.

- **Destination Interface Objects** = inside.

d) On **Translation**, configure the following:

- **Original Source** = myWebServer network object.
- **Translated Source** > **Address**= TransWebServer network object.



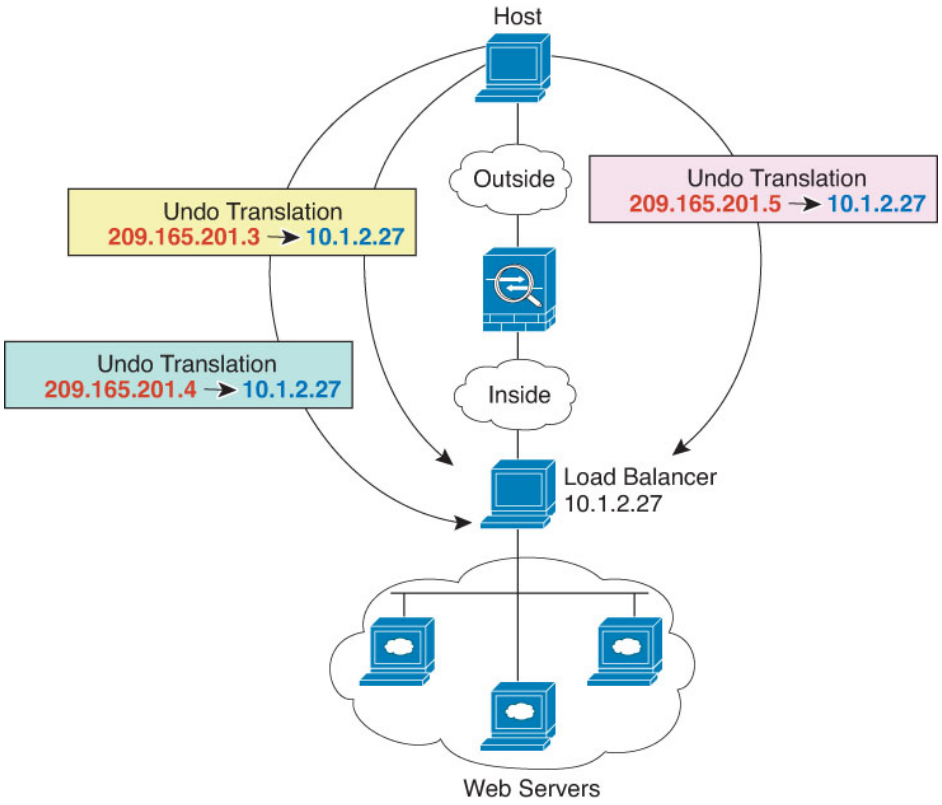
e) Click **Save**.

Step 7 Click **Save** on the NAT rule page.

Configure inside load balancer with multiple mapped addresses

The following example shows an inside load balancer that is translated to multiple IP addresses. When an outside host accesses one of the mapped IP addresses, it is untranslated to the single load balancer address. Depending on the URL requested, it redirects traffic to the correct web server.

Figure 17: Static NAT with One-to-Many for an Inside Load Balancer



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

Procedure

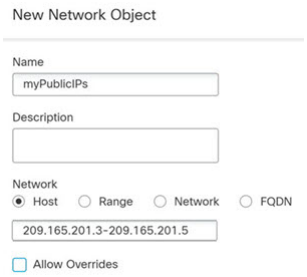
Step 1 Create a network object for the addresses to which you want to map the load balancer.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the addresses.

Name the network object (for example, myPublicIPs) and enter the network range 209.165.201.3-209.165.201.5.

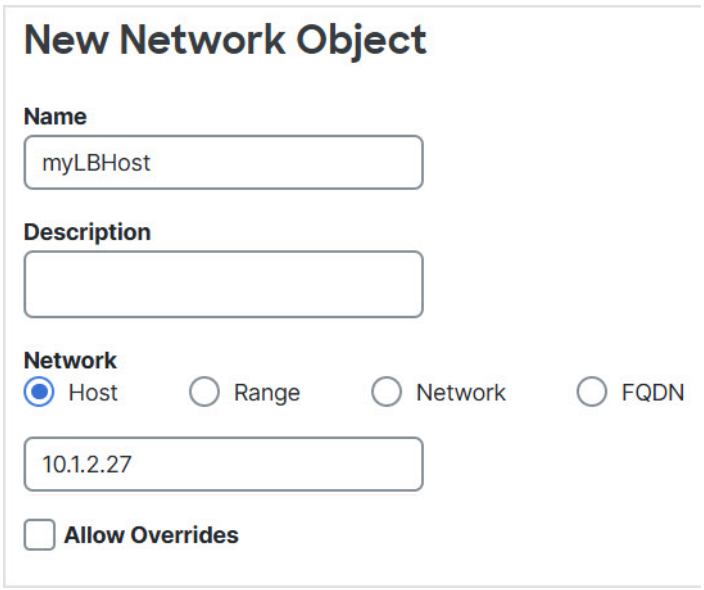


d) Click **Save**.

Step 2 Create a network object for the load balancer.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, myLBHost), enter the host address 10.1.2.27.



c) Click **Save**.

Step 3 Configure static NAT for the load balancer.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

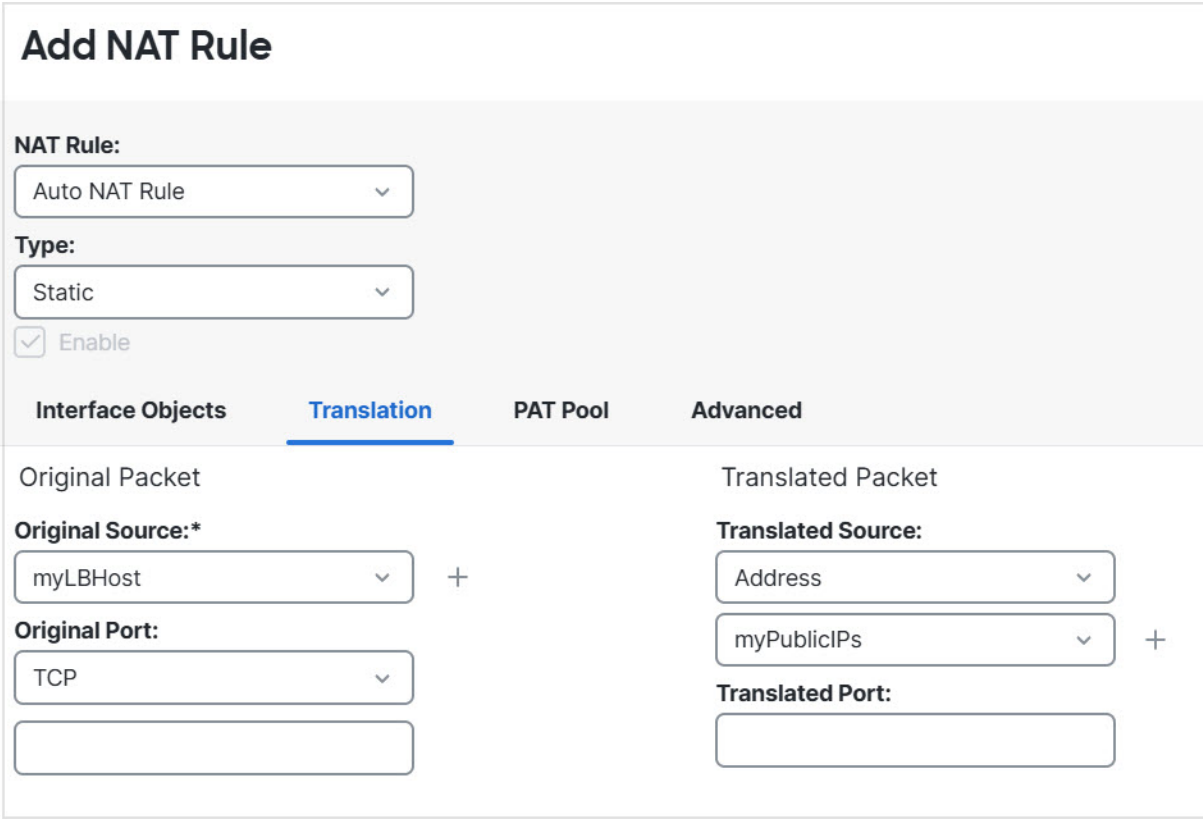
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = myLBHost network object.
- **Translated Source** > **Address** = myPublicIPs network group.



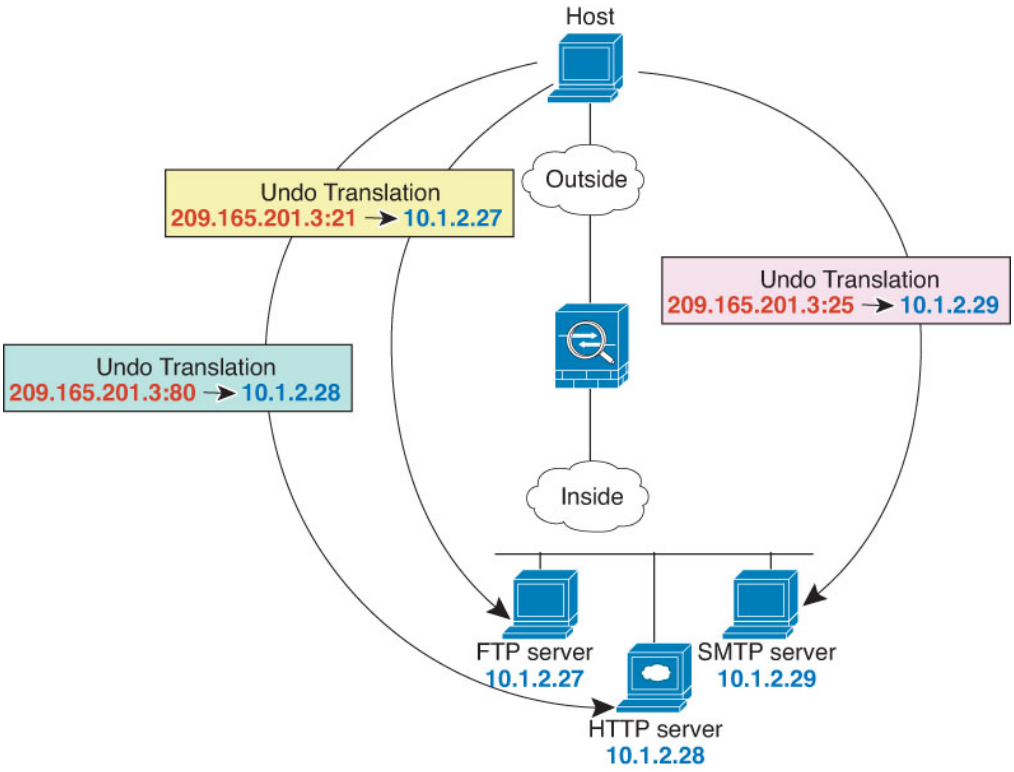
f) Click **Save**.

Step 4 Click **Save** on the NAT rule page.

Single address for FTP, HTTP, and SMTP (static auto NAT with port translation)

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

Figure 18: Static NAT-with-Port-Translation



Before you begin

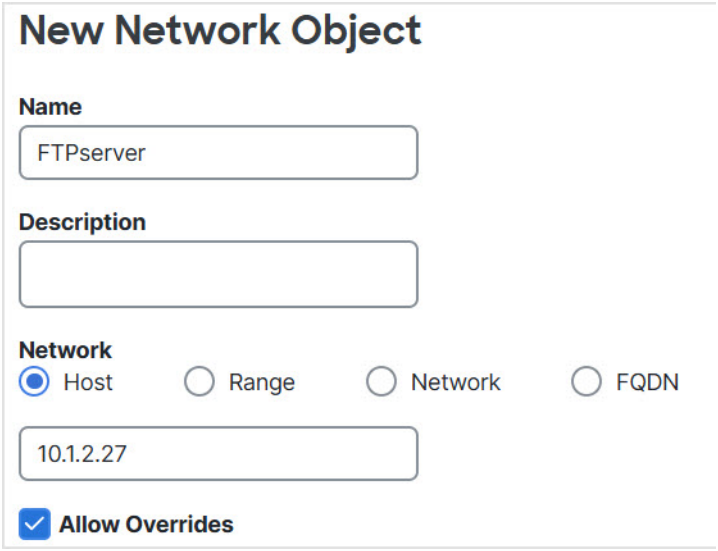
Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select and click **Add Network** > **Add GroupObjects** > **Object Management** > **Interface**.

Procedure

Step 1 Create a network object for the FTP server.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Name the network object (for example, FTPserver), and enter the real IP address for the FTP server, 10.1.2.27.

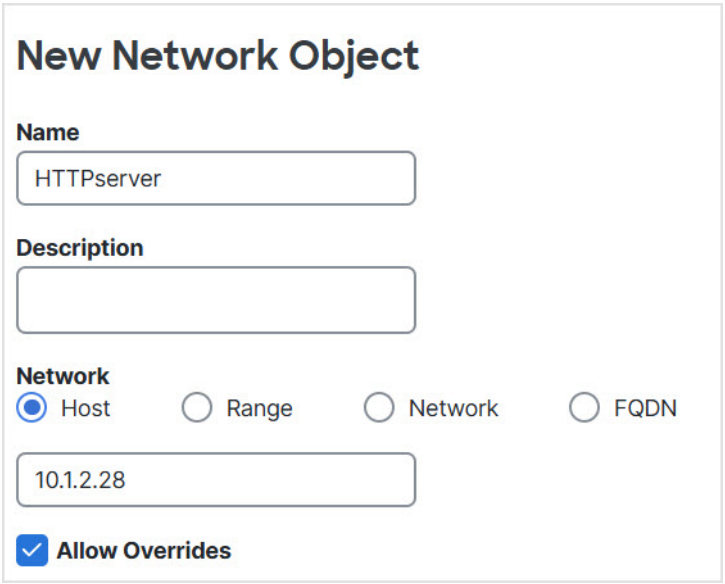


c) Click **Save**.

Step 2 Create a network object for the HTTP server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, HTTPserver), enter the host address 10.1.2.28.

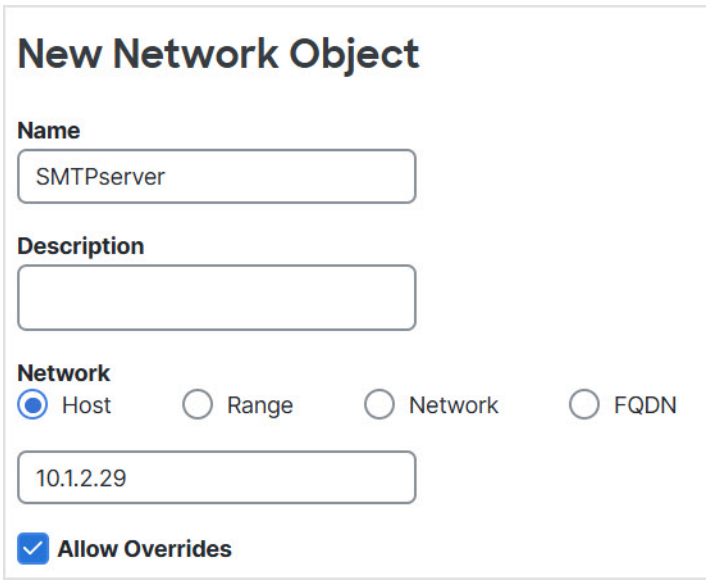


c) Click **Save**.

Step 3 Create a network object for the SMTP server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, SMTPserver), enter the host address 10.1.2.29.

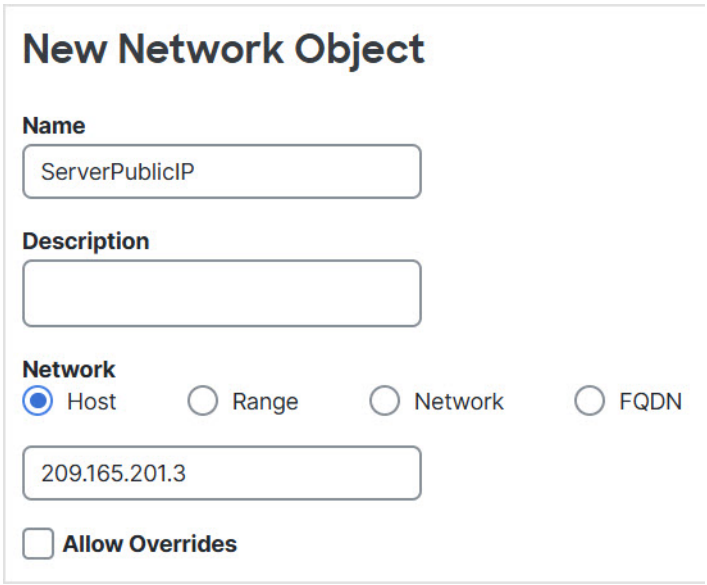


c) Click **Save**.

Step 4 Create a network object for the public IP address used for the three servers.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, ServerPublicIP) and enter the host address 209.165.201.3.



c) Click **Save**.

Step 5 Configure static NAT with port translation for the FTP server, mapping the FTP port to itself.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

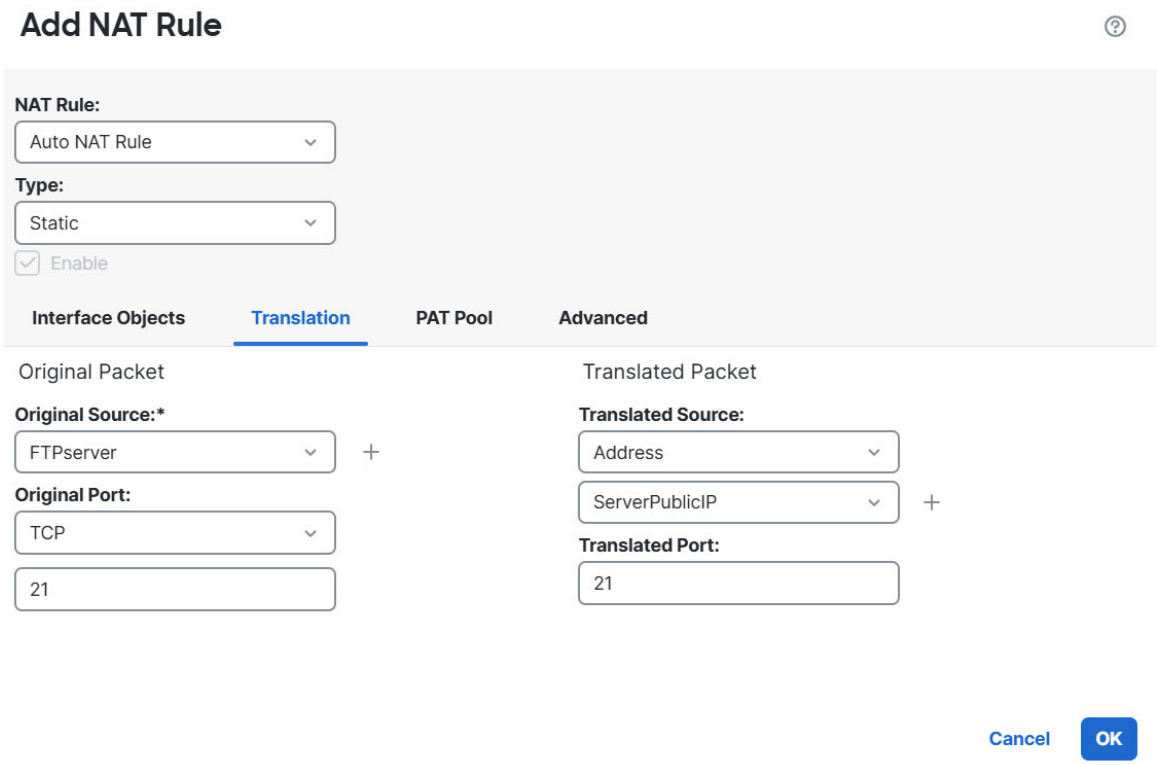
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = FTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 21.
- **Translated Port** = 21.



f) Click **Save**.

Step 6 Configure static NAT with port translation for the HTTP server, mapping the HTTP port to itself.

a) Click **Add Rule**.

b) Configure the following properties:

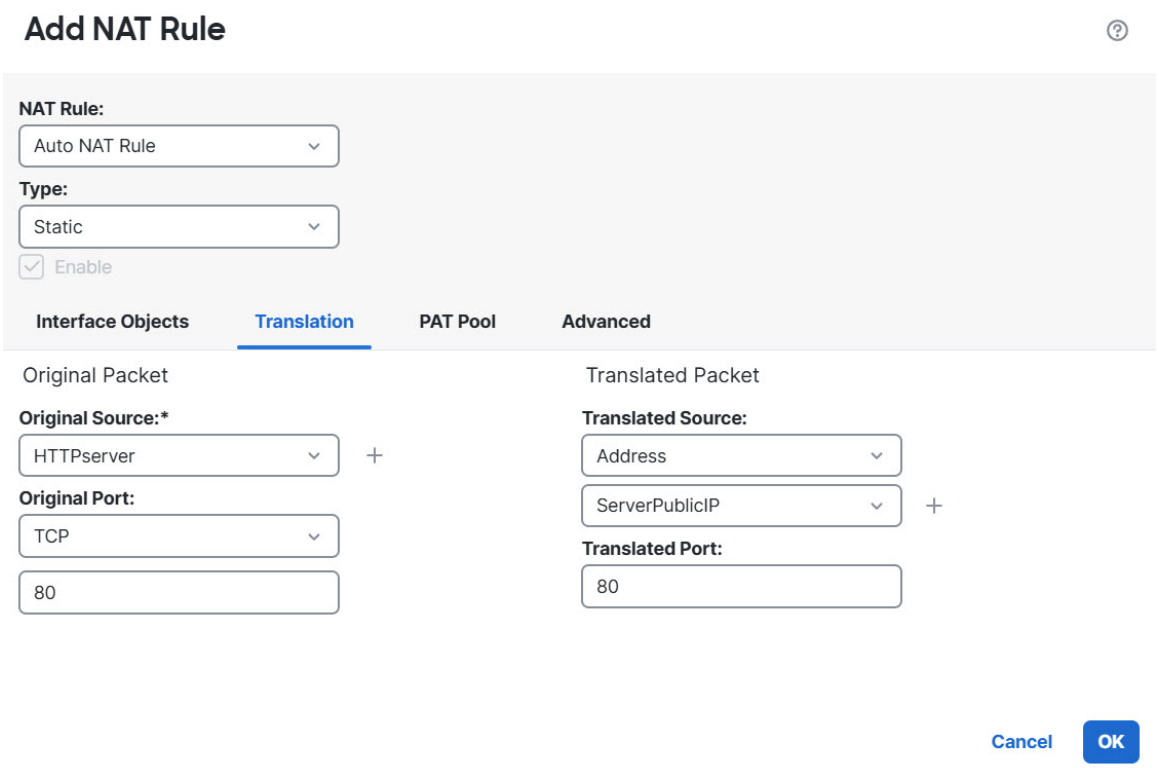
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

d) On **Translation**, configure the following:

- **Original Source** = HTTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 80.
- **Translated Port** = 80.



e) Click **Save**.

Step 7 Configure static NAT with port translation for the SMTP server, mapping the SMTP port to itself.

a) Click **Add Rule**.

b) Configure the following properties:

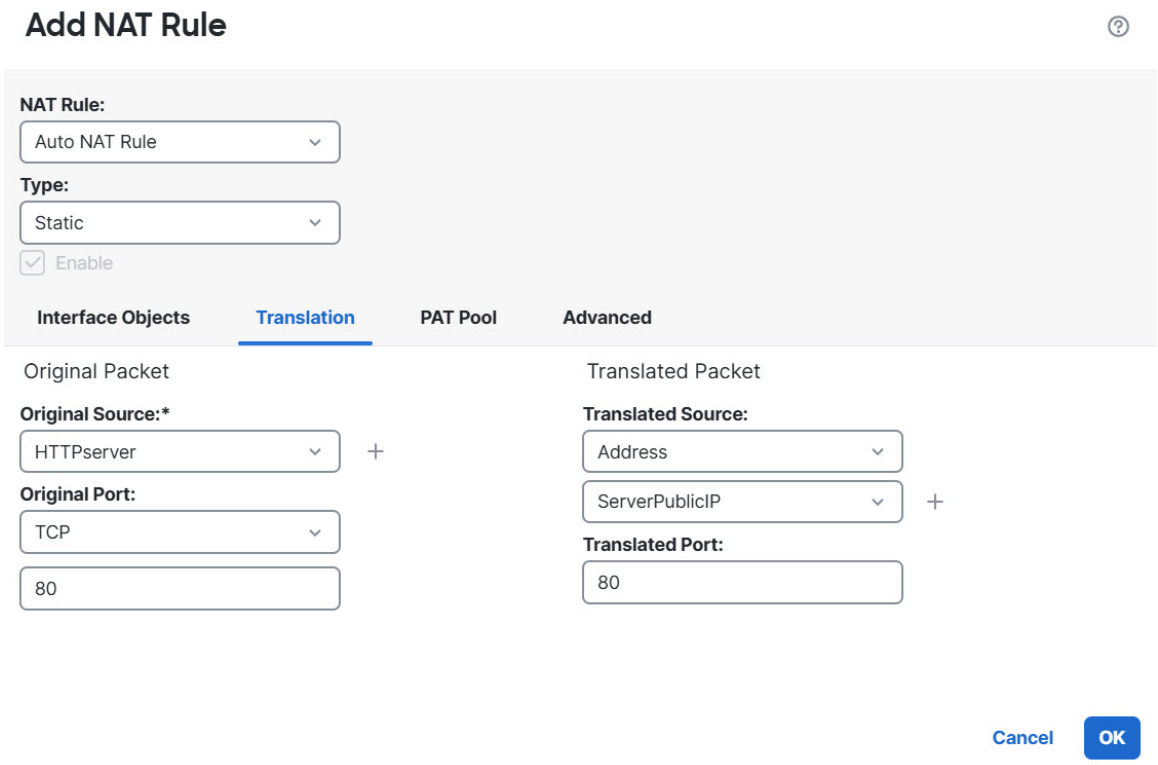
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

d) On **Translation**, configure the following:

- **Original Source** = SMTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 25.
- **Translated Port** = 25.



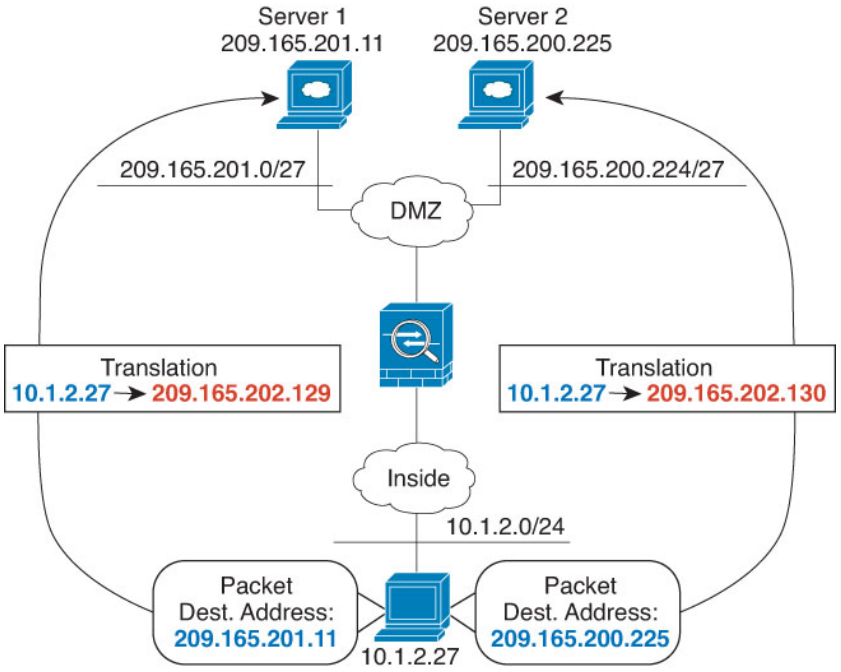
e) Click **Save**.

Step 8 Click **Save** on the NAT rule page.

Different translation depending on the destination (dynamic manual PAT)

The following figure shows a host on the 10.1.2.0/24 network accessing two different servers. When the host accesses the server at 209.165.201.11, the real address is translated to 209.165.202.129:*port*. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130:*port*.

Figure 19: Manual NAT with Different Destination Addresses



Before you begin

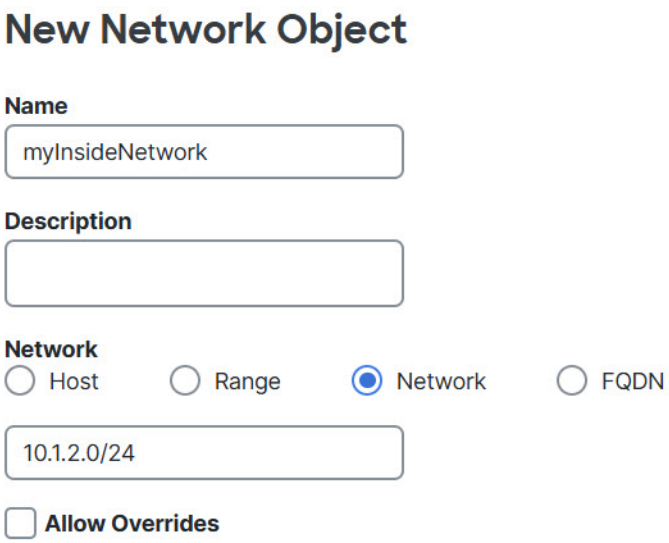
Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **dmz**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

Procedure

Step 1 Create a network object for the inside network.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Name the network object (for example, myInsideNetwork), and enter the real network address, 10.1.2.0/24.

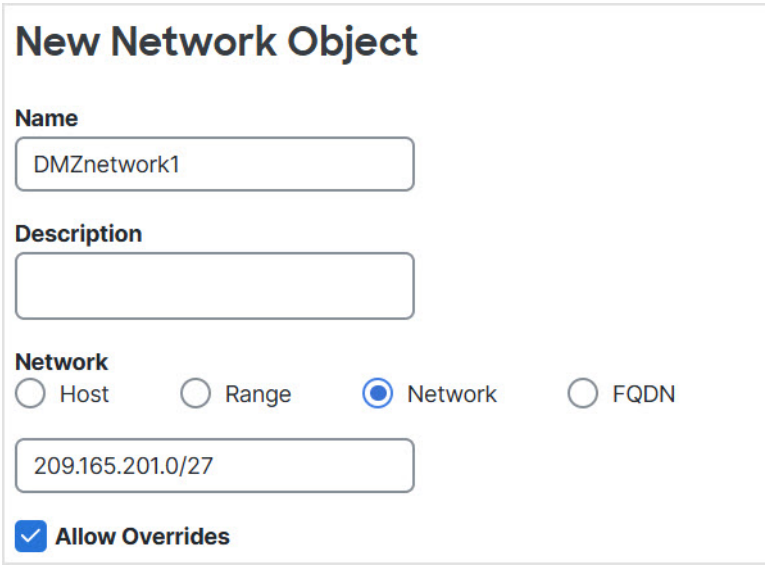


c) Click **Save**.

Step 2 Create a network object for the DMZ network 1.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, DMZnetwork1) and enter the network address 209.165.201.0/27 (subnet mask of 255.255.255.224).

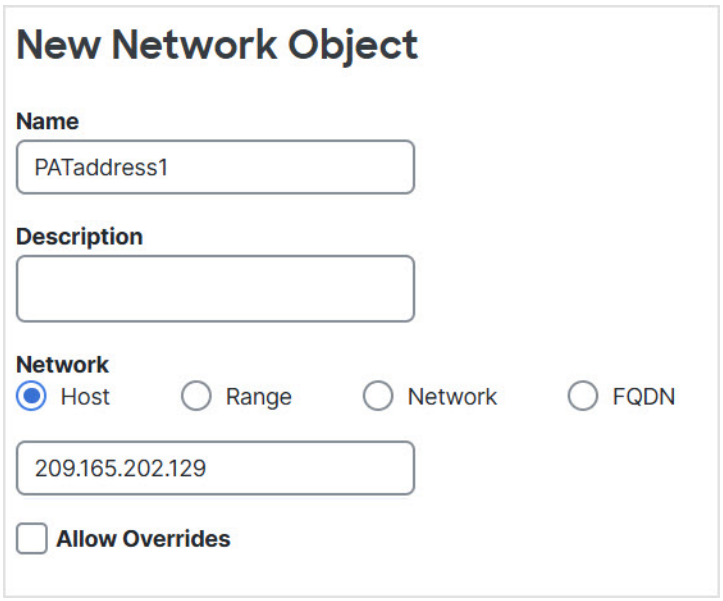


c) Click **Save**.

Step 3 Create a network object for the PAT address for DMZ network 1.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress1) and enter the host address 209.165.202.129.

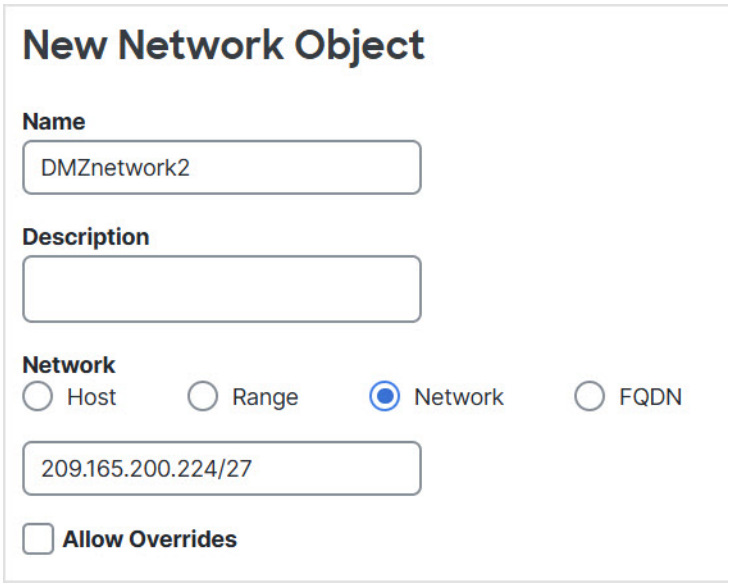


c) Click **Save**.

Step 4 Create a network object for the DMZ network 2.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, DMZnetwork2) and enter the network address 209.165.200.224/27 (subnet mask of 255.255.255.224).

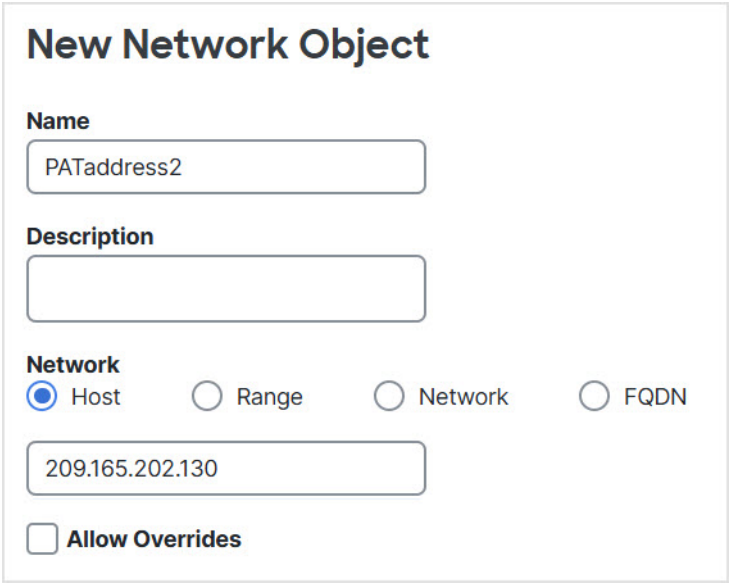


c) Click **Save**.

Step 5 Create a network object for the PAT address for DMZ network 2.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress2) and enter the host address 209.165.202.130.



c) Click **Save**.

Step 6 Configure dynamic manual PAT for DMZ network 1.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

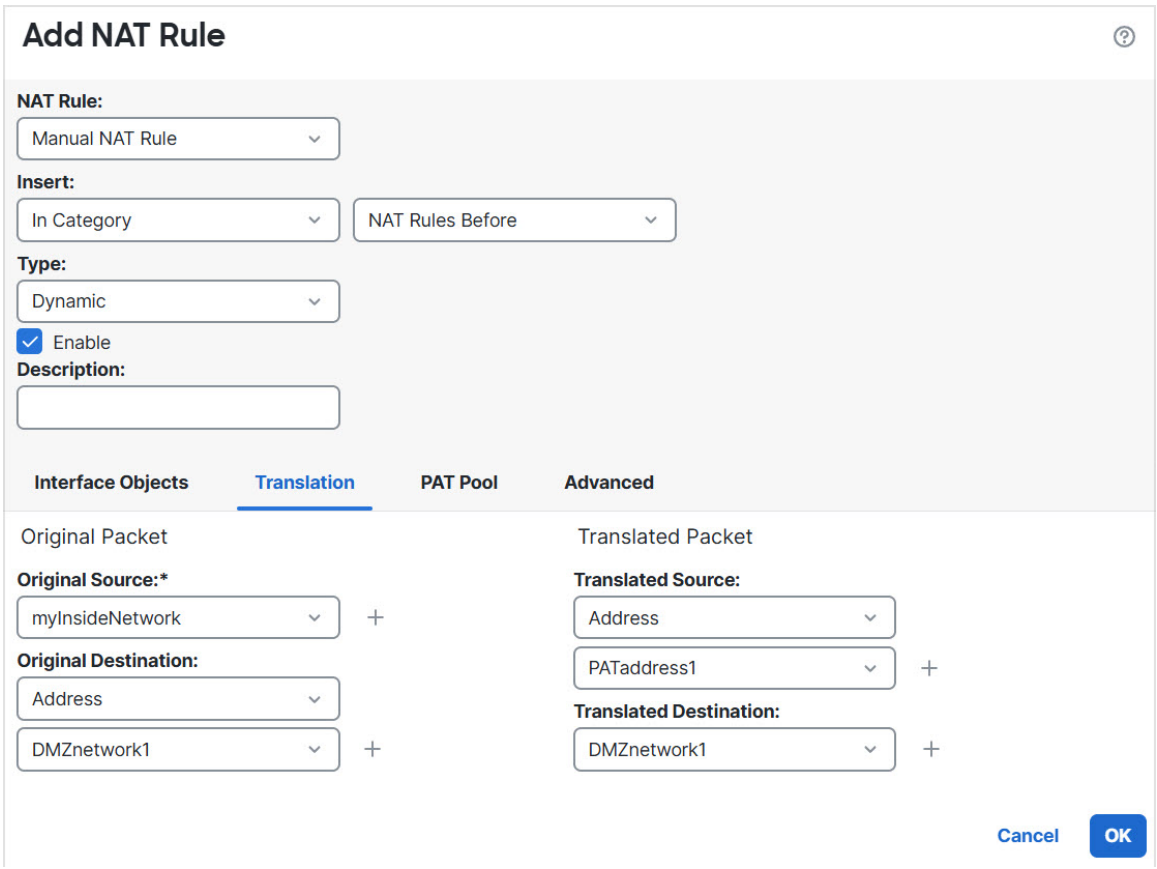
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

e) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress1 network object.
- **Original Destination** > **Address** = DMZnetwork1 network object.
- **Translated Destination** = DMZnetwork1 network object.

Note Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the original and translated destination addresses. Leave all of the port fields blank.



f) Click **Save**.

Step 7 Configure dynamic manual PAT for DMZ network 2.

a) Click **Add Rule**.

b) Configure the following properties:

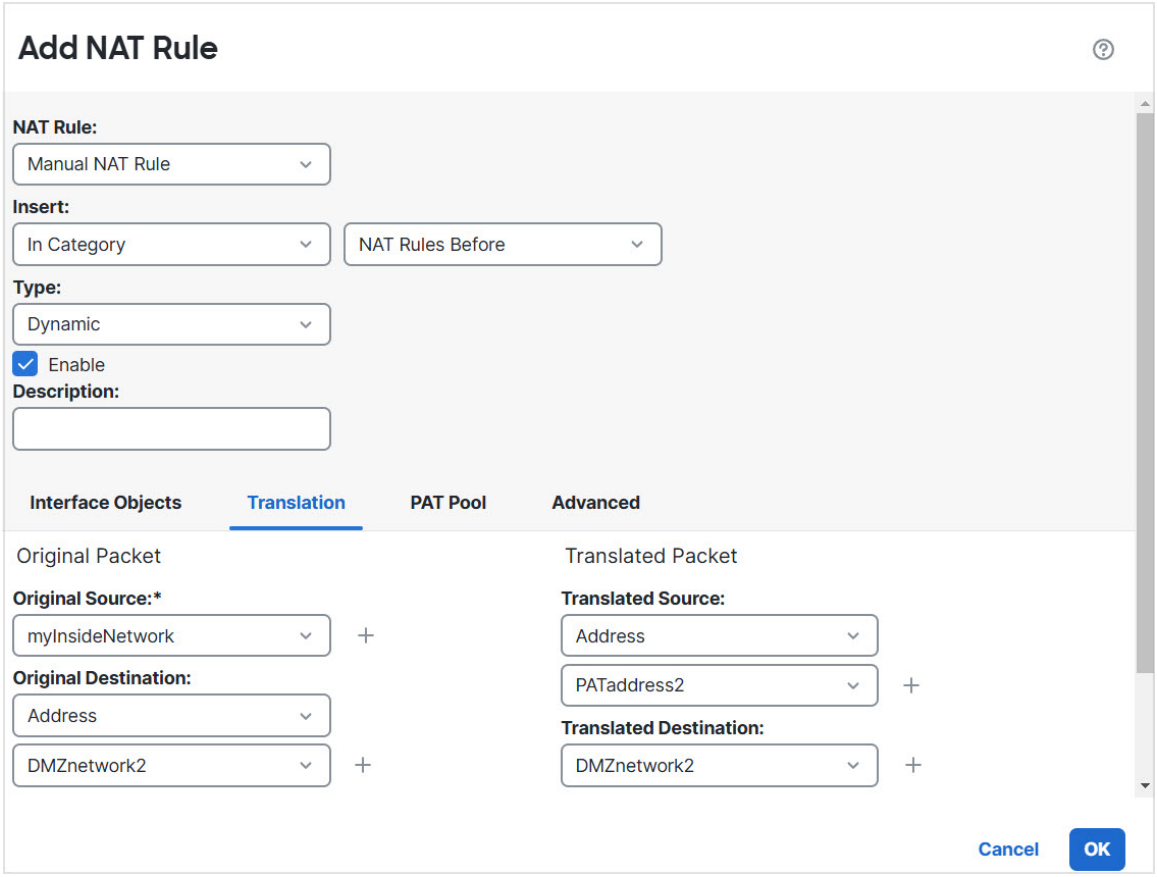
- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

d) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress2 network object.
- **Original Destination** > **Address** = DMZnetwork2 network object.
- **Translated Destination** = DMZnetwork2 network object.



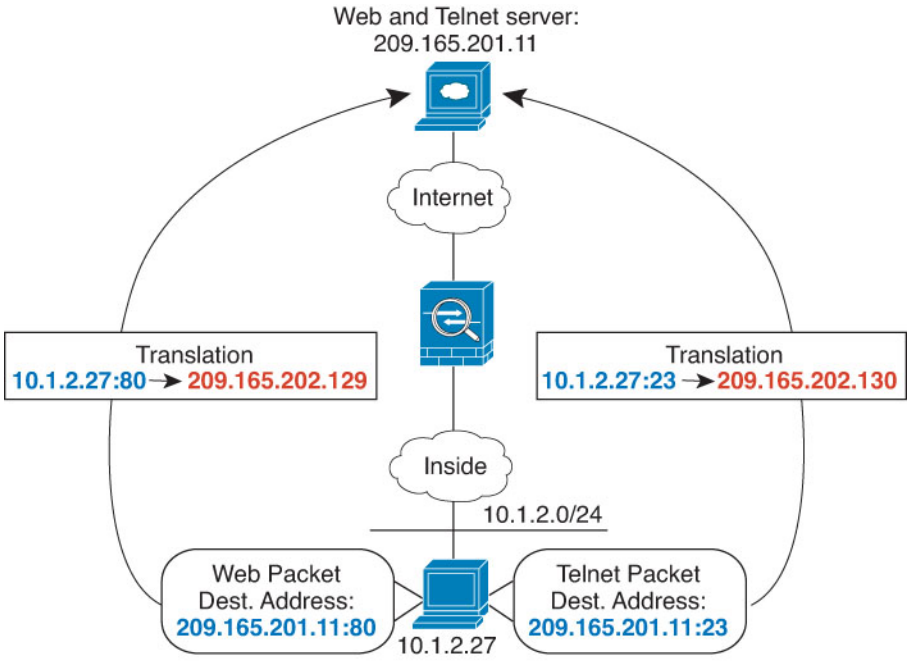
e) Click **Save**.

Step 8 Click **Save** on the NAT rule page.

Configure dynamic manual PAT with different destination ports

This illustration displays the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for Telnet services, the real address is translated to 209.165.202.129: *port*. When the host accesses the same server for web services, the real address is translated to 209.165.202.130: *port*.

Figure 20: Manual NAT with Different Destination Ports



Before you begin

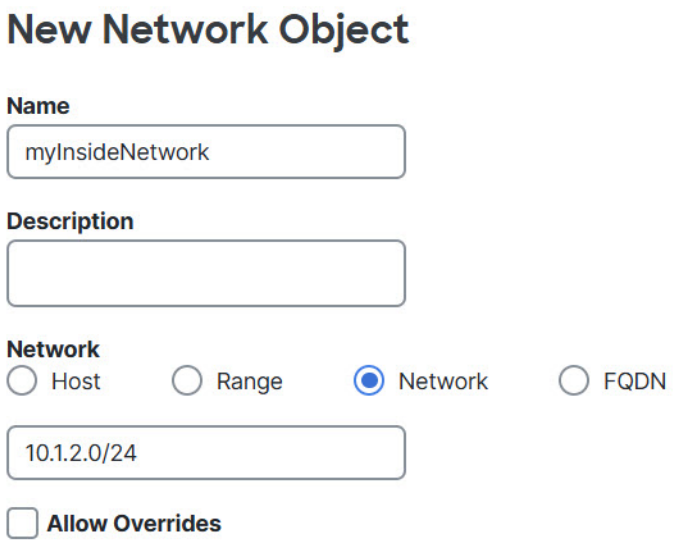
Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **dmz**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

Procedure

Step 1 Create a network object for the inside network.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Name the network object (for example, myInsideNetwork) and enter the real network address, 10.1.2.0/24.

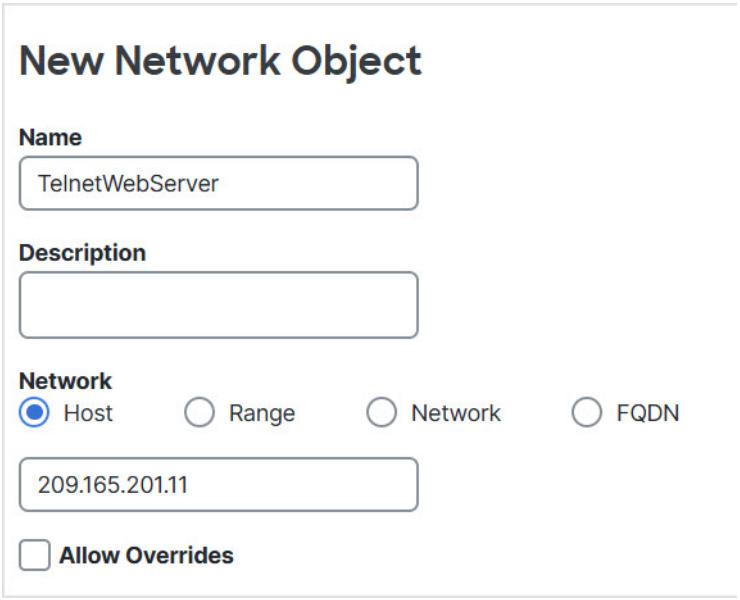


c) Click **Save**.

Step 2 Create a network object for the Telnet/Web server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, TelnetWebServer) and enter the host address 209.165.201.11.

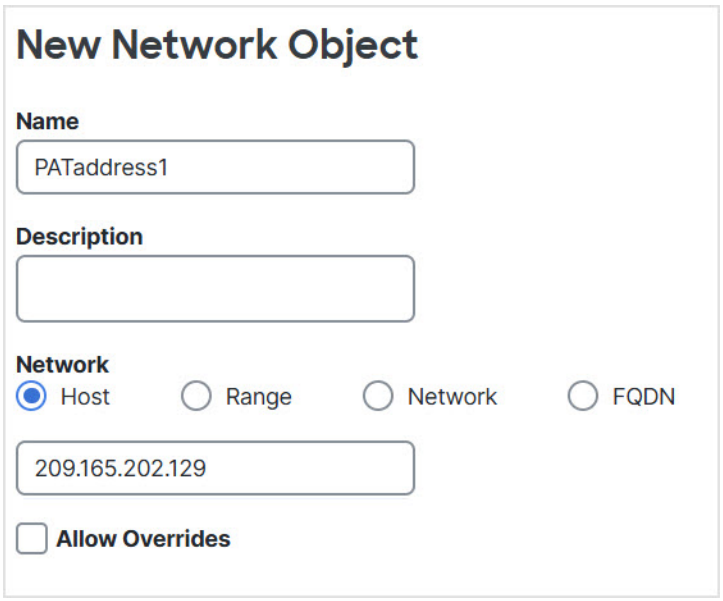


c) Click **Save**.

Step 3 Create a network object for the PAT address when using Telnet.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress1) and enter the host address 209.165.202.129.

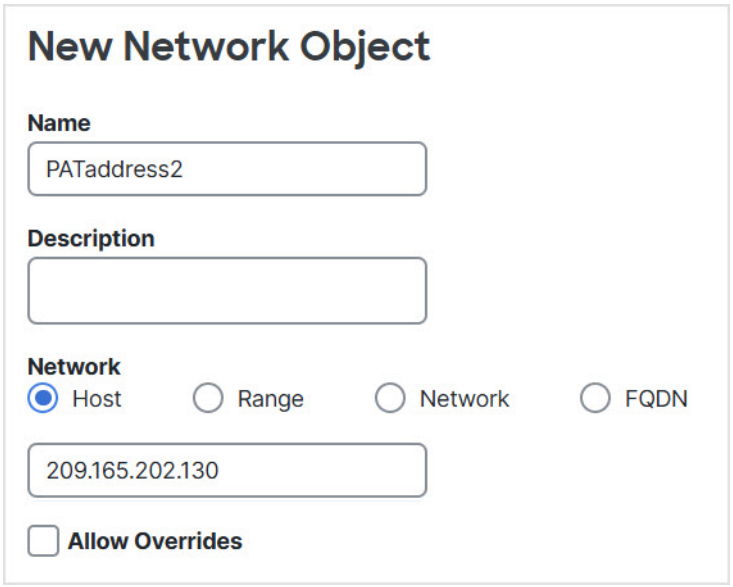


c) Click **Save**.

Step 4 Create a network object for the PAT address when using HTTP.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress2) and enter the host address 209.165.202.130.



c) Click **Save**.

Step 5 Configure dynamic manual PAT for Telnet access.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.

- **Type** = Dynamic.

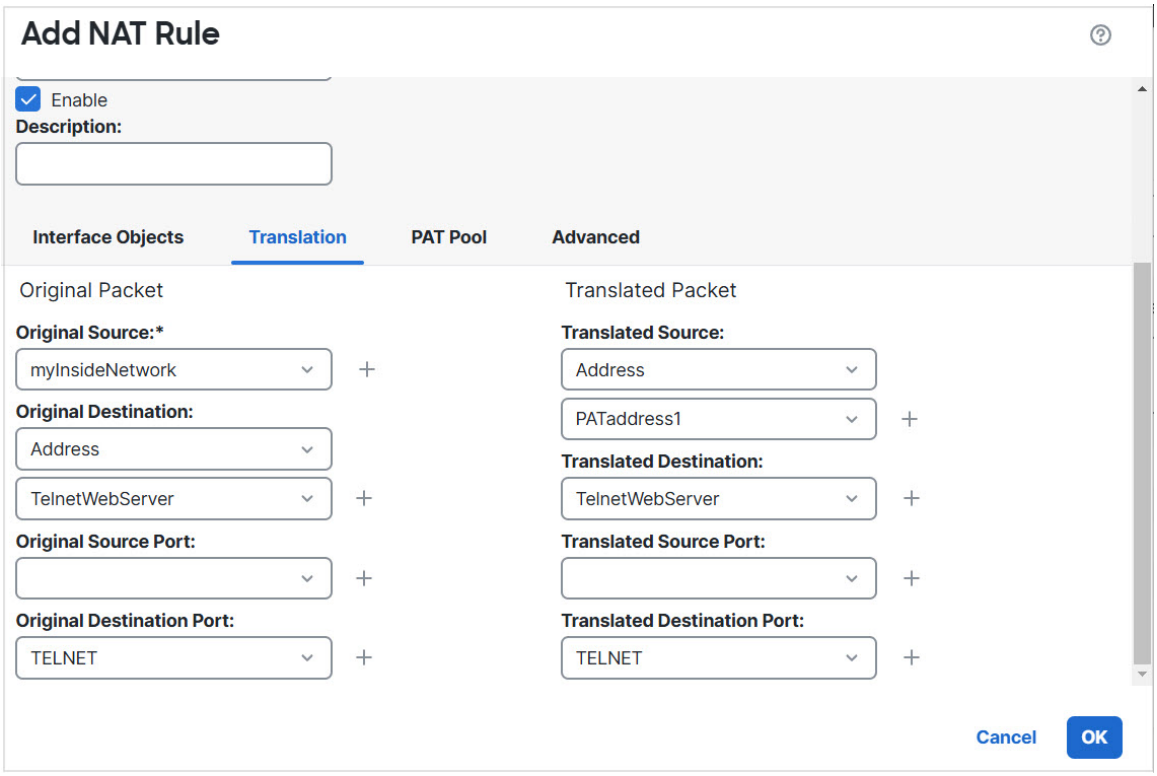
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

e) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address** = PATaddress1 network object.
- **Original Destination** > **Address** = TelnetWebServer network object.
- **Translated Destination** = TelnetWebServer network object.
- **Original Destination Port** = TELNET port object (system-defined).
- **Translated Destination Port** = TELNET port object (system-defined).

Note Because you do not want to translate the destination address or port, you need to configure identity NAT for them by specifying the same address for the original and translated destination addresses, and the same port for the original and translated port.



f) Click **Save**.

Step 6 Configure dynamic manual PAT for web access.

a) Click **Add Rule**.

b) Configure the following properties:

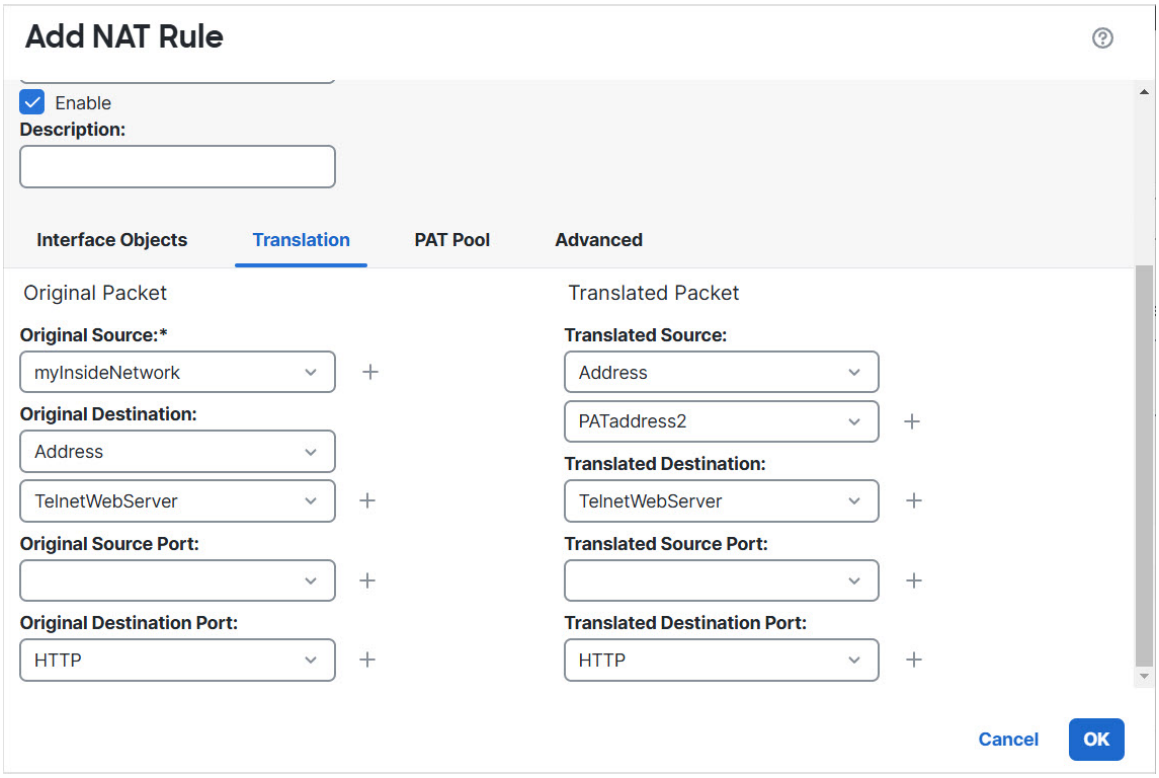
- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

d) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address** = PATaddress2 network object.
- **Original Destination** > **Address** = TelnetWebServer network object.
- **Translated Destination** = TelnetWebServer network object.
- **Original Destination Port** = HTTP port object (system-defined).
- **Translated Destination Port** = HTTP port object (system-defined).



e) Click **Save**.

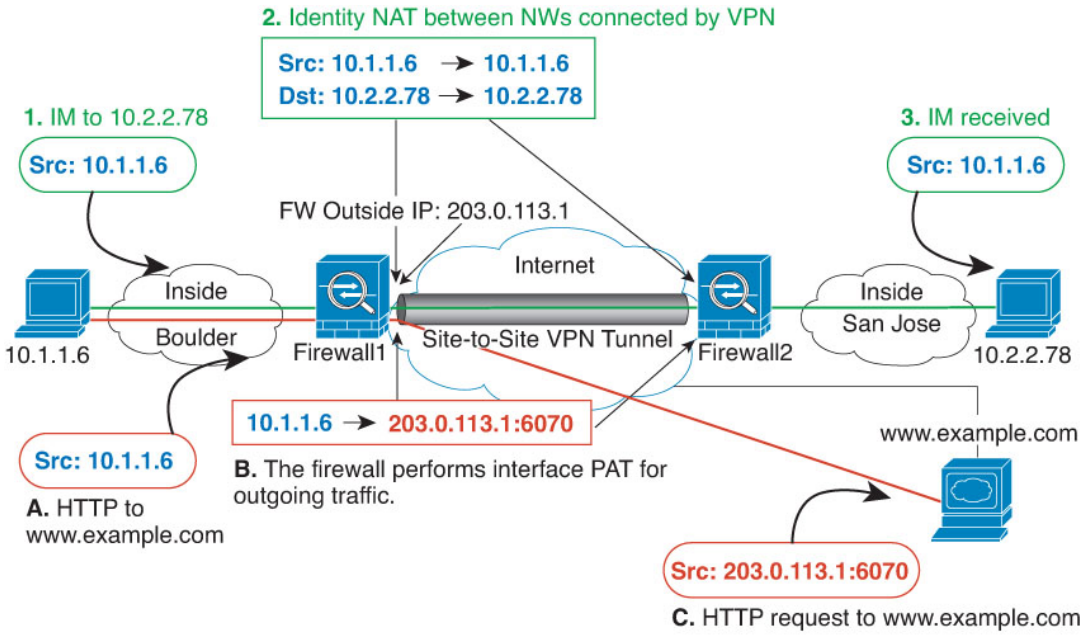
Step 7 Click **Save** on the NAT rule page.

NAT and site-to-site VPN

When you create a policy-based site-to-site VPN using the management center VPN wizard (**Device > Site To Site**), you can select the NAT Exempt option to create the rules automatically. You can view the NAT exemptions for a device in the NAT policy page (**Device > NAT > NAT Exemptions**). If you do not want to configure NAT Exempt in the VPN wizard, you can use the following procedure for NAT exemption.

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

Figure 21: Interface PAT and Identity NAT for Site-to-Site VPN



This example explains the configuration for Firewall1 (Boulder).

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the devices in the VPN. In this example, we will assume the interface objects are security zones named **inside-boulder** and **outside-boulder** for the Firewall1 (Boulder) interfaces. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

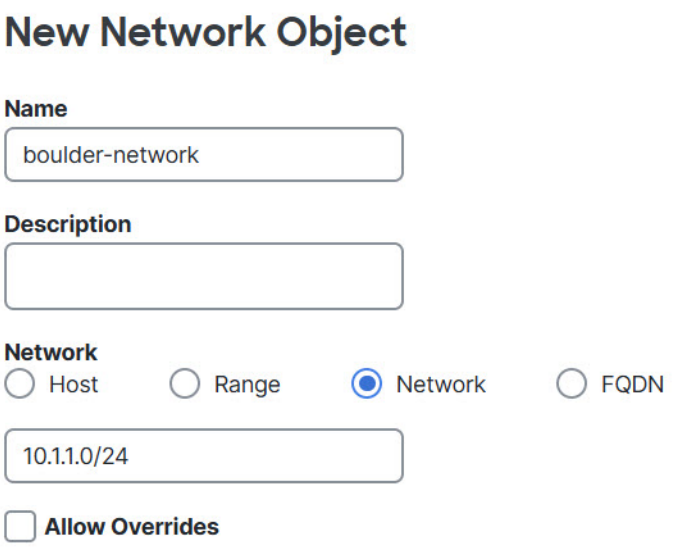
Procedure

Step 1 Create the objects to define the various networks.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Identify the Boulder inside network.

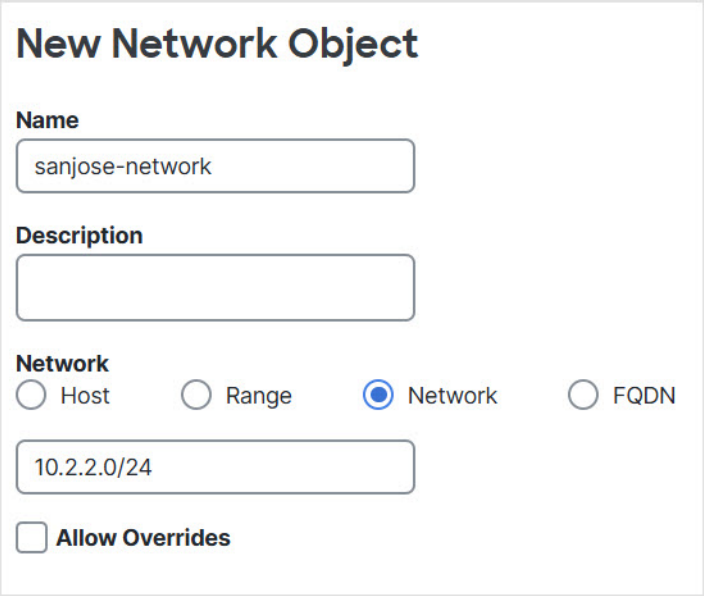
Name the network object (for example, boulder-network) and enter the network address, 10.1.1.0/24.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the inside San Jose network.

Name the network object (for example, sanjose-network) and enter the network address 10.2.2.0/24.



e) Click **Save**.

Step 2 Configure manual identity NAT for the Boulder network when going over the VPN to San Jose on Firewall1 (Boulder).

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.
b) Click **Add Rule**.
c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

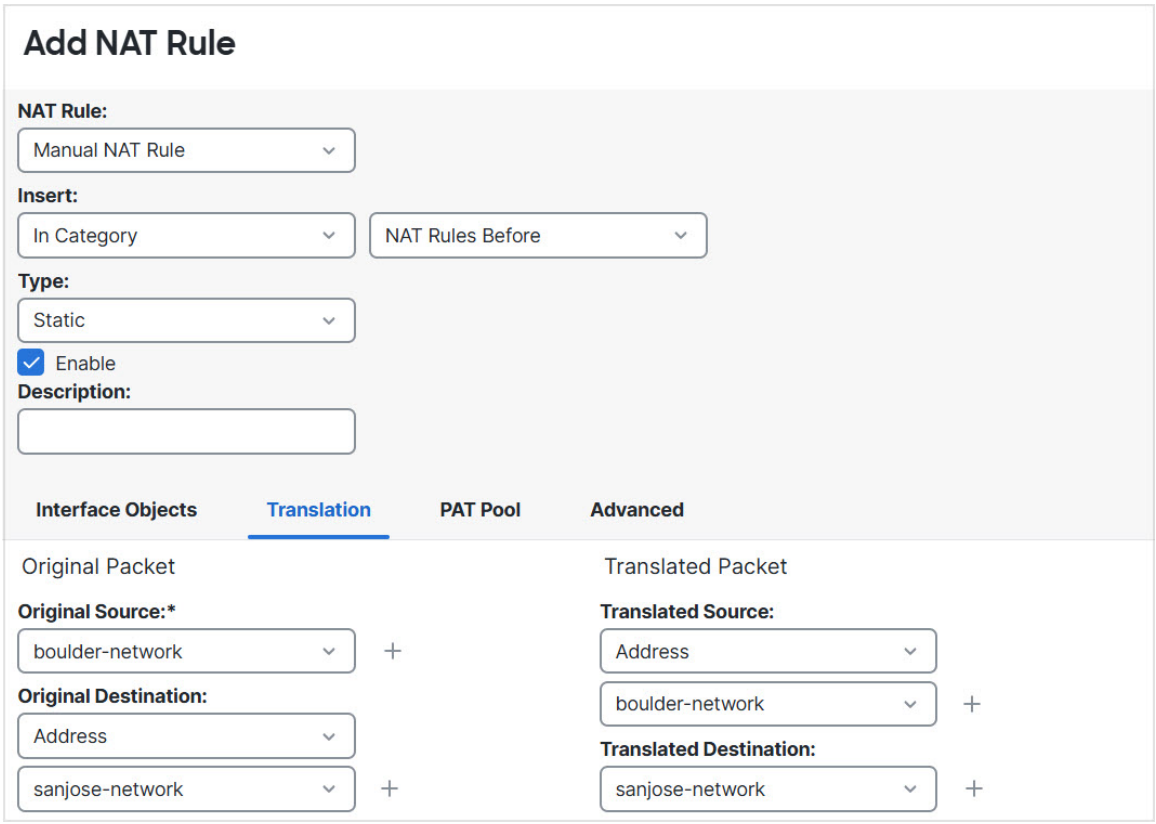
- **Source Interface Objects** = inside-boulder.
- **Destination Interface Objects** = outside-boulder.

e) On **Translation**, configure the following:

- **Original Source** = boulder-network object.
- **Translated Source** > **Address** = boulder-network object.
- **Original Destination** > **Address** = sanjose-network object.
- **Translated Destination** = sanjose-network object.

Note Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the original and translated destination addresses. Leave all of the port fields blank. This rule configures identity NAT for both source and destination.

f) On **Advanced**, select **Do not proxy ARP on Destination interface**.



g) Click **Save**.

Step 3 Configure manual dynamic interface PAT when going to the Internet for the inside Boulder network on Firewall1 (Boulder).

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.
- **Insert Rule** = any position after the first rule. Because this rule will apply to any destination address, the rule that uses sanjose-network as the destination must come before this rule, or the sanjose-network rule will never be matched. The default is to place new manual NAT rules at the end of the "NAT Rules Before Auto NAT" section.

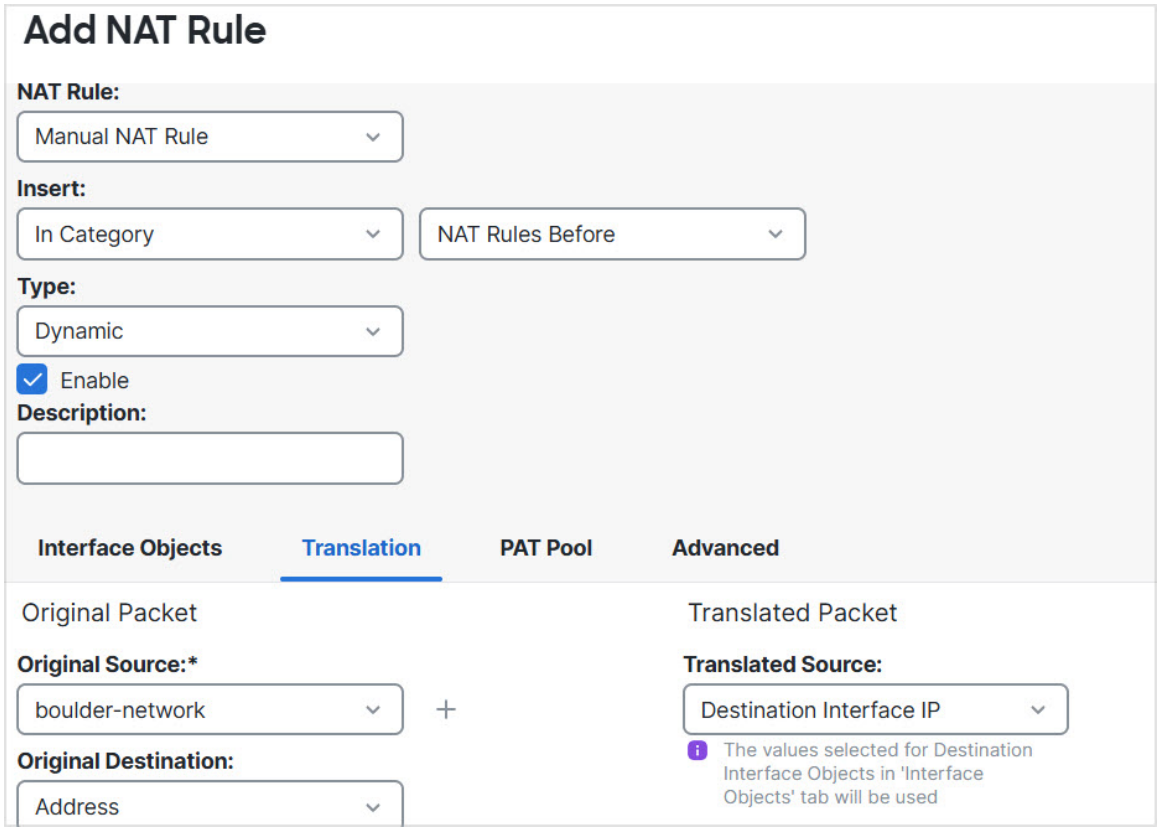
c) On **Interface Objects** , configure the following:

- **Source Interface Objects** = inside-boulder.
- **Destination Interface Objects** = outside-boulder.

d) On **Translation** , configure the following:

- **Original Source** = boulder-network object.

- **Translated Source** = **Destination Interface IP**. This option configures interface PAT using the interface contained in the destination interface object.
- **Original Destination** > **Address** = any (leave blank).
- **Translated Destination** = any (leave blank).



e) Click **Save**.

Step 4 If you are also managing Firewall2 (San Jose), you can configure similar rules for that device.

- The manual identity NAT rule would be for sanjose-network when the destination is boulder-network. Create new interface objects for the Firewall2 inside and outside networks.
- The manual dynamic interface PAT rule would be for sanjose-network when the destination is "any."

DNS query and response rewriting using NAT

DNS rewrite is the process of modifying DNS replies so that the IP address returned in the DNS response is replaced with an address that matches the NAT configuration. It is used when a translated address must be returned to the requester instead of the original address in the DNS record. This function is also known as DNS doctoring.

This feature rewrites the address in DNS queries and replies that match a NAT rule (for example, the A record for IPv4, the AAAA record for IPv6, or the PTR record for reverse DNS queries). For DNS replies traversing from a mapped interface to any other interface, the record is rewritten from the mapped value to the real value. Inversely, for DNS replies traversing from any interface to a mapped interface, the record is rewritten from the real value to the mapped value. This feature works with NAT44,NAT 66, NAT46, and NAT64.

The main circumstances when you would need to configure DNS rewrite on a NAT rule are:

- The rule is NAT64 or NAT46, and the DNS server is on the outside network. You need DNS rewrite to convert between DNS A records (for IPv4) and AAAA records (for IPv6).
- The DNS server is on the outside, clients are on the inside, and some of the fully-qualified domain names that the clients use resolve to other inside hosts.
- The DNS server is on the inside and responds with private IP addresses, clients are on the outside, and the clients access fully-qualified domain names that point to servers that are hosted on the inside.

DNS rewrite has the following limitations:

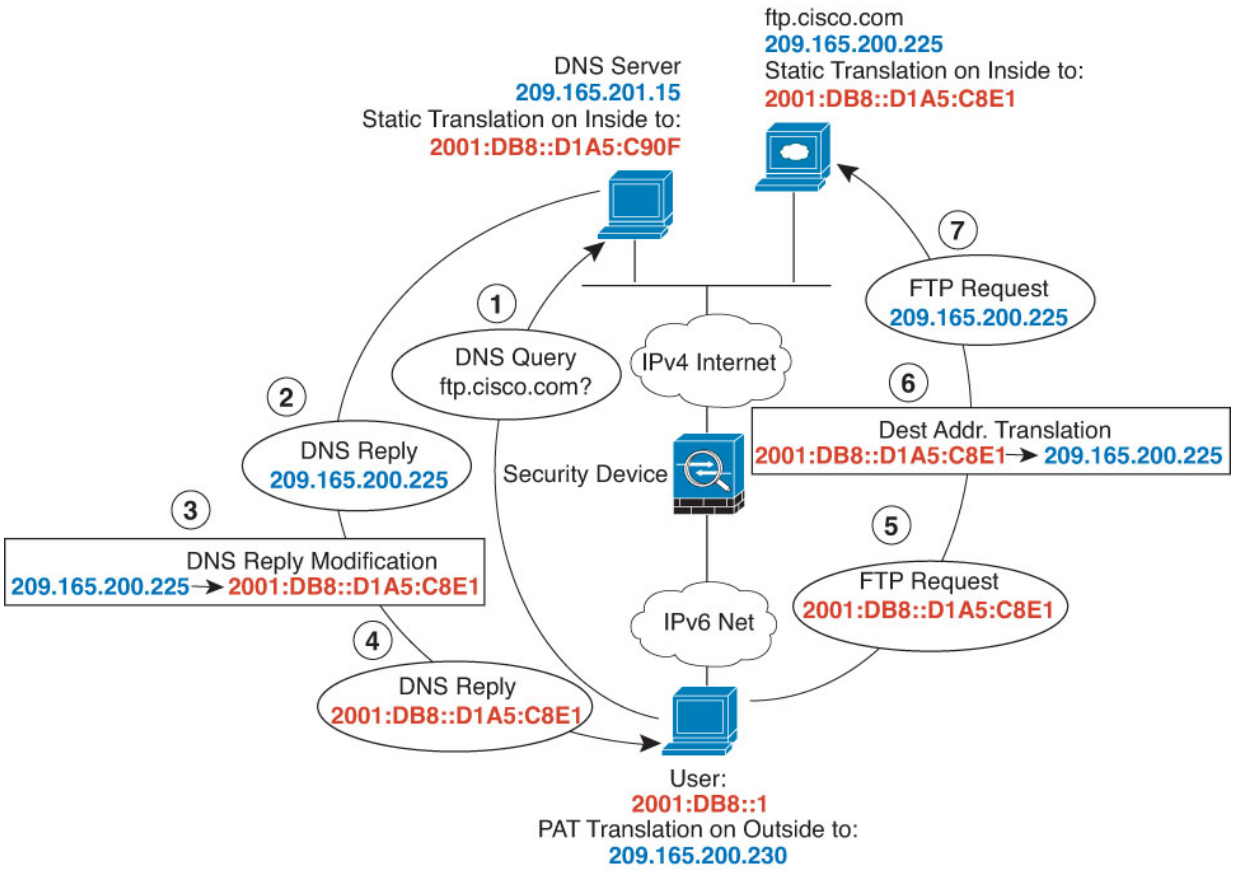
- DNS rewrite is not applicable for PAT because multiple PAT rules are applicable for each A or AAAA record, and the PAT rule to use is ambiguous.
- If you configure a manual NAT rule, you cannot configure DNS modification if you specify the destination address as well as the source address. These kinds of rules can potentially have a different translation for a single address when going to A vs. B. Therefore, the can not accurately match the IP address inside the DNS reply to the correct twice NAT rule; the DNS reply does not contain information about which source/destination address combination was in the packet that prompted the DNS request.
- You must enable DNS application inspection with DNS NAT rewrite enabled for NAT rules to rewrite DNS queries and responses. By default, DNS inspection with DNS NAT rewrite enabled is globally applied, so you probably do not need to change the inspection configuration.
- DNS rewrite is actually done on the xlate entry, not the NAT rule. Thus, if there is no xlate for a dynamic rule, rewrite cannot be done correctly. The same problem does not occur for static NAT.
- DNS rewrite does not rewrite DNS Dynamic Update messages (opcode 5).

These topics provide examples of DNS rewrite in NAT rules.

DNS64 reply modification

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

Because you want inside users to use the mapped address for ftp.cisco.com (2001:DB8::D1A5:C8E1, where D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225) you need to configure DNS reply modification for the static translation. This example also includes a static NAT translation for the DNS server, and a PAT rule for the inside IPv6 hosts.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

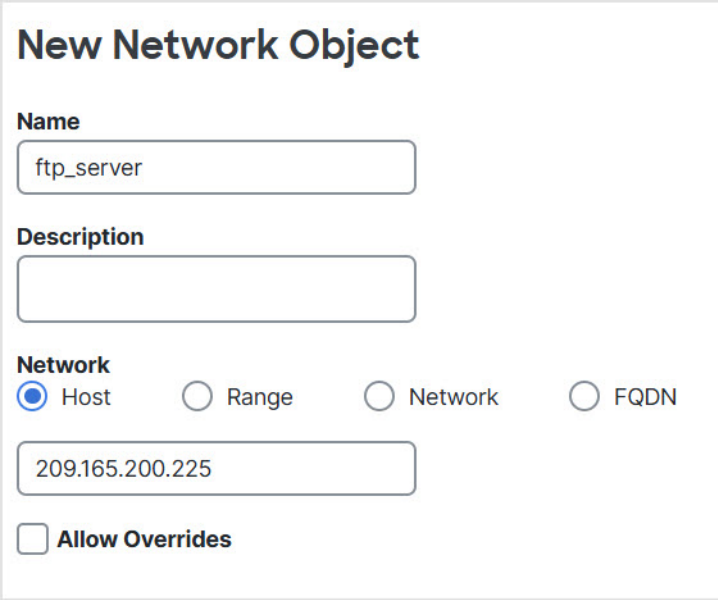
Procedure

Step 1 Create the network objects for the FTP server.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the real FTP server address.

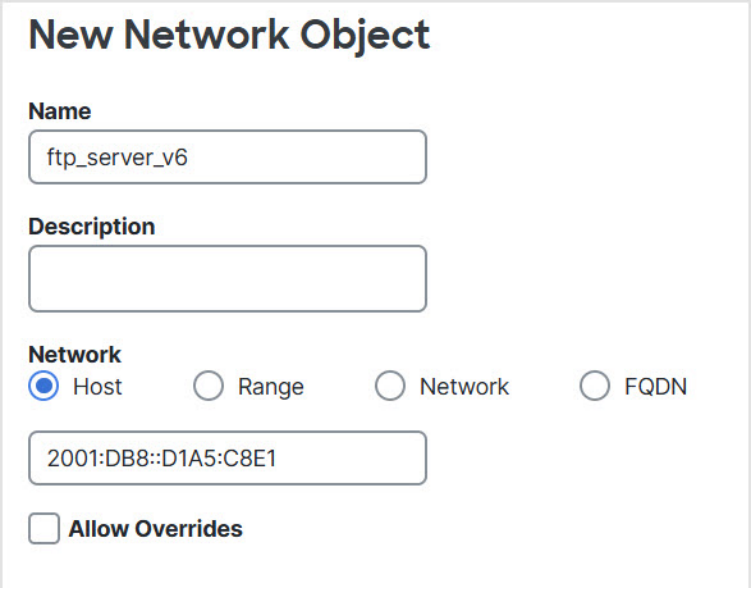
Name the network object (for example, ftp_server) and enter the host address, 209.165.200.225.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the FTP server's translated IPv6 address.

Name the network object (for example, ftp_server_v6) and enter the host address, 2001:DB8::D1A5:C8E1.

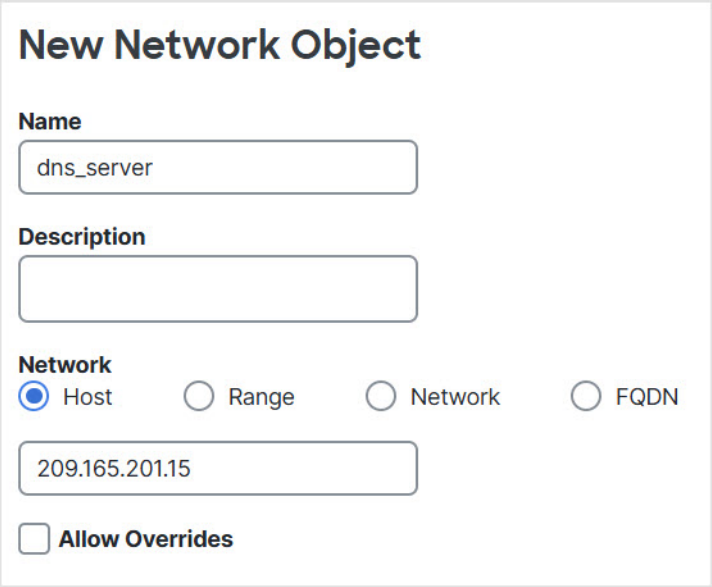


e) Click **Save**.

Step 2 Create the network objects for the DNS server.

a) Click **Add Network** > **Add Object** and define the DNS server's real address.

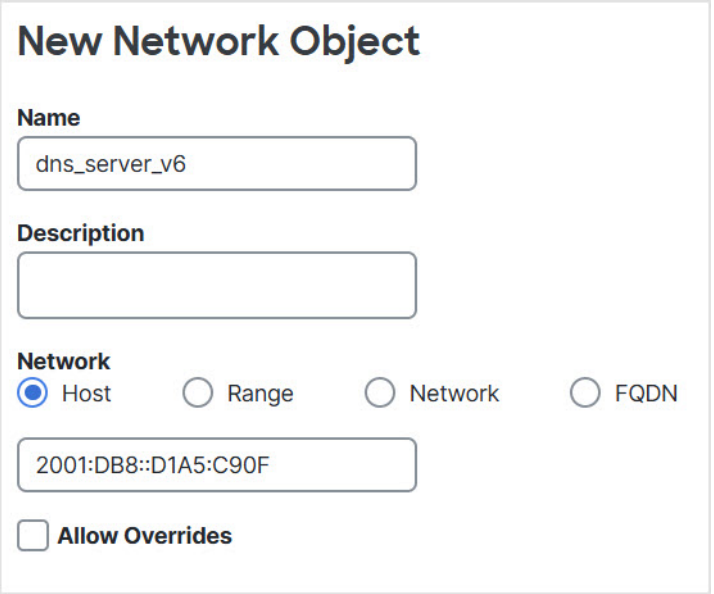
Name the network object (for example, dns_server) and enter the host address, 209.165.201.15.



b) Click **Save**.

c) Click **Add Network** > **Add Object** and define the DNS server's translated IPv6 address.

Name the network object (for example, dns_server_v6) and enter the host address, 2001:DB8::D1A5:C90F (where D1A5:C90F is the IPv6 equivalent of 209.165.201.15).

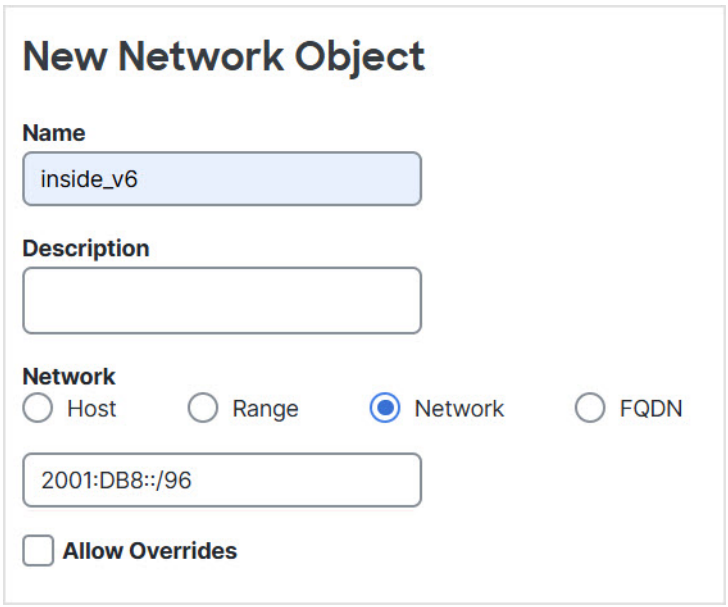


d) Click **Save**.

Step 3 Create the network objects for the inside network.

a) Click **Add Network** > **Add Object** and define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:DB8::/96.

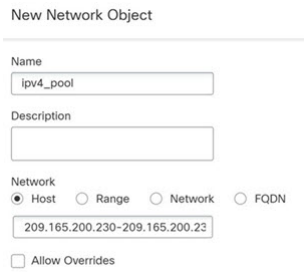


b) Click **Save**.

Step 4 Create the network objects for the PAT pool.

a) Click **Add Network** > **Add Object** and define the IPv4 PAT pool for the inside IPv6 network.

Name the network object (for example, ipv4_pool) and enter the range 209.165.200.230-209.165.200.235.



b) Click **Save**.

Step 5 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

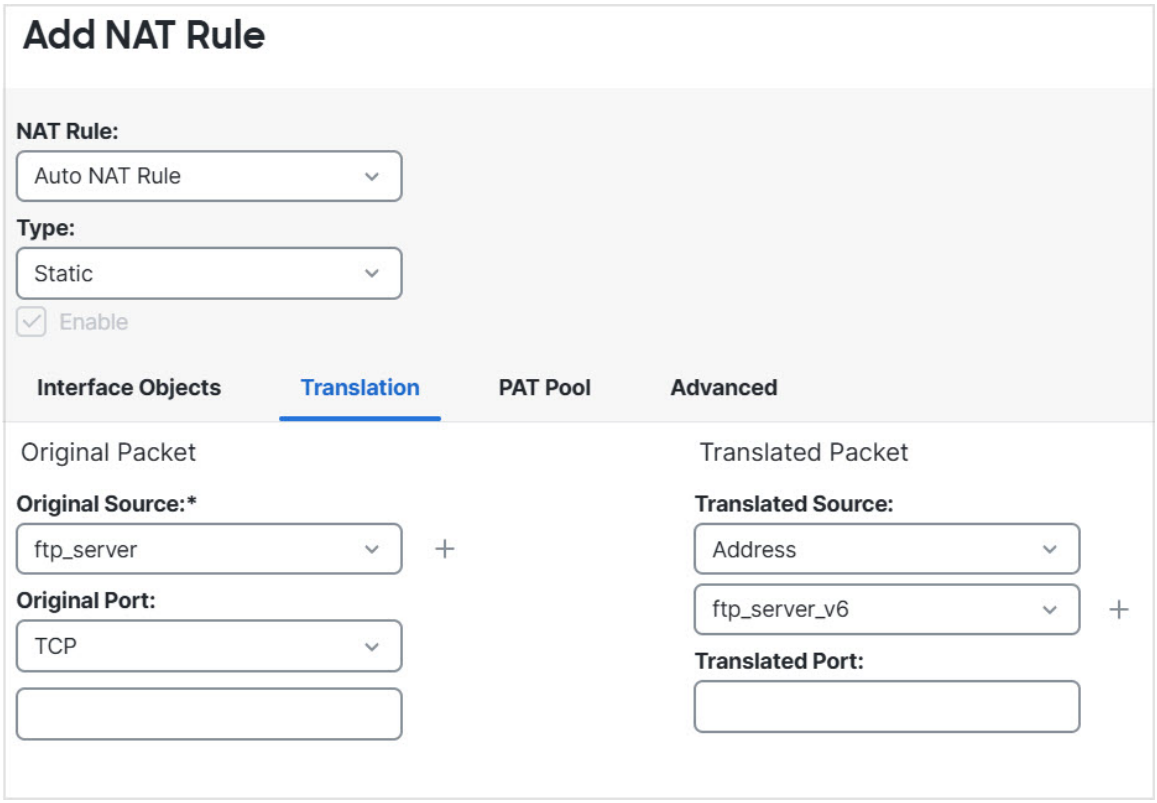
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_v6 network object.



f) On **Advanced**, select the following options:

- **Translate DNS replies that match this rule**.
- **Net to Net Mapping**, because this is a one-to-one NAT46 translation.

g) Click **OK**.

Step 6 Configure the static NAT rule for the DNS server.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

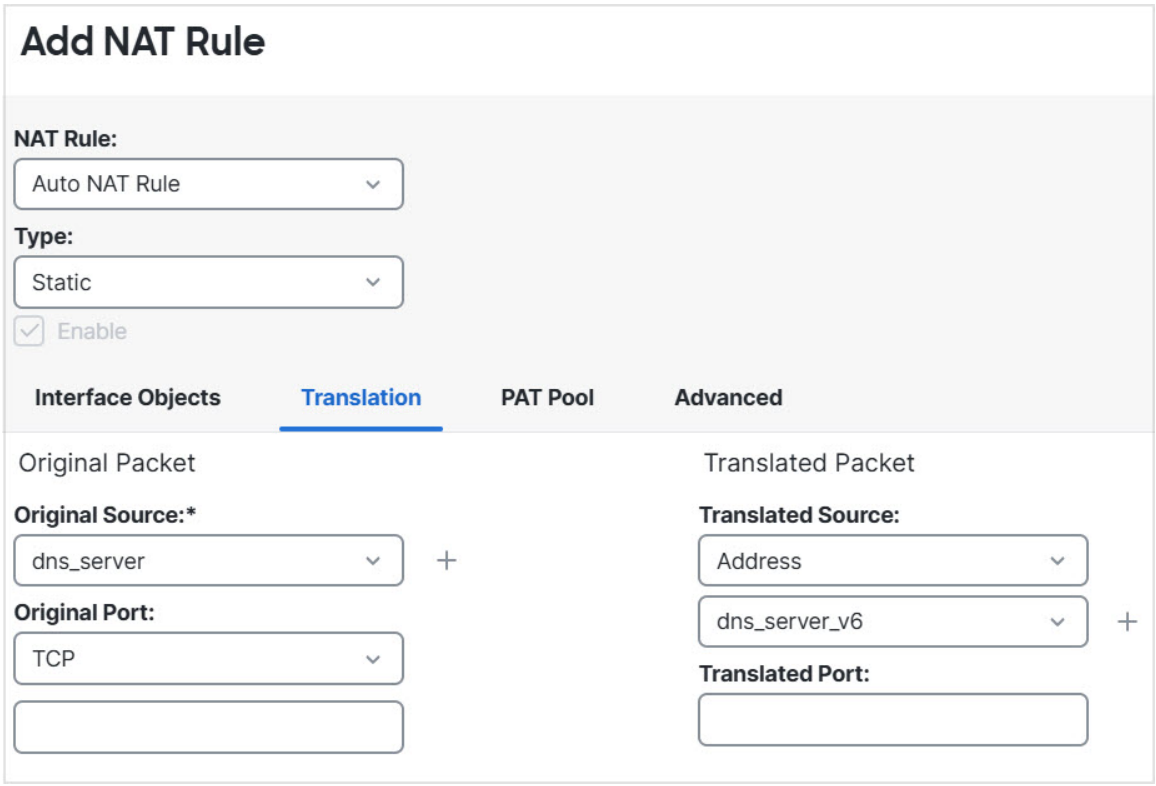
- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

d) On **Translation**, configure the following:

- **Original Source** = dns_server network object.

- **Translated Source** > **Address** = dns_server_v6 network object.

e) On **Advanced**, select **Net to Net Mapping**, because this is a one-to-one NAT46 translation.



f) Click **OK**.

Step 7 Configure the dynamic NAT with a PAT pool rule for the inside IPv6 network.

a) Click **Add Rule**.

b) Configure the following properties:

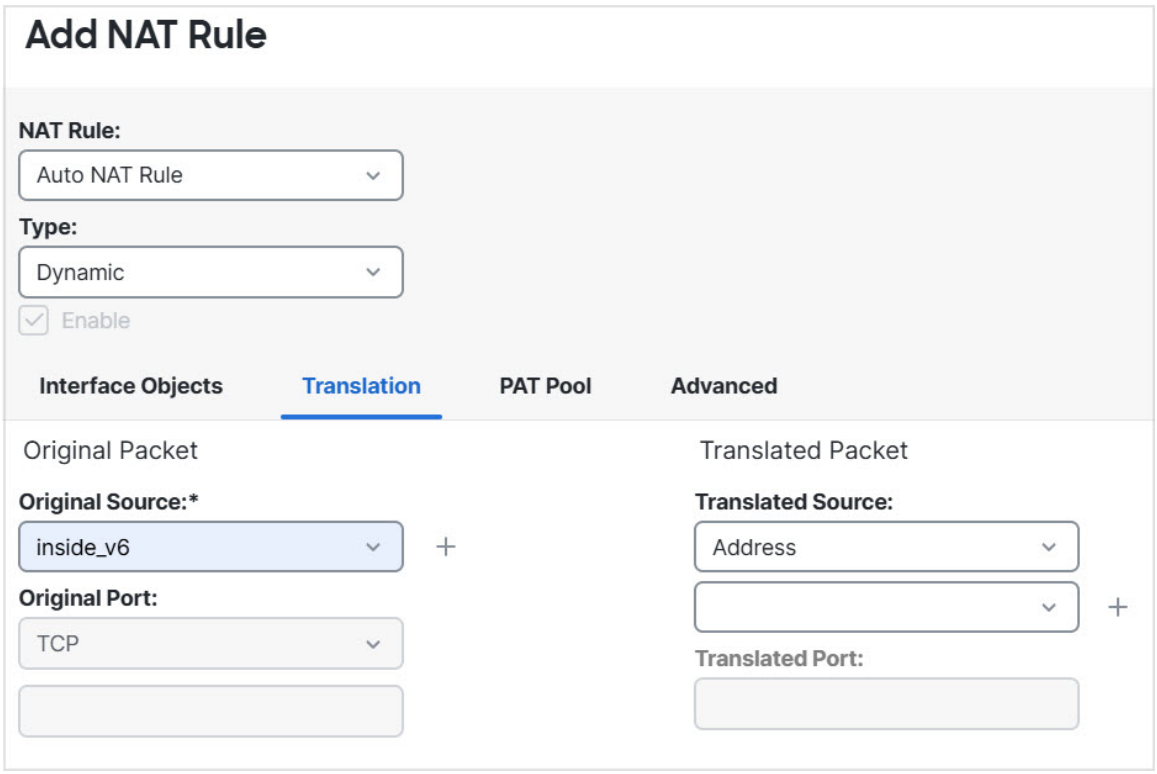
- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

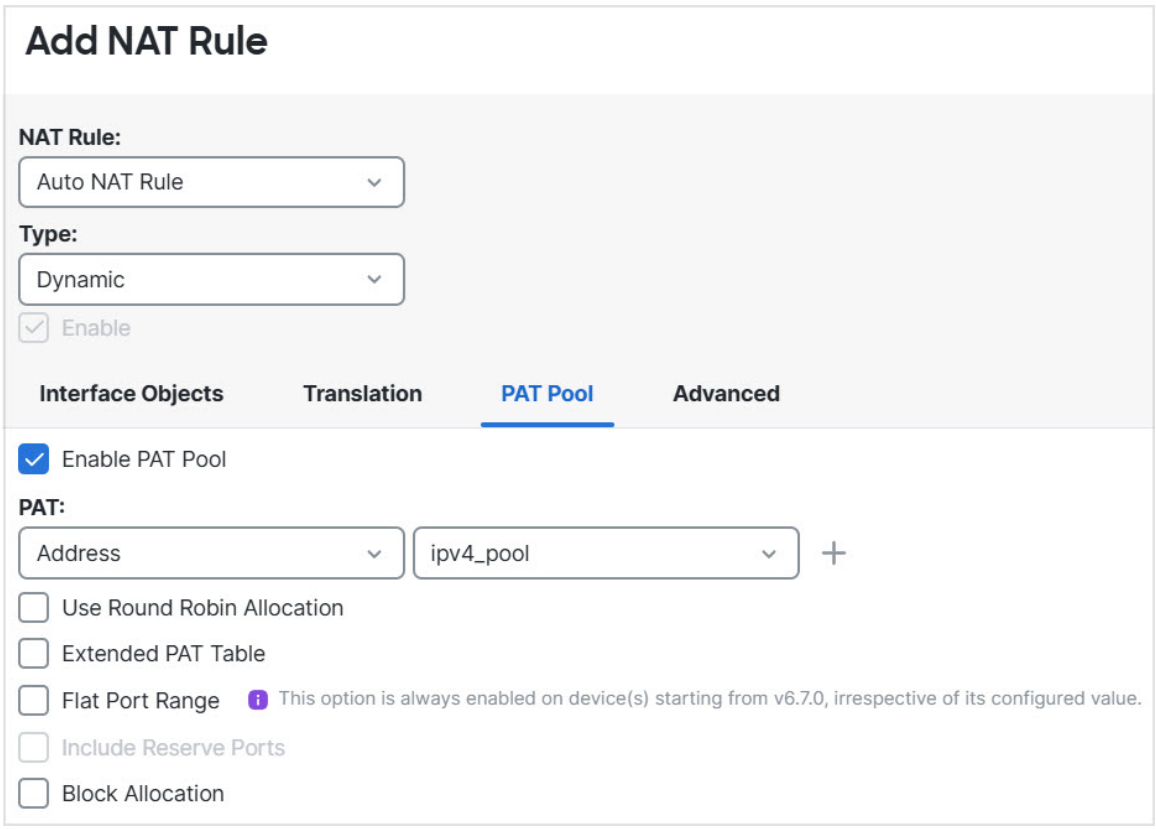
d) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** > **Address** = leave this field empty.



e) On **PAT Pool**, configure the following:

- **Enable PAT Pool** = select this option.
- **Translated Source** > **Address** = ipv4_pool network object.



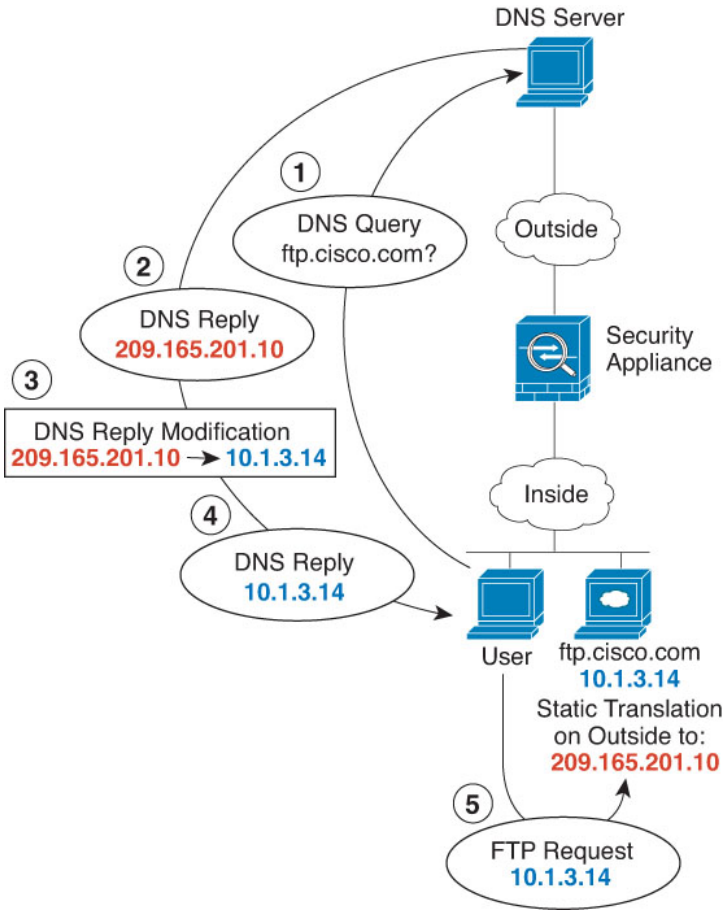
f) Click **OK**.

Create DNS reply modification with DNS server on outside

The following figure shows a DNS server that is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure NAT to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network.

In this case, you want to enable DNS reply modification on this static rule so that inside users who have access to ftp.cisco.com using the real address receive the real address from the DNS server, and not the mapped address.

When an inside host sends a DNS request for the address of ftp.cisco.com, the DNS server replies with the mapped address (209.165.201.10). The system refers to the static rule for the inside server and translates the address inside the DNS reply to 10.1.3.14. If you do not enable DNS reply modification, then the inside host attempts to send traffic to 209.165.201.10 instead of accessing ftp.cisco.com directly.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

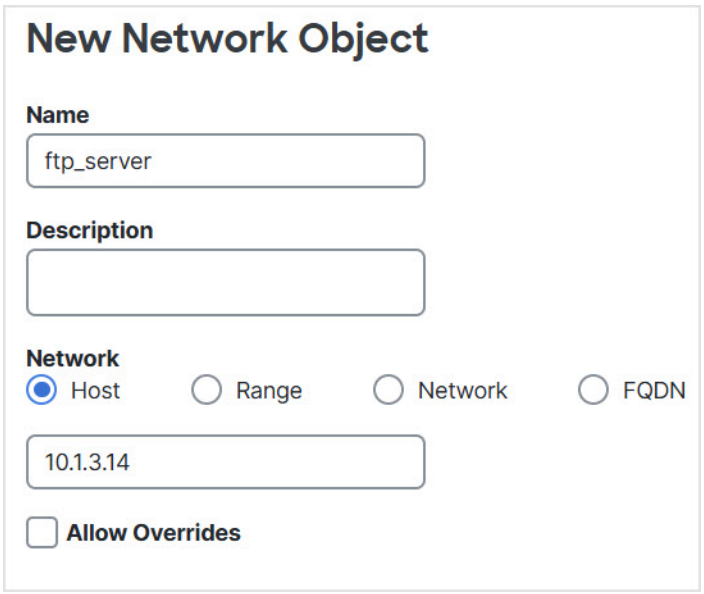
Procedure

Step 1 Create the network objects for the FTP server.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the real FTP server address.

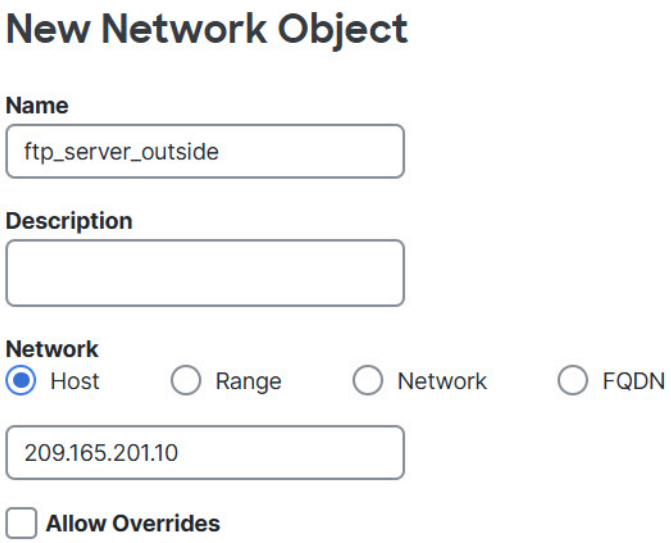
Name the network object (for example, ftp_server) and enter the host address, 10.1.3.14.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the FTP server's translated address.

Name the network object (for example, ftp_server_outside) and enter the host address, 209.165.201.10.



e) Click **Save**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

- **Type** = Static.

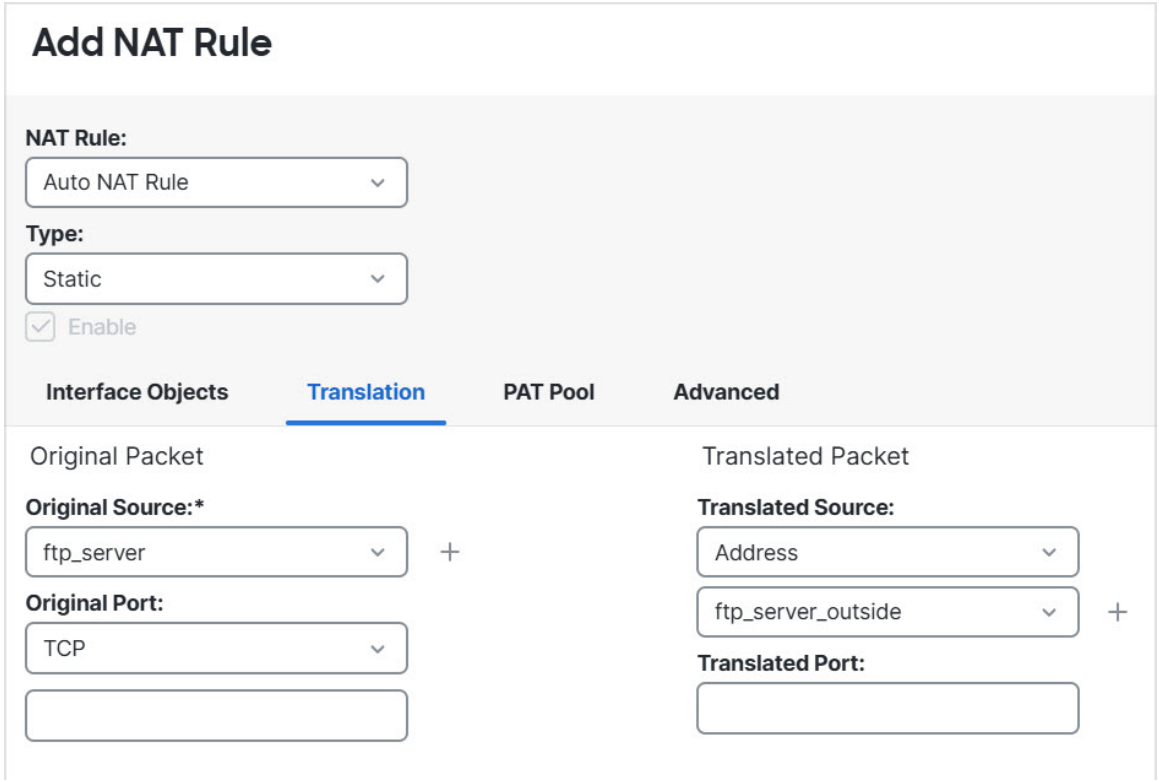
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_outside network object.

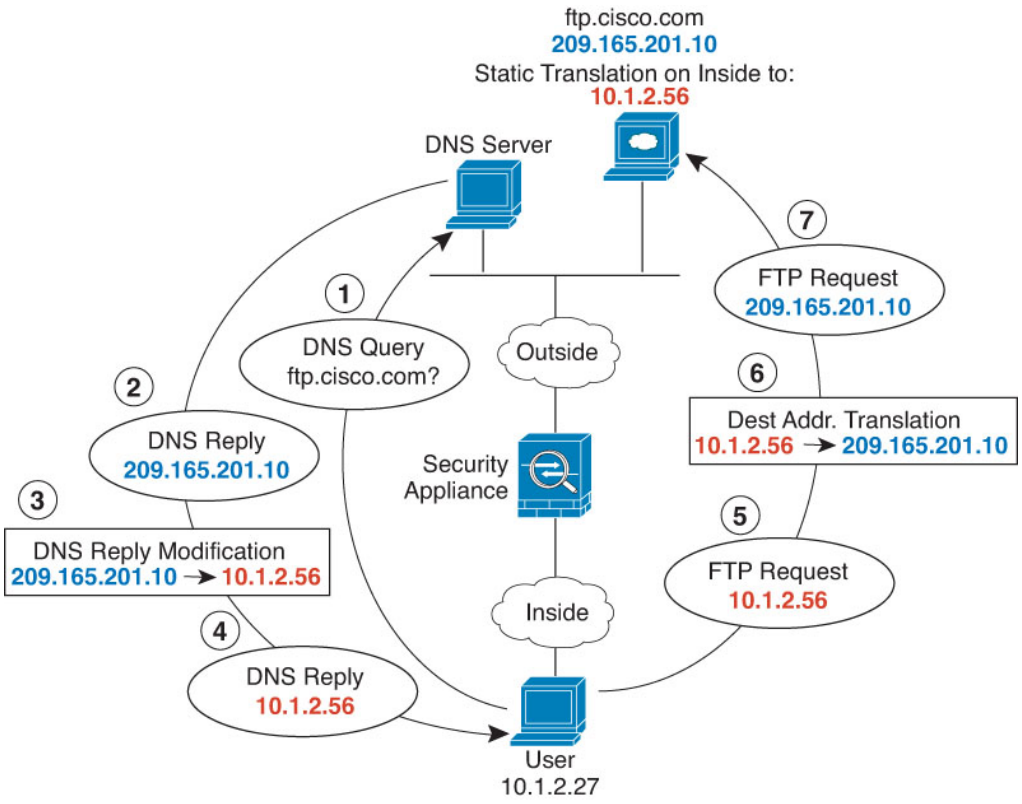
f) On **Advanced**, select **Translate DNS replies that match this rule**.



g) Click **OK**.

Configure DNS reply modification with DNS server on host network

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



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

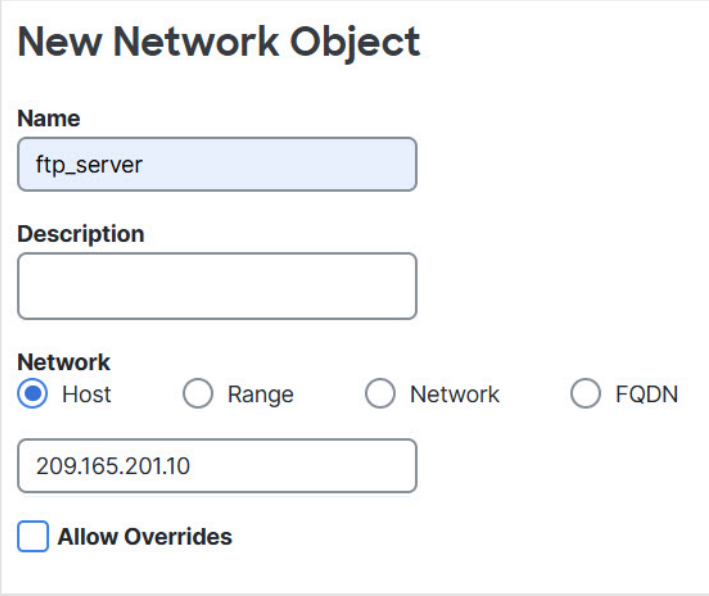
Procedure

Step 1 Create the network objects for the FTP server.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the real FTP server address.

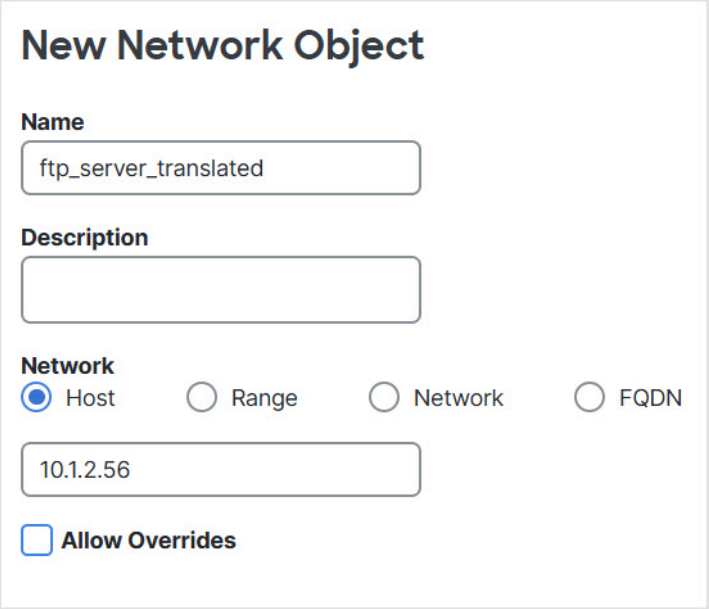
Name the network object (for example, ftp_server) and enter the host address, 209.165.201.10.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the FTP server's translated address.

Name the network object (for example, ftp_server_translated) and enter the host address, 10.1.2.56.



e) Click **Save**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

- **Type** = Static.

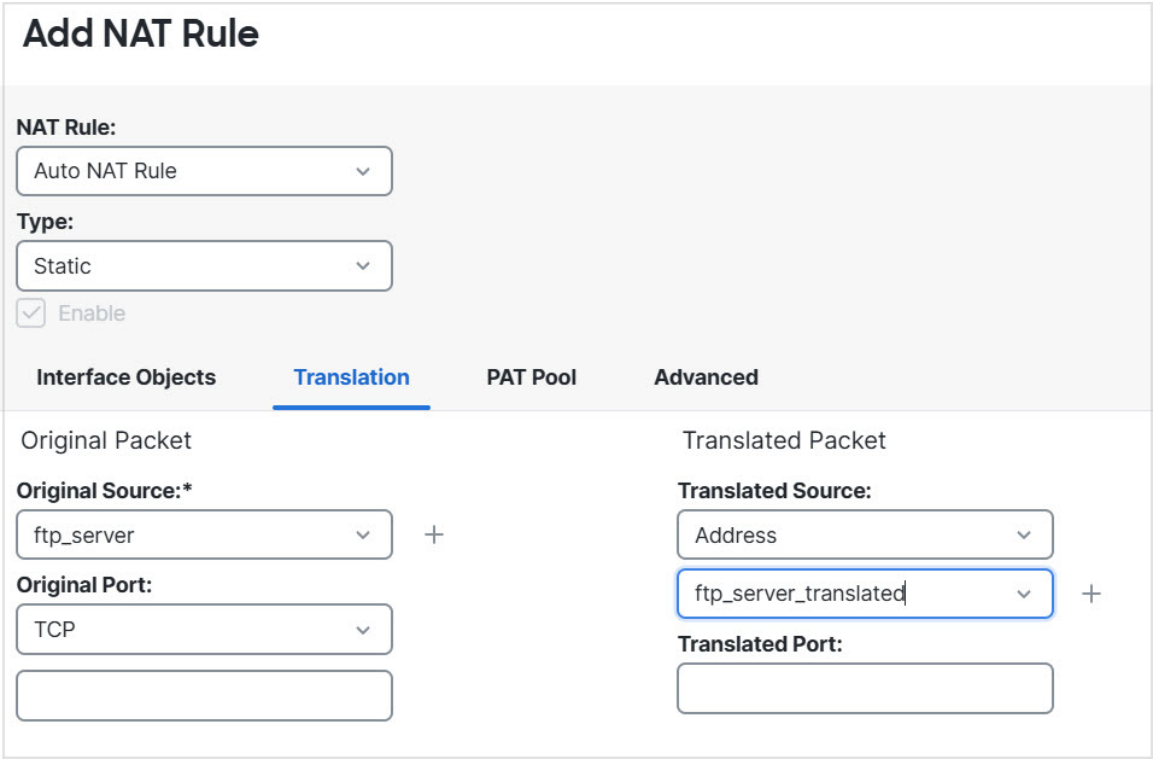
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_translated network object.

f) On **Advanced**, select **Translate DNS replies that match this rule**.



g) Click **OK**.

History for Firewall Threat Defense NAT

This table provides a chronological summary of feature changes and enhancements for the network address translation policy across different software releases.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Creation of network groups while editing NAT rules.	7.2.6 7.4.1	Any	You can create network groups in addition to network objects when you are editing a NAT rule. This feature is not supported in Version 7.3.x or 7.4.0.
Ability to enable, disable, or delete more than one NAT rule at a time.	7.2	Any	You can select multiple NAT rules and enable, disable, or delete them all at the same time. Enable and disable apply to manual NAT rules only, whereas delete applies to any NAT rule.
Manual NAT support for fully-qualified domain name (FQDN) objects as the translated destination.	7.1	Any	You can use an FQDN network object, such as one specifying www.example.com, as the translated destination address in manual NAT rules. The system configures the rule based on the IP address returned from the DNS server.
Changes to PAT address allocation in clustering. The PAT pool **Flat Port Range** option is now enabled by default and it is not configurable.	6.7	Any	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 Port Range** option in a PAT pool rule. The **Flat Port Range** option is now ignored: the PAT pool is now always flat. You can optionally select the **Include Reserved Ports** option to 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.
Ability to search and filter the Firewall Threat Defense NAT rule table.	6.7	Any	You can now search for rules in Firewall Threat Defense NAT policies to help you find rules based on IP addresses, ports, object names, and so forth. Search results include partial matches. Searching on criteria filters the rule table so only matching rules are displayed. We added a search field above the rule table when you edit Firewall Threat Defense NAT policies.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Carrier grade NAT enhancements.	6.5	Any	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). New/modified screens: We added the **Block Allocation** option to the NAT PAT Pool tab for Firewall Threat Defense NAT rules.
Support for network range objects in NAT for Firewall Threat Defense.	6.1.0	Any	You can now use network range objects in Firewall Threat Defense NAT rules where appropriate.
Network Address Translation (NAT) for Firewall Threat Defense.	6.0.1	Any	The NAT policy for Firewall Threat Defense was added. New/modified screens: Threat Defense was added as a type of NAT policy to the **Devices** > **NAT** page.