



# Network Address Translation Bindings

In Network Address Translation (NAT), the term binding describes the address binding between a local address and the global address to which the local address is translated. A binding is also called a half-entry. The different types of NAT bindings are static, dynamic, and non-PATable binds. The binding behavior of NAT is consistent across all Cisco platforms that use NAT.

This module describes the different types of NAT bindings.

- [Static NAT Binding, on page 1](#)
- [Dynamic NAT Binding, on page 2](#)
- [Non-PATable Binds, on page 2](#)
- [Recommendations on NAT Binding Configuration, on page 3](#)
- [Using VRF-Aware Software Infrastructure to Bypass NAT, on page 4](#)

## Static NAT Binding

For every static mapping in a Network Address Translation (NAT) configuration, a single static binding is created. Static bindings contain the protocol number and local and global port numbers. Sessions (with full 5-tuple entries) are created when the traffic that matches the binding passes through NAT interfaces.

Use the **ip nat inside source static** command to configure static NAT, and then configure the **show ip nat translations verbose** command to display the NAT mapping ID.

The following is sample output from the **show ip nat translations verbose** command:

```
Device(config)# ip nat inside source static 10.1.1.2 10.2.1.2
Device(config)# exit
Device# show ip nat translations verbose

Pro  Inside global      Inside local      Outside local      Outside global
---  10.1.1.2              10.2.1.2          ---                ---
create: 08/03/12 10:08:22, use: 08/03/12 10:08:22, timeout: 00:00:00
Map-Id(In): 1
Flags: static
Mac-Address: 0000.0000.0000      Input-IDB:
entry-id: 0x0, use_count:0
Total number of translations: 1
```

## Dynamic NAT Binding

Dynamic bindings are created when you configure dynamic Network Address Translation (NAT) without NAT overload configuration. In dynamic binding, the traffic matches the classification that is associated with the NAT mapping ID. Dynamic binding guarantees a one-to-one mapping between the local address and the global address. Each dynamic binding ages out when all its child sessions are aged out.

Use the **ipnat pool** command to configure dynamic NAT, and then use the **show ip nat translations verbose** command to display the mapping IDs.

The following is sample output from the **show ip nat translations verbose** command:

```
Device(config)# ip nat pool pool1 10.1.1.10 10.1.1.150 prefix-length 24
Device(config)# ip nat inside source list nat-list pool pool1
Device(config)# exit
Device# show ip nat translations verbose

Pro  Inside global      Inside local      Outside local      Outside global
---  10.1.1.10           10.2.1.2          ---                ---
    create: 08/03/12 10:10:04, use: 08/03/12 10:10:04, timeout: 23:59:51
    Map-Id(In): 1
    Mac-Address: 0000.0000.0000    Input-IDB: GigabitEthernet0/3/1
    entry-id: 0x0, use_count:2

udp  10.1.1.10:16385    10.2.1.2:16385    10.1.1.1:4003      10.1.1.1:4003
    create: 08/03/12 10:10:04, use: 08/03/12 10:10:13, timeout: 00:05:00
    Map-Id(In): 1
    Mac-Address: 0000.0000.0000    Input-IDB: GigabitEthernet0/3/1
    entry-id: 0x8bc74680, use_count:1

udp  10.1.1.10:16384    10.2.1.2:16384    10.1.1.1:4003      10.1.1.1:4003
    create: 08/03/12 10:10:03, use: 08/03/12 10:10:13, timeout: 00:05:00
    Map-Id(In): 1
    Mac-Address: 0000.0000.0000    Input-IDB: GigabitEthernet0/3/1
    entry-id: 0x8bc745b0, use_count:1

Total number of translations: 3
```

## Non-PATable Binds

Non-PATable binds are created when an application layer gateway (ALG) like Domain Name System (DNS) requests a NAT translation. When you have a dynamic NAT overload or Port-Address Translation (PAT) configuration, and have to translate a frame which is not PATable, non-PATable binds are created. A non-PATable frame is one that do not have any assigned port numbers. All IP headers are provided a protocol field; however, not every protocol is patable. Cisco IOS XE NAT only handles PAT for the following protocols:

- Internet Control Message Protocol (ICMP)
- ESP\_PROT
- Point-to-Point Tunneling Protocol (PPTP)
- TCP
- UDP



**Note** We recommend not to use non-PATable binds for overload configurations as these configurations have a one-to-one binding, which means that one local address consumes a single global address.

```
Device(config)# ip nat pool pool1 10.1.1.10 10.1.1.150 prefix-length 24
Device(config)# ip nat inside source list 1 pool pool1 overload
Device(config)# exit
Device# show ip nat translation verbose
```

Pro	Inside global	Inside local	Outside local	Outside global
---	10.1.1.10	10.1.1.2	---	---
	create: 08/03/12 10:11:53, use: 08/03/12 10:11:53, timeout: 23:59:51			
	Map-Id(In): 2			
	Mac-Address: 0000.0000.0000		Input-IDB: GigabitEthernet0/3/1	
	entry-id: 0x0, use_count:1			
---	10.1.1.10	10.2.1.2	10.1.1.1	10.1.1.1
	create: 08/03/12 10:11:54, use: 08/03/12 10:12:03, timeout: 24:00:00			
	Map-Id(In): 2			
	Mac-Address: 0000.0000.0000		Input-IDB: GigabitEthernet0/3/1	
	entry-id: 0x8bc74750, use_count:1			

Total number of translations: 2

## Recommendations on NAT Binding Configuration

When a packet arrives, a device (for example, Cisco ASR 1000 Series Aggregation Services Routers) uses the following steps to determine if the packet is subject to a Network Address Translation (NAT) rule to decide whether to use an existing translation entry, create a new translation entry, to not translate the packet. The device first checks the NAT translation table for a matching entry.

- If a matching entry is available, this entry is used for translation.
- If no matching entry is available, the device uses access control lists (ACLs) to find a match. A translation entry is created based on the configured match criteria and the IP address pool.

In the following sample dynamic Network Address Translation (NAT) configuration, the traffic that comes from the 172.16.0.0/24 network is translated by NAT and the traffic destined to 192.0.2.0/24 network is not translated.

```
Device(config)# ip nat pool NAT-POOL 10.98.198.1 10.98.198.15 netmask 255.255.255.240
Device(config)# ip nat inside source list NAT-ACL pool NAT-POOL overload
Device(config)# ip access-list extended NAT-ACL
Device(config-acl)# deny ip any 209.165.201.1 129.25.0.0 255.255.255.224
Device(config-acl)# deny ip any 192.0.2.0 144.118.0.0 255.255.255.0
Device(config-acl)# deny ip any 198.51.100.0 204.238.76.0 255.255.255.0
Device(config-acl)# deny ip any 10.0.0.0 255.0.0.0
Device(config-acl)# deny ip any 203.0.113.0 172.19.0.0 255.255.255.0
Device(config-acl)# deny ip any 192.168.0.0 255.255.0.0
Device(config-acl)# permit ip 172.16.0.0 255.240.0.0 any
```

The following is sample output from the **show ip nat translations inside** command:

```
Device# show ip nat translations inside 172.16.0.16
```

Pro	Inside global	Inside local	Outside local	Outside global
---	10.98.198.2	172.16.0.16	---	---
udp	10.98.198.2:137	172.16.0.16:137	192.0.2.4	192.0.2.4
	144.118.38.213:137			
tcp	10.98.198.2:59901	172.16.0.16:59901	192.0.2.6	192.0.2.6
	144.118.38.109:389			
udp	10.98.198.2:123	172.16.0.16:123	206.246.122.250:123	206.246.122.250:123

When the first packet arrives from 172.16.0.16 to 192.0.2.0 and a translation entry does not exist for this packet, the packet is matched against the configured ACL, and it is not translated by NAT. When the next packet arrives from 172.19.0.16 to 192.0.2.0, then that packet is matched against the NAT binding and is translated.

However, when a Domain Name System (DNS), Lightweight Directory Access Protocol (LDAP), or Netbios packet arrives from 172.19.0.16 to one of the permitted hosts, the application layer gateway (ALG) creates a binding in the translation table. When an IP address that is neither the source or the destination address is embedded in a packet payload, and the packet does not have any port numbers (for example, DNS packet), the response packet also will have an IP address that is neither the source or the destination IP address. Traffic other than Internet Control Message Protocol (ICMP), TCP, and UDP can also create NAT bindings.

## Using VRF-Aware Software Infrastructure to Bypass NAT

Use the VPN routing and forwarding (VRF)-Aware Software Infrastructure (VASI), to bypass traffic translation by Network Address Translation (NAT). In scenarios where traffic matches the deny statements in access control lists (ACLs), use VASI to bypass translation by NAT.

VASI is implemented as virtual interface pairs called Vasileft and Vasiright. These interface pairs are logically wired back-to-back and are symmetrical. For example, Vasileft1 interface and Vasiright1 interface are automatically paired. This means that a packet that enters Vasileft1 interface is internally handed over to the Vasiright1 interface without any user configuration. You can VASI for hairpinning, inter-VRF communication locally on the box, and so on. For more information about VASI, see “Configuring the VRF-Aware Software Infrastructure” module of the *Zone-Based Policy Firewall Configuration Guide*.

This section provides a sample configuration that shows how to bypass traffic translation by NAT.

- Configure an ACL with the list of subnets that must be bypassed by NAT. The traffic destined to the following subnets is not translated by NAT:
  - 209.165.201.0 255.255.255.224
  - 192.0.2.0 255.255.255.0
  - 198.51.100.0 255.255.255.0
  - 203.0.113.0 255.255.255.0
  - 192.168.0.0 255.255.0.0
- Configure policy-based routing (PBR) on the NAT inside interface to forward the traffic destined to the subnets listed above to the VASI interface. This traffic is routed from the VASI interface to the network.
- On the NAT inside interface, PBR takes precedence over NAT, and as a result, traffic that matches a PBR policy is forwarded to the Vasileft1 interface before it reaches NAT. The packet is then internally handed over to the Vasiright1 interface.
- On the NAT outside interface, the traffic appears as coming from the VASI interface that does not have a NAT configuration, and NAT translation is bypassed.

Because of traffic bypass configuration, the NAT configuration can remove all deny statements in the ACL; however, retain the permit statements:

```
ip nat inside source list NAT-ACL pool NAT-POOL overload
!
ip access-list extended NAT-ACL
 permit ip 172.19.0.0 0.0.0.255 any
```

The following is additional sample configuration required to bypass NAT translation through VASI:

```
!
interface GigabitEthernet0/0/0
 description nat outside interface
 ip address 10.2.1.1 255.255.255.0
 ip nat outside
!
interface GigabitEthernet0/0/1
 description nat inside interface
 ip address 10.2.2.1 255.255.255.0
 ip nat inside
 ip policy route-map no-NAT-rmap
!
interface vasileft1
 ip address 10.1.1.1 255.255.255.0
!
interface vasiright1
 ip address 10.1.2.1 255.255.255.0
!
ip access-list extended bypass-NAT
 permit ip any 209.165.201.0 255.255.255.224
 permit ip any 192.0.2.0 255.255.255.0
 permit ip any 198.51.100.0 255.255.255.0
 permit ip any 203.0.113.0 255.255.255.0
 permit ip any 192.168.0.0 255.255.0.0
!
route-map no-NAT-rmap permit 10
 match ip address bypass-nat
 set interface vasileft1
!
```

