

Configure NAT64

Objective

Guides the configuration of Network Address Translation for Internet Protocol version 6 to Internet Protocol version 4, including direct internet access, service-side translation, and address and port mapping techniques using encapsulation.

- Configure NAT64, on page 1
- NAT64 direct internet access, on page 1
- Service-side NAT64, on page 8
- Mapping of address and port using encapsulation with NAT64, on page 16

Configure NAT64

Describes how to configure NAT64.

A NAT64 configuration allows you to translate IPv6 addresses to IPv4 addresses for connecting IPv6 and IPv4 networks.

Traffic origination is always from the transport-side (WAN) to the service-side (LAN) in the overlay network.

NAT64 direct internet access

Describes NAT64 direct internet access.

Table 1: Feature history

Feature Name	Release Information	Description
NAT64 DIA for Cisco IOS XE Catalyst SD-WAN Devices	Cisco IOS XE Catalyst SD-WAN Release 16.12.1b Cisco vManage Release 19.2.1	The NAT64 Direct Internet Access (DIA) feature supports routing of traffic from branch sites directly to the internet instead of tunneling the internet traffic to a central site or data center for internet access. NAT64 DIA allows IPv6 clients from a branch site to access IPv4 enterprise application servers in a data center or locally in the branch. IPv6 clients can also access IPv4 servers from the branch directly using the internet.

Information about NAT64 DIA

NAT64 DIA allows an IPv4 server to access an IPv6 server either from a remote branch or from a data center.

NAT64 DIA allows an IPv4 server to access an IPv6 server either from a remote branch or from a data center.

The traffic flow for NAT64 DIA is from the LAN to DIA.

How NAT64 DIA works

Summary

Enable IPv4 and IPv6 using a Cisco VPN Interface Ethernet template.

Workflow

1. Enable IPv4 and IPv6 using a **Cisco VPN Interface Ethernet** template.
2. Configure an IPv6 route in a **Cisco VPN** template, which is the service-side VPN.

   The source and destination IPv6 addresses are translated.
3. Because NAT IPv4 DIA is configured, the source IPv4 address gets translated because of interface overload. The destination IPv4 address remains the same.

Benefits of NAT64 DIA

Summarizes the benefits of NAT64 DIA.

- Enables good application performance
- Contributes to reduced bandwidth consumption and latency
- Contributes to lower bandwidth cost
- Enables improved branch office user experience by providing DIA at remote site locations

Restrictions for NAT64 DIA

NAT64 DIA uses interface overload only.

- NAT64 DIA uses interface overload only.
- NAT DIA pool or loopback is not supported with NAT64.
- Advertisement of NAT routes into OMP is not supported.

Restrictions for a NAT64 DIA route

- You can use the following NAT64 DIA routes for installing routes in the routing table:

  Example of a NAT64 DIA route for a /128 prefix:

  ```
  nat64 route vrf 4 64:FF9B::1E00:102/128 global
  ```

  Example of a NAT64 DIA route for a /96 prefix:

  ```
  nat64 route vrf 4 64:FF9B::/96 global
  ```

- You cannot use the following NAT64 DIA route configurations for installing routes in the routing table:

  ```
  nat64 route vrf 4 64:ff9b::/64 global
  nat64 route vrf 4 ::0/0 global
  ```

Configure NAT64 DIA and a DIA route

Describes how to configure NAT64 DIA and a DIA route.

Workflow for enabling NAT64 DIA

Procedure

Step 1 Enable NAT64 using a **Cisco VPN Interface Ethernet** template for both IPv4 and IPv6.

Note NAT64 IPv4 DIA uses interface overload by default.

When configuring NAT64 for IPv6 DIA, interface overload is already configured.

A **Cisco VPN Interface Ethernet** template is a transport interface.

Step 2 Configure a NAT64 DIA IPv6 route using a **Cisco VPN** template, which is the service VPN.

NAT64 is enabled for IPv4 and IPv6, and the NAT64 DIA route is configured.

What to do next

Verify the NAT64 DIA route configuration. See Verify NAT64 DIA route configuration

Configure NAT64 DIA

Describes how to configure NAT64 DIA.

Configure NAT64 DIA with interface overload

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.

Step 2 Click **Feature Templates**.

Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

Step 3 Edit a **Cisco VPN Interface Ethernet** template by clicking **. . .** and then clicking **Edit**.

Step 4 In the **Interface Name** field, choose an interface.

Step 5 Click **NAT** and choose **IPv4**.

a) Change the scope from **Default** to **Global**.

b) Click **On** to enable NAT for IPv4.

a) In the **NAT Type** field, click **Interface** for interface overload.

Ensure that the **Interface** option is set to **On** for IPv4.

Table 2: NAT IPv4 parameters

Parameter Name	Description
NAT	Specify if NAT translation is used. The default is **Off**.
NAT Type	Specify the NAT translation type for IPv4. Available options include: **Interface**, **Pool**, and **Loopback**. The default is the **Interface** option. The **Interface** option is supported for NAT64.
UDP Timeout	Specify when NAT translations over UDP sessions time out. Range: 1–536870 seconds Default: 300 seconds (5 minutes) **Note** Starting from Cisco IOS XE Catalyst SD-WAN Release 17.6.1a and Cisco vManage Release 20.6.1, the default **UDP Timeout** value for NAT64 has been changed to 300 seconds (5 minutes).

Parameter Name	Description
TCP Timeout	Specify when NAT translations over TCP sessions time out. Enter a timeout value. Default: 3600 seconds (1 hour) **Note** Starting from Cisco IOS XE Catalyst SD-WAN Release 17.6.1a and Cisco vManage Release 20.6.1, the default **TCP Timeout** value for NAT64 has been changed to 3600 seconds (1 hour).

Step 6 Repeat Step 5, but choose **IPv6** to enable NAT for IPv6.



Note Configure both IPv4 and IPv6 for NAT64 DIA.

Step 7 In the **NAT Selection** field, click **NAT64** to enable NAT64.



Note For IPv6, interface overload is already configured.

Table 3: NAT IPv6 parameters

Parameter Name	Description
NAT	Specify if NAT translation is used. The default is **Off**.
NAT Selection	Specify NAT64. The default is the **NAT66** option.

Step 8 Click **Update**.

NAT64 DIA is configured for the selected IPv4 and IPv6 interfaces.

What to do next

Verify the NAT64 DIA route configuration. See Verify NAT64 DIA route configuration

Configure a NAT64 DIA route

Describes how to configure a NAT64 DIA route.

Configure a NAT64 DIA route using a Cisco VPN template

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.

Step 2 Click **Feature Templates**.

Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

Step 3 Edit a **Cisco VPN** feature template by clicking **. . .** and then clicking **Edit**.

Note You configure an IPv6 DIA route in a **Cisco VPN** feature template, which is the service-side VPN.

Step 4 Click **IPv6 Route**.

Step 5 Click **New IPv6 Route**.

Step 6 In the **Prefix** field, enter the well-known prefix, `64:FF9B::/96`.

Step 7 In the **Gateway** field, click **VPN**.

Step 8 In the **Enable VPN** field, change the scope from **Default** to **Global**, and click **On** to enable VPN.

Step 9 In the **NAT** field, click **NAT64**.

Step 10 Click **Update**.

The NAT64 DIA route is configured with the specified prefix and global VPN scope.

What to do next

Verify the NAT64 DIA route configuration. See Verify NAT64 DIA route configuration

Configure a NAT64 DIA route using the CLI

Describes how to configure a NAT64 DIA route using the CLI.

Procedure

Example: Configure a NAT64 DIA route

Example:

```
Device(config)# nat64 route vrf 4 64:FF9B::1E00:102/128 global
```

The NAT64 DIA route is configured.

What to do next

Verify the NAT64 DIA route configuration. See Verify NAT64 DIA route configuration

Verify NAT64 DIA route configuration

Objective

Describes how to verify NAT64 DIA route configuration.

Example 1

The following is a sample output from the **show ipv6 route vrf** command, which is for the service VPN:

```
Device# show ipv6 route vrf 4
IPv6 Routing Table - 4 - 5 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
       lp - LISP publications, a - Application, m - OMP
m   64:FF9B::/96 [251/0]
     via 172.16.255.15%default, Sdwan-system-intf%default
```

In this example, `64:FF9B::/96`, is the NAT64 well-known prefix for translating IPv6 to IPv4 addresses.

Example 2

Because NAT64 DIA is configured in the transport VPN, the routing table in the transport VPN appears as the following:

```
Device# show ipv6 route
IPv6 Routing Table - default - 2 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
       lp - LISP publications, a - Application, m - OMP, Nd - Nat-Route DIA
S   64:FF9B::/96 [1/0]
```

Configuration example for NAT64 DIA

Objective

This example shows what is configured for NAT64 DIA.

This example shows what is configured for NAT64 DIA.

```
interface GigabitEthernet1
          no shutdown
          arp timeout 1200
          ip address 10.1.15.15 10.255.255.255
          no ip redirects
```

```
          ip mtu    1500
          ip nat outside
          load-interval 30
          mtu           1500
          negotiation auto
          nat64 enable
          !
          nat64 v6v4 list nat64-global-list interface GigabitEthernet1 overload
          !
          ip nat inside source list nat-dia-vpn-hop-access-list interface GigabitEthernet1
overload
```



Note GigabitEthernet1 is a transport VPN interface.

Service-side NAT64

Describes service-side NAT64.

Table 4: Feature history

Feature Name	Release Information	Description
Service-Side NAT64 for Cisco IOS XE Catalyst SD-WAN Devices	Cisco IOS XE Catalyst SD-WAN Release 16.12.1b Cisco vManage Release 19.2.1	The service-side Network Address Translation (NAT) 64 feature translates a source IPv6 address to available IPv4 addresses in a NAT pool. The destination IPv6 address is translated to the server's actual IPv4 address since the destination IPv6 address is an IPv4 embedded IPv6 address. Service-side NAT64 allows IPv4 servers to communicate with IPv6 clients.

Information about service-side NAT64

With the dwindling IPv4 public address space and the growing need for more routable addresses, service providers and enterprises continue to build and roll out IPv6 networks.

With the dwindling IPv4 public address space and the growing need for more routable addresses, service providers and enterprises continue to build and roll out IPv6 networks. As the IPv4 internet is going to stay for a while, communication between IPv4 and IPv6 networks is an important requirement for a seamless end-user experience.

NAT IPv6 to IPv4, or NAT64, technology facilitates communication between IPv6 and IPv4 networks.

The service-side NAT64 feature translates a source IPv6 address to available IPv4 addresses in a NAT pool. The destination IPv6 address is translated to the server's actual IPv4 address since the destination IPv6 address is an IPv4-embedded IPv6 address.

Cisco IOS XE Catalyst SD-WAN devices use stateful NAT64 for translating IPv6 addresses to IPv4 addresses and IPv4 addresses to IPv6 addresses. Stateful NAT64 with NAT overload provides a 1:*n* mapping between IPv4 and IPv6 addresses.

How service-side NAT64 works

Summary

An IPv6 client attempts to connect to an IPv4 server.

Workflow

1. An IPv6 client attempts to connect to an IPv4 server.
2. The IPv6 client makes an IPv6 AAAA record DNS query, which is an IPv6 query for an IPv4 address.

   The DNS64 server responds with an IPv4-embedded IPv6 address.

   Example:

   ```
   64:ff9b::c000:0201
   ```

   which uses the NAT64 well-known prefix (WKP), `64:FF9B::/96`. The WKP is used for algorithmic mapping between address families.

   An IPv4-embedded IPv6 address consists of a variable length prefix, an embedded IPv4 address, and a variable length suffix. The last 32 bits are the hexadecimal representation of the original IPv4 address, which is 192.0.2.1 in this example.
3. The IPv6 client now tries to connect to the IPv4 server.
4. An IPv6 to IPv4 translation is performed.

   A source IPv6 address is translated to one of the available IPv4 addresses in the pool.

   A destination IPv6 address is translated to the server's actual IPv4 address since the destination IPv6 address is an IPv4-embedded IPv6 address.

Benefits of service-side NAT64

Summarizes the benefits of service-side NAT64.

- Supports communication between IPv6 clients in the service VPN with IPv4 servers on the internet
- Provides translation of IPv6 to IPv4 addresses for maintaining dual access to IPv6 and IPv4 networks
- Requires little or no changes to existing IPv4 network infrastructures when using stateful NAT64
- Provides a seamless internet experience for IPv6 users accessing IPv4 internet services, thus maintaining IPv4 business continuity
- Supports configuration of NAT64 without having to configure a data policy

Use cases for service-side NAT64

Supported traffic flow is from the IPv6 client on the remote site, in the data center, or in another branch site, to the IPv4 client or server on the local LAN.

Supported traffic flow is from the IPv6 client on the remote site, in the data center, or in another branch site, to the IPv4 client or server on the local LAN.

Note Traffic origination is always from the transport-side (WAN) to the service-side (LAN) in the overlay network.

Prerequisites for service-side NAT64

For Domain Name System (DNS) traffic to work, you must have a separate working installation of DNS64.

- For Domain Name System (DNS) traffic to work, you must have a separate working installation of DNS64.

Restrictions for service-side NAT64

Traffic must always originate on the remote branch site and go to the IPv4 server on the local LAN.

- Traffic must always originate on the remote branch site and go to the IPv4 server on the local LAN.
- Traffic cannot originate from the IPv4 server to any IPv6 client in the data center or to a remote branch site.

IPv4 address restrictions for service-side NAT64

- For more information on the usable IPv4 destination IP addresses, see the Deployment Guidelines, RFC 6052, Section 3.1.
- The well-known prefix (WKP) must not be used to represent non-global IPv4 addresses, such as those listed in the Deployment Guidelines, Section 3 of RFC 5735.

  For example, the following IPv4 prefixes are not allowed:

  - 0.0.0.0/8
  - 10.0.0.0/8
  - 127.0.0.0/8
  - 169.254.0.0/16

- You cannot use a private IPv4 address range on the service-side (LAN).

Configure service-side NAT64

Describes how to configure service-side NAT64.

The following sections provide information about configuring service-side NAT64.

Enable service-side NAT64 using a feature template

Describes how to enable service-side NAT64 using a feature template.

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.

Step 2 Click **Feature Templates**.

Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

Step 3 Edit a **Cisco VPN Interface Ethernet** template by clicking **. . .** and then clicking **Edit**.

Note The **Cisco VPN Interface Ethernet** template is a service-side interface.

Step 4 Click **NAT** and choose **IPv6** for NAT64.

Step 5 Change the scope from **Default** to **Global**.

Step 6 In the **NAT64** field, click **On** to enable NAT64.

Step 7 Click **Update**.

Service-side NAT64 is enabled using the feature template.

What to do next

Configure a service-side NAT64 pool. See Configure a service-side NAT64 pool

Configure a service-side NAT64 pool using a configuration group

Before you begin

On the **Configuration** > **Configuration Groups** page, choose **SD-WAN** as the solution type.

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.

Step 2 Create and configure a Transport VPN feature in Transport and Management profile.

a) NAT

Table 5: NAT

Field	Description
Add NAT64 v4 Pool	

Field	Description
NAT64 v4 Pool Name*	Enter a NAT pool number configured in the centralized data policy. The NAT pool name must be unique across VPNs and VRFs. You can configure up to 31 (1–32) NAT pools per router.
NAT64 Pool Range Start*	Enter a starting IP address for the NAT pool.
NAT64 Pool Range End*	Enter a closing IP address for the NAT pool.
NAT64 Overload	Enable this option to configure per-port translation. If this option is disabled, only dynamic NAT is configured on the end device. Per-port NAT is not configured. Default: Disabled

The service-side NAT64 pool is configured using a configuration group.

What to do next

Also see Deploy a Configuration Group.

Verify the service-side NAT64 configuration. See Verify configuration of service-side NAT64

Configure a service-side NAT64 pool

Describes how to configure a service-side NAT64 pool.

Before you begin

Configure a service-side NAT64 pool

Before you begin

You must have enabled service-side NAT64 using a **Cisco VPN Interface Ethernet** template before configuring a NAT64 IPv4 pool.

Procedure

Step 1 Create a new **Cisco VPN** feature template or edit an existing **Cisco VPN** feature template. The **Cisco VPN** feature template corresponds to the service-side VPN you want to configure NAT64 for.

Step 2 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.

Step 3 Click **Feature Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

Step 4 Edit a **Cisco VPN** template by clicking **. . .** and then clicking **Edit**.

Step 5 Click **NAT**.

Step 6 Click **NAT64 v4 Pool**.

Step 7 Click **New NAT64 v4 Pool**.

Step 8 In the **NAT64 Pool name** field, specify the pool name.

Note You have to specify a number for the pool name.

Step 9 In the **NAT 64 v4 Pool Range Start** field, specify the IPv4 address for the start of the pool range.

Step 10 In the **NAT 64 v4 Pool End Start** field, specify the IPv4 address for the end of the pool range.

Step 11 From the drop-down list, choose **Global**.

Step 12 Click **On** to enable **NAT 64 Overload**.

Note **NAT 64 Overload** is set to **Off** by default.

Step 13 Click **Add**.

Step 14 Click **Update** to push the configuration to the device.

The service-side NAT64 pool is configured with the specified IPv4 range and overload settings.

What to do next

Verify the service-side NAT64 configuration. See Verify configuration of service-side NAT64

Configure service-side NAT64 using the CLI

Describes how to configure service-side NAT64 using the CLI.

Table 6: Feature history

Feature Name	Release Information	Description
IPv6 Support for NAT64 Devices	Cisco IOS XE Catalyst SD-WAN Release 16.12.1b	This feature supports NAT64 to facilitate communication between IPv4 and IPv6 on Cisco IOS XE Catalyst SD-WAN devices.

Procedure

Step 1 Enable service-side NAT64 using the CLI

This section provides an example CLI configuration for enabling service-side NAT64.

Enable service-side NAT64 on the LAN interface, which is equivalent to the **Service VPN** template on Cisco SD-WAN Manager.

The IPv4 application server is on the local LAN site and the IPv6 client is in the data center or on the remote site of the LAN.

```
Device# interface GigabitEthernet 5.104
                        nat64 enable
```

Step 2 Configure a service-side NAT64 pool using the CLI

This section provides an example CLI configuration for configuring a service-side NAT64 pool.

```
Device# nat64 v4 pool pool10 192.0.2.0 192.0.2.254

nat64 v6v4 list global-list_nat64 pool pool10 vrf 4 overload
```

Service-side NAT64 is enabled and the NAT64 pool is configured.

What to do next

Verify the service-side NAT64 configuration. See Verify configuration of service-side NAT64

Verify configuration of service-side NAT64

Objective

Describes how to verify configuration of service-side NAT64.

Example - what displays in the routing table for the specified device

The following is a sample output from the **show ipv6 route vrf** command:

```
Device# show ipv6 route vrf 4
IPv6 Routing Table - 4 - 5 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
       lp - LISP publications, a - Application, m - OMP, Nd - Nat-Route DIA
Nd  64:FF9B::/96 [6/0]
     via Null0%default, directly connected
m   2001:DB8:AA:A::/64 [251/0]
     via 172.16.255.16%default, Sdwan-system-intf%default
C   2001:DB8:BB:A::/64 [0/0]
     via GigabitEthernet5.104, directly connected
L   2001:DB8:BB:A::1/128 [0/0]
     via GigabitEthernet5.104, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

In this example, the NAT64 well-known prefix, `64:FF9B::/96`, displays in the IPv6 routing table of a service VPN.

The following is a sample output from the **show ip route vrf 4** command:

```
Device# show ip route vrf 4
Routing Table: 4
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected
```

The NAT64 IPv4 pool address is installed in the routing table as `nat inside route` in the IPv4 routing table of a service VPN.

Example - what displays in the routing table on OMP

The following is a sample output from the **show ipv6 route vrf** command:

```
Device# show ipv6 route vrf 4
IPv6 Routing Table - 4 - 5 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
       lp - LISP publications, a - Application, m - OMP
m   64:FF9B::/96 [251/0]
     via 172.16.255.15%default, Sdwan-system-intf%default
C   2001:DB8:AA:A::/64 [0/0]
     via GigabitEthernet5.104, directly connected
L   2001:DB8:AA:A::1/128 [0/0]
     via GigabitEthernet5.104, receive
m   2001:DB8:BB:A::/64 [251/0]
     via 172.16.255.15%default, Sdwan-system-intf%default
L   FF00::/8 [0/0]
     via Null0, receive
```

In this example, the NAT64 well-known prefix, `64:FF9B::/96`, is received as an Overlay Management Protocol (OMP) route.

The NAT64 IPv4 pool addresses are received as an OMP route.

Configuration examples for service-side NAT64

Objective

This example shows the configuration of service-side NAT64.

This example shows the configuration of service-side NAT64.

```
nat64 v4 pool 1-4 192.0.2.0 192.0.2.254
nat64 v6v4 list nat64-list pool 1-4 vrf 4 overload
!
interface GigabitEthernet5.104
```

```
 encapsulation dot1Q 104
 vrf forwarding 4
 ip address 10.1.19.15 10.255.255.255
 ip mtu 1496
 ip ospf network broadcast
 ip ospf 4 area 0
 nat64 enable
end
```

This example shows the configuration of a NAT64 pool.

```
nat64 v4 pool 1-4 192.0.2.0 192.0.2.254
nat64 v6v4 list nat64-list pool 1-4 vrf 4 overload
!
interface GigabitEthernet5.104
 encapsulation dot1Q 104
 vrf forwarding 4
 ip address 10.1.19.15 10.255.255.255
 ip mtu 1496
 ip ospf network broadcast
 ip ospf 4 area 0
 nat64 enable
end
```

Mapping of address and port using encapsulation with NAT64

Mapping of address and port using encapsulation with NAT64

Table 7: Feature history

Feature Name	Release Information	Description
Mapping of Address and Port Using Encapsulation (MAP-E) with NAT64	Cisco IOS XE Catalyst SD-WAN Release 17.10.1a	This feature provides support for an IPv4 client to access IPv4 servers when using an IPv6-only network. IPv4 traffic is routed to the internet over an IPv6 tunnel. With this feature, you can configure a MAP-E domain and MAP-E parameters for transporting IPv4 packets over an IPv6 network using IP encapsulation. When the MAP-E customer edge (CE) device starts or when an IPv4 address changes, the device obtains the MAP-E parameters automatically from the MAP-E rule server using HTTP. This feature lets you configure a MAP-E domain and MAP-E parameters for transporting IPv4 packets over an IPv6 network using IP encapsulation. When the MAP-E customer edge (CE) device starts or when an IPv4 address changes, the device obtains the MAP-E parameters automatically from the MAP-E rule server using HTTP.

Feature Name	Release Information	Description
Mapping of Address and Port Using Encapsulation (MAP-E) with Transport Locator (TLOC)	Cisco IOS XE Catalyst SD-WAN Release 17.16.1a Cisco Catalyst SD-WAN Manager Release 20.16.1	This feature extends the support for IP over Ethernet (IPoE) MAP-E to include TLOC interfaces. The integration of MAP-E with TLOC allows the device to handle IPv4 traffic efficiently over an IPv6 network.

Information about MAP-E with NAT64

Information about MAP-E with NAT64

Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a

Mapping of Address and Port with Encapsulation (MAP-E) is an Internet Engineering Task Force (IETF) draft that describes a mechanism for transporting IPV4 packets over an IPv6-only network using encapsulation.

Within the MAP-E domain, IPv4 packets are exchanged between MAP-E CE devices and the public IPv4 internet by encapsulating and transporting the packets through the IPv6-only network.

The MAP-E with NAT64 feature supports the following configurations:

- Shared-IPv4 configuration

  MAP-E enables multiple MAP-E CE devices within a MAP-E domain to share a single IPv4 address. Each MAP-E CE device with the same IPv4 address has to use a different TCP or UDP port. MAP-E provides IPv4 connectivity using shared IPv4 addresses in an IPv6-only network.

- Fixed-IPv4 configuration

  In a fixed-IPv4 configuration, one MAP-E CE device uses a fixed-IPv4 address.

MAP-E with TLOC

Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.16.1a and Cisco Catalyst SD-WAN Manager Release 20.16.1

The IPoE MAP-E feature on Transport Locator (TLOC) interfaces enhances the Cisco Catalyst SD-WAN capabilities by providing reliable configurations for WAN services. TLOCs are identifiers representing the WAN edge router's physical or logical interfaces used for transport across the WAN. Each TLOC is characterized by three attributes: color (indicating the transport type, such as MPLS or Internet), IP address, and encapsulation type (such as IPsec or GRE).

By integrating MAP-E with TLOC, the Cisco Catalyst SD-WAN solution leverages MAP-E technology over TLOC interfaces to deliver WAN services. This integration allows efficient handling of IPv4 traffic over an IPv6 network by utilizing the TLOC's transport capabilities.

Components of a MAP-E configuration with NAT64

Each mapping of address and port using a MAP-E domain uses a different mapping rule.

Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a

Each mapping of address and port using a MAP-E domain uses a different mapping rule. A MAP-E configuration contains the following:

- Basic mapping rule (BMR)

  A BMR configures the MAP-E IPv6 address or prefix. You can configure only one BMR per IPv6 prefix. The MAP-E CE device uses the BMR to configure itself with an IPv4 address, an IPv4 prefix, or a shared IPv4 address from an IPv6 prefix. You can also use a BMR for forwarding packets in use cases where an IPv4 source address and source port are mapped to an IPv6 address or prefix. Every MAP-E CE must be provisioned with a BMR.

  The BMR IPv6 prefix along with the port parameter is used as a tunnel source address.

- One default mapping rule (DMR)

  A DMR prefix that matches with the interface address is recognized as a host and a DMR prefix with a prefix length of 128 is recognized as the tunnel source address. A border router IPv6 address is the tunnel destination address.

- Port-set ID (PSID)

  The PSID identifies the allowed ports to use.

  The border router checks whether the PSID and the port set match. If the port-set ID and port set match, the DMR matches the packet destination of the IPv6 packet. Based on the BMR, the border router constructs the IPv4 source address and extracts the IPv4 destination address from the IPv6 destination address. The IPv6 packet uses the NAT64 IPv6-to-IPv4 translation engine to construct the IPv4 packet from the IPv6 packet.

Configure the MAP-E parameters automatically from the MAP-E rule server using HTTP when the MAP-E CE device starts or whenever an IPv6 address changes.

Note It is mandatory to configure the IPv6 WAN interface as the source interface for the HTTP client. For example:

```
ip http client source-interface GigabitEthernet0/0/0
```

To configure a MAP-E domain and MAP-E parameters, use the **nat64 provisioning** command. For more information on configuring a MAP-E domain and MAP-E parameters, see the Cisco IOS XE SD-WAN Qualified Command Reference Guide.

For information on the MAP-E rule REST API specifications, see the IP Addressing: NAT Configuration Guide.

Benefits of MAP-E with NAT64

Benefits of MAP-E with NAT64

- Supports IPv4 traffic flow over an IPv6-only network.
- Supports efficient traffic delivery with lower latency without requiring additional hardware at the border router.
- Supports an easy-to-use and scalable solution for migrating to IPv6.

Restrictions for MAP-E with NAT64

Restrictions for MAP-E with NAT64

Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a

- Support for a single BMR per a MAP-E CE device. Configure different mapping rules for every address and port translation.
- No support for a BMR prefix length of 64, fragmentation, and local packet generation.
- The MAP-E with TLOC feature is supported for fixed IP mode only in Cisco IOS XE Catalyst SD-WAN Release 17.16.1a.
- The IOS-XE MAP-E solution works only when there is a single IPv6 interface connected to the service provider. Multiple WAN side IPv6 interfaces is not supported.

Workflow for MAP-E with NAT64

Summary

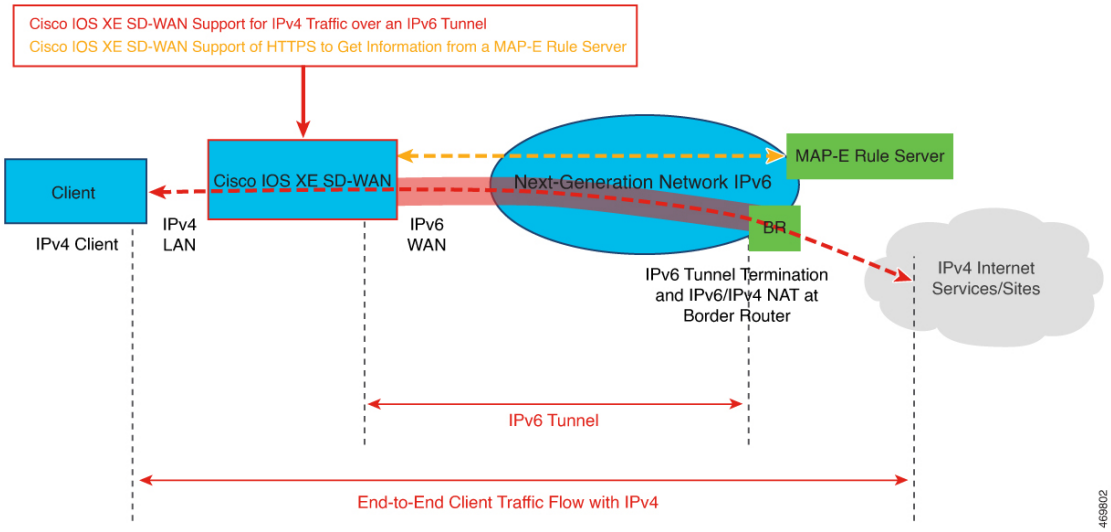
Workflow for MAP-E with NAT64

Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a

The following diagram describes the end-to-end client traffic flow for an IPv4 client to reach an IPv4 server when using an IPv6-only network with MAP-E.

Workflow

Figure 1: Workflow for MAP-E with NAT64



MAP-E interactions between a MAP-E CE device and a MAP-E rule server

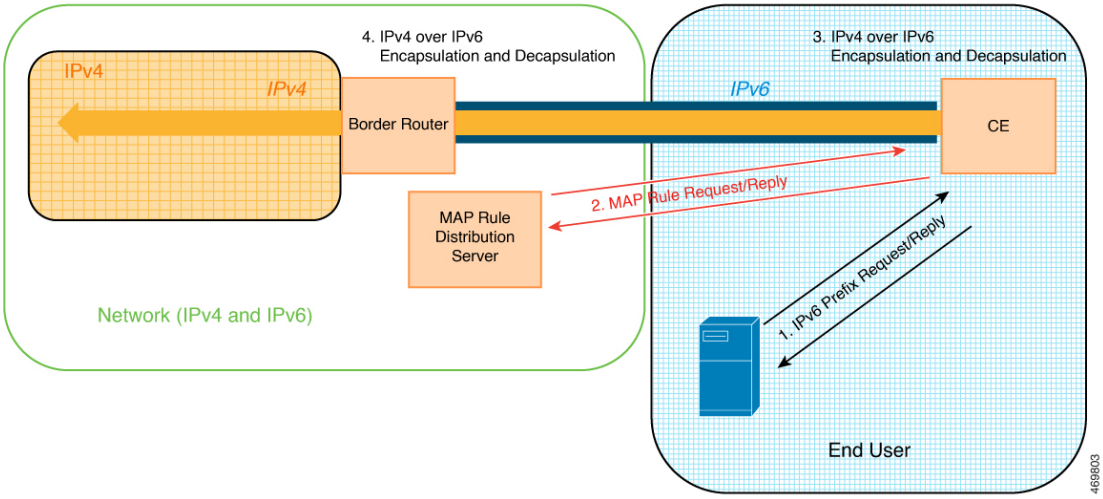
1. The MAP-E rule server acquires an IPv6 prefix from an IPv6-only network.

2. A MAP-E CE device sends an HTTP request to the MAP-E rule server and receives a response.

   MAP-E enables a Cisco IOS XE Catalyst SD-WAN device to function as a MAP-E CE device.

3. Per the MAP-E rule, a MAP-E CE device performs a translation of the incoming IPv4 packets.
4. The MAP-E CE device encapsulates the IPv4 packets into IPv6 packets and sends the IPv6 packets to the border router.
5. On receiving the encoded IPv6 packets, the IPv6 packets are decapsulated by the border router according to the MAP-E rule, and the IPv6 traffic is routed to the IPv4 public internet.

Figure 2: MAP-E workflow



Configure MAP-E with NAT64 using a CLI template

Configure MAP-E with NAT64 using a CLI template

Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a

For more information about using CLI templates, see CLI Add-On Feature Templates and CLI Templates.

Note By default, CLI templates execute commands in global config mode.

This section provides example CLI configurations to configure a MAP-E domain and MAP-E parameters.

Procedure

Step 1 Disable the NAT64 fragmentation header.

```
nat64 settings fragmentation header disable
```

Step 2 Configure a NAT64 DIA route.

```
nat64 route ip-address
                        interface-type-number
```

For more information, see Configure a NAT64 DIA Route.

Step 3 Specify the NAT64 MAP-E domain and enter MAP-E configuration mode.

```
nat64 provisioning mode jp01
```

Step 4 Configure the address resolution sever with a username and password.

```
address-resolution-server http://ipv6-prefix/directory-path
address-resolution-server 6 username encypted-user-name
address-resolution-server 6 password encypted-password
```

You can also specify an encryption type (2 or 6) for encrypting the address resolution server username and password.

Step 5 Configure a MAP-E rule server.

```
rule-server http://admin:admin@ipv6-prefix//directory-path
```

You can also specify an encryption type (2 or 6) for encrypting the rule server URL.

Step 6 (Optional) Configure a wait time in seconds for responses from the HTTP server.

```
rule-server request wait-time value-seconds
```

Step 7 Configure a host name.

Example:

```
hostname hostname
```

The host name is from the MAP-E rule server. In case you overwrite the host name, you can specify a new host name.

Step 8 Configure a tunnel interface.

```
tunnel interface Tunnelnumber
tunnel source interface-type-interface-number
```

Note Configure a tunnel interface and a tunnel source for a fixed-IPv4 configuration only.

Step 9 Configure a service prefix.

```
service-prefix ipv6-prefix
```

Note The IPv6 prefix of the MAP-E rule returned from the MAP-E rule server needs to match the IPv6 service prefix configured on the MAP-E CE device.

Note Configure a service prefix for either a fixed-IPv4 configuration or for a shared-IPv4 configuration.

The MAP-E domain and NAT64 parameters are configured.

Here is a complete configuration example for configuring a MAP-E domain and parameters.

```
nat64 settings fragmentation header disable
nat64 settings v4 tos ignore
interface GigabitEthernet1
!
nat64 settings mtu minimum 1500
nat64 provisioning mode jp01
 address-resolution-server http://2001:db8:b000:0:fe7f:6ee7:33db:5013/nic/update
 address-resolution-server password encypted-password
 address-resolution-server username encypted-username
 rule-server http://admin:admin@2001:DB8:A000::1//mape-rule.json
 rule-server request wait-time 180
 hostname hostname
 tunnel interface Tunnel1
 tunnel source GigabitEthernet2
 service-prefix 2001:DB8:b800::/48
!
nat64 route 0.0.0.0/0 GigabitEthernet1
```

What to do next

Verify MAP-E with NAT64 configuration. See Verify MAP-E with NAT64 configuration

Verify MAP-E with NAT64 configuration

Objective

Verify MAP-E with NAT64 configuration

Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a

The following example is a sample output from the **show nat64 map-e** command:

```
Device# show nat64 map-e
MAP-E Domain 9126
   Mode MAP
   Border-relay-address
      Ip-v6-address 2001:DB8::9
   Basic-mapping-rule
      Ip-v6-prefix 2001:DB8:B001:80::/60
      Ip-v4-prefix 10.1.1.0/24
      Port-parameters
         Share-ratio 4   Contiguous-ports 256   Start-port 1024
         Share-ratio-bits 2   Contiguous-ports-bits 8   Port-offset-bits 6
      Port-set-id 0
```

The output above is an example of a shared-IPv4 address configuration because the ports are shared across the MAP-E CE devices. The output displays the MAP-E parameters that were returned from the MAP-E rule server.

The following example is a sample output from the **show nat64 statistics** command:

```
Device# show nat64 statistics
NAT64 Statistics

Total active translations: 0 (0 static, 0 dynamic; 0 extended)
Sessions found: 0
Sessions created: 0
Expired translations: 0
Global Stats:
   Packets translated (IPv4 -> IPv6)
      Stateless: 0
      Stateful: 0
      nat46: 0
      MAP-T: 0
      MAP-E: 5
   Packets translated (IPv6 -> IPv4)
      Stateless: 0
      Stateful: 0
      nat46: 0
      MAP-T: 0
      MAP-E: 4

Interface Statistics
   GigabitEthernet0/0/0 (IPv4 not configured, IPv6 configured):
      Packets translated (IPv4 -> IPv6)
         Stateless: 0
         Stateful: 0
         nat46: 0
         MAP-T: 0
         MAP-E: 0
      Packets translated (IPv6 -> IPv4)
         Stateless: 0
         Stateful: 0
         nat46: 0
         MAP-T: 0
         MAP-E: 4
      Packets dropped: 0
   GigabitEthernet0/0/1 (IPv4 configured, IPv6 not configured):
      Packets translated (IPv4 -> IPv6)
         Stateless: 0
         Stateful: 0
         nat46: 0
         MAP-T: 0
         MAP-E: 5
      Packets translated (IPv6 -> IPv4)
         Stateless: 0
         Stateful: 0
         nat46: 0
         MAP-T: 0
         MAP-E: 0
      Packets dropped: 0
Dynamic Mapping Statistics
   v6v4
Limit Statistics
```