

Stateful Network Address Translation 64

The Stateful Network Address Translation 64 feature provides a translation mechanism that translates IPv6 packets into IPv4 packets and vice versa. The stateful NAT64 translator algorithmically translates the IPv4 addresses of IPv4 hosts to and from IPv6 addresses by using the configured stateful prefix. In a similar manner, the IPv6 addresses of IPv6 hosts are translated to and from IPv4 addresses through Network Address Translation (NAT). Stateful Network Address Translation 64 (NAT64) also translates protocols and IP addresses. The Stateful NAT64 translator enables native IPv6 or IPv4 communication and facilitates coexistence of IPv4 and IPv6 networks.

This document explains how Stateful NAT64 works and how to configure your network for Stateful NAT64 translation.

- Prerequisites for Configuring Stateful Network Address Translation 64, on page 1
- Restrictions for Configuring Stateful Network Address Translation 64, on page 1
- Information About Stateful Network Address Translation 64, on page 2
- How to Configure Stateful Network Address Translation 64, on page 10
- Configuration Examples for Stateful Network Address Translation 64, on page 24
- Additional References for Stateful Network Address Translation 64, on page 28
- Feature Information for Stateful Network Address Translation 64, on page 29
- Glossary, on page 30

Prerequisites for Configuring Stateful Network Address Translation 64

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

Restrictions for Configuring Stateful Network Address Translation 64

- Applications without a corresponding application-level gateway (ALG) may not work properly with the Stateful NAT64 translator.
- IP Multicast is not supported.

- The translation of IPv4 options, IPv6 routing headers, hop-by-hop extension headers, destination option headers, and source routing headers is not supported.
- Virtual routing and forwarding (VRF)-aware NAT64 is not supported.
- When traffic flows from IPv6 to IPv4, the destination IP address that you have configured must match a stateful prefix to prevent hairpinning loops. However, the source IP address (source address of the IPv6 host) must not match the stateful prefix. If the source IP address matches the stateful prefix, packets are dropped.

  Hairpinning allows two endpoints inside Network Address Translation (NAT) to communicate with each other, even when the endpoints use only each other's external IP addresses and ports for communication.

- Only TCP and UDP Layer 4 protocols are supported for header translation.
- Routemaps are not supported.
- Application-level gateways (ALGs) FTP and ICMP are not supported.
- In the absence of a pre-existing state in NAT 64, stateful translation only supports IPv6-initiated sessions.
- If a static mapping host-binding entry exists for an IPv6 host, the IPv4 nodes can initiate communication. In dynamic mapping, IPv4 nodes can initiate communication only if a host-binding entry is created for the IPv6 host through a previously established connection to the same or a different IPv4 host.

  Dynamic mapping rules that use Port-Address Translation (PAT), host-binding entries cannot be created because IPv4-initiated communication not possible through PAT.

- Both NAT44 (static, dynamic and PAT) configuration and stateful NAT64 configuration are not supported on the same interface.

Information About Stateful Network Address Translation 64

Stateful Network Address Translation 64

The Stateful NAT64 feature provides a translation mechanism that translates IPv6 packets into IPv4 packets and vice versa.

Stateful NAT64 supports TCP, and UDP traffic. Packets that are generated in an IPv6 network and are destined for an IPv4 network are routed within the IPv6 network towards the Stateful NAT64 translator. Stateful NAT64 translates the packets and forwards them as IPv4 packets through the IPv4 network. The process is reversed for traffic that is generated by hosts connected to the IPv4 network and destined for an IPv6 receiver.

The Stateful NAT64 translation is not symmetric, because the IPv6 address space is larger than the IPv4 address space and a one-to-one address mapping is not possible. Before it can perform an IPv6 to an IPv4 translation, Stateful NAT64 requires a state that binds the IPv6 address and the TCP/UDP port to the IPv4 address. The binding state is either statically configured or dynamically created when the first packet that flows from the IPv6 network to the IPv4 network is translated. After the binding state is created, packets flowing in both directions are translated. In dynamic binding, Stateful NAT64 supports communication initiated by the IPv6-only node toward an IPv4-only node. Static binding supports communication initiated by an IPv4-only node to an IPv6-only node and vice versa. Stateful NAT64 with NAT overload or Port Address Translation (PAT) provides a 1:*n* mapping between IPv4 and IPv6 addresses.

When an IPv6 node initiates traffic through Stateful NAT64, and the incoming packet does not have an existing state and the following events happen:

- The source IPv6 address (and the source port) is associated with an IPv4 configured pool address (and port, based on the configuration).

- The destination IPv6 address is translated mechanically based on the BEHAVE translation draft using either the configured NAT64 stateful prefix or the Well Known Prefix (WKP).
- The packet is translated from IPv6 to IPv4 and forwarded to the IPv4 network.

When an incoming packet is stateful (if a state exists for an incoming packet), NAT64 identifies the state and uses the state to translate the packet.

Prefixes Format for Stateful Network Address Translation 64

A set of bits at the start of an IPv6 address is called the format prefix. Prefix length is a decimal value that specifies how many of the leftmost contiguous bits of an address comprise the prefix.

When packets flow from the IPv6 to the IPv4 direction, the IPv4 host address is derived from the destination IP address of the IPv6 packet that uses the prefix length. When packets flow from the IPv4 to the IPv6 direction, the IPv4 host address is constructed using the stateful prefix.

According to the IETF address format BEHAVE draft, a u-bit (bit 70) defined in the IPv6 architecture should be set to zero. For more information on the u-bit usage, see RFC 2464. The reserved octet, also called u-octet, is reserved for compatibility with the host identifier format defined in the IPv6 addressing architecture. When constructing an IPv6 packet, the translator has to make sure that the u-bits are not tampered with and are set to the value suggested by RFC 2373. The suffix will be set to all zeros by the translator. IETF recommends that the 8 bits of the u-octet (bit range 64–71) be set to zero.

Well Known Prefix

The Well Known Prefix 64:FF9B::/96 is supported for Stateful NAT64. During a stateful translation, if no stateful prefix is configured (either on the interface or globally), the WKP prefix is used to translate the IPv4 host addresses.

Stateful IPv4-to-IPv6 Packet Flow

The packet flow of IPv4-initiated packets for Stateful NAT64 is as follows:

- The destination address is routed to a NAT Virtual Interface (NVI).

  A virtual interface is created when Stateful NAT64 is configured. For Stateful NAT64 translation to work, all packets must get routed to the NVI. When you configure an address pool, a route is automatically added to all IPv4 addresses in the pool. This route automatically points to the NVI.

- The IPv4-initiated packet hits static or dynamic binding.

  Dynamic address bindings are created by the Stateful NAT64 translator when you configure dynamic Stateful NAT64. A binding is dynamically created between an IPv6 and an IPv4 address pool. Dynamic binding is triggered by the IPv6-to-IPv4 traffic and the address is dynamically allocated. Based on your configuration, you can have static or dynamic binding.

- The IPv4-initiated packet is protocol-translated and the destination IP address of the packet is set to IPv6 based on static or dynamic binding. The Stateful NAT64 translator translates the source IP address to IPv6 by using the Stateful NAT64 prefix (if a stateful prefix is configured) or the Well Known Prefix (WKP) (if a stateful prefix is not configured).
- A session is created based on the translation information.

All subsequent IPv4-initiated packets are translated based on the previously created session.

Stateful IPv6-to-IPv4 Packet Flow

The stateful IPv6-initiated packet flow is as follows:

- The first IPv6 packet is routed to the NAT Virtual Interface (NVI) based on the automatic routing setup that is configured for the stateful prefix. Stateful NAT64 performs a series of lookups to determine whether the IPv6 packet matches any of the configured mappings based on an access control list (ACL) lookup. Based on the mapping, an IPv4 address (and port) is associated with the IPv6 destination address. The IPv6 packet is translated and the IPv4 packet is formed by using the following methods:
  - Extracting the destination IPv4 address by stripping the prefix from the IPv6 address. The source address is replaced by the allocated IPv4 address (and port).
  - The rest of the fields are translated from IPv6-to-IPv4 to form a valid IPv4 packet.



Note This protocol translation is the same for stateless NAT64.

- A new NAT64 translation is created in the session database and in the bind database. The pool and port databases are updated depending on the configuration. The return traffic and the subsequent traffic of the IPv6 packet flow will use this session database entry for translation.

IP Packet Filtering

Stateful Network Address Translation 64 (NAT64) filters IPv6 and IPv4 packets. All IPv6 packets that are transmitted into the stateful translator are filtered because statefully translated IPv6 packets consume resources in the translator. These packets consume processor resources for packet processing, memory resources (always session memory) for static configuration, IPv4 address resources for dynamic configuration, and IPv4 address and port resources for Port Address Translation (PAT).

Stateful NAT64 utilizes configured access control lists (ACLs) and prefix lists to filter IPv6-initiated traffic flows that are allowed to create the NAT64 state. Filtering of IPv6 packets is done in the IPv6-to-IPv4 direction because dynamic allocation of mapping between an IPv6 host and an IPv4 address can be done only in this direction.

Stateful NAT64 supports endpoint-dependent filtering for the IPv4-to-IPv6 packet flow with PAT configuration. In a Stateful NAT64 PAT configuration, the packet flow must have originated from the IPv6 realm and created the state information in NAT64 state tables. Packets from the IPv4 side that do not have a previously created state are dropped. Endpoint-independent filtering is supported with static Network Address Translation (NAT) and non-PAT configurations.

Differences Between Stateful NAT64 and Stateless NAT64

The table below displays the differences between Stateful NAT64 and Stateless NAT64.

Table 1: Differences Between Stateful NAT64 and Stateless NAT64

Supported Features	Stateful NAT64	Stateless NAT64
Address savings	*N*:1 mapping for PAT or overload configuration that saves IPv4 addresses.	One-to-one mapping—one IPv4 address is used for each IPv6 host).
Address space	IPv6 systems may use any type of IPv6 addresses.	IPv6 systems must have IPv4-translatable addresses (based on RFC 6052).
ALGs supported	FTP64	None
Protocols supported	ICMP, TCP, UDP	All
Standards	Draft-ieft-behave-v6v4-xlate-stateful-12	Draft-ietf-behave-v6v4-xlate-05
State creation	Each traffic flow creates a state in the NAT64 translator. The maximum number of states depends on the number of supported translations.	Traffic flow does not create any state in the NAT64 translator. Algorithmic operation is performed on the packet headers.

High-Speed Logging for NAT64

When HSL is configured, NAT64 provides a log of packets that flow through routing devices (similar to the Version 9 NetFlow-like records) to an external collector. Records are sent for each binding (binding is the address binding between the local address and the global address to which the local address is translated) and when sessions are created and destroyed. Session records contain the full 5-tuple of information (the source IP address, destination IP address, source port, destination port, and protocol). A tuple is an ordered list of elements. NAT64 also sends an HSL message when a NAT64 pool runs out of addresses (also called pool exhaustion). Because the pool exhaustion messages are rate limited, each packet that hits the pool exhaustion condition does not trigger an HSL message. Depending on your release, Stateful NAT64 supports high-speed logging (HSL) for upto 4 destinations.

Configure the **nat64 logging translations flow-export v9 udp destination** command to enable NAT64 HSL logging. The **vrf** keyword can be used to enable NAT64 HSL for a specific VRF

The table below describes the templates for HSL bind and session create or destroy. These fields (in the order they are displayed in the log) describe how the log collector must interpret the bytes in HSL records. The value for some of the fields varies based on whether the session is being created, destroyed, or modified.

Table 2: Templates for HSL Bind and Session Create or Destroy

Field	Format	ID	Value
Original IPv6 address	IPv6 address	27	Varies
Translated IPv4 address	IPv6 address	282	Varies
Translated IPv6 address	IPv4 address	225	Varies
Original IPv4 address	IPv4 address	12	Varies
Original IPv6 port	16-bit port	7	Varies

Field	Format	ID	Value
Translated IPv6 port	16-bit port	227	Varies
Translated IPv4 port	16-bit port	11	Varies
Original IPv4 port	16-bit port	228	Varies
Timestamp for an event	64 bits - milliseconds (This is a 64-bit field that holds the UNIX time, in milliseconds, when the event for the record occurred.)	323	Varies
VRF ID	32-bit ID	234	Zero
Protocol	8-bit value	4	Varies
Event	8-bit value	230	0–Invalid 1–Add event 2–Delete event

The table below describes the HSL pool exhaustion templates (in the order they are available in the template).

Table 3: Templates for HSL Pool Exhaustion

Field	Format	ID	Values
NAT pool ID	32-bit value	283	Varies
NAT event	8-bit value	230	3–Pool exhaust

How to Configure Enabling NAT64 High-Speed Logging per VRF

Enabling High-Speed Logging of NAT64 Translations

You can enable or disable high-speed logging (HSL) of all NAT64 translations or only translations for specific VPNs.

You must first use the **nat64 logging translations flow-export v9 udp destination** command to enable HSL for all VPN and non-VPN translations. The **vrf** keyword can be used to specify HSL destination address on a specific VRF. VPN translations are also known as Virtual Routing and Forwarding (VRF) translations.

After you enable HSL for all NAT translations, you can then use the **nat64 logging translations flow-export v9** *vrf-name* command to enable or disable translations for specific VPNs. When you use this command, HSL is disabled for all VPNs, except for the ones the command is explicitly enabled.

SUMMARY STEPS

1. **enable**
2. **configure terminal**

3. **nat64 logging translations flow-export v9 udp destination** *addr*|**ipv6-destination** *IPv6 address* **vrf***vrf name* **source** *interface type interface-number*
4. **nat64 logging translations flow-export v9** {*vrf-name* | **global-on**}
5. **exit**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.

Step 3	**nat64 logging translations flow-export v9 udp destination** *addr*|**ipv6-destination** *IPv6 address* **vrf***vrf name* **source** *interface type interface-number*
	Example:
	This example shows how to enable high-speed logging using an IPv4 address
	`Device(config)# nat64 logging translations flow-export v9 udp destination 10.10.0.1 1020 source GigabitEthernet 0/0/0`
	Example:
	This example shows how to enable high-speed logging using an IPv6 address
	`Device(config)# nat64 logging translations flow-export v9 udp ipv6-destination 2001::06 5050 source GigabitEthernet 0/0/0`
	Example:
	This example shows how to enable high-speed logging using an IPv6 address for a destination VRF
	`Device(config)# nat64 logging translations flow-export v9 udp ipv6-destination 2001::06 5050 vrf hslvrf source GigabitEthernet 0/0/0`
		Enables the high-speed logging of all VPN and non-VPN translations for up to four destinations. You can enable logging for a specific destination VRF using the **vrf** keyword. To specify an IPv6 address for the UDP destination, use the **ipv6-destination** keyword followed by the IPv6 address.

Step 4	**nat64 logging translations flow-export v9** {*vrf-name*	**global-on**} **Example:** `Device(config)# nat64 logging translations flow-export v9 VPN-18`	Enables or disables the high-speed logging of specific NAT VPN translations.

	Command or Action	Purpose
Step 5	**exit** **Example:** `Device(config)# exit`	(Optional) Exits global configuration mode and enters privileged EXEC mode.

FTP64 Application-Level Gateway Support

The FTP64 (or service FTP) application-level gateway (ALG) helps stateful Network Address Translation 64 (NAT64) to operate on Layer 7 data. FTP64 ALG translates IP addresses and the TCP port information embedded in the payload of an FTP control session.

NAT translates any TCP/UDP traffic that does not carry source and destination IP addresses in the application data stream. Protocols that embed the IP address information within the payload (or in the application data stream) require the support of an ALG. ALGs handle application data stream (Layer 7) protocol-specific services, such as translating embedded IP addresses and port numbers in the packet payload and extracting new connection or session information from control channels.

FTP64 is automatically enabled when Stateful NAT64 is enabled. Use the **no nat64 service ftp** command to disable the NAT64 FTP service.

Note The FTP64 ALG is not supported in Stateless NAT64 translation.

Note The FTP64 ALG does not support IPv4-compatible IPv6 addresses.

Based on *IPv6-to-IPv4 translation FTP considerations draft-ietf-behave-ftp64-02* and RFC 2228, the FTP64 ALG must switch to transparent mode (a device in a transparent mode is invisible in the network; however, this device can act as a bridge and inspect or filter packets), when commands and responses flow between the FTP client and the FTP server. When a client issues the FTP AUTH command, the FTP64 ALG transparently forwards all data on the control channel in both (ingress and egress) directions, until the end of the control channel session. Similarly, during an AUTH negotiation, the ALG must be in transparent mode, whether the negotiation is successful or not.

Based on RFC 6384, the behavior of the FTP64 ALG during a client-server communication is different. During an IPv6-to-IPv4 translation, the FTP64 ALG must transparently copy data transmitted over the control channel so that the transport layer security (TLS) session works correctly. However, the client commands and server responses are hidden from the FTP64 ALG. To ensure a consistent behavior, as soon as the initial FTP AUTH command is issued by a client, the FTP64 ALG must stop translating commands and responses and start transparently copying TCP data that is sent by the server to the client and vice versa. The FTP64 ALG must ignore the AUTH command and not go into transparent mode if the server response is in the 4*xx* or 5*xx* ranges, which comprise FTP error/warning messages.

Prior to CSCtu37975, when an IPv6 FTP client issues an FTP AUTH command, irrespective of whether the IPv4 FTP server accepts or rejects that authorization negotiation, the FTP64 ALG moves the AUTH session to transparent mode (or bypass mode). When a session is in transparent mode, NAT cannot perform translation on the packets within the session. With CSCtu37975, during a client-server communication, the FTP64 ALG's behavior is compliant with RFC 6384.

FTP64 NAT ALG Intrabox High Availability Support

Depending on your release, the FTP64 application-level gateway (ALG) adds high availability (HA) support for Stateful NAT64. The FTP64 NAT ALG Intrabox HA Support feature supports the stateful switchover between redundant Forward Processors (FPs) within a single chassis. The HA support provided by the FTP64 ALG is applicable to both intrabox HA and In-Service Software Upgrade (ISSU).

Use the **no nat64 service ftp** command to disable the NAT64 ALG service.

The FTP64 ALG synchronizes data when it receives the following messages:

- User authentication flag after 230 replies.
- ALG disable/enable flag after ALG ENABLE and ALG DISABLE messages are received.
- Fragment detection information after the first segmented packet is detected.
- Fragment detection information after the end of the segmentation is detected.



Note

- Stateful NAT64 supports only intrabox HA in some releases.
- FTP64 ALG statistics and FTP64 debug logs are not synchronized to the standby device by the FTP64 ALG.

Stateful NAT64—Intrachassis Redundancy

Depending on your release, support for the Stateful NAT64—Intrachassis Redundancy feature is available. When a second Forward Processor (FP) is available inside a single chassis, the Stateful NAT64—Intrachassis Redundancy feature enables you to configure the second FP as a standby entity. When you plug in the second FP, redundancy starts automatically with no explicit configuration. There is a short delay before the standby FP becomes the “hot standby” (which means that all sessions have been synchronized). The standby FP maintains a backup of the Stateful NAT64 session information, and when the active (first) FP fails, there is very little disruption of NAT64 sessions.

NAT64 redundancy information is sent to the standby FP in the following instances:

- When a session or a dynamic bind is created.
- When a session or a dynamic bind is deleted.
- During periodic updates. Based on the time elapsed, the active FP periodically updates the state information to the standby. Not all changes in the replicated objects are sent immediately to the standby at the time of change. The most critical updates are sent immediately, and other changes are communicated by periodic updates.

When a standby FP is inserted or when a standby FP recovers from a reload, the active FP performs a bulk synchronization to synchronize the standby FP with the active FP. NAT does an aggressive synchronization by which the active FP pushes all the state information forcefully to the standby FP.

In addition to NAT64 session information, application-specific information (application-level gateway [ALG] information) also has to be communicated to the standby FP. Each ALG has a per-session state that needs to be synchronized in the standby. The ALG triggers the sending of all ALG state information to the standby

FP. NAT provides the mechanism for actually sending the ALG state and associates the state to a particular session.

HTTP sessions are not backed up on the standby FP. To replicate HTTP sessions on the standby FP during a switchover, you must configure the **nat64 switchover replicate http enable** command.



Note The Stateful NAT64—Intrachassis Redundancy feature does not support box-to-box (B2B) redundancy or asymmetric routing.

Asymmetric Routing Support for NAT64

In Cisco IOS XE Release and later releases, Network Address Translation 64 (NAT64) supports asymmetric routing and asymmetric routing with Multiprotocol Label Switching (MPLS). In NAT 64, MPLS is enabled on the IPv4 interface. Packets coming from the IPv6 interface are switched to the IPv4 interface. No configuration changes are required to enable asymmetric routing or asymmetric routing with MPLS.

For more information, see the section “Example: Configuring Asymmetric Routing Support for NAT64”.

How to Configure Stateful Network Address Translation 64

Based on your network configuration, you can configure static, dynamic, or dynamic Port Address Translation (PAT) Stateful NAT64.



Note You need to configure at least one of the configurations described in the following tasks for Stateful NAT64 to work.

Configuring Static Stateful Network Address Translation 64

You can configure a static IPv6 address to an IPv4 address and vice versa. Optionally, you can configure static Stateful NAT64 with or without ports. Perform this task to configure static Stateful NAT64.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ipv6 unicast-routing**
4. **interface** *type number*
5. **description** *string*
6. **ipv6 enable**
7. **ipv6 address** {*ipv6-address*/*prefix-length* | *prefix-name sub-bits*/*prefix-length*}
8. **nat64 enable**
9. **exit**
10. **interface** *type number*

11. **description** *string*
12. **ip address** *ip-address mask*
13. **nat64 enable**
14. **exit**
15. **nat64 prefix stateful** *ipv6-prefix/length*
16. **nat64 v6v4 static** *ipv6-address ipv4-address*
17. **end**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ipv6 unicast-routing** **Example:** `Device(config)# ipv6 unicast-routing`	Enables the forwarding of IPv6 unicast datagrams.
Step 4	**interface** *type number* **Example:** `Device(config)# interface gigabitethernet 0/0/0`	Configures an interface type and enters interface configuration mode.
Step 5	**description** *string* **Example:** `Device(config-if)# description interface facing ipv6`	Adds a description to an interface configuration.
Step 6	**ipv6 enable** **Example:** `Device(config-if)# ipv6 enable`	Enables IPv6 processing on an interface.
Step 7	**ipv6 address** {*ipv6-address/prefix-length* | *prefix-name sub-bits/prefix-length*} **Example:** `Device(config-if)# ipv6 address 2001:DB8:1::1/96`	Configures an IPv6 address based on an IPv6 general prefix and enables IPv6 processing on an interface.
Step 8	**nat64 enable** **Example:**	Enables NAT64 translation on an IPv6 interface.

	Command or Action	Purpose
	`Device(config-if)# nat64 enable`	
Step 9	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and enters global configuration mode.
Step 10	**interface** *type number* **Example:** `Device(config)# interface gigabitethernet 1/2/0`	Configures an interface and enters interface configuration mode.
Step 11	**description** *string* **Example:** `Device(config-if)# description interface facing ipv4`	Adds a description to an interface configuration.
Step 12	**ip address** *ip-address mask* **Example:** `Device(config-if)# ip address 209.165.201.1 255.255.255.0`	Configures an IPv4 address for an interface.
Step 13	**nat64 enable** **Example:** `Device(config-if)# nat64 enable`	Enables NAT64 translation on an IPv4 interface.
Step 14	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and enters global configuration mode.
Step 15	**nat64 prefix stateful** *ipv6-prefix/length* **Example:** `Device(config)# nat64 prefix stateful 2001:DB8:1::1/96`	Defines the Stateful NAT64 prefix to be added to IPv4 hosts to translate the IPv4 address into an IPv6 address. • The Stateful NAT64 prefix can be configured at the global configuration level or at the interface level.
Step 16	**nat64 v6v4 static** *ipv6-address ipv4-address* **Example:** `Device(config)# nat64 v6v4 static 2001:DB8:1::FFFE 209.165.201.1`	Enables NAT64 IPv6-to-IPv4 static address mapping.
Step 17	**end** **Example:** `Device(config)# end`	Exits global configuration mode and enters privileged EXEC mode.

Configuring Dynamic Stateful Network Address Translation 64

A dynamic Stateful NAT64 configuration provides a one-to-one mapping of IPv6 addresses to IPv4 addresses in the address pool. You can use the dynamic Stateful NAT64 configuration when the number of active IPv6 hosts is less than the number of IPv4 addresses in the pool. Perform this task to configure dynamic Stateful NAT64.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ipv6 unicast-routing**
4. **interface** *type number*
5. **description** *string*
6. **ipv6 enable**
7. **ipv6** {*ipv6-address*/*prefix-length* | *prefix-name sub-bits*/*prefix-length*}
8. **nat64 enable**
9. **exit**
10. **interface** *type number*
11. **description** *string*
12. **ip address** *ip-address mask*
13. **nat64 enable**
14. **exit**
15. **ipv6 access-list** *access-list-name*
16. **permit ipv6** *ipv6-address* **any**
17. **exit**
18. **nat64 prefix stateful** *ipv6-prefix*/*length*
19. **nat64 v4 pool** *pool-name start-ip-address end-ip-address*
20. **nat64 v6v4 list** *access-list-name* **pool** *pool-name*
21. **end**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.

	Command or Action	Purpose
Step 3	**ipv6 unicast-routing** **Example:** `Device(config)# ipv6 unicast-routing`	Enables the forwarding of IPv6 unicast datagrams.
Step 4	**interface** *type number* **Example:** `Device(config)# interface gigabitethernet 0/0/0`	Configures an interface type and enters interface configuration mode.
Step 5	**description** *string* **Example:** `Device(config-if)# description interface facing ipv6`	Adds a description to an interface configuration.
Step 6	**ipv6 enable** **Example:** `Device(config-if)# ipv6 enable`	Enables IPv6 processing on an interface.
Step 7	**ipv6** {*ipv6-address/prefix-length* | *prefix-name sub-bits/prefix-length*} **Example:** `Device(config-if)# ipv6 2001:DB8:1::1/96`	Configures an IPv6 address based on an IPv6 general prefix and enables IPv6 processing on an interface.
Step 8	**nat64 enable** **Example:** `Device(config-if)# nat64 enable`	Enables Stateful NAT64 translation on an IPv6 interface.
Step 9	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and enters global configuration mode.
Step 10	**interface** *type number* **Example:** `Device(config)# interface gigabitethernet 1/2/0`	Configures an interface type and enters interface configuration mode
Step 11	**description** *string* **Example:** `Device(config-if)# description interface facing ipv4`	Adds a description to an interface configuration.
Step 12	**ip address** *ip-address mask* **Example:** `Device(config-if)# ip address 209.165.201.24 255.255.255.0`	Configures an IPv4 address for an interface.

	Command or Action	Purpose
Step 13	**nat64 enable** **Example:** `Device(config-if)# nat64 enable`	Enables Stateful NAT64 translation on an IPv4 interface.
Step 14	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and enters global configuration mode.
Step 15	**ipv6 access-list** *access-list-name* **Example:** `Device(config)# ipv6 access-list nat64-acl`	Defines an IPv6 access list and enters IPv6 access list configuration mode.
Step 16	**permit ipv6** *ipv6-address* **any** **Example:** `Device(config-ipv6-acl)# permit ipv6 2001:DB8:2::/96 any`	Sets permit conditions for an IPv6 access list.
Step 17	**exit** **Example:** `Device(config-ipv6-acl# exit`	Exits IPv6 access list configuration mode and enters global configuration mode.
Step 18	**nat64 prefix stateful** *ipv6-prefix/length* **Example:** `Device(config)# nat64 prefix stateful 2001:DB8:1::1/96`	Enables NAT64 IPv6-to-IPv4 address mapping.
Step 19	**nat64 v4 pool** *pool-name start-ip-address end-ip-address* **Example:** `Device(config)# nat64 v4 pool pool1 209.165.201.1 209.165.201.254`	Defines the Stateful NAT64 IPv4 address pool.
Step 20	**nat64 v6v4 list** *access-list-name* **pool** *pool-name* **Example:** `Device(config)# nat64 v6v4 list nat64-acl pool pool1`	Dynamically translates an IPv6 source address to an IPv6 source address and an IPv6 destination address to an IPv4 destination address for NAT64.
Step 21	**end** **Example:** `Device(config)# end`	Exits global configuration mode and enters privileged EXEC mode.

Configuring Dynamic Port Address Translation Stateful NAT64

A Port Address Translation (PAT) or overload configuration is used to multiplex (mapping IPv6 addresses to a single IPv4 pool address) multiple IPv6 hosts to a pool of available IPv4 addresses on a first-come first-served basis. The dynamic PAT configuration conserves the IPv4 address space while providing

connectivity to the IPv4 Internet. Configure the **nat64 v6v4 list** command with the **overload** keyword to configure PAT address translation. Perform this task to configure dynamic PAT Stateful NAT64.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ipv6 unicast-routing**
4. **interface** *type number*
5. **description** *string*
6. **ipv6 enable**
7. **ipv6** {*ipv6-address*/*prefix-length* | *prefix-name sub-bits*/*prefix-length*}
8. **nat64 enable**
9. **exit**
10. **interface** *type number*
11. **description** *string*
12. **ip address** *ip-address mask*
13. **nat64 enable**
14. **exit**
15. **ipv6 access-list** *access-list-name*
16. **permit ipv6** *ipv6-address* **any**
17. **exit**
18. **nat64 prefix stateful** *ipv6-prefix*/*length*
19. **nat64 v4 pool** *pool-name start-ip-address end-ip-address*
20. **nat64 v6v4 list** *access-list-name* **pool** *pool-name* **overload**
21. **end**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ipv6 unicast-routing** **Example:** `Device(config)# ipv6 unicast-routing`	Enables the forwarding of IPv6 unicast datagrams.

	Command or Action	Purpose
Step 4	**interface** *type number* **Example:** `Device(config)# interface gigabitethernet 0/0/0`	Configures an interface type and enters interface configuration mode.
Step 5	**description** *string* **Example:** `Device(config-if)# description interface facing ipv6`	Adds a description to an interface configuration.
Step 6	**ipv6 enable** **Example:** `Device(config-if)# ipv6 enable`	Enables IPv6 processing on an interface.
Step 7	**ipv6** {*ipv6-address/prefix-length* | *prefix-name sub-bits/prefix-length*} **Example:** `Device(config-if)# ipv6 2001:DB8:1::1/96`	Configures an IPv6 address based on an IPv6 general prefix and enables IPv6 processing on an interface.
Step 8	**nat64 enable** **Example:** `Device(config-if)# nat64 enable`	Enables Stateful NAT64 translation on an IPv6 interface.
Step 9	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and enters global configuration mode.
Step 10	**interface** *type number* **Example:** `Device(config)# interface gigabitethernet 1/2/0`	Configures an interface type and enters interface configuration mode
Step 11	**description** *string* **Example:** `Device(config-if)# description interface facing ipv4`	Adds a description to an interface configuration.
Step 12	**ip address** *ip-address mask* **Example:** `Device(config-if)# ip address 209.165.201.24 255.255.255.0`	Configures an IPv4 address for an interface.
Step 13	**nat64 enable** **Example:** `Device(config-if)# nat64 enable`	Enables Stateful NAT64 translation on an IPv6 interface.

	Command or Action	Purpose
Step 14	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and enters global configuration mode.
Step 15	**ipv6 access-list** *access-list-name* **Example:** `Device(config)# ipv6 access-list nat64-acl`	Defines an IPv6 access list and places the device in IPv6 access list configuration mode.
Step 16	**permit ipv6** *ipv6-address* **any** **Example:** `Device(config-ipv6-acl)# permit ipv6 2001:db8:2::/96 any`	Sets permit conditions for an IPv6 access list.
Step 17	**exit** **Example:** `Device(config-ipv6-acl)# exit`	Exits IPv6 access list configuration mode and enters global configuration mode.
Step 18	**nat64 prefix stateful** *ipv6-prefix/length* **Example:** `Device(config)# nat64 prefix stateful 2001:db8:1::1/96`	Enables NAT64 IPv6-to-IPv4 address mapping.
Step 19	**nat64 v4 pool** *pool-name start-ip-address end-ip-address* **Example:** `Device(config)# nat64 v4 pool pool1 209.165.201.1 209.165.201.254`	Defines the Stateful NAT64 IPv4 address pool.
Step 20	**nat64 v6v4 list** *access-list-name* **pool** *pool-name* **overload** **Example:** `Device(config)# nat64 v6v4 list nat64-acl pool pool1 overload`	Enables NAT64 PAT or overload address translation.
Step 21	**end** **Example:** `Device(config)# end`	Exits global configuration mode and enters privileged EXEC mode.

Restrictions for Enabling Stateful Network Address Conversion using VRF

- One to one static IPv6 and IPv4 mapping is not supported.
- Inter VRF translations are not supported with NAT64.
- NAT64 pools and NAT44 pools do not support IP address sharing.

Configuring VRF Aware Stateful NAT64 with Carrier Grade NAT

A Port Address Translation (PAT) or overload configuration is used to multiplex multiple IPv6 hosts to a pool of available IPv4 addresses on a first come first served basis. The dynamic PAT configuration conserves the IPv4 address space while providing connectivity to the IPv4 Internet. Configuring the **nat64 v6v4** list command with the **overload** keyword, configures the PAT address translation and the **vrf** keyword segregates the incoming and outgoing traffic in the network.

Perform this task to enable stateful NAT64 conversion within a Virtual routing and forwarding (VRF) network.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ipv6 unicast-routing**
4. **interface** *type number*
5. **description** *string*
6. **vrf forwarding***vrf-name*
7. **ipv6 enable**
8. **ipv6** {*ipv6-address*/*prefix-length* | *prefix-name sub-bits*/*prefix-length*}
9. **nat64 enable**
10. **exit**
11. **interface** *type number*
12. **description** *string*
13. **vrf forwarding***vrf-name*
14. **ip address** *ip-address mask*
15. **nat64 enable**
16. **exit**
17. **ipv6 access-list** *access-list-name*
18. **permit ipv6** *ipv6-address* **any**
19. **exit**
20. **nat64 settings mode cgn**
21. **nat64 prefix stateful** *ipv6-prefix*/*length*/*vrf*
22. **nat64 v4 pool** *pool-name start-ip-address end-ip-address*
23. **nat64 v6v4 list** *access-list-name* **pool** *pool-name* **vrf***vrf-name***overloadmatch-in-vrf**
24. **end**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.

	Command or Action	Purpose
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ipv6 unicast-routing** **Example:** `Device(config)# ipv6 unicast-routing`	Enables the forwarding of IPv6 unicast data packets.
Step 4	**interface** *type number* **Example:** `Device(config)# interface gigabitethernet 0/0/0`	Configures an interface type and enters interface configuration mode.
Step 5	**description** *string* **Example:** `Device(config-if)# description interface facing ipv6`	Adds a description to an interface configuration.
Step 6	**vrf forwarding***vrf-name* **Example:** `Device(config-if)# vrf forwarding test`	Enables IPv6 VRF routing.
Step 7	**ipv6 enable** **Example:** `Device(config-if)# ipv6 enable`	Enables IPv6 processing on an interface.
Step 8	**ipv6** {*ipv6-address*/*prefix-length* | *prefix-name sub-bits*/*prefix-length*} **Example:** `Device(config-if)# ipv6 2001:DB8:1::1/96`	Configures an IPv6 address based on an IPv6 general prefix and enables IPv6 processing on an interface.
Step 9	**nat64 enable** **Example:** `Device(config-if)# nat64 enable`	Enables Stateful NAT64 translation on an IPv6 interface.
Step 10	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and enters global configuration mode.
Step 11	**interface** *type number* **Example:** `Device(config)# interface gigabitethernet 1/2/0`	Configures an interface type and enters interface configuration mode
Step 12	**description** *string* **Example:**	Adds a description to an interface configuration.

	Command or Action	Purpose
	`Device(config-if)# description interface facing ipv4`	
Step 13	**vrf forwarding***vrf-name* **Example:** `Device(config-if)# vrf forwarding test`	Enables IPv4 VRF routing.
Step 14	**ip address** *ip-address mask* **Example:** `Device(config-if)# ip address 209.165.201.24 255.255.255.0`	Configures an IPv4 address for an interface.
Step 15	**nat64 enable** **Example:** `Device(config-if)# nat64 enable`	Enables Stateful NAT64 translation on an IPv6 interface.
Step 16	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and enters global configuration mode.
Step 17	**ipv6 access-list** *access-list-name* **Example:** `Device(config)# ipv6 access-list nat64-acl`	Defines an IPv6 access list and enters IPv6 access list configuration mode.
Step 18	**permit ipv6** *ipv6-address* **any** **Example:** `Device(config-ipv6-acl)# permit ipv6 2001:DB8:2::/96 any`	Sets permit conditions for an IPv6 access list.
Step 19	**exit** **Example:** `Device(config-ipv6-acl)# exit`	Exits IPv6 access list configuration mode and enters global configuration mode.

Step 20	**nat64 settings mode cgn**
	Example:
	`Device(config)# nat64 settings mode cgn`
	Example:
	`Device(config)# no nat64 settings mode cgn`
		Enables the CGN operating mode.
		• To disable the CGN operating mode, use the **no** form of this command.

Step 21	**nat64 prefix stateful** *ipv6-prefix/length/vrf* **Example:** `Device(config)# nat64 prefix stateful 2001:DB8:1::1/96 vrf 1`	Enables NAT64 IPv6-to-IPv4 address mapping.
Step 22	**nat64 v4 pool** *pool-name start-ip-address end-ip-address* **Example:**	Defines the stateful NAT64 IPv4 address pool.

	Command or Action	Purpose
	`Device(config)# nat64 v4 pool pool1 209.165.201.1 209.165.201.254`	
Step 23	**nat64 v6v4 list** *access-list-name* **pool** *pool-name* **vrf***vrf-name***overloadmatch-in-vrf** **Example:** `Device(config)# nat64 v6v4 list nat64-acl pool pool1 vrf 1 overload match-in-vrf`	Enables NAT64 conversion for a VRF network.
Step 24	**end** **Example:** `Device(config)# end`	Exits global configuration mode and enters privileged EXEC mode.

Verifying VRF Aware Stateful NAT64 with Carrier Grade NAT (CGN)

SUMMARY STEPS

1. **show nat64 translations**
2. **show ip route vrf***vrf*
3. **show ipv6 route vrf***vrf*
4. **show platform hardware qfp active feature nat64 data statistics**
5. **show platform hardware qfp active feature nat64 data sess-dump**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	**show nat64 translations** **Example:** `Device# show nat64 translations`	Displays all NAT 64 protocol information.
Step 2	**show ip route vrf***vrf* **Example:** `Device# show ip route vrf VRF1`	Displays the IP routing table associated with a VRF.
Step 3	**show ipv6 route vrf***vrf* **Example:** `Device(config)# show ipv6 route vrf VRF1`	Displays IPv6 routing table information associated with a VPN routing and forwarding (VRF) instance.
Step 4	**show platform hardware qfp active feature nat64 data statistics** **Example:**	Displays all NAT64 translation global statistics.

	Command or Action	Purpose
	`Device(config)# show platform hardware qfp active feature nat64 data statistics`	
Step 5	**show platform hardware qfp active feature nat64 data sess-dump** **Example:** `Device(config)# show platform hardware qfp active feature nat64 data sess-dump`	Displays all NAT64 session dump information.

Monitoring and Maintaining a Stateful NAT64 Routing Network

Use the following commands in any order to display the status of your Stateful Network Address Translation 64 (NAT64) configuration.

SUMMARY STEPS

1. **show nat64 aliases** [*lower-address-range upper-address-range*]
2. **show nat64 logging**
3. **show nat64 prefix stateful** {**global** | {**interfaces** | **static-routes**} [**prefix** *ipv6-address/prefix-length*]}
4. **show nat64 timeouts**

DETAILED STEPS

Procedure

Step 1 **show nat64 aliases** [*lower-address-range upper-address-range*]

This command displays the IP aliases created by NAT64.

Example:

```
Device# show nat64 aliases

Aliases configured: 1
Address   Table ID  Inserted   Flags   Send ARP   Reconcilable   Stale   Ref-Count
10.1.1.1  0         FALSE      0x0030  FALSE      TRUE           FALSE   1
```

Step 2 **show nat64 logging**

This command displays NAT64 logging.

Example:

```
Device# show nat64 logging

NAT64 Logging Type

Method       Protocol   Dst. Address   Dst. Port   Src. Port
translation
flow export  UDP        10.1.1.1       5000        60087
```

Step 3 **show nat64 prefix stateful** {**global** | {**interfaces** | **static-routes**} [**prefix** *ipv6-address/prefix-length*]}

This command displays information about NAT64 stateful prefixes.

Example:

```
Device# show nat64 prefix stateful interfaces

Stateful Prefixes

Interface               NAT64   Enabled   Global Prefix
GigabitEthernet0/1/0    TRUE    TRUE      2001:DB8:1:1/96
GigabitEthernet0/1/3    TRUE    FALSE     2001:DB8:2:2/96
```

Step 4 **show nat64 timeouts**

This command displays statistics for NAT64 translation session timeout.

Example:

```
Device# show nat64 timeouts

NAT64 Timeout

Seconds    CLI Cfg    Uses 'All'    all flows
86400      FALSE      FALSE         udp
300        FALSE      TRUE          tcp
7200       FALSE      TRUE          tcp-transient
240        FALSE      FALSE         icmp
60         FALSE      TRUE
```

Configuration Examples for Stateful Network Address Translation 64

Example: Configuring Static Stateful Network Address Translation 64

```
Device# configure terminal
Device(config)# ipv6 unicast-routing
Device(config)# interface gigabitethernet 0/0/0
Device(config-if)# description interface facing ipv6
Device(config-if)# ipv6 enable
Device(config-if)# ipv6 address 2001:DB8:1::1/96
Device(confif-if)# nat64 enable
Device(config-fi)# exit
Device(config)# interface gigabitethernet 1/2/0
Device(config-if)# description interface facing ipv4
Device(config-if)# ip address 209.165.201.1 255.255.255.0
Device(config-if)# nat64 enable
Device(config-if)# exit
Device(config)# nat64 prefix stateful 2001:DB8:1::1/96
Device(config)# nat64 v6v4 static 2001:DB8:1::FFFE 209.165.201.1
Device(config)# end
```

Example: Configuring Dynamic Stateful Network Address Translation 64

```
Device# configure terminal
Device(config)# ipv6 unicast-routing
Device(config)# interface gigabitethernet 0/0/0
Device(config-if)# description interface facing ipv6
Device(config-if)# ipv6 enable
Device(config-if)# ipv6 2001:DB8:1::1/96
Device(config-if)# nat64 enable
Device(config-if)# exit
Device(config)# interface gigabitethernet 1/2/0
Device(config-if)# description interface facing ipv4
Device(config-if)# ip address 209.165.201.24 255.255.255.0
Device(config-if)# nat64 enable
Device(config-if)# exit
Device(config)# ipv6 access-list nat64-acl
Device(config-ipv6-acl)# permit ipv6 2001:db8:2::/96 any
Device(config-ipv6-acl)# exit
Device(config)# nat64 prefix stateful 2001:db8:1::1/96
Device(config)# nat64 v4 pool pool1 209.165.201.1 209.165.201.254
Device(config)# nat64 v6v4 list nat64-acl pool pool1
Device(config)# end
```

Example: Configuring Dynamic Port Address Translation Stateful NAT64

```
enable
 configure terminal
  ipv6 unicast-routing
  interface gigabitethernet 0/0/0
   description interface facing ipv6
   ipv6 enable
   ipv6 2001:DB8:1::1/96
   nat64 enable
   exit
  interface gigabitethernet 1/2/0
   description interface facing ipv4
   ip address 209.165.201.24 255.255.255.0
   nat64 enable
   exit
  ipv6 access-list nat64-acl
   permit ipv6 2001:db8:2::/96 any
   exit
  nat64 prefix stateful 2001:db8:1::1/96
  nat64 v4 pool pool1 209.165.201.1 209.165.201.254
  nat64 v6v4 list nat64-acl pool pool1 overload
  end
```

Example: Configuring Asymmetric Routing Support for NAT64

The following example shows how to configure asymmetric routing for Network Address Translation 64 (NAT64):

```
!RouterA Configuration
```

```
Device(config)# ipv6 unicast-routing
Device(config)# nat64 prefix stateful 2001:db8:2::/96
Device(config)# nat64 v6v4 static 2001:db8:1::5 150.0.0.1 redundancy 1 mapping-id 150
Device(config)# nat64 v6v4 static 2001:db8:1::6 150.0.0.2 redundancy 1 mapping-id 151
Device(config)# nat64 switchover replicate http enable port 80
!
Device(config)# redundancy
Device(config-red)# application redundancy
Device(config-red-app)# group 1
Device(config-red-app-grp)# name RG1
Device(config-red-app-grp)# data gigabitethernet 1/1/0
Device(config-red-app-grp)# control gigabitethernet 1/1/1 protocol 1
Device(config-red-app-grp)# asymmetric-routing interface gigabitethernet 1/1/2
Device(config-red-app-grp)# priority 150 failover threshold 140
Device(config-red-app-grp)# asymmetric-routing always-divert enable
Device(config-red-app-grp)# exit
Device(config-red-app)# exit
Device(config-red)# exit
!
Device(config)# interface gigabitethernet 1/1/0
Device(config-if)# ip address 10.10.10.1 255.255.255.0
Device(config-if)# no shutdown
Device(config-if)# exit
!
Device(config)# interface gigabitethernet 1/1/1
Device(config-if)# ip address 172.16.0.1 255.240.0.0
Device(config-if)# no shutdown
Device(config-if)# exit
!
Device(config)# interface gigabitethernet 1/1/2
Device(config-if)# ip address 192.168.0.1 255.255.0.0
Device(config-if)# no shutdown
Device(config-if)# exit
!
Device(config)# interface gigabitethernet 1/1/3
Device(config-if)# ipv6 enable
Device(config-if)# no shutdown
Device(config-if)# nat64 enable
Device(config-if)# ipv6 address 2001:db8:1::2/96
Device(config-if)# redundancy rii 100
Device(config-if)# redundancy group 1 ipv6 2001:db8:1::1/96 exclusive decrement 15
Device(config-if)# exit
!
Device(config)# interface gigabitethernet 1/1/4
Device(config-if)# ip address 192.0.2.1 255.255.255.0
Device(config-if)# nat64 enable
Device(config-if)# no shutdown
Device(config-if)# redundancy rii 101
Device(config-if)# redundancy asymmetric-routing enable
Device(config-if)# exit
!
Device(config)# router ospf 90
Device(config-router)# network 192.0.2.0 255.255.255.0 area 0
Device(config-router)# end


! Router B Configuration

Device(config)# ipv6 unicast-routing
Device(config)# nat64 prefix stateful 2001:db8:2::/96
Device(config)# nat64 v6v4 static 2001:db8:1::5 150.0.0.1 redundancy 1 mapping-id 150
Device(config)# nat64 v6v4 static 2001:db8:1::6 150.0.0.2 redundancy 1 mapping-id 151
```

```
Device(config)# nat64 switchover replicate http enable port 80
!
Device(config)# redundancy
Device(config-red)# application redundancy
Device(config-red-app)# group 1
Device(config-red-app-grp)# name RG1
Device(config-red-app-grp)# data gigabitethernet 1/2/0
Device(config-red-app-grp)# control gigabitethernet 1/2/1 protocol 1
Device(config-red-app-grp)# asymmetric-routing interface gigabitethernet 1/2/2
Device(config-red-app-grp)# priority 140 failover threshold 135
Device(config-red-app-grp)# asymmetric-routing always-divert enable
Device(config-red-app-grp)# exit
Device(config-red-app)# exit
Device(config-red)# exit
!
Device(config)# interface gigabitethernet 1/2/0
Device(config-if)# ip address 10.10.10.2 255.255.255.0
Device(config-if)# no shutdown
Device(config-if)# exit
!
Device(config)# interface gigabitethernet 1/2/1
Device(config-if)# ip address 172.16.0.2 255.240.0.0
Device(config-if)# no shutdown
Device(config-if)# exit
!
Device(config)# interface gigabitethernet 1/2/2
Device(config-if)# ip address 192.168.0.2 255.255.0.0
Device(config-if)#no shutdown
Device(config-if)# exit
!
Device(config-if)# interface gigabitethernet 1/2/3
Device(config-if)# ipv6 enable
Device(config-if)# no shutdown
Device(config-if)# nat64 enable
Device(config-if)# ipv6 addr 2001:db8:1::3/96
Device(config-if)# redundancy rii 100
Device(config-if)# redundancy group 1 ipv6 2001:db8:1::1/96 exclusive decrement 15
Device(config-if)# exit
!
Device(config)# interface gigabitethernet 1/2/4
Device(config-if)# ip address 198.51.100.1 255.255.255.0
Device(config-if)# nat64 enable
Device(config-if)# no shutdown
Device(config-if)# redundancy rii 101
Device(config-if)# redundancy asymmetric-routing enable
Device(config-if)# exit
!
Device(config)# router ospf 90
Device(config-router)# network 198.51.100.0 255.255.255.0 area 0
Device(config-router)# end
```

Additional References for Stateful Network Address Translation 64

Related Documents

Related Topic	Document Title
Cisco IOS commands	Master Command List, All Releases
NAT commands	IP Addressing Services Command Reference

Standards and RFCs

Standard/RFC	Title
Framework for IPv4/IPv6 Translation	Framework for IPv4/IPv6 Translation draft-ietf-behave-v6v4-framework-06
FTP ALG for IPv6-to-IPv4 translation	An FTP ALG for IPv6-to-IPv4 translation draft-ietf-behave-ftp64-06
IP/ICMP Translation Algorithm	IP/ICMP Translation Algorithm draft-ietf-behave-v6v4-xlate-10
IPv6 Addressing of IPv4/IPv6 Translators	IPv6 Addressing of IPv4/IPv6 Translators draft-ietf-behave-address-format-07
RFC 2228	FTP Security Extensions
RFC 2373	IP Version 6 Addressing Architecture
RFC 2464	Transmission of IPv6 Packets over Ethernet Networks
RFC 2765	Stateless IP/ICMP Translation Algorithm (SIIT)
RFC 2766	Network Address Translation - Protocol Translation (NAT-PT)
RFC 4787	Network Address Translation (NAT) Behavioral Requirements for Unicast UDP
RFC 4966	Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status
RFC 6384	An FTP Application Layer Gateway (ALG) for IPv6-to-IPv4 Translation
Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers	Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers draft-ietf-behave-v6v4-xlate-stateful-12

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Stateful Network Address Translation 64

Table 4: Feature Information for Stateful Network Address Translation 64

Feature Name	Releases	Feature Information
Asymmetric Routing Support for NAT64	Cisco IOS XE Release 3.16S	In Cisco IOS XE Release and later releases, Network Address Translation 64 (NAT64) supports asymmetric routing and asymmetric routing with Multiprotocol Label Switching (MPLS).
FTP64 NAT ALG Intrabox HA Support	Cisco IOS XE Release 3.5S	In Cisco IOS XE Release 3.5S, the FTP64 ALG adds HA support for Stateful NAT64. The FTP64 NAT ALG Intrabox HA Support feature supports the stateful switchover between redundant FPs within a single chassis. The HA support provided by the FTP64 ALG is applicable to both intrabox and interbox HA and In-Service Software Upgrade (ISSU).
Stateful NAT64 ALG—Stateful FTP64 ALG Support	Cisco IOS XE Release 3.4S	Cisco IOS XE Release 3.4S and later releases support FTP64 (or service FTP) ALGs. The FTP64 ALG helps Stateful NAT64 operate on Layer 7 data. An FTP ALG translates IP addresses and the TCP port information embedded in the payload of an FTP control session. The following commands were introduced or modified: **nat64 service ftp**.

Feature Name	Releases	Feature Information
Stateful NAT64—Intra-Chassis Redundancy	Cisco IOS XE Release 3.5S Cisco IOS XE Release 3.10S	Cisco IOS XE Release 3.5S and later releases support the Stateful NAT64—Intra-Chassis Redundancy feature. When a second Forward Processor (FP) is available inside a single chassis, the Stateful NAT64 Intra-Chassis Redundancy feature enables you to configure the second FP as a standby entity. The standby FP maintains a backup of the stateful NAT64 session information and when the active (first) FP fails, there is no disruption of NAT64 sessions. The following commands were introduced or modified: **nat64 switchover replicate http port**.
Stateful Network Address Translation 64	Cisco IOS XE Release 3.4S	The Stateful Network Address Translation 64 feature provides a translation mechanism that translates IPv6 packets into IPv4 packets and vice versa. The Stateful NAT64 translator, algorithmically translates the IPv4 addresses of IPv4 hosts to and from IPv6 addresses by using the configured stateful prefix. In a similar manner, the IPv6 addresses of IPv6 hosts are translated to and from IPv4 addresses through NAT. The following commands were introduced or modified: **clear nat64 statistics**, **debug nat64**, **nat64 logging**, **nat64 prefix stateful**, **nat64 translation**, **nat64 v4**, **nat64 v4v6**, **nat64 v6v4**, **show nat64 aliases**, **show nat64 limits**, **show nat64 logging**, **show nat64 mappings dynamic**, **show nat64 mappings static**, **show nat64 services**, **show nat64 pools**, **show nat64 prefix stateful**, **show nat64 statistics**, **show nat64 timeouts**, and **show nat64 translations**.

Glossary

ALG—application-layer gateway or application-level gateway.

FP—Forward Processor.

IPv4-converted address—IPv6 addresses used to represent the IPv4 hosts. These have an explicit mapping relationship to the IPv4 addresses. This relationship is self-described by mapping the IPv4 address in the IPv6 address. Both stateless and stateful translators use IPv4-converted IPv6 addresses to represent the IPv4 hosts.

IPv6-converted address—IPv6 addresses that are assigned to the IPv6 hosts for the stateless translator. These IPv6-converted addresses have an explicit mapping relationship to the IPv4 addresses. This relationship is

self-described by mapping the IPv4 address in the IPv6 address. The stateless translator uses the corresponding IPv4 addresses to represent the IPv6 hosts. The stateful translator does not use IPv6-converted addresses, because the IPv6 hosts are represented by the IPv4 address pool in the translator via dynamic states.

NAT—Network Address Translation.

RP—Route Processor.

stateful translation—In stateful translation a per-flow state is created when the first packet in a flow is received. A translation algorithm is said to be stateful if the transmission or reception of a packet creates or modifies a data structure in the relevant network element. Stateful translation allows the use of multiple translators interchangeably and also some level of scalability. Stateful translation is defined to enable the IPv6 clients and peers without mapped IPv4 addresses to connect to the IPv4-only servers and peers.

stateless translation—A translation algorithm that is not stateful is called stateless. A stateless translation requires configuring a static translation table, or may derive information algorithmically from the messages it is translating. Stateless translation requires less computational overhead than stateful translation. It also requires less memory to maintain the state, because the translation tables and the associated methods and processes exist in a stateful algorithm and do not exist in a stateless one. Stateless translation enables the IPv4-only clients and peers to initiate connections to the IPv6-only servers or peers that are equipped with IPv4-embedded IPv6 addresses. It also enables scalable coordination of IPv4-only stub networks or ISP IPv6-only networks. Because the source port in an IPv6-to-IPv4 translation may have to be changed to provide adequate flow identification, the source port in the IPv4-to-IPv6 direction need not be changed.