

NPTv6 Support

The NPTv6 feature supports translating IPv6 packet headers and source address prefixes in both directions, from inside to outside and vice versa. A router that implements an NPTv6 prefix translation function is referred to as an NPTv6 Translator.

To support inter-VRF communication, you can use VRF-Aware Software Infrastructure Scale feature. The VRF-Aware Software Infrastructure (VASI) Scale feature allows you to apply services such as access control lists (ACLs), Network Address Translation (NAT), policing, and zone-based firewalls to MPLS traffic or IPv4 and IPv6 traffic that is flowing across two different Virtual Routing and Forwarding (VRF) instances. The VASI interfaces support redundancy of the Route Processor (RP) and Forwarding Processor (FP).

- Information About NPTv6 support, on page 1
- Configuring NPTv6 Support on VASI, on page 4
- Additional References for NPTv6 support, on page 9

Information About NPTv6 support

The IPv6-to-IPv6 Network Prefix Translation (NPTv6) serves as a useful mechanism for implementing address independence in an IPv6 environment. A major benefit associated with NPTv6 is the fact that it avoids the requirement for an NPTv6 Translator to rewrite the transport layer headers which reduces the load on network devices. NPTv6 also does not interfere with encryption of the full IP payload.

The NPTv6 support allows for greater reliability as it provides support for load balancing and achieves the translation without breaking the end-to-end reachability at the network layer.

Interconnect Different Networks

The NPTv6 support allows you to redirect or forward packets from one network to another in an IPV6 environment. The NPTv6 support on is an algorithmic translation function which provides a 1:1 relationship between the addresses within the inside and outside network. When NPTv6 is used, you can interconnect different networks and support multihoming, load balancing, peer-to-peer networking.

Stateless Support

The NPTv6 does not create any state in the date plane and hence, can operate using minimal memory and supports High Availability (HA) by default.

Improved Support and Scaling

The NPTv6 supports prefix longer than 64 bits and supports static IPv6 host to host translations. You can configure IPv4 and IPv6 translations on the same interface using NPTv6 support and scaling is supported. The NPTv6 feature also supports packet tracing and conditional debugging.

Access to Services Hosted on a Global Network

Implementing VASI by using virtual interface pairs, where each of the interfaces in the pair is associated with a different VRF lets you access different services on the internet. The VASI virtual interface is the next hop interface for any packet that needs to be switched between these two VRFs. VASI interfaces provide the framework necessary to configure a firewall or a NAT between VRF instances.

Pairing of Interfaces

Each interface pair is associated with two different VRF instances. The two virtual interfaces, called vasileft and vasiright, in a pair are logically wired back-to-back and are completely symmetrical. Each interface has an index. The association of the pairing is done automatically based on the two interface indexes such that vasileft automatically gets paired to vasiright.

Static or Dynamic Routing

You can configure either static routing or dynamic routing with Border Gateway Protocol (BGP), Enhanced Interior Gateway Routing Protocol (EIGRP), or Open Shortest Path First (OSPF). BGP dynamic routing protocol restrictions and configuration are valid for BGP routing configurations between VASI interfaces.

Benefits of Using NPTv6 support

- When NPTv6 is used, you can interconnect different networks and support multihoming, load balancing,peer-to-peer networking. The NPTv6 does not create any state in the date plane and hence can operate using minimal memory and supports High Availability (HA) by default.
- You can configure IPv4 and IPv6 translations on the same interface using NPTv6 support and scaling is supported. The NPTv6 feature also supports Packet tracing and conditional debugging.

Restrictions for NPTv6 support

- Multicast is not supported.
- Firewall is not supported.
- High Speed Logging (HSL) and syslog is not supported..

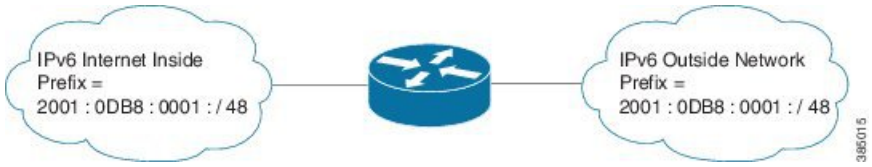
Deployment Scenarios for NPTv6 Support

Single Inside and Outside Network

You can use an NPTv6 Translator to interconnect two network links, one which is an internal network linked to a leaf network which is within a single administrative domain and the other which is external network with connectivity to a global network like the Internet. All hosts on the internal network use addresses from a single prefix which is routed locally. The addresses will be translated to and from the addresses in a globally routable prefix when the IP datagrams transit the NPTv6 Translator. The lengths of these two prefixes will be functionally the same and if the prefix lengths are different, the longer of the two prefixes limits the ability to use subnets in the shorter prefix.

The figure below illustrates NPTv6 deployment having a single inside and outside network.

Figure 1: NPTv6 using Single Inside and Outside Network

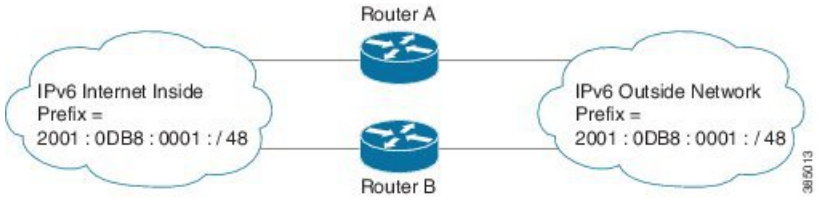


Redundancy and Load Sharing

When more than one NPTv6 Translator is attached to a network, the NPTv6 Translators are configured with the same internal and external prefixes. Since the translation is algorithmic, even though there are multiple translators, they map only one external address to the internal address.

The figure below illustrates NPTv6 deployment in redundancy and load-sharing network.

Figure 2: NPTv6 in Redundancy and Loadsharing Network

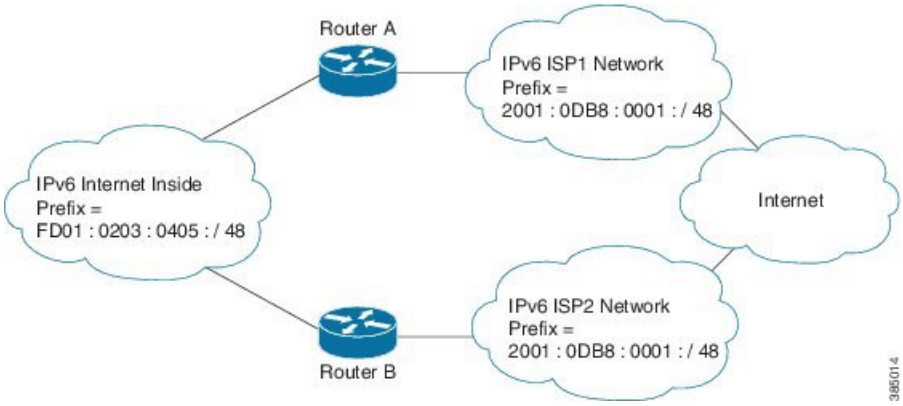


Multihoming

In a multihomed network the NPTv6 Translators are attached to an internal network, but are connected to different external networks. The NPTv6 Translators are configured with the same internal prefix but different external prefixes. Since there are multiple translations, the NPTv6 Translator maps multiple external addresses to the common internal address.

The figure below illustrates NPTv6 deployment in multihoming network.

Figure 3: NPTv6 in Multihoming Network



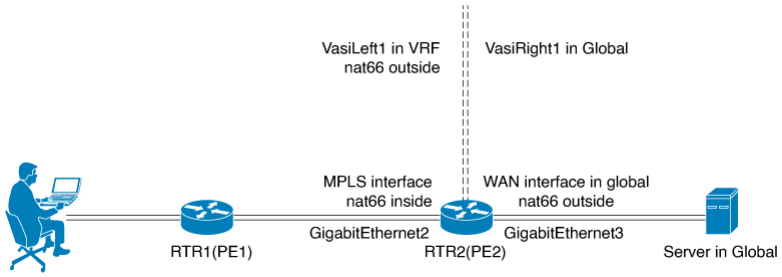
NPTv6 Support on VASI

VPN customers on 6vPE deployment could access services in global network like internet using NPTv6 translator on VASI interfaces (or by configuring NPTv6 on VASI interfaces). VASI allows applying NPTv6 translator to the traffic between VRFs/VPNs.

To support inter-VRF communication, you can use VRF-Aware Software Infrastructure Scale feature. The VRF-Aware Software Infrastructure (VASI) Scale feature allows you to apply services such as access control lists (ACLs), Network Address Translation (NAT), policing, and zone-based firewalls to MPLS traffic or IPv4 and IPv6 traffic that is flowing across two different Virtual Routing and Forwarding (VRF) instances. The VASI interfaces support redundancy of the Route Processor (RP) and Forwarding Processor (FP).

The figure below illustrates VPN customer in 6vPE deployment accessing services in global network using NPTv6 and VASI on PE2:

Figure 4: NPTv6 Support on VASI



Configuring NPTv6 Support on VASI

Configuring NPTv6 Support on VASI involves the following steps:

- Configure 6VPE for PE1
- Configure 6VPE for PE2
- Configure Virtual Interfaces on PE2
- Configure NPTv6 on PE2

Configure 6VPE for PE 1

To configure 6VPE for PE 1:

```
vrf definition client_vpn
rd 1:1
route-target export 1:1
route-target import 1:1
!
address-family ipv4
exit-address-family
!
address-family ipv6
exit-address-family
!
!
interface GigabitEthernet2
```

```
vrf forwarding client_vpn
ipv6 address 1001:1:2::2/64
!
interface GigabitEthernet3
ip address 10.2.0.2 255.255.255.0
!
interface Loopback 100
ip address 100.0.0.2 255.255.255.255
!
router ospf 1
network 100.0.0.2 0.0.0.0 area 0.1.0.0
network 10.2.0.2 0.0.0.0 area 0.1.0.0
!
interface GigabitEthernet3
ip ospf network point-to-point
!
mpls ldp router-id Loopback100 force
interface GigabitEthernet3
mpls ip
mpls label protocol ldp
!
router bgp 65002
bgp router-id 2.2.2.2
bgp log-neighbor-changes
neighbor 100.0.0.3 remote-as 65003
neighbor 100.0.0.3 ebgp-multihop 255
neighbor 100.0.0.3 update-source Loopback100
address-family ipv4
no neighbor 100.0.0.3 activate
exit-address-family
!
address-family vpnv6
neighbor 100.0.0.3 activate
neighbor 100.0.0.3 send-community both
exit-address-family
!
address-family ipv6 vrf client_vpn
redistribute connected
exit-address-family
!
```

Configure 6VPE for PE2

To Configure 6VPE for PE2:

```
ipv6 unicast-routing
vrf definition client_vpn
rd 1:1
route-target export 1:1
route-target import 1:1
!
address-family ipv4
exit-address-family
!
address-family ipv6
exit-address-family
!
!
interface GigabitEthernet2
ip address 10.2.0.3 255.255.255.0
!
interface GigabitEthernet3
ipv6 address 1001:3:4::3/64
!
```

```
interface Loopback 100
ip address 100.0.0.3 255.255.255.255
!
router ospf 1
network 100.0.0.3 0.0.0.0 area 0.1.0.0
network 10.2.0.3 0.0.0.0 area 0.1.0.0
!
interface GigabitEthernet2
ip ospf network point-to-point
!
mpls ldp router-id Loopback100 force
interface GigabitEthernet2
mpls ip
mpls label protocol ldp
!
router bgp 65003
bgp router-id 3.3.3.3
bgp log-neighbor-changes
neighbor 100.0.0.2 remote-as 65002
neighbor 100.0.0.2 ebgp-multihop 255
neighbor 100.0.0.2 update-source Loopback100
address-family ipv4
no neighbor 100.0.0.2 activate
exit-address-family
!
address-family vpnv6
neighbor 100.0.0.2 activate
neighbor 100.0.0.2 send-community both
exit-address-family
!
address-family ipv6 vrf client_vpn
exit-address-family
!
```

Configure Virtual Interfaces on PE2

To configure Virtual Interfaces on PE2:

```
interface vasileft1
vrf forwarding client_vpn
ipv6 address 1003:3:3::1/120
ipv6 address FE80:1:1:1::1 link-local
interface vasiright1
ipv6 address 1003:3:3::2/120
ipv6 address FE80:1:1:1::2 link-local
!
ipv6 prefix-list DENY_BGP_ROUTES_v6 deny 1001:1:2::/64
router bgp 65003
neighbor 1003:3:3::1 remote-as 60001
neighbor 1003:3:3::1 local-as 60002 no-prepend replace-as
neighbor 1003:3:3::1 description PEERING to the VASI left Interface
!
address-family ipv6
network 1001:3:4::/64
neighbor 1003:3:3::1 activate
neighbor 1003:3:3::1 send-community
neighbor 1003:3:3::1 next-hop-self
exit-address-family
!
address-family ipv6 vrf client_vpn
bgp router-id 5.5.5.5
neighbor 1003:3:3::2 remote-as 60002
neighbor 1003:3:3::2 local-as 60001 no-prepend replace-as
neighbor 1003:3:3::2 description Peer to VASI in Global
```

```
neighbor 1003:3:3::2 activate
neighbor 1003:3:3::2 send-community
neighbor 1003:3:3::2 prefix-list DENY_BGP_ROUTES_v6 out
exit-address-family
```

Configure NPTv6 on PE2

To configure NPTv6 on PE2:

```
interface GigabitEthernet2
nat66 inside
!
interface vasileft1
nat66 outside
!
interface GigabitEthernet3
nat66 outside
!
nat66 prefix inside 1001:1:2::/120 outside 2001:2001:2001::/120 vrf client_vpn
```

Verifying NPTv6 Configuration

To verify the various functions under the overall NPTv6 feature, refer to the below list of commands:

Command	Description

show nat66 prefix
Example:
`Device# show nat66 prefix`
`Prefixes configured: 1`
`NAT66 Prefixes`
`Id: 1 Inside 2002:AB01::/64 Outside 2002:AB02::/64`
	Verify stateless NAT66 prefix configuration.

show nat66 statistics
Example:
`Device# show nat66 statistics`
`NAT66 Statistics`
`Global Stats:`
`Packets translated (In -> Out)`
`: 7`
`Packets translated (Out -> In)`
`: 7`
	Verify NAT66 translation statistics.

show platform hardware qfp active feature nat66 datapath basecfg
 	
Example:
 	
`Device# show platform hardware qfp active feature`
`nat66 datapath basecfg`
`nat66 cfg_flags 0x00000001, dbg_flags 0x00000000`
`nat66_prefix_hash_table_entries 2048, nat66_prefix_hash_table 0x89628400`
`prefix hasht 0x89628400 max 2048 chunk 0x8c392bb0`
`hash_salt 719885386`
	Verify global stateless NPTv6 prefix in the data plane and other base configuration information.

show platform hardware qfp active feature nat66 datapath prefix
 	
Example:
 	
`Device# show platform hardware qfp active feature`
`nat66 datapath prefix`
`prefix hasht 0x89628400 max 2048 chunk 0x8c392bb0`
`hash_salt 719885386`
`NAT66 hash[1] id(1) len(64) vrf(0) in: 2002:ab01:0000:0000:0000:0000:0000:0000 out: 2002:ab02:0000:0000:0000:0000:0000:0000`
`in2out: 7`
`out2in: 7`
	Verify the stateless NPTv6 prefix configuration on passed interfaces.

show platform hardware qfp active feature nat66 datapath statistics
 	
Example:
 	
`Device# show platform hardware qfp act feat nat66`
`data statistics`
 	
`in2out xlated pkts 7 out2in xlated pkts 7`
`NAT66_DROP_SC_INVALID_PKT 0`
`NAT66_DROP_SC_BAD_DGLEN 0`
`NAT66_DROP_SC_PLU_FAIL 22786`
`NAT66_DROP_SC_PROCESS_V6_ERR 0`
`NAT66_DROP_SC_INVALID_EMBEDDED 0`
`NAT66_DROP_SC_SRC_RT 0`
`NAT66_DROP_SC_NOT_ENABLED 0`
`NAT66_DROP_SC_NO_GPM 0`
`NAT66_DROP_SC_LOOP 0`
`in2out_pkts 22768 out2in_pkts 22793`
`in2out_pkts_untrans 22761 out2in_pkts_untrans 22786`
`in2out_lookup_pass 7 out2in_lookup_pass 7`
`in2out_lookup_fail 0 out2in_lookup_fail 22786`
`mem_alloc_fail 0 prefix_fail 0`
`total prefix count 1`
	Verify global NPTv6 statistics.

Troubleshooting Tips

You must make sure that the inside and outside interfaces are configured.

Use the following debug commands if you have any configuration issues:

debug platform hardware qfp active feature nat66 datapath detailed	Provides detailed debugging information about the data plane layer.
debug platform hardware qfp active feature nat66 datapath all	Displays debugging information about the data plane layer.
debug platform condtion feature nat66 datapath submode detailed	Provides data plane layer debugging information using buginf_cond. ACL filter can be supplied via the debug condition infrastructure.

Additional References for NPTv6 support

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Command List, All Releases
IP Addressing Services commands	Cisco IOS IP Addressing Services Command Reference
VASI (VRF-Aware Software Infrastructure)	Configuring the VASI (VRF-Aware Software Infrastructure) Scale

Standards and RFCs

Standard/RFC	Title
RFC 6296	*IPv6-to-IPv6 Network Prefix Translation*

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