Implementing IPv6 VPN Provider Edge Transport over MPLS

IPv6 Provider Edge or IPv6 VPN Provider Edge (6PE/6VPE) uses the existing MPLS IPv4 core infrastructure for IPv6 transport. 6PE/6VPE enables IPv6 sites to communicate with each other over an MPLS IPv4 core network using MPLS label switched paths (LSPs).

This feature relies heavily on multiprotocol Border Gateway Protocol (BGP) extensions in the IPv4 network configuration on the provider edge (PE) router to exchange IPv6 reachability information (in addition to an MPLS label) for each IPv6 address prefix. Edge routers are configured as dual-stack, running both IPv4 and IPv6, and use the IPv4 mapped IPv6 address for IPv6 prefix reachability exchange.

Familiarity with MPLS and BGP4 configuration and troubleshooting is required for implementing 6PE/6VPE.

Overview of 6PE/6VPE

Table 1: Feature History Table

Feature Name	Release	Description
6PE/6VPE	Release 25.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature is supported on Cisco 8711-48Z-M routers.

Feature Name	Release	Description
6PE/6VPE	Release 25.1.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) This feature support is now supported on Cisco 8011-4G24Y4H-I routers.
6PE/6VPE	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: K100])(select variants only*); Modular Systems (8800 [LC ASIC: Q100, P100])(select variants only*) The router now enable seamless integration of IPv6 networks over an MPLS backbone, allowing service providers to offer IPv6 services without changing the core MPLS infrastructure. Facilitates IPv6 routing by leveraging existing IPv4 MPLS paths, ensuring efficient and scalable network expansion. IPv6 prefixes are advertised as VPNv4 routes, which simplifies deployment and reduces operational complexity. This approach requires no IPv6 capabilities in the MPLS core, leading to cost savings and a smoother transition to IPv6. *Previously this feature was supported on Q200 and Q100. It is now extended to: • 8712-MOD-M • 8212-48FH-M • 8711-32FH-M • 88-LC1-52Y8H-EM • 88-LC1-12TH24FH-E • 88-LC1-36EH

Multiple techniques are available to integrate IPv6 services over service provider core backbones:

- Dedicated IPv6 network running over various data link layers

- Dual-stack IPv4-IPv6 backbone

• Existing MPLS backbone leverage

These solutions are deployed on service providers' backbones when the amount of IPv6 traffic and the revenue generated are in line with the necessary investments and the agreed-upon risks. Conditions are favorable for the introduction of native IPv6 services, from the edge, in a scalable way, without any IPv6 addressing restrictions and without putting a well-controlled IPv4 backbone in jeopardy. Backbone stability is essential for service providers that have recently stabilized their IPv4 infrastructure.

Service providers running an MPLS/IPv4 infrastructure follow similar trends because several integration scenarios that offer IPv6 services on an MPLS network are possible. Cisco Systems has specially developed Cisco 6PE or IPv6 Provider Edge Router over MPLS, to meet all those requirements.

Inter-AS support for 6PE requires support of Border Gateway Protocol (BGP) to enable the address families and to allocate and distribute PE and ASBR labels.

Note   Cisco IOS XR displays actual IPv4 next-hop addresses for IPv6 labeled-unicast and VPNv6 prefixes. IPv4-mapped-to-IPv6 format is not supported.

Benefits of 6PE/6VPE

Service providers who currently deploy MPLS experience these benefits of Cisco 6PE/6VPE:

• Minimal operational cost and risk—No impact on existing IPv4 and MPLS services.

• Provider edge routers upgrade only—A 6PE/6VPE router can be an existing PE router or a new one dedicated to IPv6 traffic.

• No impact on IPv6 customer edge routers—The ISP can connect to any customer CE running Static, IGP or EGP.

• Production services ready—An ISP can delegate IPv6 prefixes.

• IPv6 introduction into an existing MPLS service—6PE/6VPE routers can be added at any time

Deploying IPv6 over MPLS Backbones

Backbones enabled by 6PE (IPv6 over MPLS) allow IPv6 domains to communicate with each other over an MPLS IPv4 core network. This implementation requires no backbone infrastructure upgrades and no reconfiguration of core routers, because forwarding is based on labels rather than on the IP header itself. This provides a very cost-effective strategy for IPv6 deployment.

IPv6 on the Provider Edge and Customer Edge Routers

Service Provider Edge Routers

6PE is particularly applicable to service providers who currently run an MPLS network. One of its advantages is that there is no need to upgrade the hardware, software, or configuration of the core network, and it eliminates

the impact on the operations and the revenues generated by the existing IPv4 traffic. MPLS is used by many service providers to deliver services to customers. MPLS as a multiservice infrastructure technology is able to provide layer 3 VPN, QoS, traffic engineering, fast re-routing and integration of ATM and IP switching.

Customer Edge Routers

Using tunnels on the CE routers is the simplest way to deploy IPv6 over MPLS networks. It has no impact on the operation or infrastructure of MPLS and requires no changes to the P routers in the core or to the PE routers. However, tunnel meshing is required as the number of CEs to connect increases, and it is difficult to delegate a global IPv6 prefix for an ISP.

The following figure illustrates the network architecture using tunnels on the CE routers.

Figure 1: IPv6 Using Tunnels on the CE Routers
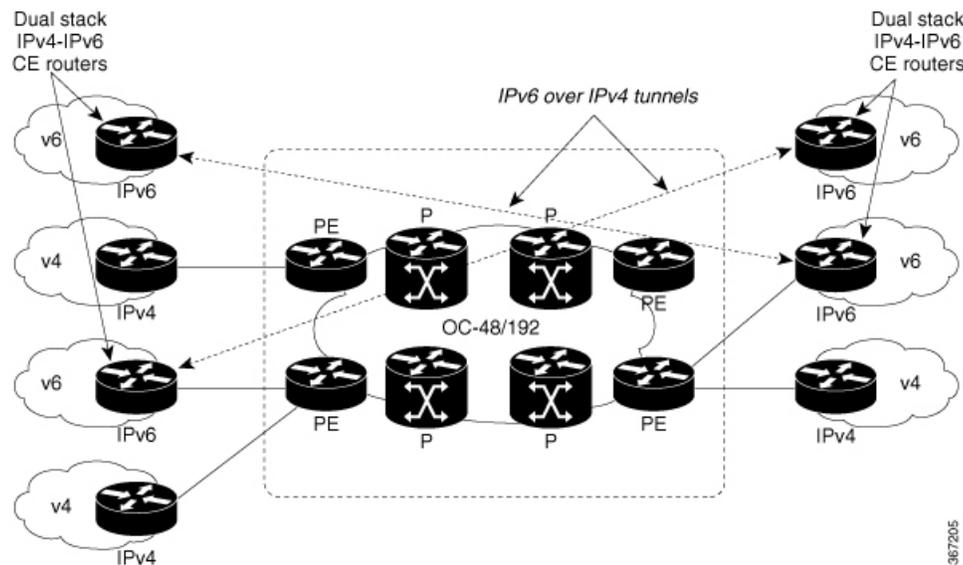
IPv6 Provider Edge Multipath

Internal and external BGP multipath for IPv6 allows the IPv6 router to balance load between several paths (for example, the same neighboring autonomous system (AS) or sub-AS, or the same metrics) to reach its destination. The 6PE multipath feature uses multiprotocol internal BGP (MP-IBGP) to distribute IPv6 routes over the MPLS IPv4 core network and to attach an MPLS label to each route.

When MP-IBGP multipath is enabled on the 6PE router, all labeled paths are installed in the forwarding table with available MPLS information (label stack). This functionality enables 6PE to perform load balancing.

OSPFv3 (CE to PE)

The Open Shortest Path First version 3 (OSPFv3) IPv6 VPN Provider Edge (6VPE) feature adds VPN routing and forwarding (VRF) and provider edge-to-customer edge(PE-CE) routing support to Cisco IOS XR OSPFv3 implementation. This feature allows:

- Multiple VRF support per OSPFv3 routing process

- OSPFV3 PE-CE extensions

Multiple VRF Support

OSPFv3 supports multiple VRFs in a single routing process that allows scaling to tens and hundreds of VRFs without consuming too much route processor (RP) resources. Multiple OSPFv3 processes can be configured on a single router. In large-scale VRF deployments, this allows partition VRF processing across multiple RPs. It is also used to isolate default routing table or high impact VRFs from the regular VRFs. It is recommended to use a single process for all the VRFs. If needed, a second OSPFv3 process must be configured for IPv6 routing.

Note   A maximum of four OSPFv3 processes are supported.

OSPFv3 PE-CE Extensions

IPv6 protocol is being vastly deployed in today's customer networks. Service Providers (SPs) need to be able to offer Virtual Private Network (VPN) services to their customers for supporting IPv6 protocol, in addition to the already offered VPN services for IPv4 protocol.

In order to support IPv6, routing protocols require additional extensions for operating in the VPN environment. Extensions to OSPFv3 are required in order for OSPFv3 to operate at the PE-CE links.

VRF Lite

VRF lite feature enables VRF deployment without BGP or MPLS based backbone. In VRF lite, the PE routers are directly connected using VRF interfaces. For OSPFv3, the following needs to operate differently in the VRF lite scenario, as opposed to the deployment with BGP or MPLS backbone:

  • DN bit processing—In VRF lite environment, the DN bit processing is disabled.

  • ABR status—In VRF context (except default VRF), OSPFv3 router is automatically set as an ABR, regardless to it's connectivity to area 0. This automatic ABR status setting is disabled in the VRF lite environment.

Note   To enable VRF Lite, issue the **capability vrf-lite** command in the OSPFv3 VRF configuration submode.

Configuring 6PE/6VPE

Configuration Example

This example shows how to configure 6PE on PE routers to transport the IPv6 prefixes across the IPv4 cloud. Ensure that you configure 6PE on PE routers participating in both the IPv4 cloud and IPv6 clouds. Pointers:

  • For 6PE, you can use all routing protocols supported on Cisco IOS XR software such as BGP, OSPF, IS-IS, and Static to learn routes from both clouds. However, for 6VPE, you can use only the BGP, and Static routing protocols to learn routes. Also, 6VPE supports OSPFv3 routing protocol between PE and CE routers.

  • While configuring 6PE/6VPE on the router, it is mandatory to configure label allocation mode, per-vrf for all routers including peer routers.

- Route policies must be configured prior to configuring 6PE/6VPE.

- BGP uses the **per-vrf** label mode for transporting local and redistributed IP prefixes. Before IOS XR Release 7.5.3, BGP assigned a random label for the prefixes. Starting from Release 7.5.3, BGP assigns a label value of **2**, the IPv6 Explicit NULL Label, for the same prefixes.

```
Router#configure
Router(config)#router bgp 10
Router(config-bgp)#bgp router-id 11.11.11.11
Router(config-bgp)#graceful-restart
Router(config-bgp)#log neighbor changes detail
Router(config-bgp)#address-family ipv6 unicast
Router(config-bgp-af)#redistribute connected
Router(config-bgp-af)#redistribute ospfv3 7
Router(config-bgp-af)#allocate-label all
Router(config-bgp-af)#commit
Router(config-bgp)#neighbor 66:1:2::2
Router(config-bgp-nbr)#remote-as 102
Router(config-bgp-nbr)#address-family ipv6 unicast
Router(config-bgp-nbr-af)#route-policy pass-all in
Router(config-bgp-nbr-af)#route-policy pass-all out
Router(config-bgp-nbr-af)#commit
Router(config-bgp)#neighbor 13.13.13.13
Router(config-bgp-nbr)#remote-as 10
Router(config-bgp-nbr)#update-source Loopback0
Router(config-bgp-nbr)#address-family vpnv4 unicast
Router(config-bgp-nbr-af)#address-family ipv6 labeled-unicast
Router(config-bgp-nbr-af)#address-family vpnv6 unicast
Router(config-bgp-nbr-af)#commit
Router(config-bgp-nbr-af)#exit
Router(config-bgp-nbr)#exit
Router(config-bgp)#vrf red
Router(config-bgp-vrf)#rd 500:1
Router(config-bgp-vrf)#address-family ipv4 unicast
Router(config-bgp-vrf-af)#label mode per-vrf
Router(config-bgp-vrf-af)#redistribute connected
Router(config-bgp-vrf-af)#redistribute static
Router(config-bgp-vrf-af)#exit
Router(config-bgp-vrf)#address-family ipv6 unicast
Router(config-bgp-vrf-af)#label mode per-vrf
Router(config-bgp-vrf-af)#redistribute connected
Router(config-bgp-vrf-af)#redistribute static
Router(config-bgp-vrf-af)#commit
Router(config-bgp-vrf-af)#!
!
Router(config)#interface HundredGigE0/0/1/0
Router(config-if)#vrf red
Router(config-if)#ipv6 address 4002:110::1/128
Router(config-if)#exit
Router(config)#vrf red
Router(config-vrf)#address-family ipv4 unicast
Router(config-vrf-af)#label mode per-vrf
Router(config-vrf-af)#import route-target
Router(config-vrf-import-rt)#500:1
Router(config-vrf-import-rt)#!
Router(config-vrf-import-rt)#export route-target
Router(config-vrf-export-rt)#500:1
Router(config-vrf-export-rt)#!
Router(config-vrf-export-rt)#!
Router(config-vrf-export-rt)#address-family ipv6 unicast
Router(config-vrf-af)#label mode per-vrf
```

```
Router(config-vrf-af)#import route-target
Router(config-vrf-import-rt)#500:1
Router(config-vrf-import-rt)#!
Router(config-vrf-import-rt)#export route-target
Router(config-vrf-export-rt)#500:1
Router(config-vrf-export-rt)#commit
```

Running Configuration

```
router bgp 10
bgp router-id 11.11.11.11
bgp graceful-restart
bgp log neighbor changes detail
!
address-family ipv6 unicast
 redistribute connected
  redistribute ospfv3 7
  allocate-label all
!
!
neighbor 66:1:2::2
  remote-as 201
  address-family ipv6 unicast
   route-policy pass-all in
   route-policy pass-all out
  !
!
neighbor 13.13.13.13
  remote-as 10
  update-source Loopback0
  address-family vpnv4 unicast
  !
  address-family ipv6 labeled-unicast
  !
  address-family vpnv6 unicast
!
vrf red
  rd 500:1
  address-family ipv4 unicast
   label mode per-vrf
   redistribute connected
   redistribute static
  !
  address-family ipv6 unicast
   label mode per-vrf
   redistribute connected
   redistribute static
  !
 !
!
interface HundredGigE0/0/1/0
vrf red
Ipv6 address 4002:110::1/128
!
exit
vrf red
address-family ipv4 unicast
import route-target
500:1
!
export route-target
```

```
500:1
!
!
address-family ipv6 unicast
import route-target
500:1
!
export route-target
500:1
!
```

Verification

```
Router# show route ipv6
Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
       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, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local, G  - DAGR, l - LISP
       A - access/subscriber, a - Application route
       M - mobile route, r - RPL, (!) - FRR Backup path
Gateway of last resort is not set

L    ::ffff:127.0.0.0/104
      [0/0] via ::, 02:10:49
C    66:1:2::/64 is directly connected,
      02:09:39, TenGigE0/0/0/10.2
L    66:1:2::1/128 is directly connected,
      02:09:39, TenGigE0/0/0/10.2
C    66:1:3::/64isdirectlyconnected,
[20/0] via fe80::200:2cff:fe64:99e2, 02:07:38, TenGigE0/0/0/10.2
B    2000:0:0:1c::/64
      [20/0] via fe80::200:2cff:fe64:99e2, 02:07:38, TenGigE0/0/0/10.2
B    2000:0:0:1d::/64
```

Local PE :
```
Router# show bgp ipv6 labeled-unicast 2000:0:0:1c::/64
BGP routing table entry for 2000:0:0:1c::/64
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker             5033        5033
    Local Label: 66313
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.1
  Advertised to peers (in unique update groups):
    13.13.13.13
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0.1
  Advertised to peers (in unique update groups):
    13.13.13.13
  201
    66:1:2::2 from 66:1:2::2 (39.229.0.1)
      Origin IGP, localpref 100, valid, external, best, group-best
      Received Path ID 0, Local Path ID 0, version 5033
      Origin-AS validity: not-found
```

Remote PE
```
Router# show bgp ipv6 labeled-unicast 2000:0:0:1c::/64
BGP routing table entry for 2000:0:0:1c::/64
Versions:
  Process           bRIB/RIB  SendTblVer
```

```
  Speaker              139679      139679
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.2
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0.2
  201
    11.11.11.11 (metric 5) from 13.13.13.13 (11.11.11.11)
      Received Label 66313
      Origin IGP, localpref 100, valid, internal, best, group-best, labeled-unicast
      Received Path ID 0, Local Path ID 0, version 139679
      Originator: 11.11.11.11, Cluster list: 5.5.5.5
```

Configuring OSPFv3 as the Routing Protocol Between the PE and CE Routers

Configuration Example

This example shows how to configure provider edge (PE)-to-customer edge (CE) routing sessions that use Open Shortest Path First version 3 (OSPFv3).

```
Router#config
Router(config)#router ospfv3 7
Router(config-ospfv3)#router-id 10.200.1.7
Router(config-ospfv3)#vrf vrf1
Router(config-ospfv3-vrf)#area 7
Router(config-ospfv3-vrf-ar)#interface Loopback7
Router(config-ospfv3-vrf-ar-if)#!
Router(config-ospfv3-vrf-ar-if)#interface TenGigE0/7/0/0/3.7
Router(config-ospfv3-vrf-ar-if)#
```

Running Configuration

```
router ospfv3 7
router-id 10.200.1.7
vrf vrf1
  area 7
   interface Loopback7
   !
   interface TenGigE0/7/0/0/3.7
   !
  !
!
```

Verification

```
Router#show ospfv3 7 vrf vrf1 neighbor
# Indicates Neighbor awaiting BFD session up

Neighbors for OSPFv3 7, VRF vrf1

Neighbor ID    Pri   State          Dead Time   Interface ID    Interface
10.201.7.1     0     FULL/DROTHER   00:00:36    0               TenGigE0/7/0/0/3.7
    Neighbor is up for 1w0d

Total neighbor count: 1
```

Configure BGP as the Routing Protocol Between the PE and CE Routers

BGP distributes reachability information for VPN-IPv6 prefixes for each VPN. PE to PE or PE to route reflector (RR) sessions are iBGP sessions, and PE to CE sessions are eBGP sessions. PE to CE eBGP sessions can be directly or indirectly connected (eBGP multihop).

Configuration Example

This example lists the steps to configure BGP as the routing protocol between the PE and CE routers. The route policy, *pass-all* in this example, must be configured before it can be attached.

PE1:

```
Router-PE1#configure
Router-PE1(config)#router bgp 2001
Router-PE1(config-bgp)#bgp router-id 13.13.13.1
Router-PE1(config-bgp)#address-family ipv6 unicast
Router-PE1(config-bgp-af)#exit
Router-PE1(config-bgp)#address-family vpnv6 unicast
Router-PE1(config-bgp-af)#exit
/* VRF configuration */
Router-PE1(config-bgp)#vrf vrf1601
Router-PE1(config-bgp-vrf)#rd 2001:1601
Router-PE1(config-bgp-vrf)#address-family ipv6 unicast
Router-PE1(config-bgp-vrf-af)#label mode per-vrf
Router-PE1(config-bgp-vrf-af)#redistribute connected
Router-PE1(config-bgp-vrf-af)#exit
Router-PE1(config-bgp-vrf)#neighbor 2002:1::3
Router-PE1(config-bgp-vrf-nbr)#remote-as 7501
Router-PE1(config-bgp-vrf-nbr)#address-family ipv6 unicast
Router-PE1(config-bgp-vrf-nbr-af)#route-policy pass-all in
Router-PE1(config-bgp-vrf-nbr-af)#route-policy pass-all out
Router-PE1(config-bgp-vrf-nbr-af)#commit
```

CE1:

```
Router-CE1#configure
Router-CE1(config)#router bgp 2001
Router-CE1(config-bgp)#bgp router-id 8.8.8.1
Router-CE1(config-bgp)#address-family ipv6 unicast
Router-CE1(config-bgp-af)#exit
Router-CE1(config-bgp)#address-family vpnv6 unicast
Router-CE1(config-bgp-af)#exit
Router-CE1(config-bgp)#neighbor  2001:1::1
Router-CE1(config-bgp-nbr)#remote-as 2001
Router-CE1(config-bgp-nbr)#address-family ipv6 unicast
Router-CE1(config-bgp-nbr-af)#route-policy pass-all in
Router-CE1(config-bgp-nbr-af)#route-policy pass-all out
Router-CE1(config-bgp-nbr-af)#commit
```

Running Configuration

PE1:

```
router bgp 2001
 bgp router-id 13.13.13.1
 address-family ipv6 unicast
 !
 address-family vpnv6 unicast
 !
 vrf vrf1601
  rd 2001:1601
  address-family ipv6 unicast
   label mode per-vrf
   redistribute connected
  !
  neighbor 2002:1::3
   remote-as 7501
   address-family ipv6 unicast
    route-policy pass-all in
    route-policy pass-all out
   !
  !
 !
```

CE1:

```
router bgp 7501
 bgp router-id 8.8.8.1
 address-family ipv6 unicast
 !
 address-family vpnv6 unicast
 !
 neighbor 2002:1::1
 remote-as 2001
  address-family ipv6 unicast
   route-policy pass-all in
   route-policy pass-all out
 !
!
```

Verification

- **PE1:**

```
Router-PE1#show bgp neighbor
BGP neighbor is 2002:1::3
 Remote AS 6553700, local AS 2001, external link
 Administratively shut down
 Remote router ID 2002:1::2
  BGP state = Established
  NSR State: None
  Last read 00:00:04, Last read before reset 00:00:00
  Hold time is 60, keepalive interval is 20 seconds
  Configured hold time: 60, keepalive: 30, min acceptable hold time: 3
  Last write 00:00:16, attempted 19, written 19
  Second last write 00:00:36, attempted 19, written 19
  Last write before reset 00:00:00, attempted 0, written 0
  Second last write before reset 00:00:00, attempted 0, written 0
  Last write pulse rcvd  Apr 12 10:31:20.739 last full not set pulse count 27939
  Last write pulse rcvd before reset 00:00:00
  Socket not armed for io, armed for read, armed for write
  Last write thread event before reset 00:00:00, second last 00:00:00
```

```
        Last KA expiry before reset 00:00:00, second last 00:00:00
        Last KA error before reset 00:00:00, KA not sent 00:00:00
        Last KA start before reset 00:00:00, second last 00:00:00
        Precedence: internet
        Non-stop routing is enabled
        Graceful restart is enabled
        Restart time is 120 seconds
        Stale path timeout time is 360 seconds
        Enforcing first AS is enabled
        Multi-protocol capability not received
        Received 0 messages, 0 notifications, 0 in queue
        Sent 0 messages, 0 notifications, 0 in queue
        Minimum time between advertisement runs is 30 secs
        Inbound message logging enabled, 3 messages buffered
        Outbound message logging enabled, 3 messages buffered

       For Address Family: IPv6 Unicast
        BGP neighbor version 0
        Update group: 0.2 Filter-group: 0.0  No Refresh request being processed
        Inbound soft reconfiguration allowed
        AF-dependent capabilities:
          Outbound Route Filter (ORF) type (128) Prefix:
            Send-mode: advertised
            Receive-mode: advertised
          Graceful Restart capability advertised
            Local restart time is 120, RIB purge time is 600 seconds
            Maximum stalepath time is 360 seconds
        Route refresh request: received 0, sent 0
        Policy for incoming advertisements is pass-all
        Policy for outgoing advertisements is pass-all
        0 accepted prefixes, 0 are bestpaths
        Cumulative no. of prefixes denied: 0.
        Prefix advertised 0, suppressed 0, withdrawn 0
        Maximum prefixes allowed 1048576
        Threshold for warning message 75%, restart interval 0 min
        An EoR was not received during read-only mode
        Last ack version 1, Last synced ack version 0
        Outstanding version objects: current 0, max 0
        Additional-paths operation: None
       Advertise VPNv6 routes enabled with defaultReoriginate,disable Local with stitching-RT
      option
        Advertise VPNv6 routes is enabled with default option


        Connections established 1; dropped 0
        Local host: 2002:1::3, Local port: 23456, IF Handle: 0x00000000
        Foreign host: 2002:1::1, Foreign port: 179
        Last reset 03:12:58, due to Admin. shutdown (CEASE notification sent - administrative
      shutdown)
        Time since last notification sent to neighbor: 03:12:58
        Notification data sent:
          None
       External BGP neighbor not directly connected.
```

- **CE1:**

```
 Router-CE1#show bgp neighbor
 BGP neighbor is 2001:1::1
  Remote AS 2001, local AS 6553700, external link
  Remote router ID 2002:1::1
   BGP state = Established
   NSR State: None
   Last read 00:00:04, Last read before reset 00:00:00
```

```
        Hold time is 60, keepalive interval is 20 seconds
        Configured hold time: 60, keepalive: 30, min acceptable hold time: 3
        Last write 00:00:16, attempted 19, written 19
        Second last write 00:00:36, attempted 19, written 19
        Last write before reset 00:00:00, attempted 0, written 0
        Second last write before reset 00:00:00, attempted 0, written 0
        Last write pulse rcvd  Apr 12 10:31:20.739 last full not set pulse count 27939
        Last write pulse rcvd before reset 00:00:00
        Socket not armed for io, armed for read, armed for write
        Last write thread event before reset 00:00:00, second last 00:00:00
        Last KA expiry before reset 00:00:00, second last 00:00:00
        Last KA error before reset 00:00:00, KA not sent 00:00:00
        Last KA start before reset 00:00:00, second last 00:00:00
        Precedence: internet
        Non-stop routing is enabled
        Graceful restart is enabled
        Restart time is 120 seconds
        Stale path timeout time is 360 seconds
        Enforcing first AS is enabled
        Multi-protocol capability not received
        Received 0 messages, 0 notifications, 0 in queue
        Sent 0 messages, 0 notifications, 0 in queue
        Minimum time between advertisement runs is 30 secs
        Inbound message logging enabled, 3 messages buffered
        Outbound message logging enabled, 3 messages buffered

    For Address Family: IPv6 Unicast
     BGP neighbor version 0
     Update group: 0.1 Filter-group: 0.0  No Refresh request being processed
     Inbound soft reconfiguration allowed
     AF-dependent capabilities:
        Outbound Route Filter (ORF) type (128) Prefix:
          Send-mode: advertised
          Receive-mode: advertised
        Graceful Restart capability advertised
          Local restart time is 120, RIB purge time is 600 seconds
          Maximum stalepath time is 360 seconds
     Route refresh request: received 0, sent 0
     Policy for incoming advertisements is pass-all
     Policy for outgoing advertisements is pass-all
     0 accepted prefixes, 0 are bestpaths
     Cumulative no. of prefixes denied: 0.
     Prefix advertised 0, suppressed 0, withdrawn 0
     Maximum prefixes allowed 1048576
     Threshold for warning message 75%, restart interval 0 min
     An EoR was not received during read-only mode
     Last ack version 1, Last synced ack version 0
     Outstanding version objects: current 0, max 0
     Additional-paths operation: None

    Connections established 0; dropped 0
    Local host: 2002:1::1, Local port: 179, IF Handle: 0x00000000
    Foreign host: 2001:1::3, Foreign port: 23456
    Last reset 00:00:00
    External BGP neighbor not directly connected.
```