CONFIDENTIAL INFORMATION
DESTROY BY SHREDDER ONLY

ISOCORE

MPLS Interoperability

IPv6 Migration Techniques- 6PE/VPE

DRAFT

August 2012

Document ID:ITD-10345v1
Status: Draft

TABLE OF CONTENTS

1 Introduction

This document describes the test plan for Isocore proposed IPv6 to IPv4 transitional methodologies and test plans. These methods have been discussed with various industry leaders and have been identified based on the input from Isocore Internetworking Lab members. The testing that would be included under this umbrella will allow vendors to interwork with the implementations, which either only support single IPv6 stack or dual IP stack.

This test documents proposes test scenario for evaluating implementations in a multivendor environment for deploying IPv6 islands over MPLS enabled network. This technology enables enterprise customers to deployed IPv6 over existing IPv4/MPLS services. Following areas are classified to cover the capability verification of devices supporting IPv6 over a packet switched IPv4 network:

1. Connecting IPv6 islands over MPLS/IPv4 core
2. Evaluating MP-Extensions in the IPv4 network to exchange IPv6 reachability information
3. Associating MPLS label for each IPv6 address prefix announced
4. Dual-stack capability of the PE routers
5. This test plan also extends to the scenario when the core devices are pure IPv6 capable
6. Evaluating the capability of the PE devices to support VPNIPv6 Address families

This is a “leading edge” interoperability event. It is understood that some of the systems under test are in beta or even alpha stages in the development process. As such, it is NOT imperative to implement all the functions defined in the test cases in order to participate in the upcoming test activity.

The Isocore IPv6/IPv4 integration program is driven by pragmatic considerations. One of the goals of this program is to validate various IPv6 to IPv4 migration strategies and offer the stability of these transitional paths in the commercial market. The results of these tests will be used to provide input to the IETF to correct any discrepancies in the specifications. Any interoperability problems that may be discovered will be addressed individually between the vendors involved in full confidentiality.

1.1 Goals and objectives

The objective of this project is to assess various migration strategies for IPv4 to IPv6 transition.
This testing objective will evaluate the capability of the dual stack devices to interconnect the IPv6 clouds to IPv4 networks. An effort will be made to deploy extensions to Layer 3 VPNs to forward traffic from IPv6 enabled enterprise customers. Also, it will be considered to evaluate the capability of the multiservice platforms by enabling layer 2 services in conjunction with advanced IP layer 3 VPNs interoperability of different implementations of VPLS using test cases derived from Carrier operational requirements.

This test is driven by pragmatic considerations:

- One of the goals of the Isocore interoperability program is to promote rapid adoption of leading edge network technologies in the commercial market.

- The IETF standards process also benefits because inconsistencies and ambiguities that are identified in the protocol specifications will be fed back into the standards process sothat required modifications can be made to strengthen the final specifications.

1.2 Preliminary Configuration

The final test topology with all the interface and loopback IP addresses assigned will be provided on the day of the testing. Figure 1 and Figure 2 refer to the test setup and is referred to in all the test cases. It is suggested that participating DUTs can be configured for:

1. ISIS must be configured on all PE router interfaces and IP reachability must be verified**.**
2. For CE devices connectivity either static routes will be configured or OSPFv3/ISISv6/RIPng should be used for the IPv6 route distribution in the IPv6 site
3. PE devices should support dual IP stack (IPv4 and IPv6 support)
4. PE should be capable of being configured for IPv4 addresses on the interfaces facing core network, and IPv6 capable interfaces facing sites [RFC 3513]
5. PE device should support one of the tunneling techniques that include MPLS, IPsec, and IPinIP.
6. PE devices with dual stack should have capability to exchange routing information using IGP, IPv6 EBGP, and static routes
7. PE routers to advertise IPv6 reachability information and distribute aggregate IPv6 labels to other neighboring VPN capable PE devices
8. If MPLS is used as the tunneling mechanism, LDP must be configured on all PE router interfaces and targeted LDP sessions must be established or RSVPTE tunnels may be used.
9. BGP must be configured on all PE router interfaces and BGP sessions must be established in full mesh between PE routers
10. RSVPTE must be configured on all PE and P router interfaces,
11. MPLS Traffic Engineering (TE) tunnels must be established in full mesh between PE routers using RSVPTE,
12. CE devices should be capable of setting up IPv6 addresses on the interfaces.

2 Control/Data Plane Verification for IPv6 over MPLS Interworking

In this section we provide the test topologies and detailed test specification for pairwise and integrated network testing of Interconnection of IPv6 across IPv4/MPLS network.

2.1 IPv6/IPv4 Integration with VPN context topological reference

Figure 1 shows the sample network diagram that may represent a typical deployment scenario of the BGP/MPLS VPN extensions for IPv6 and dual stack environment for PE devices supporting IPv4 core and IPv6 sites. MPLS decoupling of the control plane and data plane allows the integration of IPv4, IPv6 and VPN services. The method relies on the BGP extensions in the IPv4 network Provider Edge routers to exchange IPv6 reachability information along with the MPLS label for each IPv6 address prefix announced. All the PE devices shown in the figure 1 may be dual stack (IPv4 and IPv6).The label hierarchy helps in improving the scalability of this model, the top label provide connectivity in the MPLS core (or if any other tunneling mechanism is used), LDP or RSVPTE is used to distribute the labels. The next label or inner label is used at each PE device for IPv6 forwarding and it is distributed using [MPLS BGP] and [RFC 2858].

The CE devices can be customer routers or Ethernet switches. The CE-PE links can either be direct physical links such as 100BaseT, GigE, or logical links such as ATM, Frame Relay, and MPLS based VCs. A CE device can be connected with multiple links to one PE router or it can be multi-homed.

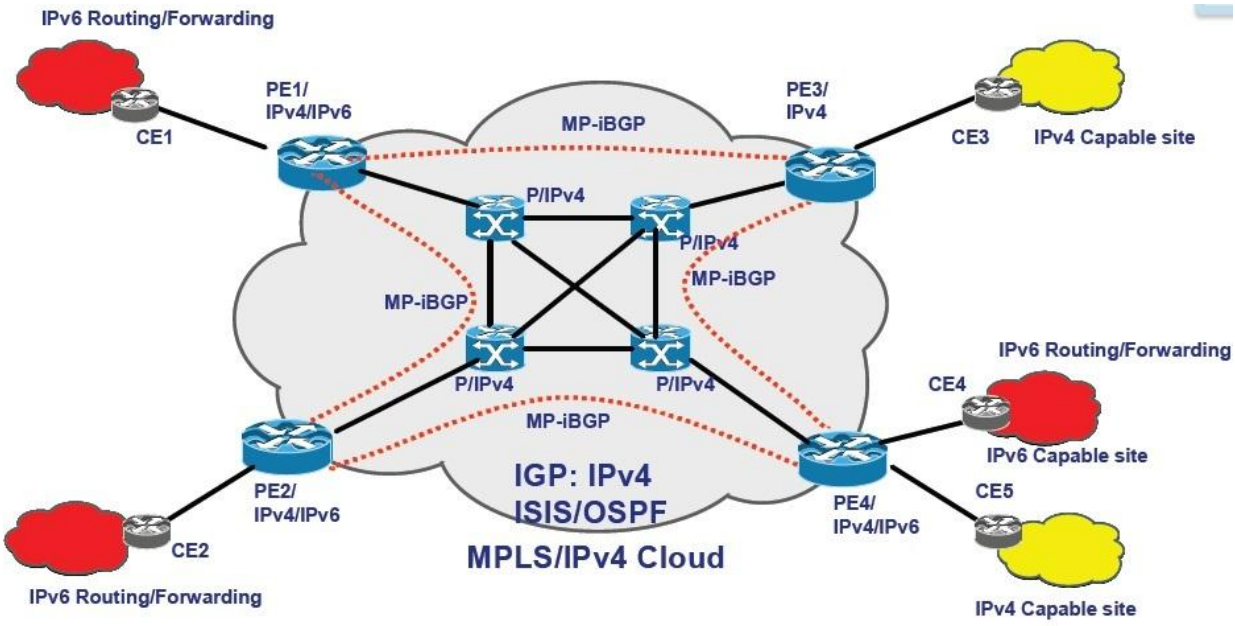

Figure 1: Topological Model for IPv6/IPv4 Integration

2.2 BGP/MPLS IPv6/IPv4 Integration Control Plane Interworking (IPv6 over MPLS Tunnels)

Overview: This section presents the test scenarios that evaluate the capability of the PE devices to interconnect multiple IPv6 islands across existing IPv4 network. The main approach that is being discussed in this test document is MPBGP over IPv4. The main description of the approached is given in [BGP Tunnel]. It is required that the PE devices participating in this capability testing supports dual Stack (DS-BGP) MPBGP. The dual-stack MPBGP routers provide access to IPv6 customers and may provide access to IPv4 customers in addition. Interconnecting the IPv6 islands over an IPv4 cloud requires capability to exchange IPv6 reachability information among DSBGP routers and Tunneling IPv6 packets from Ingress to egress PE nodes

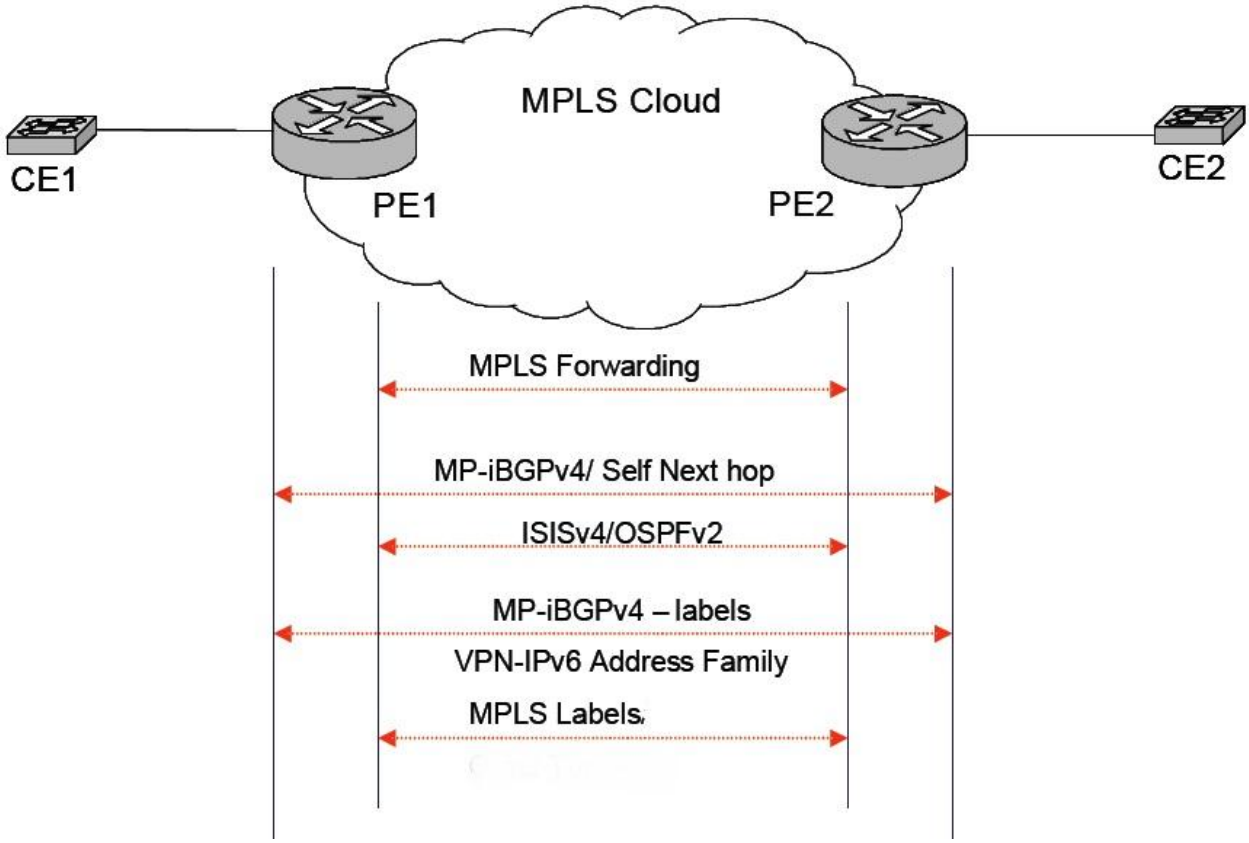

Figure 2: Pairwise testing topology

2.2.1 Test Case V6INT 1: Verify the capability of the PE1 and PE2 (figure 1) to exchange IPv6 reachability information

Test Setup:

Refer the figure for reference components of the topology and section 1.2 for basic setup requirements

Procedure/Expected Behavior:

1. Configure the PE 1, 2, 3 and 4 for IBGP sessions
2. Ensure that all the participating PE routers have setup IBGP sessions successfully and are exchanging the routing information
3. Ensure that PE2, PE3 and PE4 advertises its IPv4 address as the BGP next hop address
4. Verify the capability of the PE routers in figure 1 to support IPv4 embedded IPv6 address:
    a) Format of the IPv4mapped IPv6 address is as follows, and is also mentioned in [RFC 3513]
    IPv4 address: 192.168.16.1
    IPv4mapped IPv6 address: 0:0:0:0:0: FFFF: 192.168.16.1 >: FFFF: 192.168.16.1
5. Ensure that the PE correctly maps the BGP next hop address to be of the same address family as NLRI
6. Verify that the BGP sessions are maintained the BGP routes are correctly installed in the FIB.

2.2.2 Test Case V6INT 2: Verify the capability of dual-stack PE devices to correctly maps the IPv6 data to its peers

Procedure/Expected Behavior:

1. Verify that the PE (1...4) establish MP-BGP sessions
2. Verify that the DUT correctly exchange the IPv6 reachability information
3. Verify that the PE1 (for traffic forwarding from CE1 to CE3) is able to resolve the IPv4 address of the remote PE attached to CE3 to reach a destination IPv6 prefix
4. Ensure that PE1 extracts the IPv4 address of the PE attached to CE3 contained in the IPv4 IPv6 address and gets the label associated to the LSP for this destination
5. IPv4 label is stored with the BGP label for this destination IPv6 subnet forwarding table of PE1
6. PE1 tunnels the data over IPv4 cloud towards the egress PE

2.2.3 Test Case V6INT 3: Verify the capability of the DUT to automatically determine the IPv4 tunnel end point by parsing the MPBGP routing info

Procedure/Expected Behavior:

1. Configure the PE devices to setup the MPLS LSP mesh in the IPv4 network using the PE destination IPv4 address
2. Verify that the MPLS LSPs are established between PE1 and all other PE devices interconnecting the IPv6 islands MPLS LSP could be established by either using RSVPTE or LDP
3. Verify that PE1 distributed label associated with the IPv6 sites connected to PE1 to PE2/PE3/PE4 and vice versa

4. The labels are distributed via MPBGP and [MPLS BGP]
5. Ensure that P routers correctly execute the PHOP behavior before the PE2 receives the single label packet from PE1
6. Verify that the PE1 correctly encode the SAFI value to be 4 when the labeled IPv6 prefixes are exchanged between the PE devices (for VPN context the SAFI value is set to 128)
7. AFI value is still set to 2

3 Control/Data Plane Verification for IPv6 over MPLS

This section defines the test scenarios covering the Service Provider requirements to use its current packet switched network (PSN) to provide Virtual Private Network (VPN) services to its IPv6 customers [BGP/MPLS IPv6]. This is extensions to [IP VPN] method for supporting IPv6. In order to support IPv6 sites in the BGP/VPNs, it is required that the PE devices support VPNIPV6 address family and it should be capable of distributing this information via Multiprotocol BGP. A VPN that is capable of supporting IPv6 sites is said to be IPv6VPN. PEs uses VRFs to isolate the IPv6 sites thus ensuring that address separation is maintained and address duplication could be used.

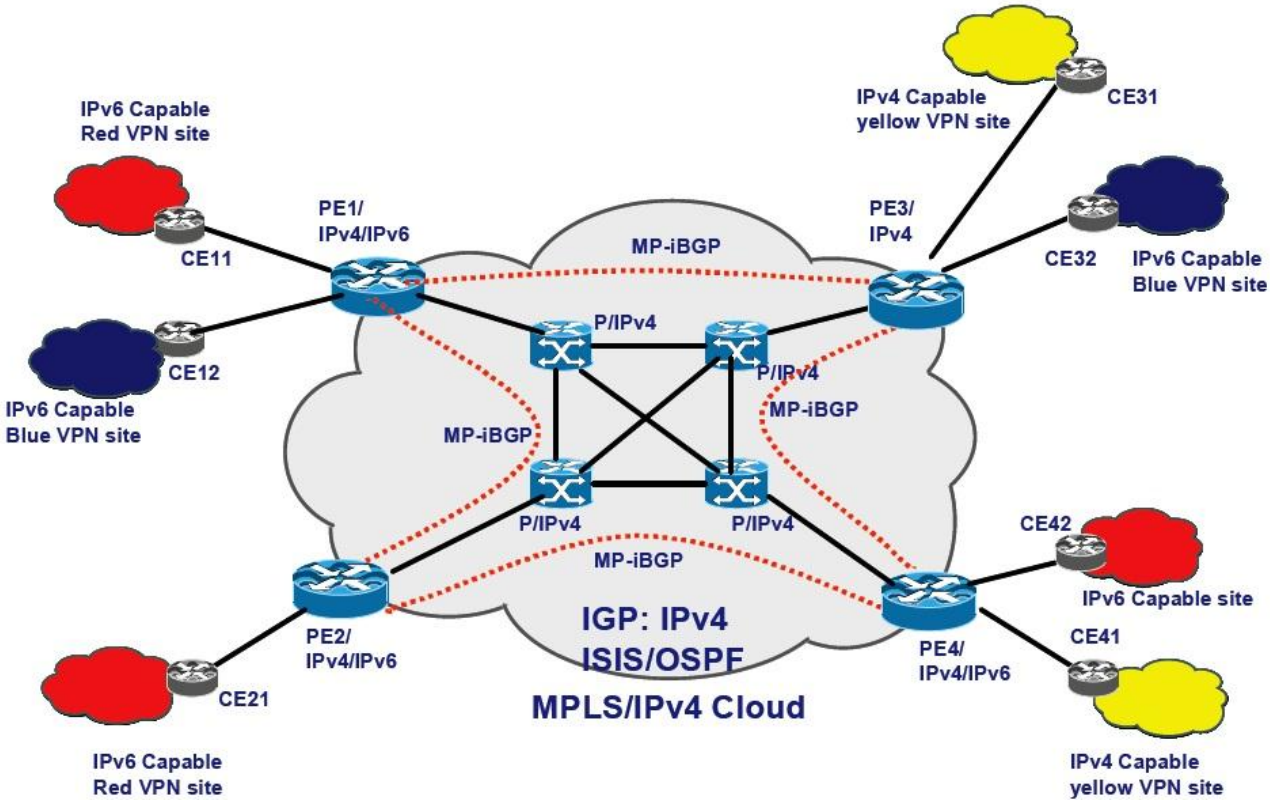

Figure 3: IPv6 VPN reference Test Setup

3.1.1 Test Case V6VPN 1: Verify the capability of the PE devices to support IPV6 address families

Overview (section 2.1, [BGP/MPLS IPv6]):
[MPBGP] allows the usage of address family in the MP_REACH_NLRI attribute. Two fields are specified in this attribute to let the other PE devices know the type of the NLRI expected and next hop field. VPNIPv6 address family is introduced in the [BGP/MPLS IPv6] and is considered similar to VPMIPv4 address family. VPNIPv6 is a 24byte quantity as compared to the 12byte of VPNIPv4. VPNIPv6 address consists of 8byte of RD and 16byte of IPv6 address.

Procedure/Expected Behavior:

1. Configure the PE devices with IBGP sessions
2. Ensure that there is full iBGP mesh in the network (figure 1)
3. Ensure that IPv4 routes are being exchanged between PE1/PE2/PE4
4. Configure the PE1 connected to CE11 for IPv6 address
5. Ensure that routes are being exchanged between CE11 and PE1
6. Configure the VRF associated with the interface connected to CE11
7. Configure the import and export route-targets at PE1, PE2 and PE4 to be same initially, and call it as “red” VPN
8. Verify that the routes learnt from CE11 are being shown in the “red” VRF at PE2
9. If the routes are correctly installed in the red VRF at PE2, ensure that ICMP response works
10. Once the routes have been correctly exchanged, use a test device connected to CE2 to forward IPv6 packets to CE11
11. Verify that the packets are correctly delivered

3.1.2 Test Case V6VPN 2: Verify the capability of the DUT to configure the duplicate IPv6 addresses in two VRFs

Usage of Route Distinguishers in the VPN context allows the Service Provider to use duplicate addresses in the in different sites belonging to different VPNs (refer figure 3). The purpose of the RD is to allow PE device to create distinct routes to a common IPv6 prefix. RDs can also be used to create multiple routes to the same destination; this is achieved by creating two VPNs with different RDs. This allows BGP to install multiple routes to the same address prefix.

Procedure/Expected Behavior::

1. Configure the PE devices with IBGP sessions
2. Ensure that there is full iBGP mesh in the network (figure 1)
3. Ensure that IPv4 routes are being exchanged between PE1/PE2/PE4
4. Configure the PE1 connected to CE11 for IPv6 address

5. Ensure that routes are being exchanged between CE11 and PE1
6. Configure the VRF (with export/import route targets to be 110:0) associated with the interface connected to CE11
7. Configure the End station attached to the CE11 with the following configuration:
   a. IPv6 address 2001::200.1.1.1/64
8. Configure the VPN blue associated with PE3 with import/export targets set to 210:0
   a. IPv6 address to be configured on the end system attached to CE32 should be set to 2001::200.1.1.1/64
9. Ensure that the PE1 correctly forward the packets received from CE12 (blue VPN) to CE32 attached to blue VPN at PE3
10. Ensure that PE3 correctly forwards traffic received from C31 for CE11 (red VPN).

3.1.3 Test Case V6VPN 3: Verify the capability of the PE devices to correctly advertise the reachability information for the IPv6 VPNs and BGP next hop

Route redistribution among PEs by BGP is achieved by establishing mesh sessions of iBGP between all the PE devices participating in the IPv6/IPv4 VPNs. The PE routers are expected to exchange the reachability information via MPBGP, and advertise themselves as the BGP next hop. BGP next hop field in the MP_REACH_NLRI [MPLS BGP] should contain the VPNIPv6 address whose RD is set to 0 and whose 16byte IPv6 address is encoded as IPv4mapped Ipv6 address

Procedure/Expected Behavior:

1. Configure the PE devices with IBGP sessions
2. Ensure that there is full iBGP mesh in the network (figure 1)
3. Ensure that IPv4 routes are being exchanged between PE1/PE2/PE4
4. Ensure that PE1 advertise its IPv4mapped IPv6 address as the BGP next hop to PE3.

3.1.4 Test V6VPN 4: Verify the capability of the PE devices to correctly advertise the labeled IPv6 VPN routes to its VPN peers

The PE routers should advertise and distribute MPLS labels with IPv6 VPN routes. When the PE device receives a packet it forwards the packet directly based on the label or perform a lookup in the corresponding IPv6VPN context. The BGP Multiprotocol extensions are used to encode MP_REACH_NLRI. The AFI and SAFI fields are set to be:

AFI: 2 for IPv6 address family

SAFI: 128 for MPLS labeled VPN –IPv6 routes

Procedure/Expected Behavior:

1. Configure the PE devices with IBGP sessions
2. Ensure that there is full iBGP mesh in the network (figure 1)
3. Ensure that IPv4 routes are being exchanged between PE1/PE2/PE4
4. Configure the PE1 connected to CE11 for IPv6 address
5. Ensure that routes are being exchanged between CE11 and PE1
6. Verify that at CE42/PE4 the routes have two labels associated with the prefixes advertised by PE1 for Red VPN
7. Verify that the BGP capability is negotiated correctly before the labeled VPN IPv6 routes are exchanged
8. The outer label is being used for forwarding the packets to the BGP destinations and inner lab is bound and distributed by the MPBGP extensions and [MPLS BGP]
9. Verify that the PE4 correctly forwards the packets to PE1 for CE11 and the packets leave at the right IPv4 interface
10. Verify that the routes are withdrawn when the link between PE1 and CE11 is brought down
11. Verify that the VPN sessions between PE1 and PE4 disappear on flapping the iBGP sessions between the two

3.1.5 Test Case V6VPN 5: Verify the capability of the PE devices to use various tunneling techniques to tunnel the IPv6 packets to the remote PE devices

The tunneling type to be used in the IPv6 VPN is determined by the PE configurations. When PE receives a packet from a CE, it looks up the packet's IPv6 destination address in the VRF corresponding to that CE. If this route is found in the VRF, it will have a MPLS label associated with it and BGP next hop. Before the packet is forwarded the MPLS label is pushed on the packet and it is then encapsulated in the tunnel for transport to the egress PE identified by the BGP next hop.

Procedure/Expected Behavior:

1. Configure the PE devices with IBGP sessions
2. Ensure that there is full iBGP mesh in the network (figure 1)
3. Ensure that IPv4 routes are being exchanged between PE1/PE2/PE4
4. Enable LDP/RSVPTE on all the P router interfaces participating in the core network to be used for the IPv6 VPN
5. Configure the PE devices to setup MPLS LSPs to the BGP next hops for the VPN routes
6. Ensure that the PE1 establishes the LSP with destination address to be the BGP next hop (that is embedded IPv4 address in the IPv6 format)

4 Adhoc Test Scenarios

Overview: Additional test scenarios pertaining to the performance aspects may be included if time and resources permit. These cover various aspects of the IPv6 VPN. The areas are specified in this section. Based on the input from the vendors participating in the test event, the listed test cases may be moved to the regular testing agenda.

1. Verifying the capability of the PE devices to support IPv4 and IPv6 sites in parallel
2. Verifying the failure scenarios in the core and monitor the effects on the BGP route withdrawal and updates
3. Verify that there are no stale entries in the BGP routing tables when CE devices withdraw the routes due to:
    a. Node failures
    b. Link failures
    c. Adjacencies flapping
    d. Route withdrawal due to end systems failure
4. Monitoring the convergence of IPv6 routes learnt via BGP
5. Application of route policies in the IPv6 VPN context and verifying the PE device capabilities to correctly select routes
    a. Using weight, AS path, MED, route origin, local preference etc.
6. Other test scenarios may be added

References

[BGP/MPLS IPv6]	"BGP-MPLS IP VPN extension for IPv6 VPN", RFC 4659
[RFC 3513]	"Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513
[BGP Tunnel]	"Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", RFC 4798
[RFC 3036]	"LDP Specification", RFC 3036
[RFC 2858]	"Multiprotocol Extensions for BGP-4", RFC 2858
[MPLS BGP]	"Carrying Label Information in BGP-4", RFC 3107