Contents

Introduction

This document describes the configuration and verification of Inter-As layer 3 mpls vpn, option B feature. IOS and IOS-XR platform are used for explanation and verification. It shows a sample network scenario and its configuration and outputs for better understanding.

Prerequisites

Requirements

There are no such requirements, however basic understanding of MPLS (Multi Protocol Label Switching) and working knowledge of IOS-XR platform would certainly help.

Components Used

This document is not restricted to specific software and hardware versions. The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Background Information

MPLS is widely deployed across ISPs (Internet Service Providers) worldwide. One such service is MPLS layer 3 VPN (Virtual Private Network). MPLS Layer 3 VPNS mainly stretch the routing boundaries of a customer from one geographical location to another, ISP is mainly used as a transit. Peering with ISP on one geographical location and on the other geographical location is done, then the customer specific routes are received on the CE (Customer Edge) device from the PE (Provider Edge/ISP) device.

Now if the requirement is to stretch routing boundaries for a customer, for two different geographical locations where two different ISPs have presence. Then the two ISPs need to coordinate so that the MPLS layer 3 VPN is provided to the end customer. Such a solution is called as Inter-As layer 3 MPLS VPN.

Inter-As MPLS layer 3 VPNS can be deployed in 4 different ways, called as Option A, Option B, Option C and Option D.

Implementation using Option B is explained in this document.

Configure

Network Diagram

The topology for the Inter-As Option B exchange is shown below.

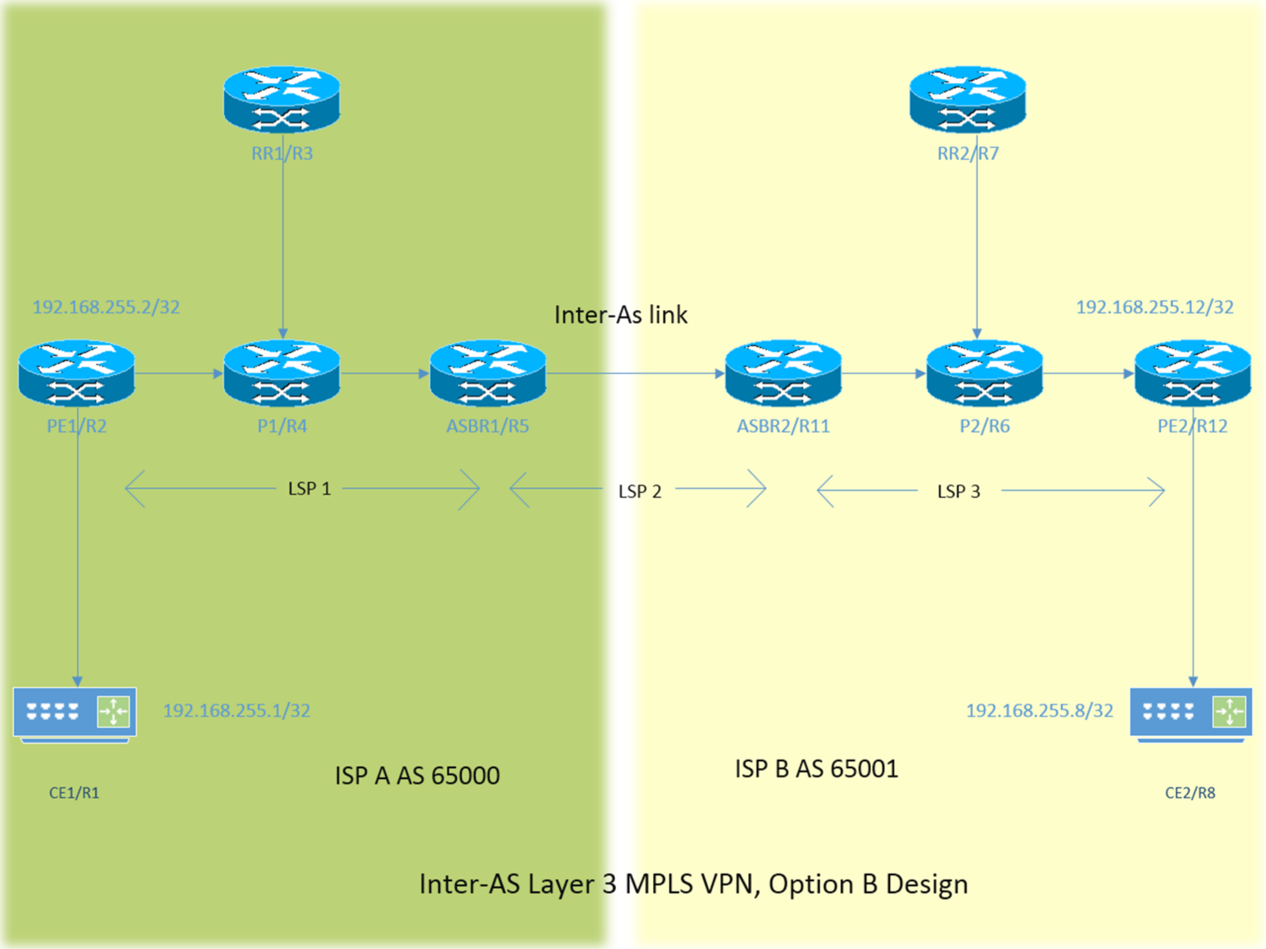

The addressing scheme is very simple. Every router has loopback1 interface described as 192.168.255.X where is X=1 when router 1 is under concern. The interface addressing is of the type 192.168.XY.X . Suppose R1 and R2 are in under consideration, configuration of the interface under the router R1 is 192.168.12.1 (here X =1 , Y = 2).

CE - Customer Edge

PE - Provider Edge

RR - Route Reflector

ASBR - Autonomous System Boundary Router

Throughout the document, the term CE denotes to both the Customer Edge devices, if specific reference has to be made for a particular device then it will be referenced as CE1. This applies to PE, RR and ASBR as well.

All the devices run IOS, however ASBR2/R11 and PE2/R12 run IOS-XR.

Two ISPs are being referenced with AS (Autonomous System) 65000 and AS 65001. ISP with AS 65000 is on the left side of the topology and is referenced as ISP A and ISP with AS 65001 is on the right side of the topology and is referenced as ISP B.

Configurations

The configurations of the devices are described below.

CE1

PE1

P1

RR1

ASBR1

ASBR2

P2

RR2

PE2

CE2

Explanation

- EIGRP as the PE-CE routing protocol is being deployed.
- OSPF is used as the IGP for the ISP core. On both the ISPs on all the physical links LDP + IGP is deployed. LDP + IGP is not configured on the Inter-As link between ASBR1 and ASBR2.
- Redistribution of EIGRP under vrf A into BGP and vice versa is performed on PE.
- Only VPNv4 address-family on PE is activated with the route reflector. The command "no bgp default ipv4-unicast" disables the default ipv4 address family peering in IOS. For IOS-XR no such command is required as it only forms the peering, with respect to the address family under which neighbor is configured.
- These redistributed routes are advertised as VPNv4 routes to the route reflector (RR).
- The route reflector reflects these routes to the ASBR device. Since reflecting the vpn4 routes

is needed, so only vpnv4 address family is activated. The route reflector will not lie in the transit path.

- The P device is just switching the labels and lies the transit path of the traffic.
- On the ASBR device "no bgp default route-target filter" for IOS and "retain route-target all" for the IOS-XR has been configured. This is important as the ABBR devices are not route-reflectors and they do not have any vrfs with RT (route target) configured, so they will implicitly drop the routing update sent to them from the route-reflectors. This is an expected behaviour as IOS and IOS-XR tend optimize the routing table information and drop the updates for those vrfs with RTs which are not locally configured.
- On the ASBRs the eBGP VPNv4 peering is configured. MPLS is not enabled with ldp on the link connecting the ASBRs.
- When the eBGP VPNv4 peering comes up on the ASBR1 (IOS) with the IOS-XR device, automatically the "mpls bgp forwarding" is configured on the Inter-As link. Exchange of the labels with ASBR2, is accomplished not via ldp but via BGP. IOS also automatically adds static /32 route to ASBR2's interface so that mpls label is bound to a /32 route and label switching is properly done.
- For IOS-XR over Inter-As link there is a different logic as compared to that of IOS. It is required to configure a static /32 route to ASBR1's interface, so that mpls label is bound for a /32 prefix. If this is not done then control plane will come up but the traffic will not be forwarded.
- IOS-XR does not sends or receives routing updates with EBGP peers unless a route policy is configured. A route policy is configured with the name DEFAULT. The action is to "pass" which means to send/receive all updates.

Verify

Ping from CE1 to CE2 and vice versa

The output of ping from CE1 to CE2 using the loopback1 interface as the source is shown below.

The output of ping from CE2 to CE1 using the loopback1 interface as the source is shown below.

Explanation of Updates Exchanged and MPLS Labels

- On CE1 show ip route gives the route for loopback1 of the CE2 on the other end.

- The traffic flow with mpls labels imposed/disposed along the path CE1 to CE2 is discussed here, i.e. how reachability is obtained when going from source loopback1 of CE1 to loopback1 of CE2. Similar information regarding the return path i.e. from CE2 loopback1 to CE1 loopback1 is also discussed.

- In MPLS layer 3 vpn designs, it should be remembered that during the label switch operation the transport label is swapped and the vpn label is untouched. The VPN label is exposed when PHP (Penutimate Hop Popping) occurs and traffic reaches PE or when a LSP (label switch path) is terminated.

- On PE1 the loopback1 of CE2 is learned via BGP VPNv4 and redistributed into to vrf aware

EIGRP. The loopback1 learned via CE1 via EIGRP is redistributed into BGP and it also becomes a VPNv4 route.

- From the above output it can be understood that, to reach to 192.168.255.8/32 prefix a vpn label of 27 learned. This output also indicates that label 23 is vpn label allocated by the BGP to advertise the reachability to the 192.168.255.1/32. The next hop for the VPNv4 prefix decides the transport label as well as the label switch path. So "show mpls forwarding-table" for the next hop 192.168.255.5 gives the transport label information to reach 192.168.255.8/32.

- The outgoing label is 21 and hence it can be concluded that to reach 192.168.255.8/32, a transport label of 21 and vpn label of 27 will be used by PE1.

- It can be also concluded that return traffic coming to 192.168.255.1/32 will be PHP'd already by the P1 router and hence will hit PE1 with vpn label of 23 and mpls forwarding table sends that traffic to Fa0/0 i.e. the CE1 after popping out the vpn label .

- The output on the route reflector gives the confirmation of the information discussed so far.

- The real interesting part is the ASBR1, here label to reach 192.168.255.1/32 is sent to ASBR2 and ASBR2 advertises the label information to reach 192.168.255.8/32. As described earlier, the next hop in the bgp vpnv4 update decides the transport label, keeping that in mind, the next hop 192.168.255.5 (for the 192.168.255.8/32 prefix learnt on PE1) belongs to the loopback1 of ASBR1. So as per the process of PHP (Penultimate hop popping) the transport label will have already been removed by P1 when the traffic destined to 192.168.255.8 hits ASBR1. So the traffic which hits the ASBR1 will hit with a vpn label of 27.The output on ASBR1 is shown below.

- It can now be clearly observed that the traffic destined to 192.168.255.8/32 when hits ASBR1 with a label of 27 will be forwarded to ASBR2 with a label of 24009 to the next hop of ASBR2 192.168.115.11. In the similar fashion, traffic destined to 192.168.255.1/32 from ASBR2 will come with label 25 and label will be swapped to 23 (vpn label) and then proper transport label will be encapsulated to forward the traffic to next-hop 192.168.255.2 (PE1).

- So the return traffic will take label 19 as the transport label and 23 as the vpn label to reach PE1 from ASBR1.

- It is important to understand that when traffic is traversing the Inter-As link, there is only single mpls label, mainly the vpn label. When the traffic is within an AS, two mpls labels are observed.

- On ASBR2 i.e. the IOS-XR device similar labels are observed.

- Here it is observed that ASBR2 advertises label 24009 to ASBR1 for the prefix 192.168.255.8/32. This output also shows that to reach 192.168.255.1/32 prefix ASBR1 has advertised label 25. Now since it is seen that to reach 192.168.255.8/32 next hop is 192.168.255.12 (PE2). The mpls forwarding table will have the LDP label or the transport label to reach the next-hop.

- To reach the 192.168.255.12 outgoing label of 19 is being used. So traffic from ASBR2 to PE2 will have two mpls labels, 19 as the transport label and 24001 as the vpn label.

- In similar way as discussed above the return traffic, i.e. from CE2 to CE1 will hit ASBR2 with a vpn label of 24007 as the transport label would already have been PHP'd by the P2 router. The label swap operation occurs and label is swapped to 25 and it is sent to next hop 192.168.115.5 i.e. the ASBR1 Inter-As link.

- PE2 is itself the next hop for the prefix 192.168.255.8/32, so PHP will be performed by the P2 router and traffic destined for 192.168.255.8/32 will hit PE2 with single mpls label i.e. the VPN label 24001.

- Hence, when the traffic hits PE2 with vpn label 24001 it is forwarded to CE2 over the link Gi0/0/0/1 and the vpn label is also removed. Also, to send traffic to 192.168.255.1/32 a vpn label of 24007 and transport label of 20 will be used by PE2.

Verification via Traceroutes

Traceroute from CE1 to CE2.

- The labels can be seen the traceroute and are exactly the same as discussed above.
- It was already mentioned that next hop the vpnv4 update controls the label switch path and hence the transport label.
- The next hop for a prefix in an Option B Inter-As design, changes 3 times and hence 3 LSPs exist.
- The prefix 192.168.255.8/32 is originated from PE2, so in the AS 65001 PE2 is the next hop for the vpnv4 update.
- This update reaches ASBR2 and now ASBR2 advertises this update to ASBR1 over the Inter-As link and hence the ASBR2 now becomes the next hop for the vpnv4 update.
- Again the same prefix is now advertised in AS 65000 via ASBR1 as vpnv4 update and so for AS 65000 ASBR1 is the next hop for the vpnv4 update.
- Since the next hop determines the LSP and it changes 3 times, 3 distinct LSPs are highlighted in the traceroute.
- It should be observed that for a distinct LSP the vpn label remains intact and does not change.

Traceroute from CE2 to CE1.

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.