Implementing Secure Converged Wide Area Networks (ISCW)

Multi-Protocol Label Switching (MPLS)

http://www.INE.com

What is MPLS?

- Multiprotocol Label Switching
  - Open standard per RFC 3031
  - Previously Cisco proprietary *Tag Switching*
- Multiprotocol
  - Can transport different payloads
    - Ethernet, Frame Relay, ATM, IPv4, IPv6, etc.
- Label Switching
  - Labels used to switch traffic instead of normal IPv4/IPv6 route
  - Similar to how Frame Relay DLCI or ATM VPI/VCI works

What is MPLS? (cont.)

- MPLS is not a WAN protocol in the traditional sense such as…
  - PoS (PPP/HDLC)
  - ATM
  - Frame Relay
  - Ethernet
- Instead MPLS runs on top of these protocols at "Layer 2.5"

Ethernet	**IP**	**TCP**	**Data**

Ethernet	**MPLS**	**IP**	**TCP**	**Data**

Why Use MPLS?

- Simplifies Service Provider (SP) network equipment and resources
  - Interoperate PPP/FR/ATM/Ethernet over IP
  - Reduce routing table lookups
  - Load distribution and high availability
- Value add services to customer
  - Provisioning quick and easy
  - SLA enforcements
  - L2/L3 VPN

How MPLS Works

- Layer 2 (e.g. Ethernet) or Layer 3 (e.g. IPv4) traffic enters SP network
- MPLS “label” is added to incoming traffic at SP edge
- SP core switches traffic towards exit point using MPLS label
- MPLS label is removed as traffic exits SP network
- Traffic is “tunneled” from SP entry to SP exit since SP core does not inspect payload

MPLS Terms

- Customer Edge (CE) Router
  - CPE equipment that (typically) does not run MPLS
  - Can be layer 2 only or layer 3 aware
- Provider Edge (PE) Router
  - Connects to CE
  - Adds MPLS label on incoming traffic from CE
  - Removes MPLS label on outgoing traffic to CE
  - AKA Label Edge Router (LER)
- Provider (P) Router
  - Connects only to PEs and other P routers
  - Switches traffic between interfaces based on MPLS labels
  - AKA Label Switch Router (LSR)

Copyright © 2009 Internetwork Expert, Inc
www.INE.com

MPLS Terms (cont.)

- MPLS Label
  - Used to choose which interface to switch traffic out
  - When more than one label exists, called the "Label Stack"
- Label Push
  - Process of adding a label to a packet
    - e.g. packet received from customer on PE
  - AKA "Label Imposition"
- Label Swap
  - Process of changing a label before switching
- Label Pop
  - Process of removing a label from a packet
    - e.g. packet sent out to customer on PE
  - AKA "Label Disposition"

MPLS Terms (cont.)

- Label Distribution Protocol
  - Dynamic protocol used to advertise IP prefix to label associations
- Forwarding Equivalency Class (FEC)
  - Actual prefix to label association
- Label Forwarding Information Based (LFIB)
  - Table where FECs are stored
  - Equivalent of the IP routing table, but for labels

MPLS Label Format

- 4 byte header with 4 fields
  - Label
    - 20 bits
    - labels 0 – 15 reserved
  - Experimental Bits (MPLS EXP)
    - 3 bits
    - QoS marking similar to IP Precedence
  - Bottom of Stack Indicator
    - 1 bit
    - Value of 1 means last label in the stack
  - Time to Live (TTL)
    - 8 bits

How Labels Work

- IP routing table populated with IP prefixes via IGP
- Labels are created to represents each prefix
  - Label to prefix binding is called a Forwarding Equivalency Class (FEC)
- Labels are dynamically advertised using a *Label Distribution Protocol*
  - Tag Distribution Protocol (TDP)
    - Legacy and Cisco Proprietary
  - Label Distribution Protocol (LDP)
    - Open Standard
  - BGP
    - Multiprotocol Extensions
  - Resource Reservation Protocol (RSVP)
    - Used for MPLS TE

Creating Label Bindings

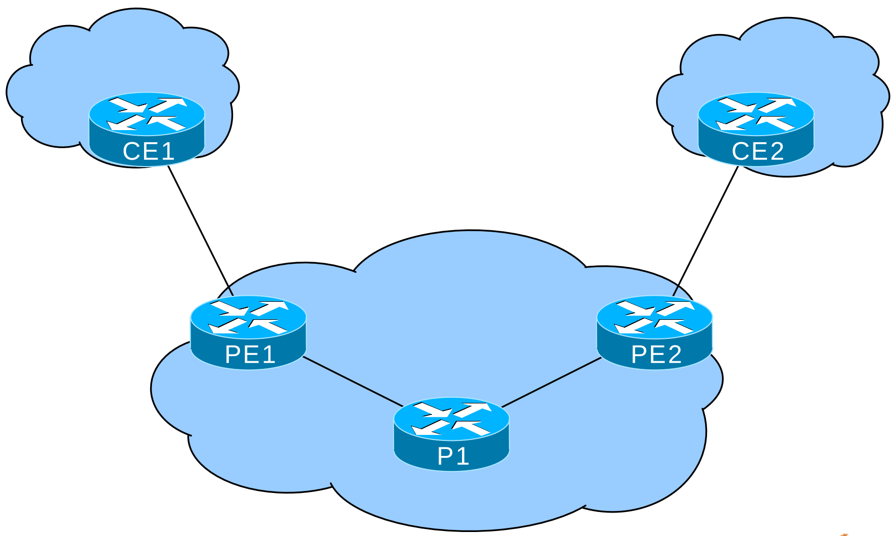

How Switching Works

- Label + Prefix binding forms Label Forwarding Information Base (LFIB)
- When traffic is received, LFIB is consulted in order to perform one of three operations
  - Label push
    - Add a label to an incoming packet
    - AKA label imposition
  - Label swap
    - Replace the label on an incoming packet
  - Label pop
    - Remove the label from an outgoing packet
    - AKA label disposition

Label Switching

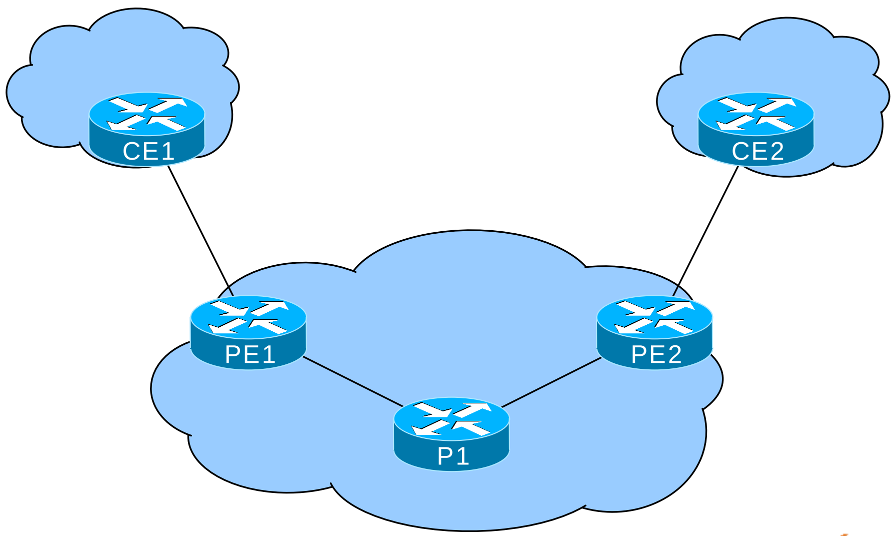

Penultimate Hop Popping (PHP)

- PE routers normally perform two lookups for traffic going to CE
  - Find label in LFIB (1)
  - Remove Label
  - Find route in routing table (2)
  - Forward packet
- Second lookup can be avoided by the next-to-last (penultimate) hop performing the pop operation
  - PE receives packet as unlabeled, no LFIB lookup required
- Accomplished by PE advertising *implicit-null* label for destination
- Default operation for Cisco IOS

Label Switching with PHP

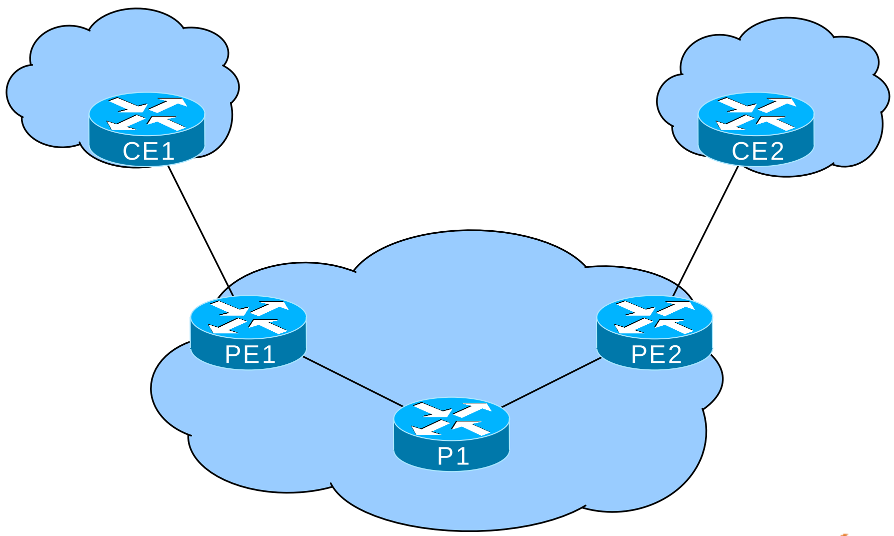

Why MPLS - Revisited

- To switch based on labels, we had to…
  - Run IGP to exchange prefixes
  - Run LDP to exchange labels
- Now we have the overhead of IGP *plus* LDP tables, how does that help?
  - LFIB lookup not faster than FIB lookup, both are done with CEF
- Answer…
  - Route only on the edge
  - Switch in the core

MPLS Tunnels

- One large advantage for SPs is that MPLS can provide "BGP free" core
  - P routers only need IGP information for internal SP routes
  - Routes outside the SP network can be label switched based on the BGP next-hop
- Logic is from P's perspective is…
  - PE sent P_1, P_2… P_n to me with BGP via next-hop X
  - I know X via IGP, and my label for it is Y
  - I don't need P_1, P_2… P_n, I'll just label switch using Y
- SP core now has tens or hundreds of routes as opposed to hundreds of thousands of routes

MPLS Tunnel Example

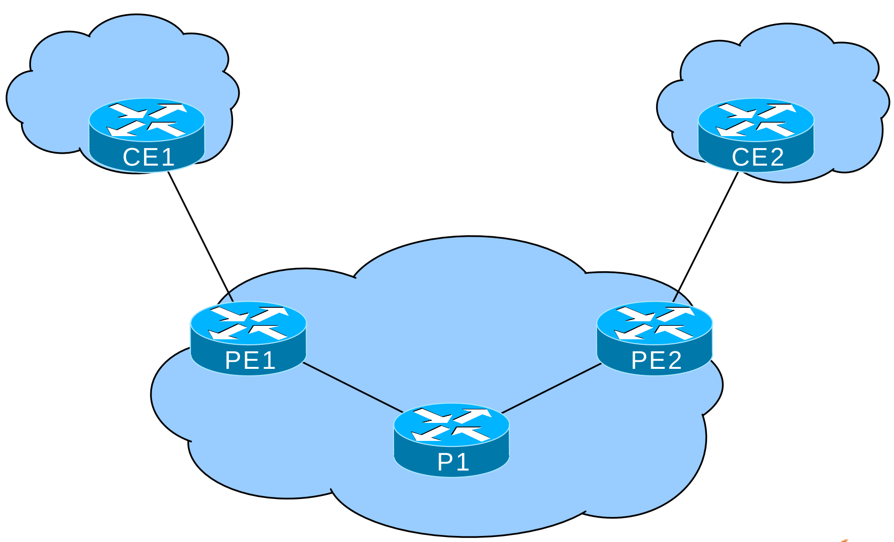

VPN Models: Overlay vs Peer

- Overlay VPNs
  - Service Provider does not participate in customer routing
- Must be provisioned prior to communication
  - Frame Relay & ATM PVCs
  - Leased lines
  - GRE & IPSec Tunnels
- Overlay suffers from (n*(n-1))/2 scalability issues
- Allows customers to use flexible addressing scheme

VPN Models: Overlay vs Peer

- Peer-to-Peer VPNs
  - Service Provider ***does*** participate in customer routing
- No static provisioning required
- Service Provider required to keep customer traffic separate through route filtering and access-lists
- Does not allow customers to use flexible addressing
- Problems with default routing

MPLS VPNs

- Best of both worlds from overlay and peer-to-peer VPNs
- No static provisioning required
  - Adding new sites doesn't necessarily require reconfiguration of other sites
- Service provider separates control plane and data plane per customer
  - Allows flexibility in customer addressing
  - Manual route and ACL filtering not required in SP
  - Customers can use default routing as needed

MPLS VPN Types

- By design, MPLS payload does not have to be IPv4 packet
- Payload could be
  - IPv4
  - IPv6
  - Ethernet
  - ATM
  - Frame Relay
  - PPP
  - …
- Result is that MPLS can offer both layer 2 and layer 3 VPN services

MPLS Layer 2 VPNs

- AKA L2VPN
- Works by encapsulating MPLS on top of full layer 2 frame before forwarding to P routers
  - e.g.

PPP	MPLS	Ethernet	IP	TCP	Data

- Traffic label switched in P network without looking at payload
- Customer does not need to route with P network
  - Customer endpoints effectively in same broadcast domain
- Called Any Transport over MPLS (AToM) and Virtual Private LAN Services (VPLS)

MPLS Layer 3 VPNs

- AKA L3VPN
- Works by encapsulating MPLS on top of IPv4/IPv6 packet before forwarding to P routers
  - e.g.

Ethernet	MPLS	IP	TCP	Data

- Traffic label switched in P network without looking at payload
- Customer must run routing protocol with P network
  - Customer endpoints on different subnets
  - Called VPNv4 routing inside P network

How L3VPN Works

- Layer 3 MPLS VPNs have two basic components…
  - Separation of customer routing information
    - Virtual Routing and Forwarding Instance (VRF)
    - Different customers have different "virtual" routing tables on the PEs
  - Exchange of customer routing information
    - MP-BGP advertises VRF routes between PEs
    - Traffic is label switched towards PEs' BGP next-hops
- Result is that P routers do not know or need to know customer routes
  - Inherently secure unless misconfigured

Virtual Routing and Forwarding

- VRFs create "virtual routers" inside IOS
- Each VRF has its own routing table
  - **show ip route vrf [*name* | *]**
- PE-CE routing protocol exchanges information
- Addressing ***can*** overlap in different VRFs
- Interfaces not in a VRF and routes learned from them are in the *global table*
- One VRF can't talk to an interface in another VRF or in the global table

Multiprotocol BGP

- RFC 4364 - BGP/MPLS IP Virtual Private Networks (VPNs)
- PE-PE routing protocol to exchange VRF routes
- Three key fields encoded as BGP extended community values
  - Route Distinguisher (RD)
  - Route Target (RT)
  - MPLS Label

Route Distinguisher (RD)

- 8 byte value included in MP-BGP update
  - Format *ASN:nn* or *w.x.y.z:nn*
- Unique per VPN or per VPN site
- Combination of unique RD and IP prefix gives global significance to route in MPLS network
  - Called VPNv4 route

Route Target (RT)

- 8 byte value included in MP-BGP update
  - Format similar to RD
  - *ASN:nn* or *w.x.y.z:nn*
- Controls which routes enter and leave the VRF
  - "export" route-target
    - Controls which routes go from VRF into BGP
  - "import" route-target
    - Controls which routes go from BGP into VRF
- Allows granular control over which sites have what routes
  - If all sites in a VPN import and export the same RT, connectivity is a full mesh

MPLS L3VPN Example

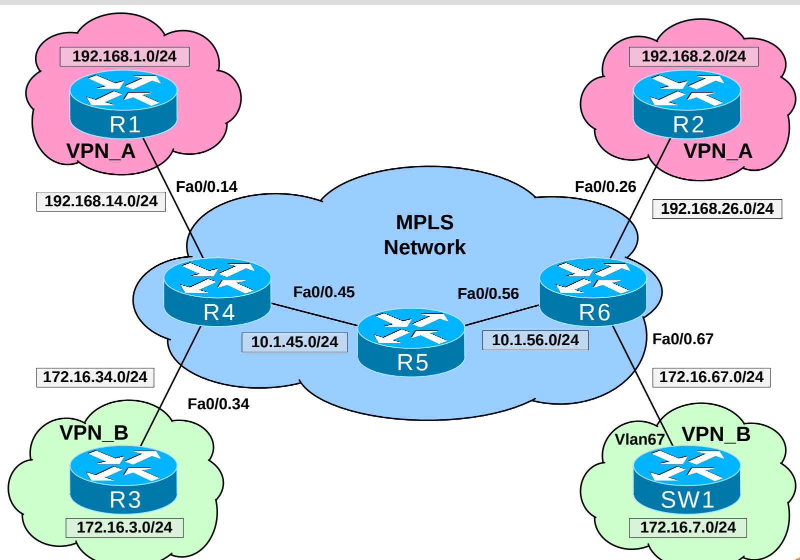

CE Configuration – VPN_A

```
R1#
interface Loopback0
 ip address 192.168.1.1 255.255.255.0
!
interface FastEthernet0/0.14
 encapsulation dot1Q 14
 ip address 192.168.14.1 255.255.255.0
!
router rip
 version 2
 network 192.168.1.0
 network 192.168.14.0
 no auto-summary

R2#
interface Loopback0
 ip address 192.168.2.2 255.255.255.0
!
interface FastEthernet0/0.26
 encapsulation dot1Q 26
 ip address 192.168.26.2 255.255.255.0
!
router rip
 version 2
 network 192.168.2.0
 network 192.168.26.0
 no auto-summary
```


CE Configuration – VPN_B

```
R3#
interface Loopback0
 ip address 172.16.3.3 255.255.255.0
!
interface FastEthernet0/0.34
 encapsulation dot1Q 34
 ip address 172.16.34.3 255.255.255.0
!
router eigrp 1
 no auto-summary
 network 172.16.0.0

SW1#
ip routing
!
interface Loopback0
 ip address 172.16.7.7 255.255.255.0
!
interface Vlan67
 ip address 172.16.67.7 255.255.255.0
!
router eigrp 1
 no auto-summary
 network 172.16.0.0
```

P Network IGP/MPLS Configuration

```
R4#
ip cef
!
mpls label protocol ldp
!
interface Loopback0
 ip address 10.1.4.4 255.255.255.255
 ip ospf 1 area 0
!
interface FastEthernet0/0.45
 encapsulation dot1Q 45
 ip address 10.1.45.4 255.255.255.0
 ip ospf 1 area 0
 mpls ip

R6#
ip cef
!
mpls label protocol ldp
!
interface Loopback0
 ip address 10.1.6.6 255.255.255.255
 ip ospf 1 area 0
!
interface FastEthernet0/0.56
 encapsulation dot1Q 56
 ip address 10.1.56.6 255.255.255.0
 ip ospf 1 area 0
 mpls ip
```

```
R5#
ip cef
!
mpls label protocol ldp
!
interface Loopback0
 ip address 10.1.5.5 255.255.255.255
 ip ospf 1 area 0
!
interface FastEthernet0/0.45
 encapsulation dot1Q 45
 ip address 10.1.45.5 255.255.255.0
 ip ospf 1 area 0
 mpls ip
!
interface FastEthernet0/0.56
 encapsulation dot1Q 56
 ip address 10.1.56.5 255.255.255.0
 ip ospf 1 area 0
 mpls ip
```


PE VRF Configuration

```
R4#
ip vrf VPN_A
 rd 100:1
 route-target export 100:1
 route-target import 100:1
!
ip vrf VPN_B
 rd 100:2
 route-target export 100:2
 route-target import 100:2
!
interface FastEthernet0/0.14
 encapsulation dot1Q 14
 ip vrf forwarding VPN_A
 ip address 192.168.14.4 255.255.255.0
!
interface FastEthernet0/0.34
 encapsulation dot1Q 34
 ip vrf forwarding VPN_B
 ip address 172.16.34.4 255.255.255.0
```

```
R6#
ip vrf VPN_A
 rd 100:1
 route-target export 100:1
 route-target import 100:1
!
ip vrf VPN_B
 rd 100:2
 route-target export 100:2
 route-target import 100:2
!
interface FastEthernet0/0.26
 encapsulation dot1Q 26
 ip vrf forwarding VPN_A
 ip address 192.168.26.6 255.255.255.0
!
interface FastEthernet0/0.67
 encapsulation dot1Q 67
 ip vrf forwarding VPN_B
 ip address 172.16.67.6 255.255.255.0
```

PE VRF Routing Configuration

```
R4#
router eigrp 65535
 !
 address-family ipv4 vrf VPN_B
  network 172.16.0.0
  no auto-summary
  autonomous-system 1
 exit-address-family
!
router rip
 !
 address-family ipv4 vrf VPN_A
  network 192.168.14.0
  no auto-summary
  version 2
 exit-address-family
```

```
R6#
router eigrp 65535
 !
 address-family ipv4 vrf VPN_B
  network 172.16.0.0
  no auto-summary
  autonomous-system 1
 exit-address-family
!
router rip
 !
 address-family ipv4 vrf VPN_A
  network 192.168.26.0
  no auto-summary
  version 2
 exit-address-family
```


PE MP-BGP Configuration

```
R4#
router bgp 100
 neighbor 10.1.6.6 remote-as 100
 neighbor 10.1.6.6 update-source Loopback0
 !
 address-family vpnv4
  neighbor 10.1.6.6 activate
  neighbor 10.1.6.6 send-community extended
 exit-address-family

R6#
router bgp 100
 neighbor 10.1.4.4 remote-as 100
 neighbor 10.1.4.4 update-source Loopback0
 !
 address-family vpnv4
  neighbor 10.1.4.4 activate
  neighbor 10.1.4.4 send-community extended
 exit-address-family
```

PE Redistribution Configuration

```
R4#
router eigrp 65535
 !
 address-family ipv4 vrf VPN_B
  redistribute bgp 100 metric 100000
100 255 1 1500
 exit-address-family
!
router rip
 !
 address-family ipv4 vrf VPN_A
  redistribute bgp 100 metric
transparent
 exit-address-family
!
router bgp 100
 !
 address-family ipv4 vrf VPN_B
  redistribute eigrp 1
 exit-address-family
 !
 address-family ipv4 vrf VPN_A
  redistribute rip
 exit-address-family
```

```
R6#
router eigrp 65535
 !
 address-family ipv4 vrf VPN_B
  redistribute bgp 100 metric 100000
100 255 1 1500
 exit-address-family
!
router rip
 !
 address-family ipv4 vrf VPN_A
  redistribute bgp 100 metric
transparent
 exit-address-family
!
router bgp 100
 !
 address-family ipv4 vrf VPN_B
  redistribute eigrp 1
 exit-address-family
 !
 address-family ipv4 vrf VPN_A
  redistribute rip
 exit-address-family
```


P Network Global Routing Table Verification

```
R4#show ip route
<output omitted>

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O       10.1.6.6/32 [110/3] via 10.1.45.5, 00:17:37, FastEthernet0/0.45
O       10.1.5.5/32 [110/2] via 10.1.45.5, 00:17:37, FastEthernet0/0.45
C       10.1.4.4/32 is directly connected, Loopback0
C       10.1.45.0/24 is directly connected, FastEthernet0/0.45
O       10.1.56.0/24 [110/2] via 10.1.45.5, 00:17:37, FastEthernet0/0.45

R5#show ip route
<output omitted>

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O       10.1.6.6/32 [110/2] via 10.1.56.6, 00:18:39, FastEthernet0/0.56
C       10.1.5.5/32 is directly connected, Loopback0
O       10.1.4.4/32 [110/2] via 10.1.45.4, 00:17:39, FastEthernet0/0.45
C       10.1.45.0/24 is directly connected, FastEthernet0/0.45
C       10.1.56.0/24 is directly connected, FastEthernet0/0.56

R6#show ip route
<output omitted>

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C       10.1.6.6/32 is directly connected, Loopback0
O       10.1.5.5/32 [110/2] via 10.1.56.5, 00:07:50, FastEthernet0/0.56
O       10.1.4.4/32 [110/3] via 10.1.56.5, 00:07:50, FastEthernet0/0.56
O       10.1.45.0/24 [110/2] via 10.1.56.5, 00:07:50, FastEthernet0/0.56
C       10.1.56.0/24 is directly connected, FastEthernet0/0.56
```

P Network LDP Neighbor Verification

```
R4#show mpls ldp neighbor
    Peer LDP Ident: 10.1.5.5:0; Local LDP Ident 10.1.4.4:0
        TCP connection: 10.1.5.5.50583 - 10.1.4.4.646
        State: Oper; Msgs sent/rcvd: 103/102; Downstream
        Up time: 01:23:26
        LDP discovery sources:
          FastEthernet0/0.45, Src IP addr: 10.1.45.5
        Addresses bound to peer LDP Ident:
          10.1.45.5       10.1.5.5        10.1.56.5

R5#show mpls ldp neighbor
    Peer LDP Ident: 10.1.6.6:0; Local LDP Ident 10.1.5.5:0
        TCP connection: 10.1.6.6.47231 - 10.1.5.5.646
        State: Oper; Msgs sent/rcvd: 104/105; Downstream
        Up time: 01:24:25
        LDP discovery sources:
          FastEthernet0/0.56, Src IP addr: 10.1.56.6
        Addresses bound to peer LDP Ident:
          10.1.56.6       10.1.6.6
    Peer LDP Ident: 10.1.4.4:0; Local LDP Ident 10.1.5.5:0
        TCP connection: 10.1.4.4.646 - 10.1.5.5.50583
        State: Oper; Msgs sent/rcvd: 102/103; Downstream
        Up time: 01:23:27
        LDP discovery sources:
          FastEthernet0/0.45, Src IP addr: 10.1.45.4
        Addresses bound to peer LDP Ident:
          10.1.45.4       10.1.4.4

R6#show mpls ldp neighbor
    Peer LDP Ident: 10.1.5.5:0; Local LDP Ident 10.1.6.6:0
        TCP connection: 10.1.5.5.646 - 10.1.6.6.47231
        State: Oper; Msgs sent/rcvd: 105/104; Downstream
        Up time: 01:24:26
        LDP discovery sources:
          FastEthernet0/0.56, Src IP addr: 10.1.56.5
        Addresses bound to peer LDP Ident:
          10.1.45.5       10.1.5.5        10.1.56.5
```


P Network LFIB Verification

```
R4#show mpls forwarding-table
Local  Outgoing      Prefix            Bytes Label   Outgoing   Next Hop
Label  Label or VC   or Tunnel Id      Switched      interface
16     No Label      192.168.1.0/24[V] 1298          Fa0/0.14   192.168.14.1
17     No Label      172.16.3.0/24[V]  590           Fa0/0.34   172.16.34.3
18     No Label      192.168.14.0/24[V]   \
                                       0             aggregate/VPN_A
19     17            10.1.6.6/32       0             Fa0/0.45   10.1.45.5
20     Pop Label     10.1.5.5/32       0             Fa0/0.45   10.1.45.5
21     Pop Label     10.1.56.0/24      0             Fa0/0.45   10.1.45.5
22     No Label      172.16.34.0/24[V] 0             aggregate/VPN_B

R5#show mpls forwarding-table
Local  Outgoing      Prefix            Bytes Label   Outgoing   Next Hop
Label  Label or VC   or Tunnel Id      Switched      interface
16     Pop Label     10.1.4.4/32       5167          Fa0/0.45   10.1.45.4
17     Pop Label     10.1.6.6/32       6600          Fa0/0.56   10.1.56.6

R6#show mpls forwarding-table
Local  Outgoing      Prefix            Bytes Label   Outgoing   Next Hop
Label  Label or VC   or Tunnel Id      Switched      interface
16     Pop Label     10.1.5.5/32       0             Fa0/0.56   10.1.56.5
17     16            10.1.4.4/32       0             Fa0/0.56   10.1.56.5
18     Pop Label     10.1.45.0/24      0             Fa0/0.56   10.1.56.5
19     No Label      192.168.2.0/24[V] 1866          Fa0/0.26   192.168.26.2
20     No Label      192.168.26.0/24[V]   \
                                       0             aggregate/VPN_A
21     No Label      172.16.7.0/24[V]  590           Fa0/0.67   172.16.67.7
22     No Label      172.16.67.0/24[V] 0             aggregate/VPN_B
```

PE MP-BGP Peering Verification

```
R4#show ip bgp vpnv4 all summary
BGP router identifier 10.1.4.4, local AS number 100
BGP table version is 17, main routing table version 17
8 network entries using 1248 bytes of memory
8 path entries using 544 bytes of memory
9/8 BGP path/bestpath attribute entries using 1512 bytes of memory
3 BGP extended community entries using 144 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 1 (at peak 2) using 32 bytes of memory
BGP using 3480 total bytes of memory
BGP activity 8/0 prefixes, 8/0 paths, scan interval 15 secs

Neighbor        V          AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.6.6        4        100      22      24       17    0    0 00:08:31        4

R5#show ip bgp vpnv4 all summary
% BGP not active

R6#show ip bgp vpnv4 all summary
BGP router identifier 10.1.6.6, local AS number 100
BGP table version is 17, main routing table version 17
8 network entries using 1248 bytes of memory
8 path entries using 544 bytes of memory
9/8 BGP path/bestpath attribute entries using 1512 bytes of memory
3 BGP extended community entries using 144 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 1 (at peak 1) using 32 bytes of memory
BGP using 3480 total bytes of memory
BGP activity 16/8 prefixes, 16/8 paths, scan interval 15 secs

Neighbor        V          AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.4.4        4        100      24      22       17    0    0 00:08:34        4
```


PE VPNv4 MP-BGP Table Verification

```
R4#show ip bgp vpnv4 all
BGP table version is 17, local router ID is 10.1.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf VPN_A)
*> 192.168.1.0      192.168.14.1             1         32768 ?
*>i192.168.2.0      10.1.6.6                 1    100      0 ?
*> 192.168.14.0     0.0.0.0                  0         32768 ?
*>i192.168.26.0     10.1.6.6                 0    100      0 ?
Route Distinguisher: 100:2 (default for vrf VPN_B)
*> 172.16.3.0/24    172.16.34.3         156160         32768 ?
*>i172.16.7.0/24    10.1.6.6            156160    100      0 ?
*> 172.16.34.0/24   0.0.0.0                  0         32768 ?
*>i172.16.67.0/24   10.1.6.6                 0    100      0 ?

R5#show ip bgp vpnv4 all
% BGP not active

R6#show ip bgp vpnv4 all
BGP table version is 17, local router ID is 10.1.6.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf VPN_A)
*>i192.168.1.0      10.1.4.4                 1    100      0 ?
*> 192.168.2.0      192.168.26.2             1         32768 ?
*>i192.168.14.0     10.1.4.4                 0    100      0 ?
*> 192.168.26.0     0.0.0.0                  0         32768 ?
Route Distinguisher: 100:2 (default for vrf VPN_B)
*>i172.16.3.0/24    10.1.4.4            156160    100      0 ?
*> 172.16.7.0/24    172.16.67.7         156160         32768 ?
*>i172.16.34.0/24   10.1.4.4                 0    100      0 ?
*> 172.16.67.0/24   0.0.0.0                  0         32768 ?
```

PE VRF Routing Table Verification - VPN_A

```
R4#show ip route vrf VPN_A

Routing Table: VPN_A
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       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
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    192.168.14.0/24 is directly connected, FastEthernet0/0.14
B    192.168.26.0/24 [200/0] via 10.1.6.6, 00:06:22
R    192.168.1.0/24 [120/1] via 192.168.14.1, 00:00:27, FastEthernet0/0.14
B    192.168.2.0/24 [200/1] via 10.1.6.6, 00:06:22

R5#show ip route vrf VPN_A
% IP routing table VPN_A does not exist

R6#show ip route vrf VPN_A

Routing Table: VPN_A
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       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
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

B    192.168.14.0/24 [200/0] via 10.1.4.4, 00:06:16
C    192.168.26.0/24 is directly connected, FastEthernet0/0.26
B    192.168.1.0/24 [200/1] via 10.1.4.4, 00:06:16
R    192.168.2.0/24 [120/1] via 192.168.26.2, 00:00:21, FastEthernet0/0.26
```


PE VRF Routing Table Verification - VPN_B

```
R4#show ip route vrf VPN_B

Routing Table: VPN_B
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       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
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 4 subnets
C       172.16.34.0 is directly connected, FastEthernet0/0.34
B       172.16.7.0 [200/156160] via 10.1.6.6, 00:06:41
D       172.16.3.0 [90/156160] via 172.16.34.3, 00:12:38, FastEthernet0/0.34
B       172.16.67.0 [200/0] via 10.1.6.6, 00:06:41

R5#show ip route vrf VPN_B
% IP routing table VPN_B does not exist

R6#show ip route vrf VPN_B

Routing Table: VPN_B
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       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
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 4 subnets
B       172.16.34.0 [200/0] via 10.1.4.4, 00:06:36
D       172.16.7.0 [90/156160] via 172.16.67.7, 00:13:38, FastEthernet0/0.67
B       172.16.3.0 [200/156160] via 10.1.4.4, 00:06:36
C       172.16.67.0 is directly connected, FastEthernet0/0.67
```

CE Routing Table Verification - VPN_A

```
R1#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       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
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    192.168.14.0/24 is directly connected, FastEthernet0/0.14
R    192.168.26.0/24 [120/1] via 192.168.14.4, 00:00:27, FastEthernet0/0.14
C    192.168.1.0/24 is directly connected, Loopback0
R    192.168.2.0/24 [120/2] via 192.168.14.4, 00:00:27, FastEthernet0/0.14

R2#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       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
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

R    192.168.14.0/24 [120/1] via 192.168.26.6, 00:00:13, FastEthernet0/0.26
C    192.168.26.0/24 is directly connected, FastEthernet0/0.26
R    192.168.1.0/24 [120/2] via 192.168.26.6, 00:00:13, FastEthernet0/0.26
C    192.168.2.0/24 is directly connected, Loopback0
```


CE Routing Table Verification - VPN_B

```
R3#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       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
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 4 subnets
C       172.16.34.0 is directly connected, FastEthernet0/0.34
D       172.16.7.0 [90/158720] via 172.16.34.4, 00:04:14, FastEthernet0/0.34
C       172.16.3.0 is directly connected, Loopback0
D       172.16.67.0 [90/30720] via 172.16.34.4, 00:04:14, FastEthernet0/0.34

SW1#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       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
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 4 subnets
D       172.16.34.0 [90/28416] via 172.16.67.6, 00:04:07, Vlan67
C       172.16.7.0 is directly connected, Loopback0
D       172.16.3.0 [90/156416] via 172.16.67.6, 00:04:07, Vlan67
C       172.16.67.0 is directly connected, Vlan67
```

LFIB & CEF Detailed Verification

```
R4#show ip cef vrf VPN_A 192.168.2.0 detail
192.168.2.0/24, epoch 0
  recursive via 10.1.6.6 label 19
    nexthop 10.1.45.5 FastEthernet0/0.45 label 17

R4#show arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.1.45.4               -   001e.7a9e.417c  ARPA   FastEthernet0/0.45
Internet  10.1.45.5             113   001a.6ddc.2654  ARPA   FastEthernet0/0.45

R4#show mpls forwarding-table vrf VPN_A 192.168.2.0 detail
Local  Outgoing      Prefix            Bytes Label   Outgoing   Next Hop
Label  Label or VC   or Tunnel Id      Switched      interface
None   19            192.168.2.0/24[V] 0             Fa0/0.45   10.1.45.5
        MAC/Encaps=18/26, MRU=1496, Label Stack{17 19}
        001A6DDC2654001E7A9E417C8100002D8847 0001100000013000
        VPN route: VPN_A
        No output feature configured

R4#traceroute vrf VPN_A 192.168.2.2 source 192.168.14.4

Type escape sequence to abort.
Tracing the route to 192.168.2.2

  1 10.1.45.5 [MPLS: Labels 17/19 Exp 0] 4 msec 0 msec 4 msec
  2 192.168.26.6 [MPLS: Label 19 Exp 0] 0 msec 4 msec 0 msec
  3 192.168.26.2 4 msec *  0 msec
```


For More Information on MPLS

- Beginner
  - MPLS Fundamentals (Cisco Press) ISBN 1-58705-197-4
- Intermediate
  - MPLS Configuration on Cisco IOS Software (Cisco Press) ISBN 1-58705-199-0
- Advanced
  - MPLS-Enabled Applications: Emerging Developments and New Technologies (Wiley) ISBN 0-47001-453-9

MPLS Q&A