# Multi Protocol Label Switching (MPLS)

Presenters:

Daryl Wan

Yash

# Agenda

1. Overview
2. Use Cases
3. Details and Caveats
4. Configuration
5. Best Practices
6. Troubleshooting
7. Demo
8. Additional Resources

# Overview

# MPLS Overview

- 10.9 adds support for MPLS and L3 VPNs
- MPLS switches packets based on labels

- Supported platform:
  - 8360

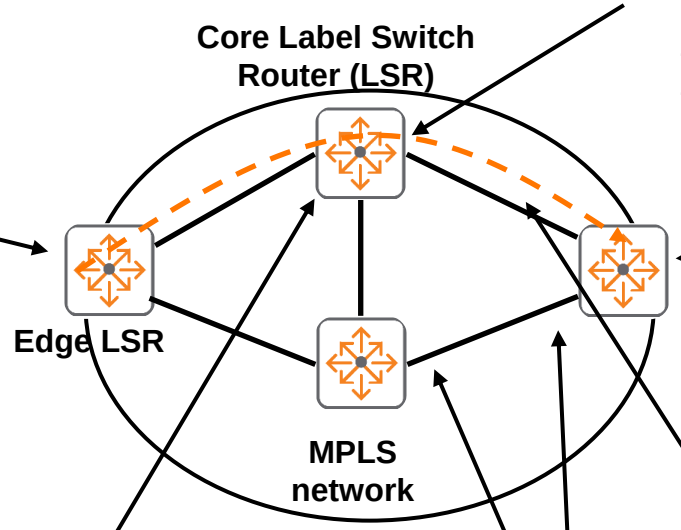


1. At Ingress Edge:

**Label imposition** (PUSH)

– Classify and label packets

2. In the Core:

**Label switching** (SWAP)

– Forward using labels (not IP)

– Label indicates service class and destination

3. At Egress Edge:

**Label disposition** (POP)

– Remove labels and forward packets

**Label Switch Path (LSP)**

– The chain of labels that are swapped at each hop to get from one LSR to another

– Unidirectional

Penultimate Hop Popping (PHP):

– Optimizes MPLS performance by reducing 1 less lookup at Edge LSR

– Edge LSR will just need to perform IP lookup instead of removing label and performing IP lookup

– Refer to caveats section

Label Distribution Protocol (LDP):

– LDP sessions are established between peers to distribute labels

# MPLS L3 VPN Overview

- MPLS L3 VPNs are deployed by service providers to provide L3 network connectivity and multi-tenant traffic isolation using an MPLS network
  - Different VRFs may utilize overlapping IP subnets if desired
  - Core LSR = P/RR
  - Edge LSR = PE
  - CE routers do not need to understand MPLS or VPNv4
- Multi Protocol BGP (MP-BGP) is used between PE routers to exchange VPNv4 addresses, extended community and labels
- Full mesh MP-IBGP peering is avoided between PEs when P routers function as VPNv4 RR
  - PEs peer to dual RRs for redundancy

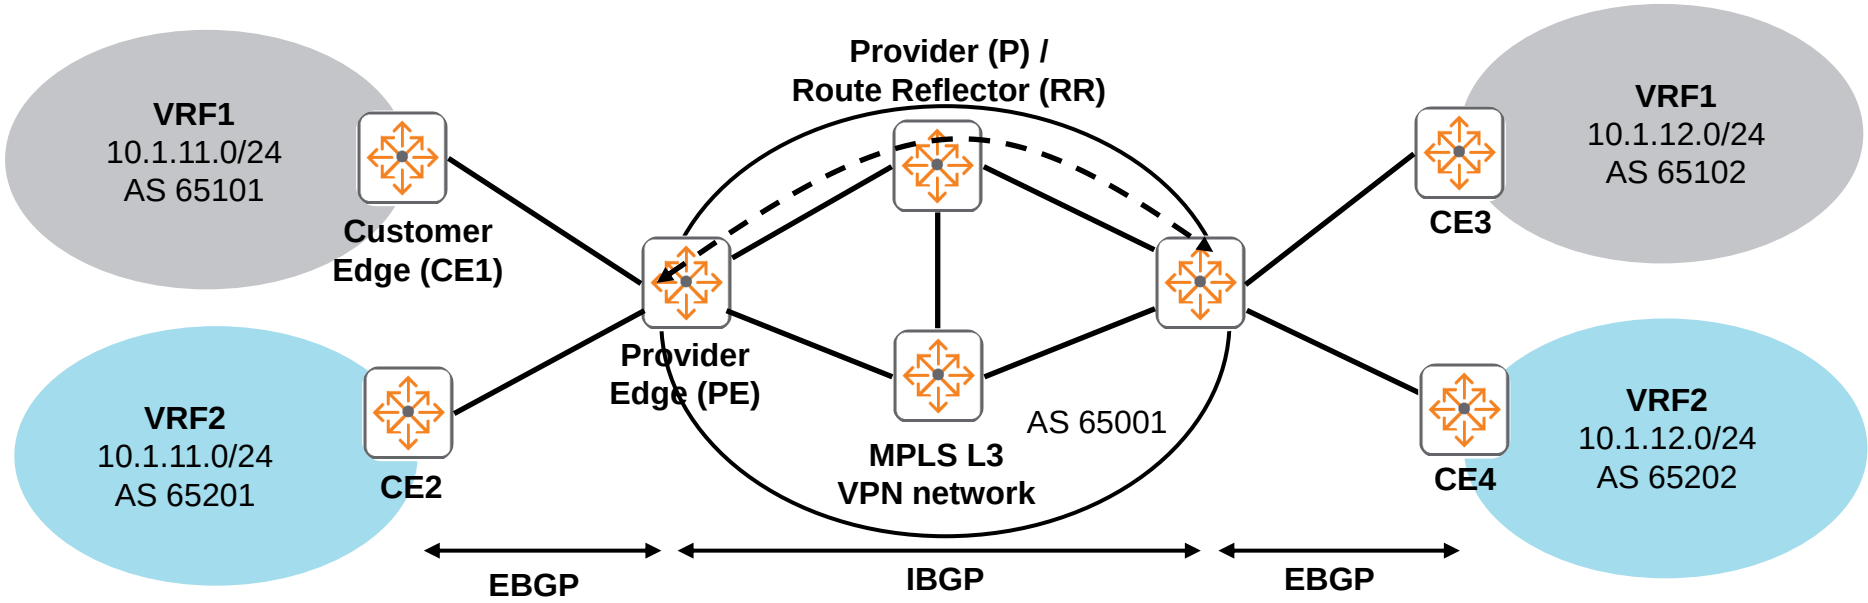

# Use Cases

# Existing MPLS L3 VPN network expansion with 8360 PEs

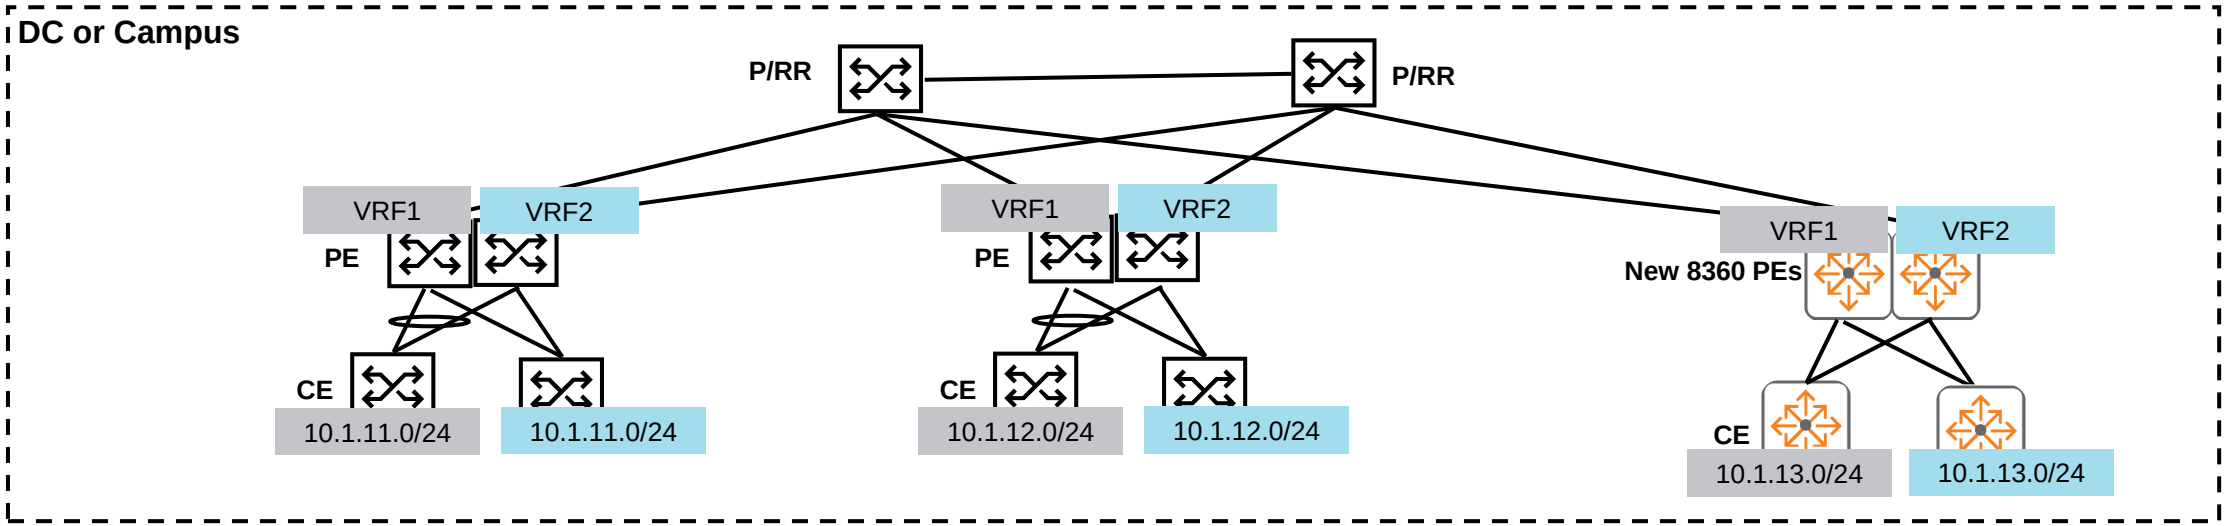


- Customer not looking to change network architecture
- Existing MPLS L3 VPN network could be Comware or other 3rd party vendor
- Willing to deploy AOS-CX MPLS for expansion

# New MPLS L3 VPN network with 8360s

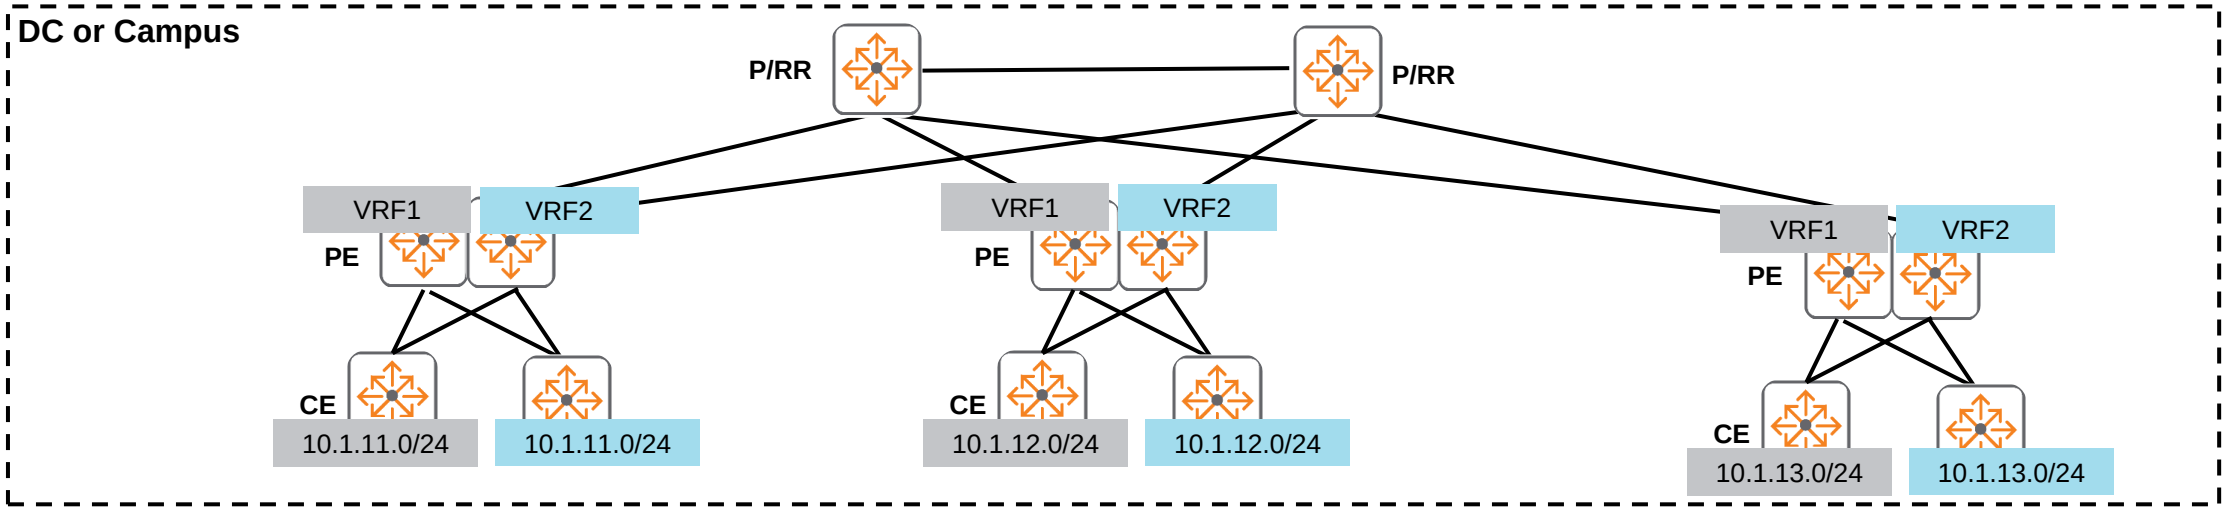


- Isolated tenant VRFs
  - Overlapping IP subnets are possible
- New MPLS L3 VPN network
  - Size of MPLS network limited to port density of 8360 P/RRs

| | |
|---|---|
| 8360-12C Port to Power Bundle | 12 QSFP28 100GbE ports.<br>[ JL708A $22,995.00 ] |
| 8360-12C Power to Port Bundle | 12 QSFP28 100GbE ports.<br>[ JL709A $22,995.00 ] |
| 8360-32Y4C MACsec Port to Power Bundle | 32 SFP28 25GbE ports, 4 QSFP28 100GbE ports.<br>[ JL700A $26,995.00 ] |
| 8360-32Y4C MACsec Power to Port Bundle | 32 SFP28 25GbE ports, 4 QSFP28 100GbE ports.<br>[ JL701A $26,995.00 ] |

# Hub & Spoke MPLS L3 VPN network with 8360s

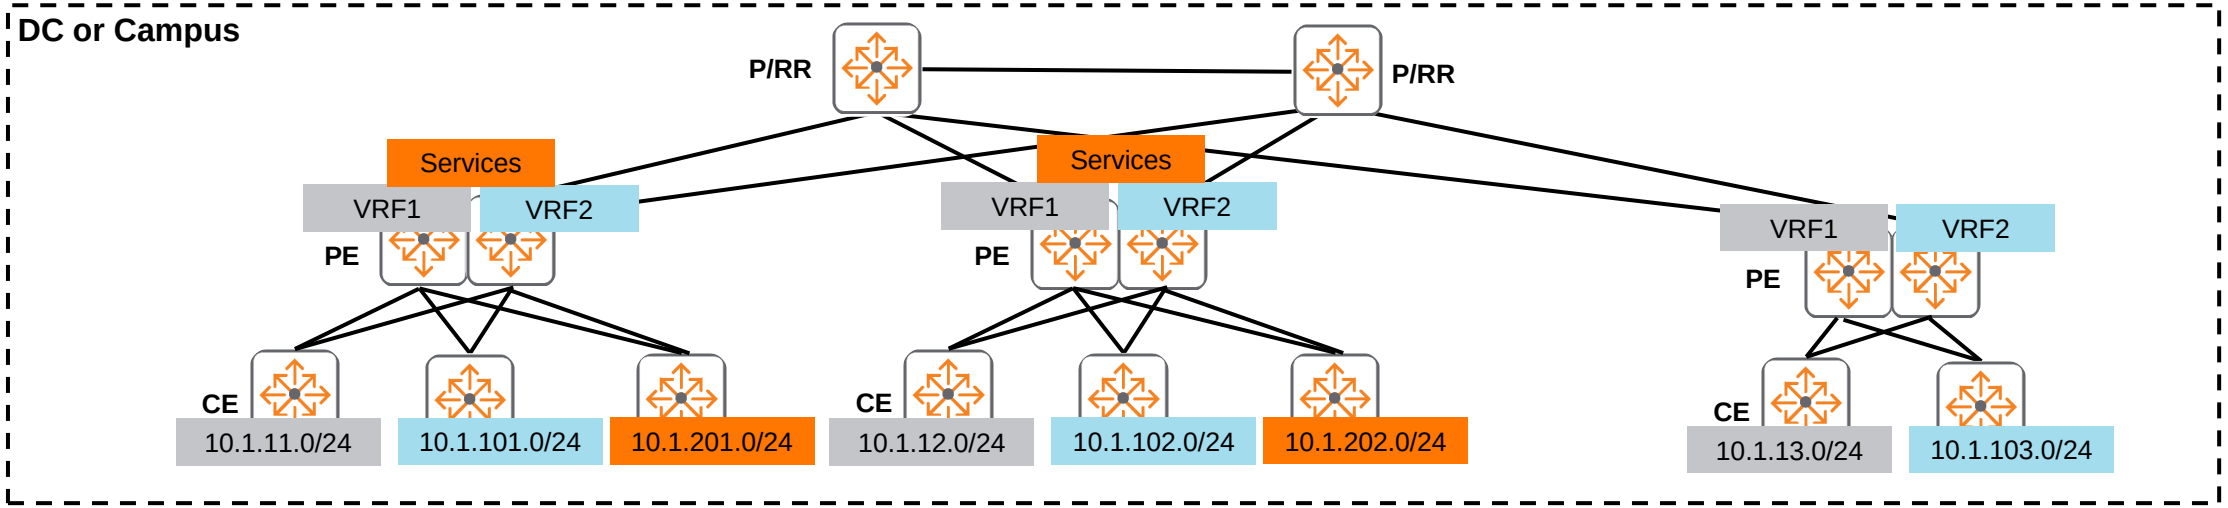


- Isolated tenant VRFs
  - No communication between subnets in VRF1 and VRF2
  - No overlapping IP subnets
- Shared “Services” VRF
  - VRF1/VRF2 has access to shared VRF subnets
  - Shared VRF has access to VRF1/VRF2 subnets

# Details

# LDP and Downstream Label Allocation

– LDP sessions are established between peers to distribute labels in the downstream direction

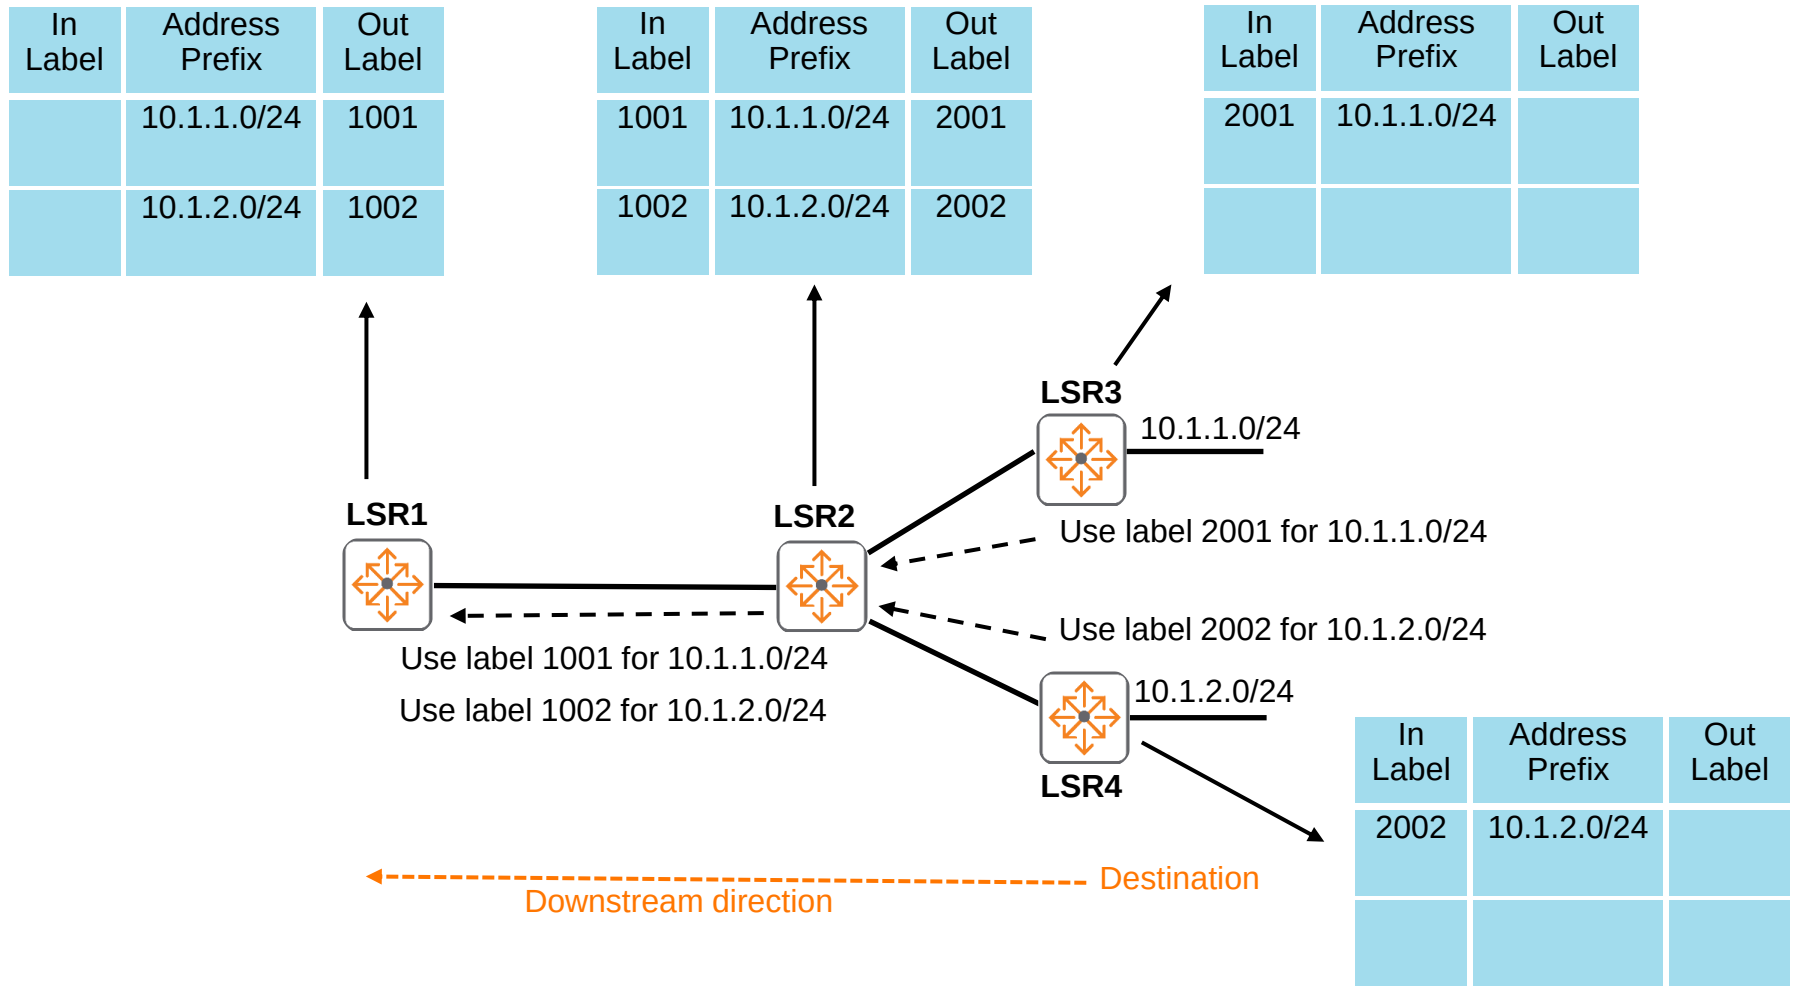

# Label Based Forwarding

– Labels are pushed, swapped and popped in the upstream direction

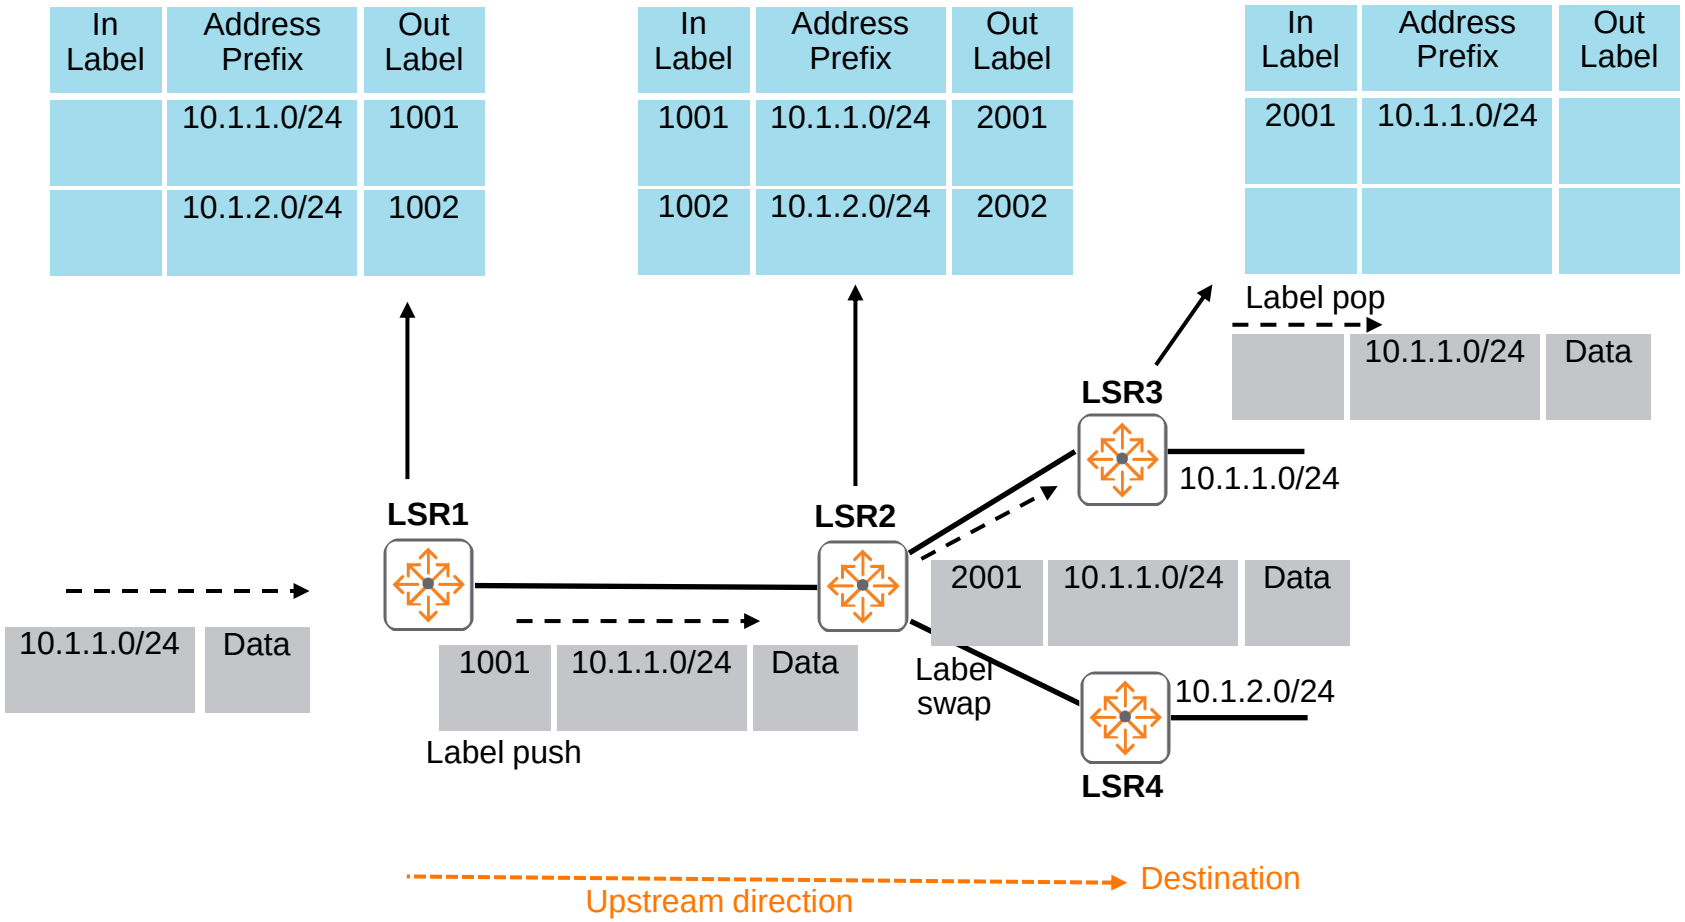

# MPLS L3 VPN Label Based Forwarding

- Label stacking required for MPLS L3 VPNs
- Outer labels (based on LDP) are used to forward traffic towards PE loopbacks
- Inner labels (based on VPNv4) are used to forward traffic towards CE subnets
  - VPNv4 label lookup only required on PEs

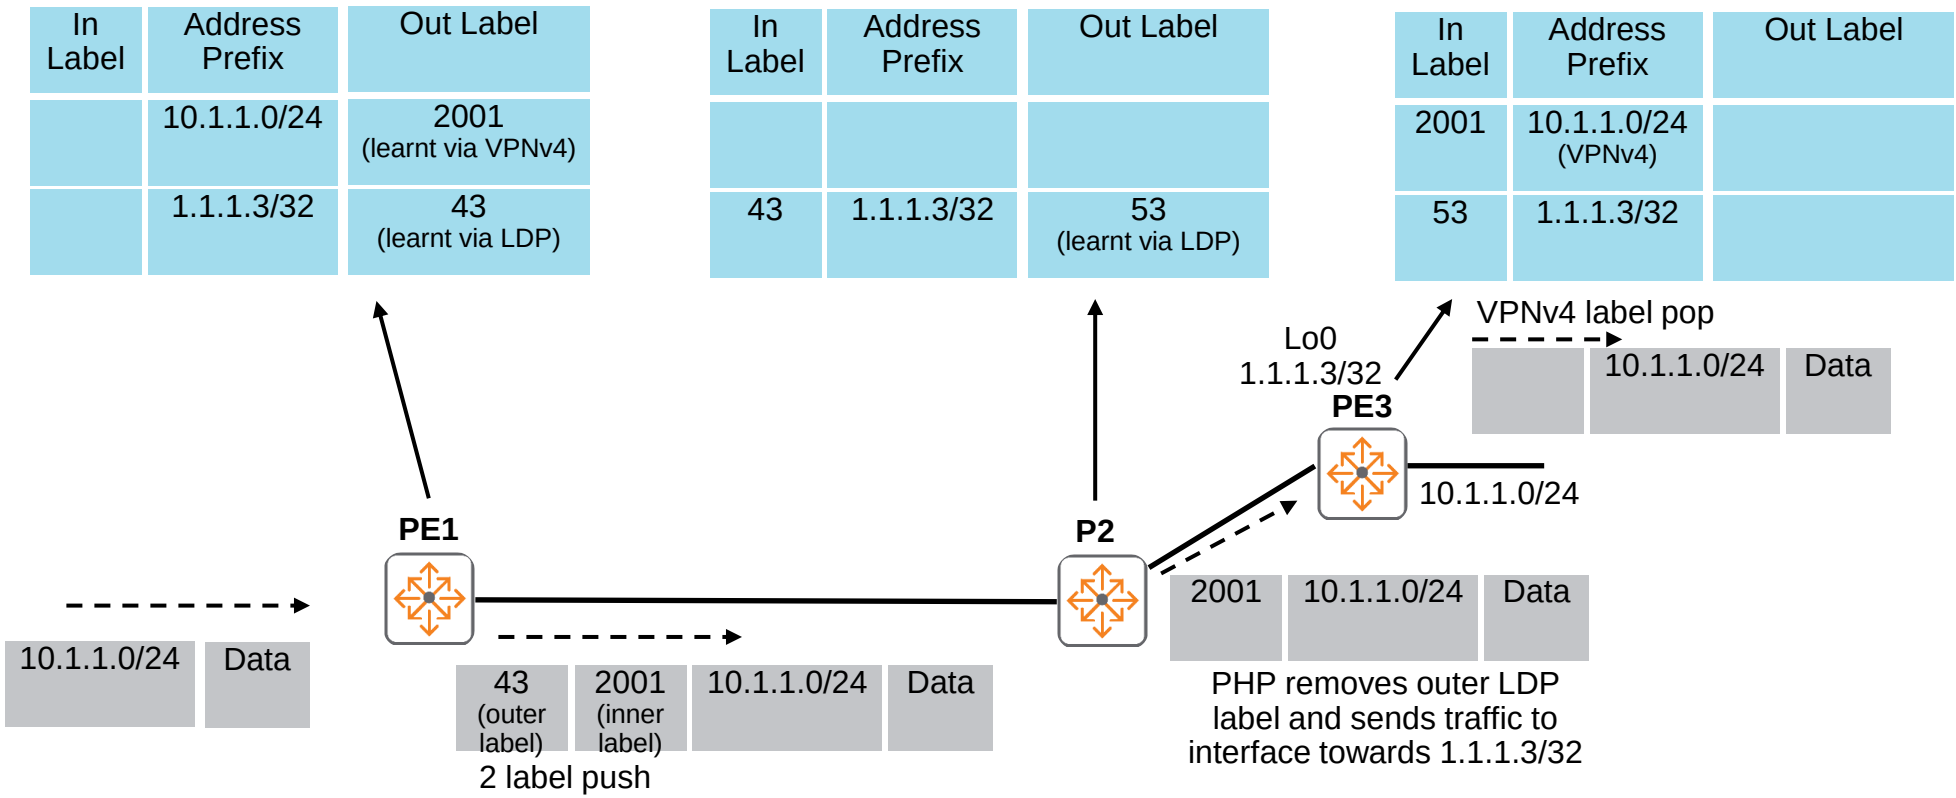

# Outer LDP Label Packet Capture

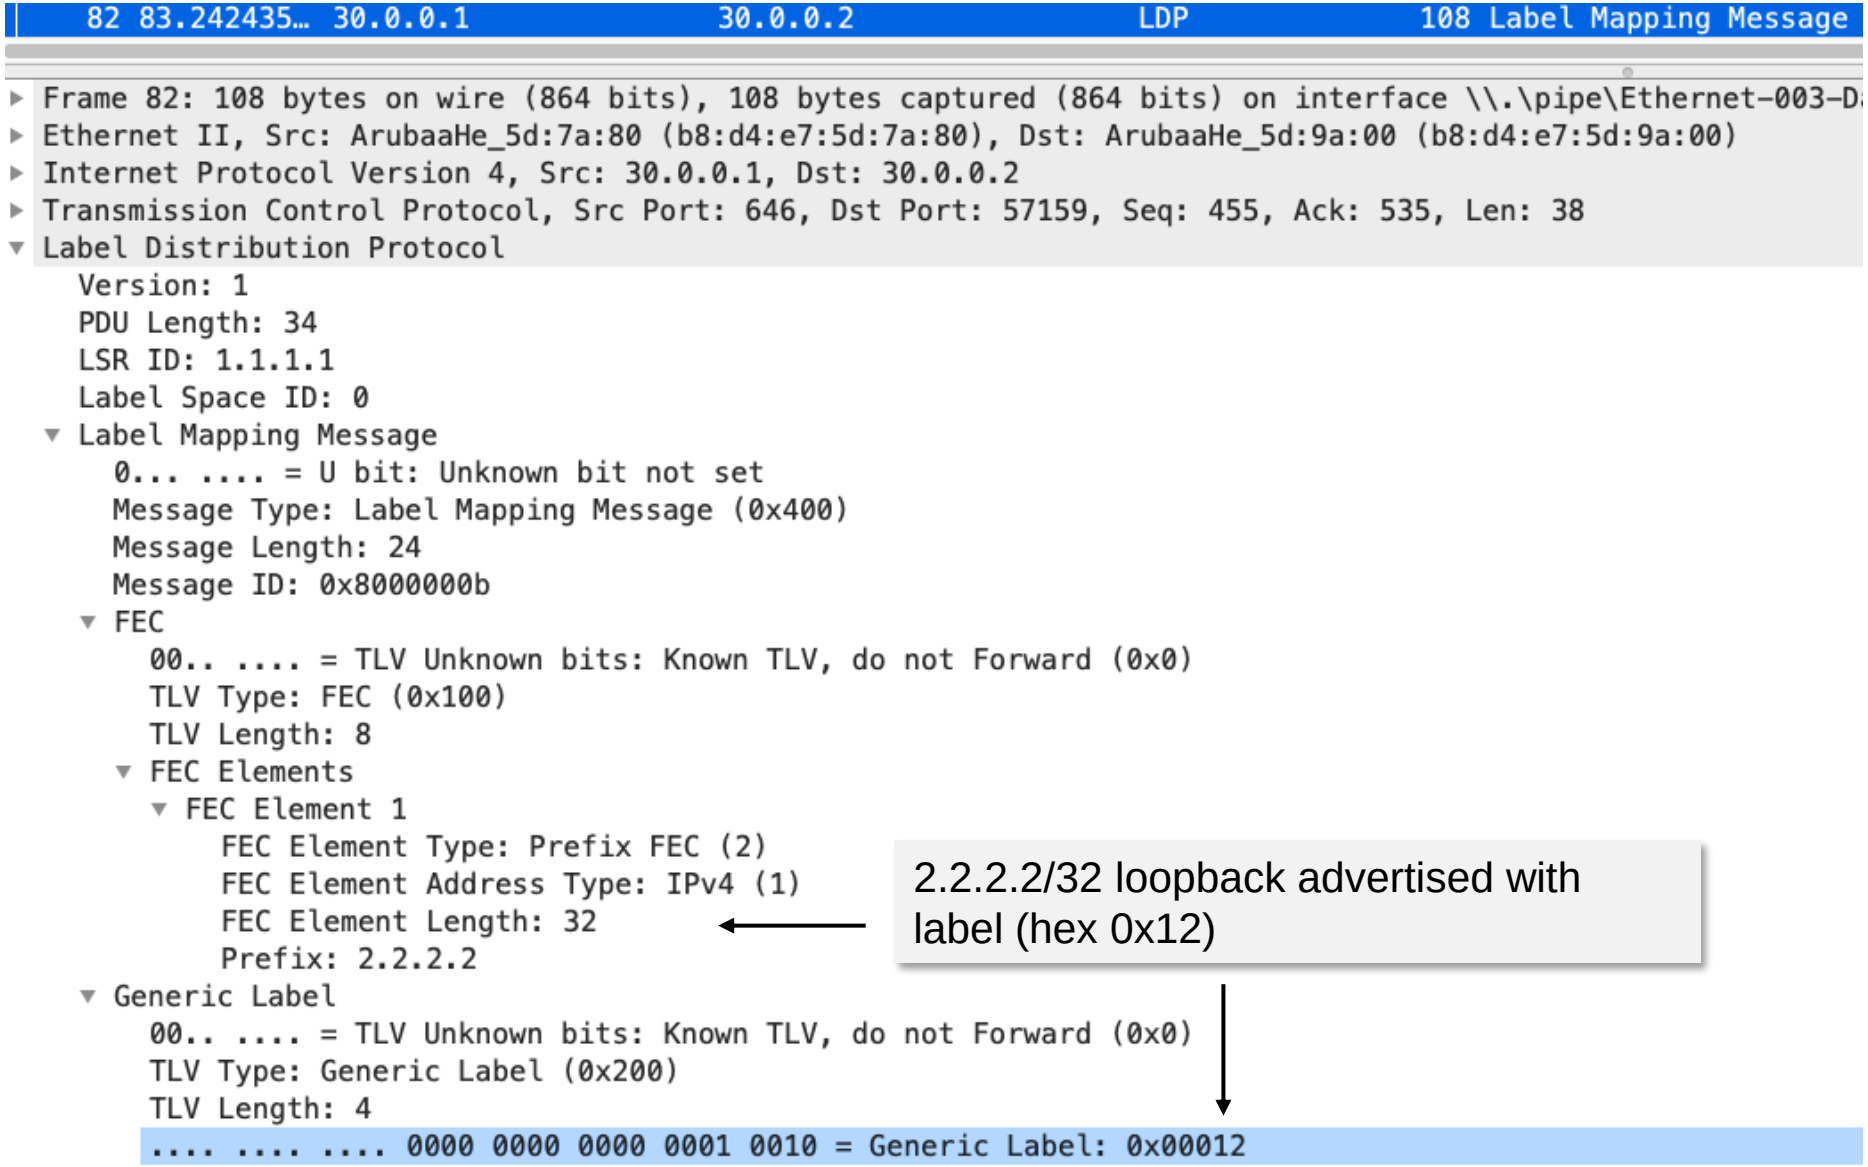

# Inner VPNv4 Label Packet Capture

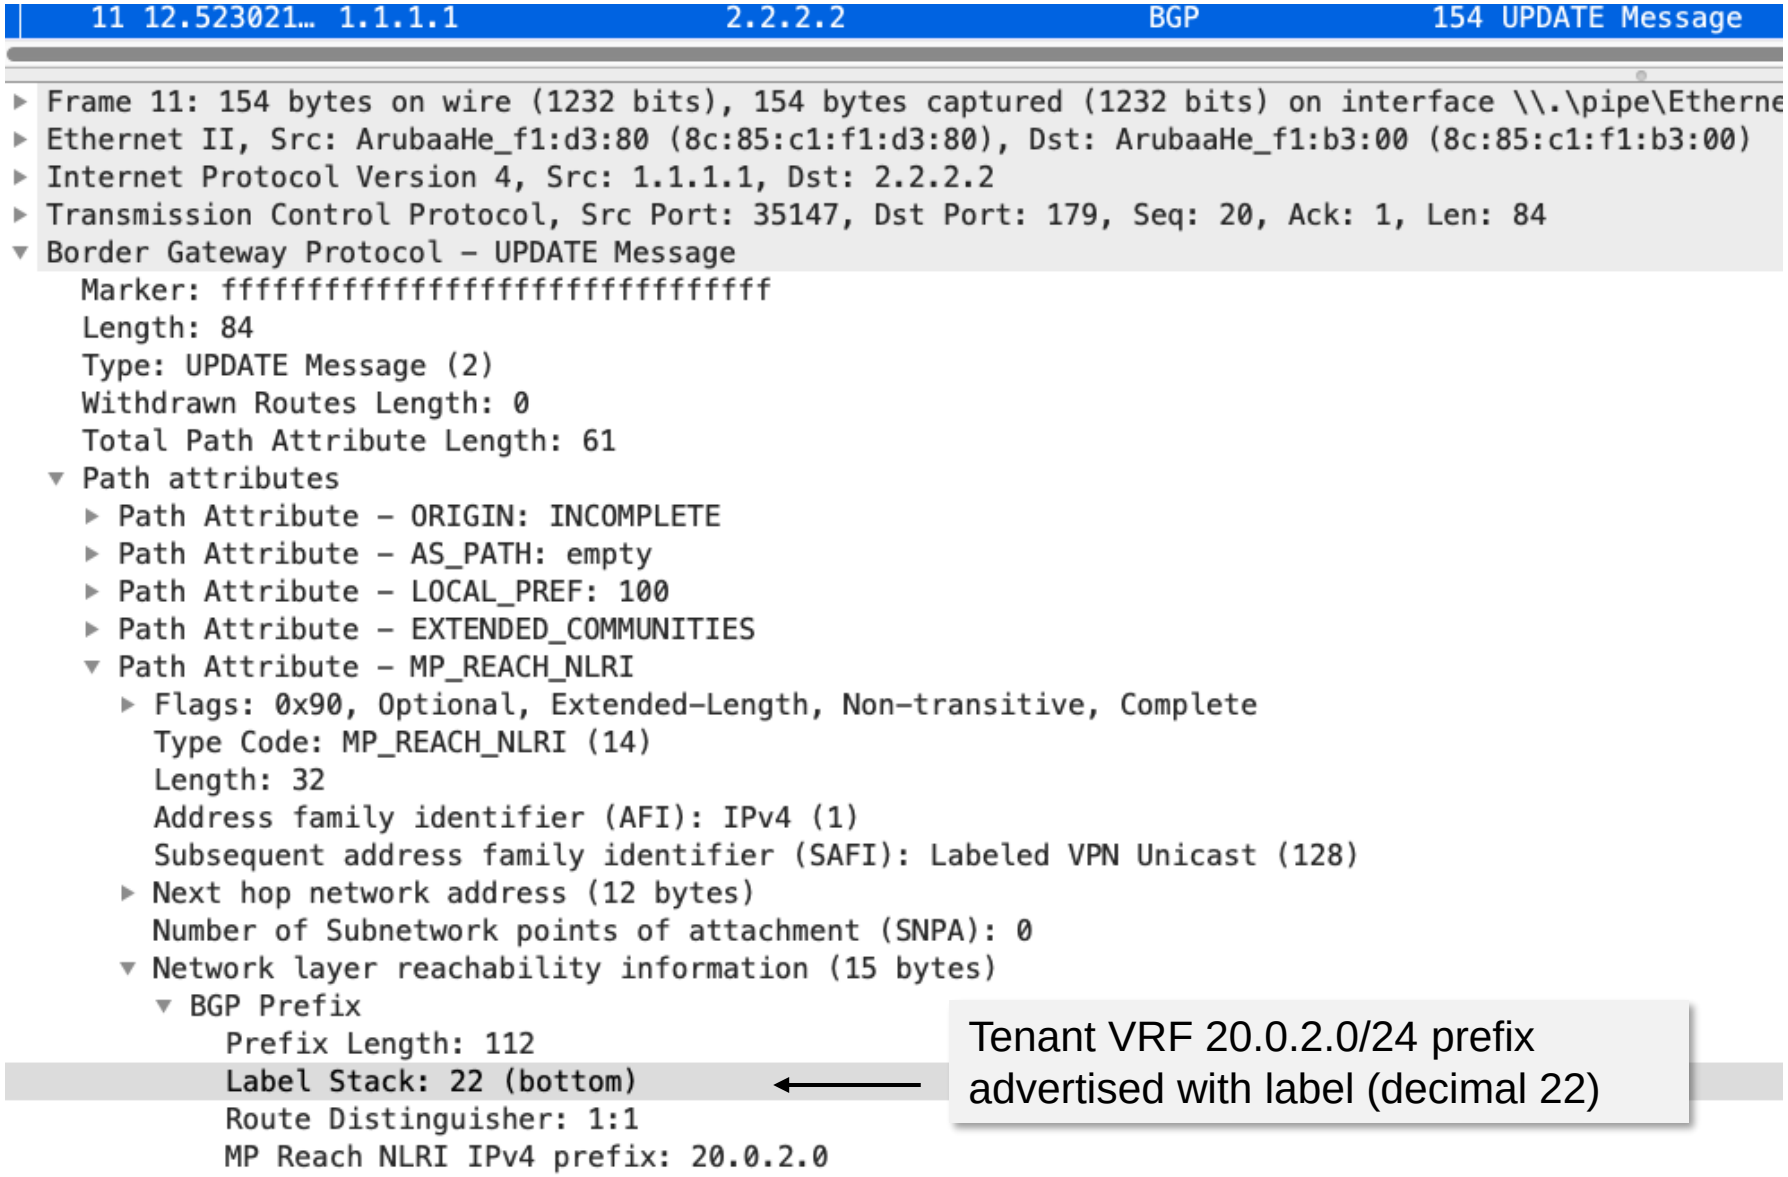


Tenant VRF 20.0.2.0/24 prefix advertised with label (decimal 22)

# MPLS L3 VPN Deployment Requirements and Recommendations

- High availability recommendations:
  - Dual standalone P/RRs
  - Dual standalone PEs
  - Dual CEs (could be VSX or 2 x standalone switches)
  - Redundant ECMP links
- OSPF required to exchange loopback IPs between PE/P
- LDP required to exchange outer labels
- IBGP MP-BGP VPNv4 peering between PE/P(RR) loopbacks to exchange inner VPNv4 labels
- EBGP IPv4 peering recommended between CE and PE

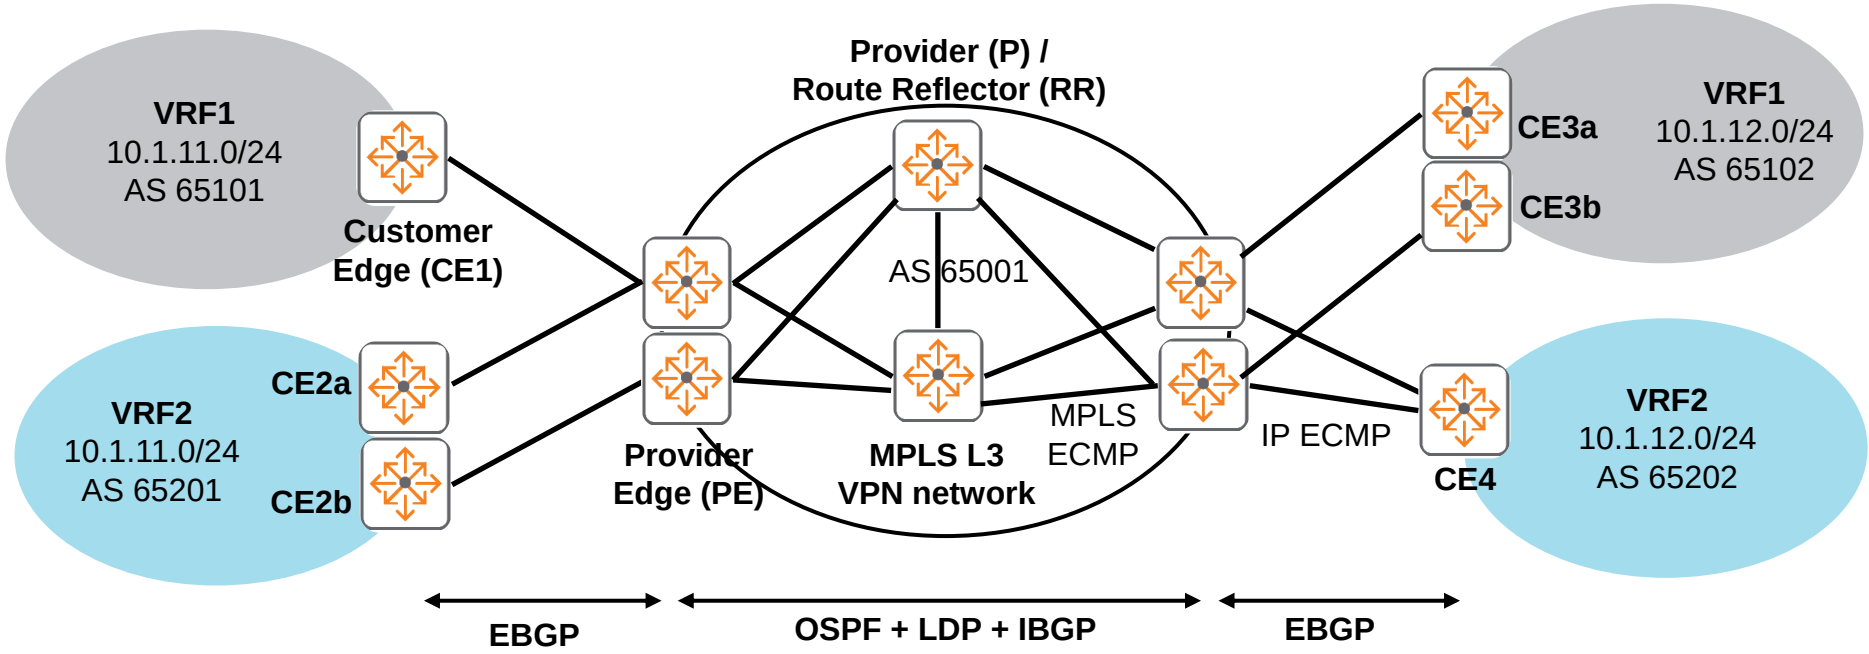

# MPLS L3 VPN and QoS Uniform Mode

- With implicit null/PHP enabled on egress PE
  - If CE sends traffic with DSCP markings
  - DSCP values are automatically copied to EXP bits on ingress PE
  - EXP bits are used within the MPLS network
  - Remote CE will receive DSCP marked traffic

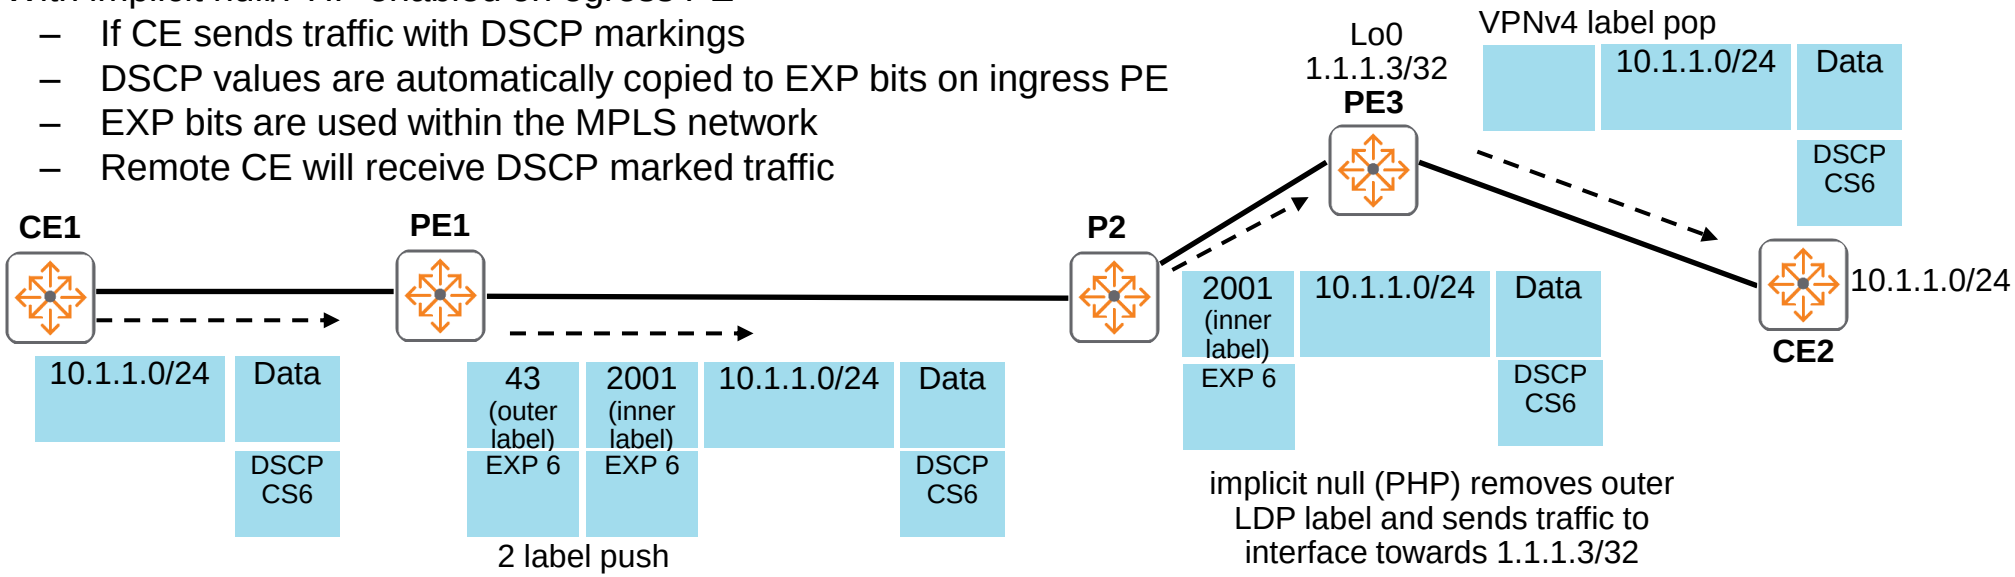


---

- With explicit null enabled on egress PE
- Allows the outer EXP value to be modified and sends modified DSCP value to remote CE

- exp-null sent by egress PE3 to P2
- Copies top label EXP value to DSCP

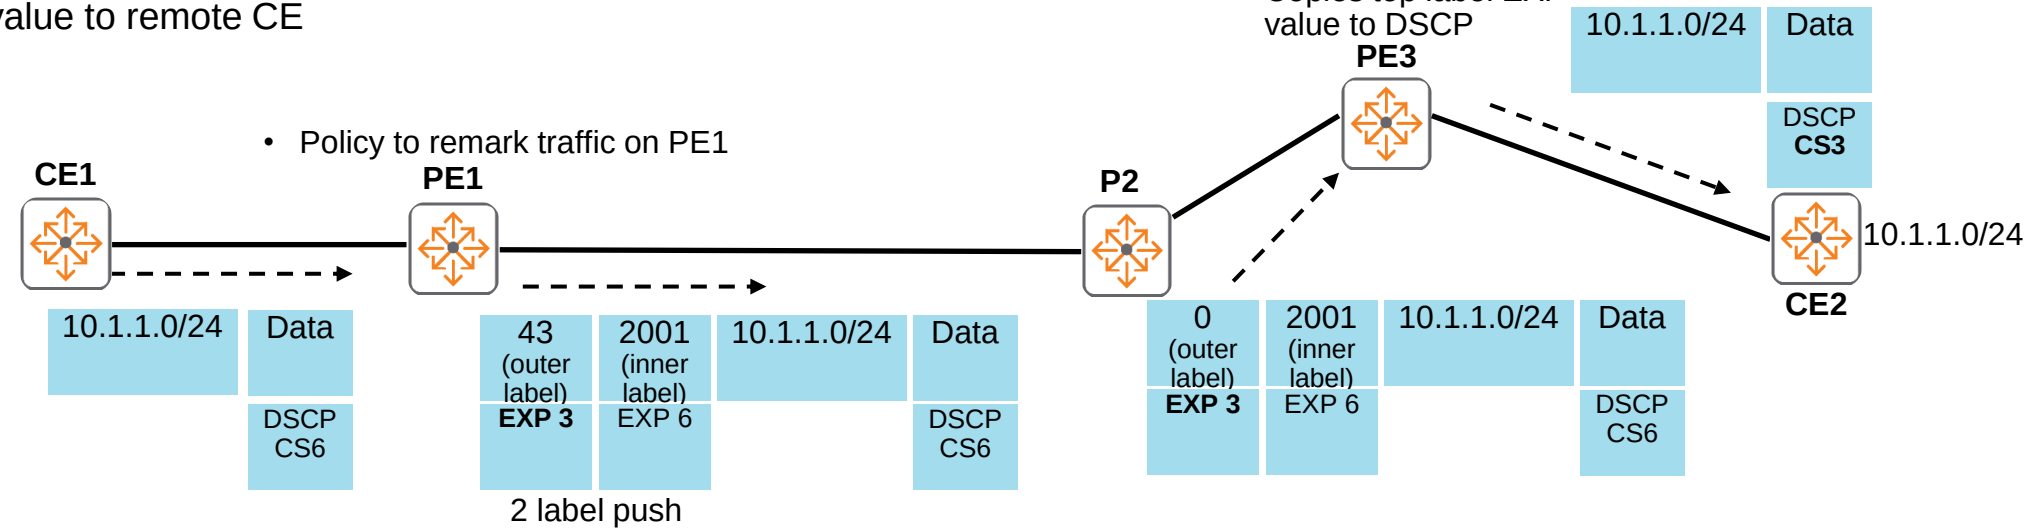

# MPLS ECMP

- MPLS Equal Cost Multi Pathing (ECMP) provides load balancing across Link Aggregation Group (LAG) and links over multiple paths that have the same cost
- As shown in the example below, ingress PE would see 2 different labels towards the egress PE Lo0 (2.2.2.2/32) and load share traffic across both P routers

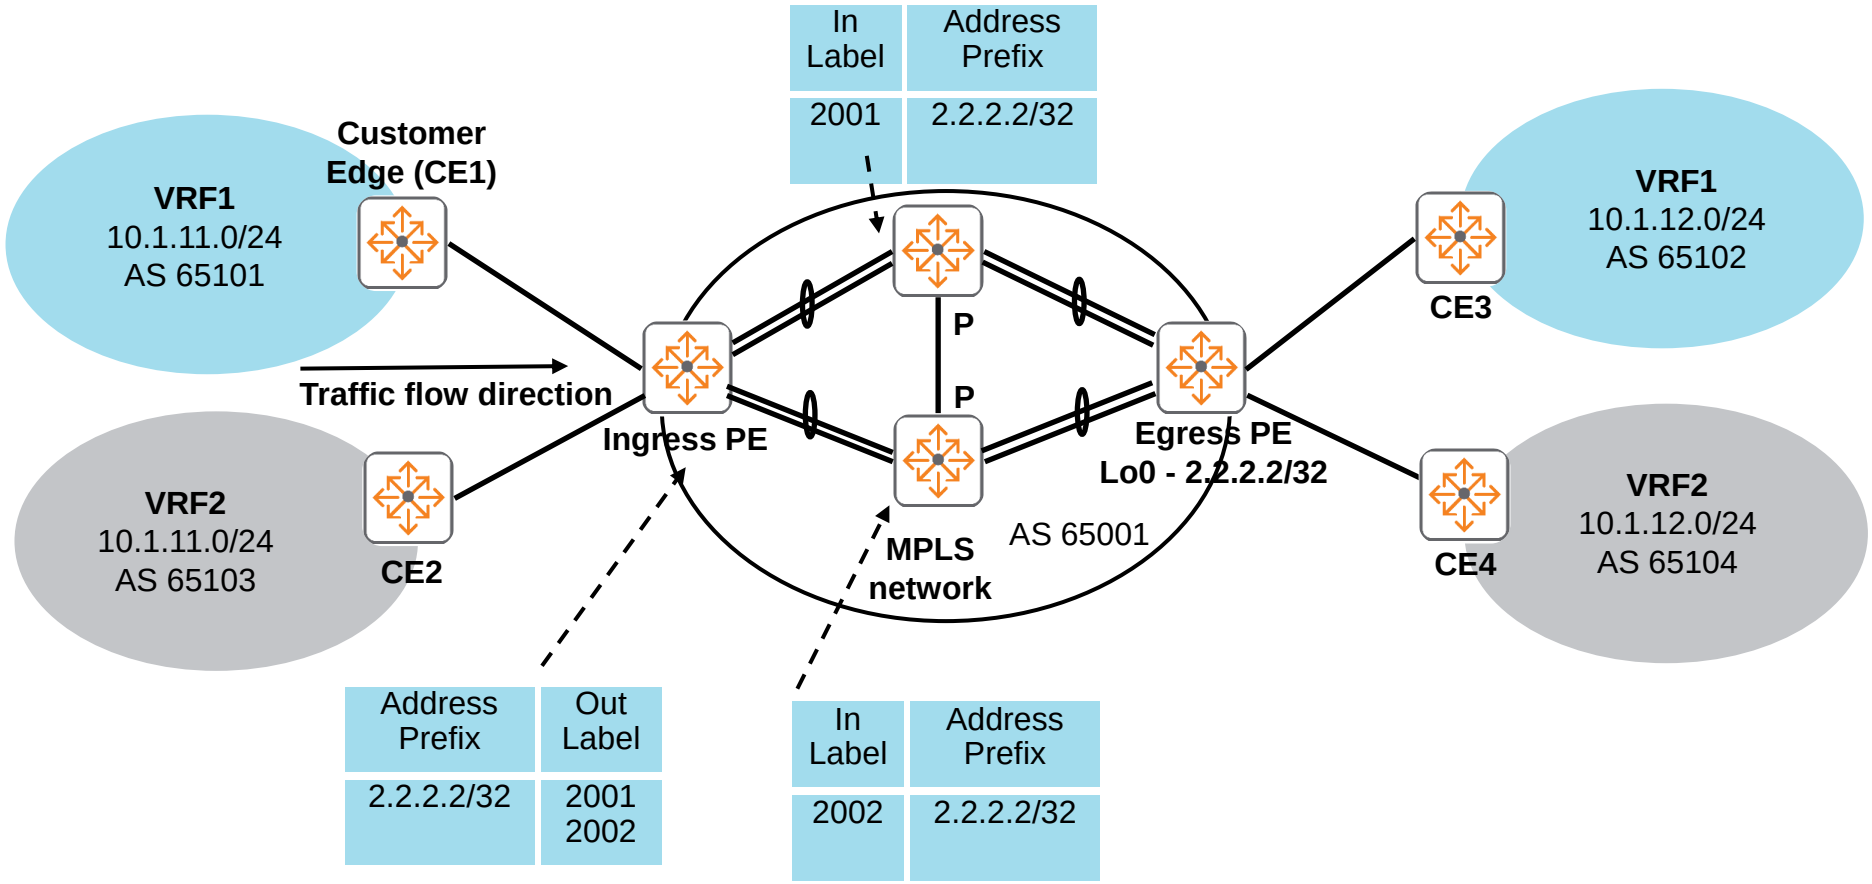

# Same CE AS#

– “allow-as-loop” on CE is required to support this

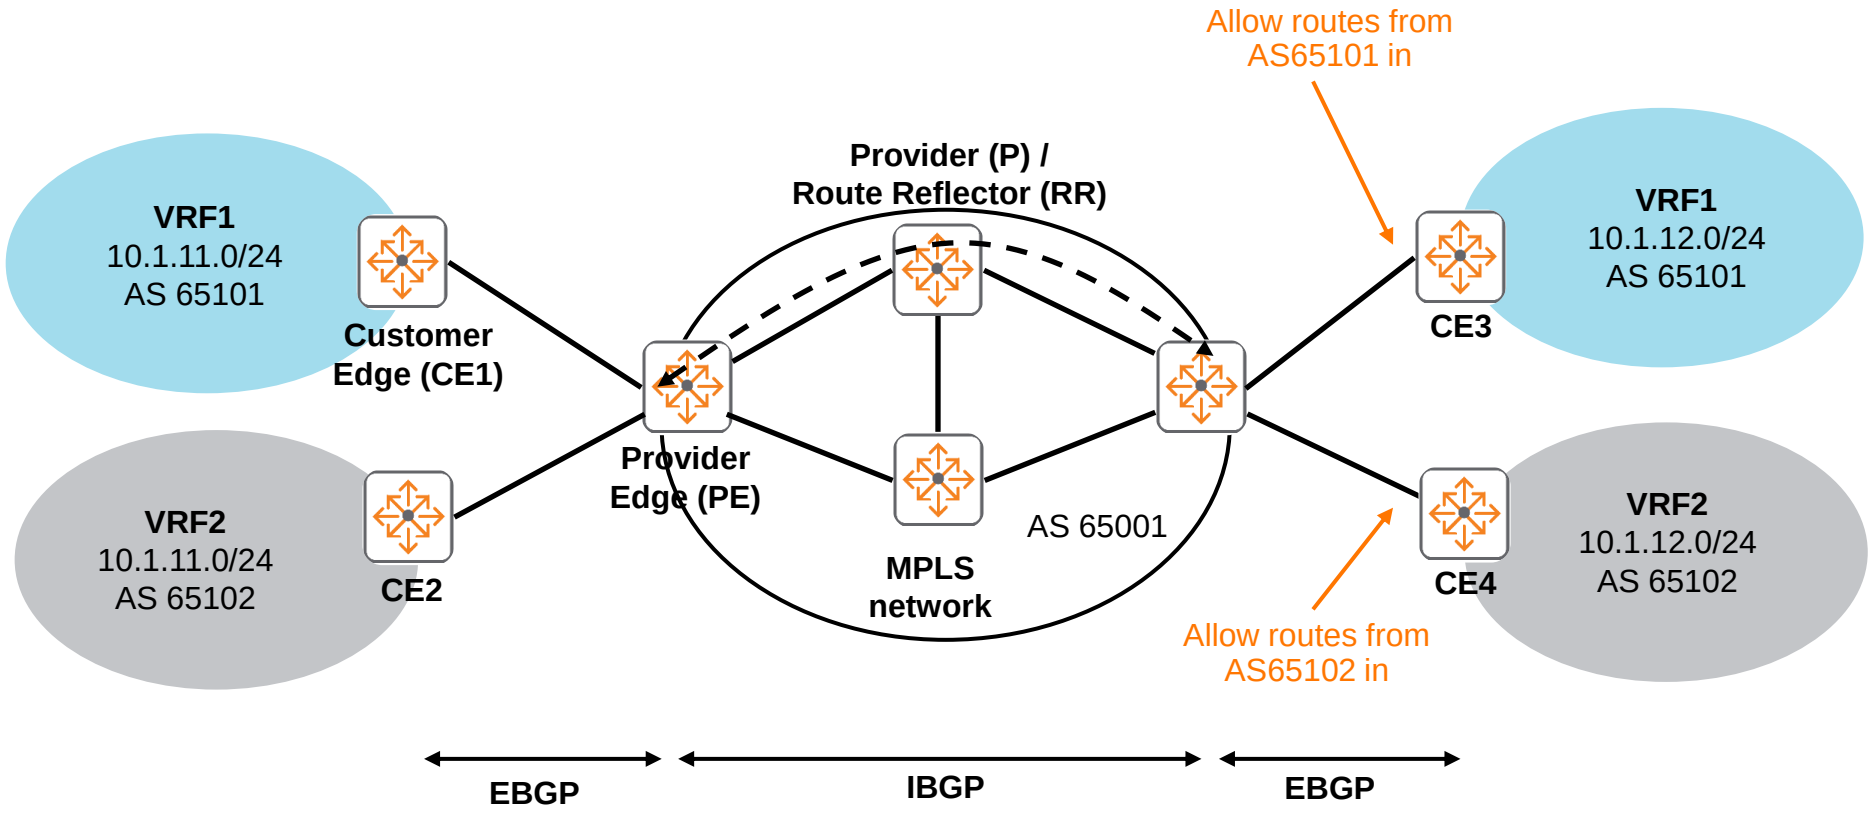

# 10.9 MPLS supported features

- MPLS L3 VPN
  - IPv4 unicast
- Static LSP
- OSPF/LDP between P/RR and PE
- MP-BGP VPNv4 RR on P
- MP-BGP VPNv4 RR client on PE
- QoS uniform mode
- VRFs (isolated or shared VRF with route leaking)
- PE-CE connectivity
  - EBGP or static routing (e.g static default route on CE)
  - PE redistribute connected routes (e.g. PE-CE subnet)
- 3rd party MPLS VPN interop
- MPLS ECMP (targeted for 10.9.10 CPE)
- Active/active egress PEs (targeted for 10.9.10 CPE)
- LSP ping and traceroute (targeted for 10.9.1000 CPE)

# 10.9 MPLS caveats

- Only supported on 8360
- VSX PEs are not supported
  - Use dual standalone PEs for redundancy
- No multicast between MPLS PEs
  - PIM multicast over GRE tunnels is also not supported
  - Use multicast over VXLAN tunnels between CEs if multicast is required as a workaround
- No IPv6
- No FRR/GR
- No entropy label support
- No implicit null/PHP
- No MPLS SNMP or REST
- VXLAN EVPN cannot be enabled when MPLS L3 VRF is enabled
- Sub-interfaces or split-ports, e.g. 1/1/1.10 or 1/1/1:1 are not supported

# Configuration

# MPLS L3 VPN Configuration – P/RR example

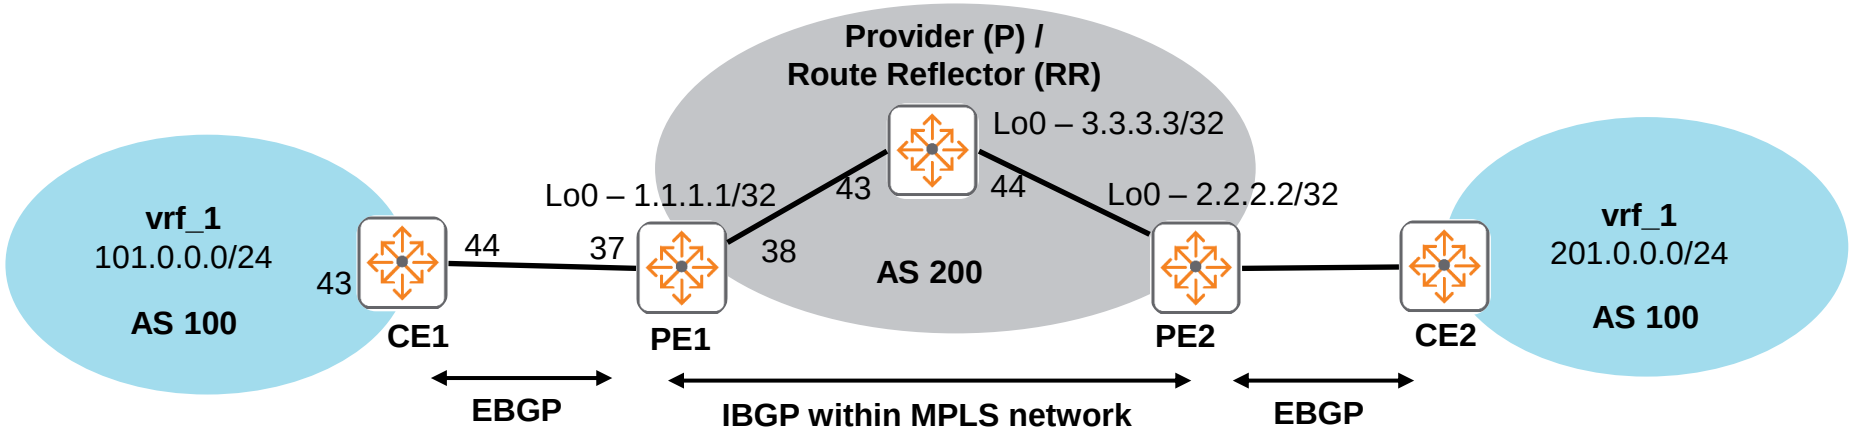


```
hostname P
!
interface 1/1/43
    no shutdown
    description PE1
    ip address 30.0.0.2/24
    ip ospf 1 area 0.0.0.0
    mpls enable
    mpls ldp enable
interface 1/1/44
    no shutdown
    description PE2
    ip address 30.1.0.1/24
    ip ospf 1 area 0.0.0.0
    mpls enable
    mpls ldp enable
interface loopback 0
    ip address 3.3.3.3/32
    ip ospf 1 area 0.0.0.0
```

```
router ospf 1
    area 0.0.0.0
router bgp 200
    neighbor 1.1.1.1 remote-as 200
    neighbor 1.1.1.1 update-source loopback 0
    neighbor 2.2.2.2 remote-as 200
    neighbor 2.2.2.2 update-source loopback 0
    address-family vpnv4 unicast
        neighbor 1.1.1.1 activate
        neighbor 1.1.1.1 send-community extended
        neighbor 1.1.1.1 route-reflector-client
        neighbor 2.2.2.2 activate
        neighbor 2.2.2.2 send-community extended
        neighbor 2.2.2.2 route-reflector-client
```

```
mpls
    enable
    label-protocol ldp
        enable
        router-id loopback0
```

# MPLS L3 VPN Configuration – PE example

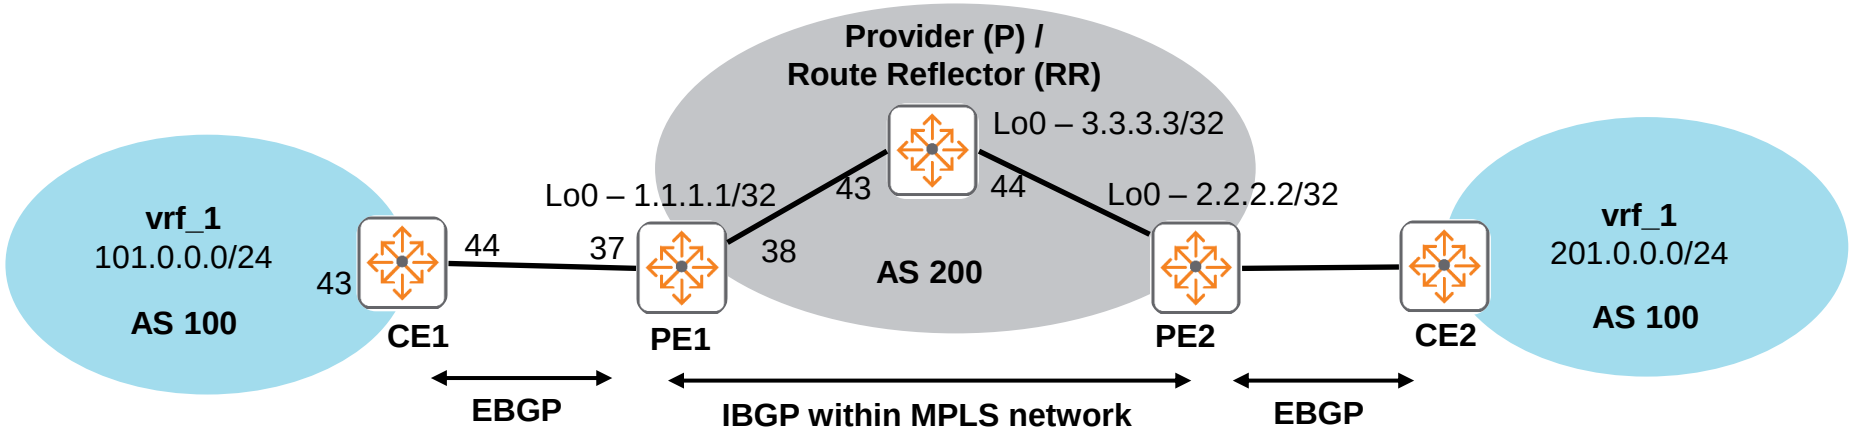


```
hostname PE1
!
vrf vrf_1
    rd 1.1.1.1:1
    l3vpn-only
    address-family ipv4 unicast
        route-target export 100:1
        route-target import 100:1
interface 1/1/37
    no shutdown
    vrf attach vrf_1
    description CE1
    ip address 20.1.0.2/24
interface 1/1/38
    no shutdown
    description P1
    ip address 30.0.0.1/24
    ip ospf 1 area 0.0.0.0
    mpls enable
    mpls ldp enable
```

```
interface loopback 0
    ip address 1.1.1.1/32
    ip ospf 1 area 0.0.0.0
!
router ospf 1
    area 0.0.0.0
router bgp 200
    neighbor 3.3.3.3 remote-as 200
    neighbor 3.3.3.3 update-source loopback 0
    address-family vpnv4 unicast
        neighbor 3.3.3.3 activate
        neighbor 3.3.3.3 send-community extended
    vrf vrf_1
        neighbor 20.1.0.1 remote-as 100
        address-family ipv4 unicast
            neighbor 20.1.0.1 activate
            neighbor 20.1.0.1 send-community both
            redistribute connected
```

```
mpls
    enable
    label-protocol ldp
        enable
        router-id loopback0
```

**Note:**

- **l3vpn-only under VRF context is required for VPNv4 address family to be active**

# MPLS L3 VPN Configuration – CE example

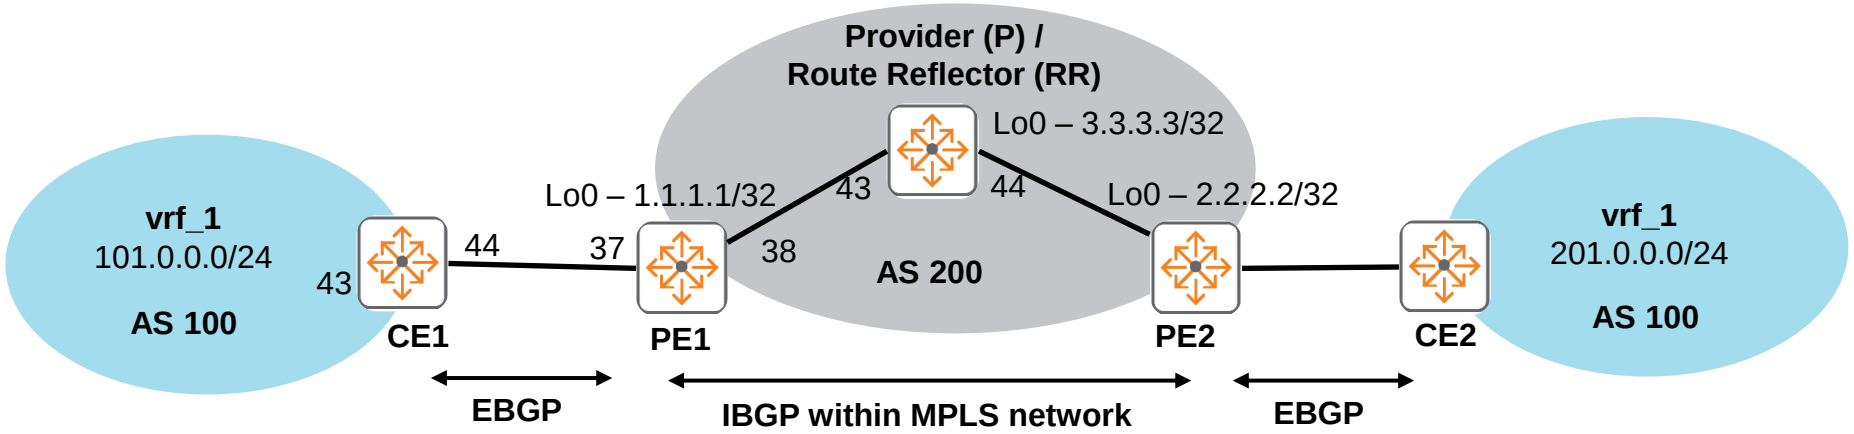


```
hostname CE1
!
interface 1/1/43
    no shutdown
    description LAN
    ip address 20.0.0.2/24
interface 1/1/44
    no shutdown
    description PE1
    ip address 20.1.0.1/24
```

```
router bgp 100
    neighbor 20.1.0.2 remote-as 200
    address-family ipv4 unicast
        neighbor 20.1.0.2 activate
        neighbor 20.1.0.2 allowas-in 100
        neighbor 20.1.0.2 send-community both
        neighbor 20.1.0.2 soft-reconfiguration inbound
```

# Best Practices

# Best Practices

- Use redundant VPNv4 RRs to avoid full mesh IBGP peers
- Use peer groups to simplify configs on both P/RR and PE

```
router bgp 65001
 neighbor RR-client peer-group
    neighbor RR-client remote-as 65001
    neighbor RR-client description all RR clients
    neighbor RR-client update-source loopback 0
    neighbor 1.1.1.1 peer-group RR-client
    neighbor 2.2.2.2 peer-group RR-client
    address-family vpnv4 unicast
        neighbor RR-client route-reflector-client
        neighbor RR-client send-community both
        neighbor 1.1.1.1 activate
        neighbor 2.2.2.2 activate
```

- Utilize EBGP between CE and PE for dynamic routing, use unique CE AS# if possible
- Remember
  - To add “l3vpn-only” under VRF context
  - Implicit null is not supported, explicit null has to be configured on 3rd party egress PE towards AOS-CX P
  - If egress PE and P are AOS-CX switches, changing it to explicit null is not required as explicit null is automatically sent

# Troubleshooting

# MPLS Troubleshooting

- Have a topology diagram ready
- Ensure IPs, interface, AS details are included
- Check physical cabling and generate “show tech” when opening a TAC case
- Check network: show LLDP neighbor, ensure unicast network works using ping and traceroute between loopbacks and interfaces, fix any issues found

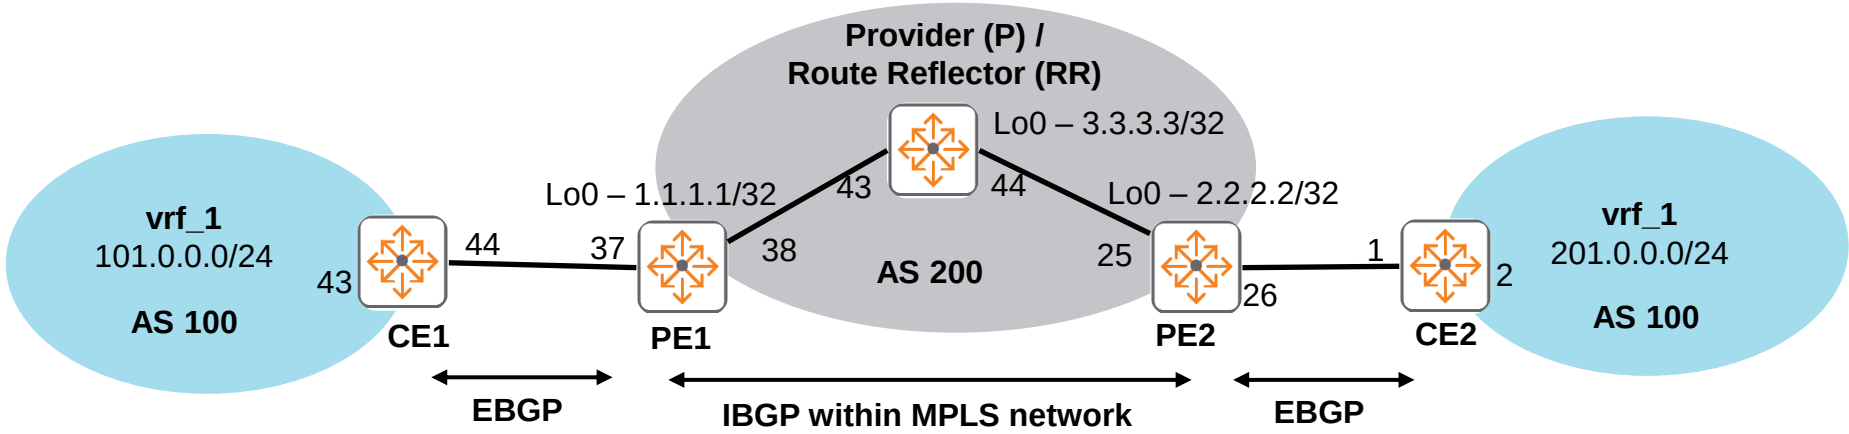


- Mirror traffic and packet capture if required

```
mirror session 1
    enable
    destination interface 1/1/40
    source interface 1/1/51 both
```

- If requested by TAC engineer

```
diagnostics
diag-dump mpls basic

debug mpls all
sh debug buffer
no debug all
```

# MPLS Troubleshooting

– Recommended troubleshooting flow

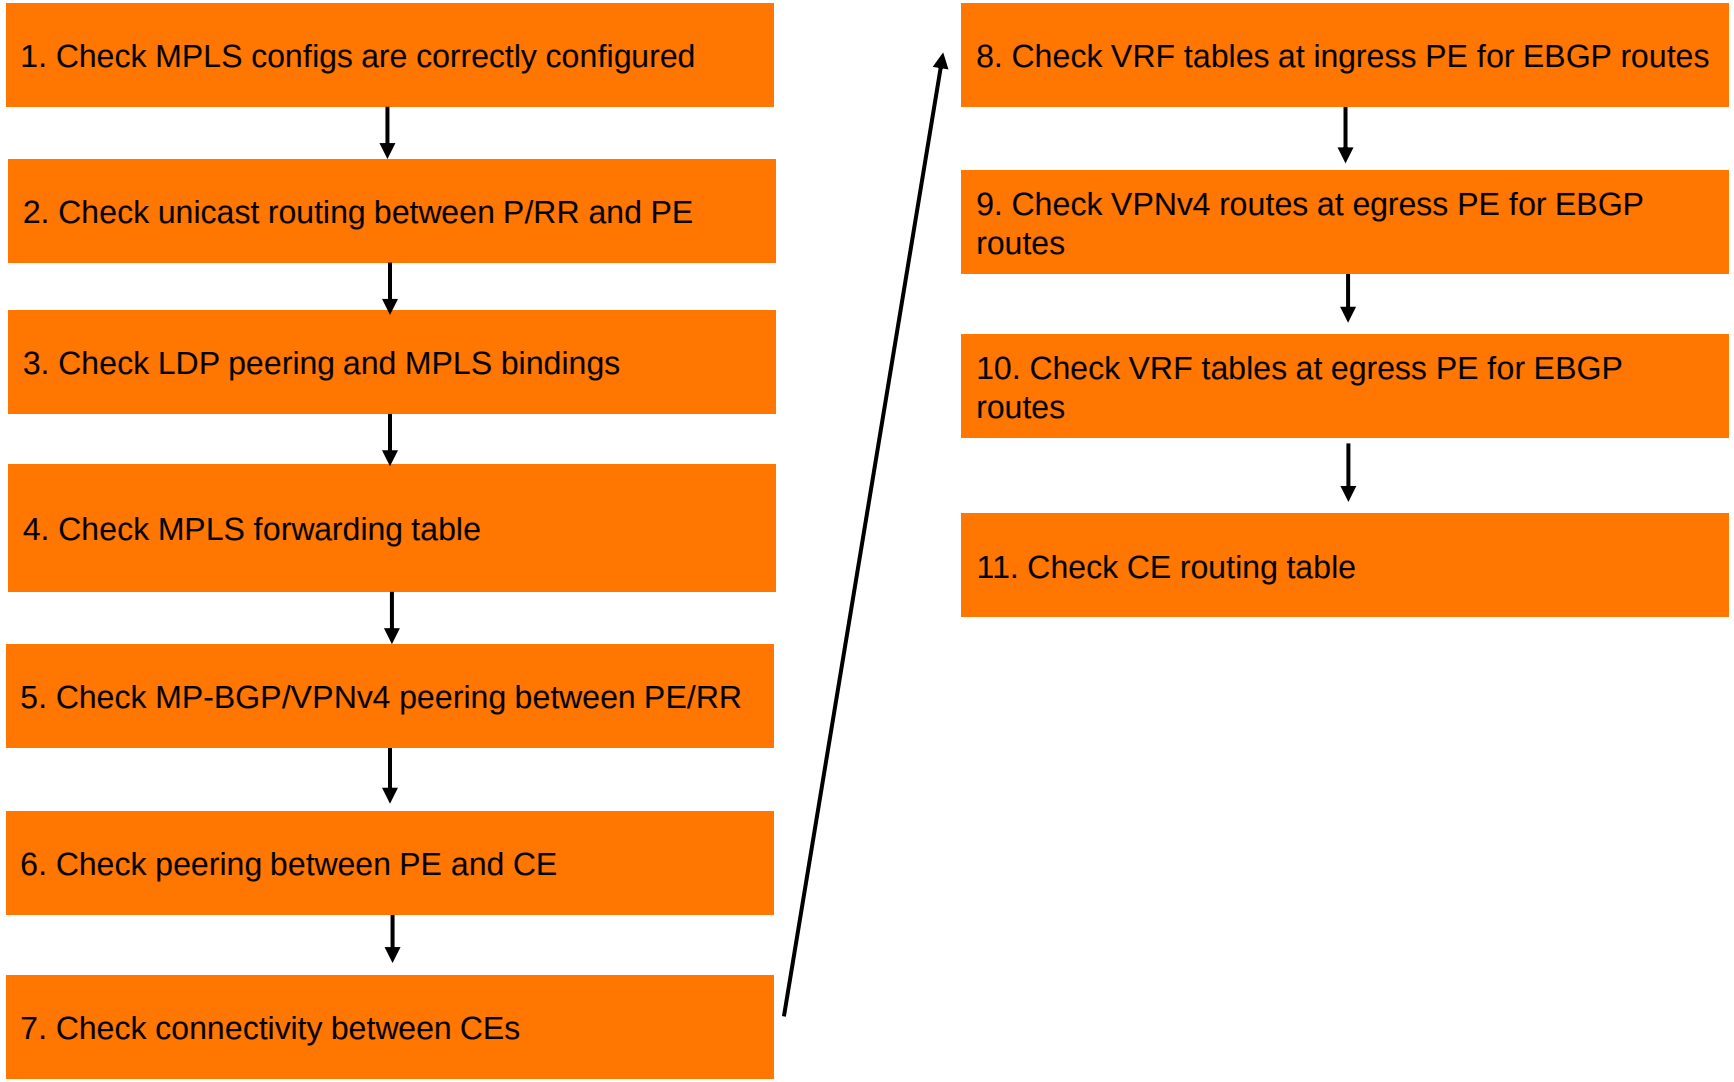

# 1. Check MPLS configs are correctly configured

– Refer to config section for CE, PE and P sample configs

# 2. Check unicast routing between P/RR and PE

- On PE, verify loopbacks from remote P and PE are learnt
- If not seen, verify OSPF is enabled and advertising loopbacks on remote P and PE

```
PE1# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix              Nexthop          Interface     VRF(egress)   Origin/  Distance/    Age
                                                                   Type     Metric
-------------------------------------------------------------------------------------------
1.1.1.1/32          -                loopback0     -               L         [0/0]        -
2.2.2.2/32          30.0.0.2         1/1/38        -               O         [110/20]     00h:12m:32s
3.3.3.3/32          30.0.0.2         1/1/38        -               O         [110/10]     00h:12m:52s
30.0.0.0/24         -                1/1/38        -               C         [0/0]        -
30.0.0.1/32         -                1/1/38        -               L         [0/0]        -
30.1.0.0/24         30.0.0.2         1/1/38        -               O         [110/20]     00h:12m:52s

Total Route Count : 6
```

# 3. Check LDP peering and MPLS bindings

– Verify LDP neighbor is up

```
PE1#   show mpls ldp neighbor
Local LDP Identifier: 1.1.1.1:0, Peer LDP
Identifier: 3.3.3.3:0
    Graceful Restart: No
    Session Holdtime: 40 sec
    Up time: 00:13:49
    State: operational
    LDP Discovery Sources: 1/1/38
    Addresses bound to this peer:
         3.3.3.3 30.0.0.2 30.1.0.1
```

– Verify labels and MPLS bindings for remote P and PE exist

```
PE1# show mpls ldp bindings
30.1.0.1/32
        No local binding
        remote binding: lsr:3.3.3.3:0 label: exp-null
1.1.1.1/32
        local binding: label: exp-null
        remote binding: lsr:3.3.3.3:0 label:16
30.0.0.1/32
        local binding: label: exp-null
        No remote binding
30.0.0.2/32
        No local binding
        remote binding: lsr:3.3.3.3:0 label: exp-null
3.3.3.3/32
        local binding: label: 17
        remote binding: lsr:3.3.3.3:0 label: exp-null
2.2.2.2/32
        local binding: label: 18
        remote binding: lsr:3.3.3.3:0 label:17
```

– If not seen, verify remote PE has LDP established

# 4. Check MPLS forwarding table

– Verify nexthop + labels for remote VPNv4 and PE and routes exist

```
PE1# sh mpls forwarding

MPLS Bindings

Entry Bindings   : 4
Exit Bindings    : 3
Transit Bindings : 2

Entry Bindings:
Origin   Prefix              Ingress           Nexthop           Outgoing  Egress        Egress       Status
                              VRF               Address           Label     Interface     VRF
-------------------------------------------------------------------------------------------------------------
LDP      2.2.2.2/32           default           30.0.0.2          17        1/1/38        default      operational
LDP      3.3.3.3/32           default           30.0.0.2          exp-null  1/1/38        default      operational
BGP      20.2.0.0/24          vrf_1             2.2.2.2           18        1/1/38        default      operational
BGP      201.0.0.0/24         vrf_1             2.2.2.2           18        1/1/38        default      operational

Exit Bindings:

Origin   Prefix              Incoming  Service   Egress            Status
                             Label     Label     VRF
-------------------------------------------------------------------------
static   n/a                 exp-null  -         default           operational
BGP      n/a                 imp-null  16        vrf_1             operational
static   n/a                 7         -         default           operational
```

# 5. Check MP-BGP VPNv4 peering between PE/RR

- Verify MP-BGP VPNv4 peering is established
- If it's down, verify and fix network reachability between source and destination peering IPs

```
PE1# sh bgp vpnv4 un neighbors
Codes: ^ Inherited from peer-group

VRF : default

BGP Neighbor 3.3.3.3 (Internal)
    Description          :
    Peer-group           :

    Remote Router Id     : 3.3.3.3          Local Router Id     : 1.1.1.1
    Remote AS            : 200              Local AS            : 200
    Remote Port          : 43923            Local Port          : 179
    State                : Established      Admin Status        : Up
    Conn. Established    : 2                Conn. Dropped       : 1
    Passive              : No               Update-Source       : loopback0
```

# 6. Check peering between PE and CE

– On PE, verify EBGP peering to CE

```
PE1# show bgp vrf vrf_1 ipv4 unicast neighbors
Codes: ^ Inherited from peer-group

VRF : vrf_1

BGP Neighbor 20.1.0.1 (External)
    Description          :
    Peer-group           :

    Remote Router Id     : 20.1.0.1           Local Router Id     : 20.1.0.2
    Remote AS            : 100                Local AS            : 200
    Remote Port          : 51588              Local Port          : 179
    State                : Established        Admin Status        : Up
    Conn. Established    : 1                  Conn. Dropped       : 0
    Passive              : No                 Update-Source       :
    Cfg. Hold Time       : 180                Cfg. Keep Alive     : 60
    Neg. Hold Time       : 180                Neg. Keep Alive     : 60
    Up/Down Time         : 00h:20m:28s        Connect-Retry Time  : 120
    Local-AS Prepend     : No                 Alt. Local-AS       : 0
    BFD                  : Disabled
    Password             :
    Last Err Sent        : Cease
    Last SubErr Sent     : Administrative Shutdown
    Last Err Rcvd        : No Error
    Last SubErr Rcvd     : No Error
```

# 7. Check connectivity between PEs and between CEs

– Ping from CE to remote CE WAN and LAN to verify MPLS L3 VPN connectivity

```
CE1# ping 20.2.0.2
PING 20.2.0.2 (20.2.0.2) 100(128) bytes of data.
108 bytes from 20.2.0.2: icmp_seq=1 ttl=61 time=0.307 ms
108 bytes from 20.2.0.2: icmp_seq=2 ttl=61 time=0.328 ms
108 bytes from 20.2.0.2: icmp_seq=3 ttl=61 time=0.340 ms
108 bytes from 20.2.0.2: icmp_seq=4 ttl=61 time=0.335 ms
108 bytes from 20.2.0.2: icmp_seq=5 ttl=61 time=0.354 ms
```

# 8. Check VRF tables at ingress PE for EBGP routes

- On ingress PE, verify routes from both connected CE and remote PE are learnt
- If the expected route is not learned, check if CE and remote PE are learning and advertising them out

```
PE1# show ip rib vrf vrf_1

Displaying ipv4 routes in RIB

Origin Codes: R - RIP, O - OSPFv2, B - BGP
              C - connected, S - static, L - local
Type Codes:   E - External BGP, I - Internal BGP
              IA - OSPF inter area, ia - OSPF intra area
              E1 - OSPF external type 1, E2 - OSPF external type 2
              EV - BGP EVPN, V - BGP VPN
* indicates selected for forwarding

VRF: vrf_1

Prefix                Nexthop           Interface     VRF(egress)   Origin/  Distance/    Age
                                                                    Type     Metric
-------------------------------------------------------------------------------------------
*20.1.0.0/24          -                 1/1/37        -             C        [0/0]        -
*20.1.0.2/32          -                 1/1/37        -             L        [0/0]        -
*20.2.0.0/24          2.2.2.2           -             -             B/V      [200/0]      00h:24m:27s
*101.0.0.0/24         20.1.0.1          -             -             B/E      [20/0]       00h:14m:53s
*201.0.0.0/24         2.2.2.2           -             -             B/V      [200/0]      00h:14m:50s

Total Route Count : 5
```

Route from remote PE

Route from directly connected CE

# 8. Check VRF tables at ingress PE for EBGP routes

- On ingress PE, verify routes from both connected CE and remote PE are learnt
- If the expected route is not learned, check if CE and remote PE are learning and advertising them out

```
PE1# sh ip route vrf vrf_1

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: vrf_1

Prefix              Nexthop          Interface     VRF(egress)   Origin/  Distance/    Age
                                                                 Type     Metric
-------------------------------------------------------------------------------------------
20.1.0.0/24         -                1/1/37        -             C          [0/0]        -
20.1.0.2/32         -                1/1/37        -             L          [0/0]        -
20.2.0.0/24         2.2.2.2          1/1/38        default       B/V        [200/0]      00h:26m:51s
101.0.0.0/24        20.1.0.1         1/1/37        -             B/E        [20/0]       00h:11m:42s
201.0.0.0/24        2.2.2.2          1/1/38        default       B/V        [200/0]      00h:11m:42s

Total Route Count : 5
```

Route from remote PE

Route from directly connected CE

# 9. Check VPNv4 routes at egress PE for EBGP routes

– Verify CE LAN routes appear here

```
PE2# show bgp vpnv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 2.2.2.2

    Network            Nexthop          Metric     LocPrf     Weight Path

Route Distinguisher: 1:1                   (Label 16)
*>i 20.1.0.0/24        1.1.1.1          0          100        0       ?

Route Distinguisher: 1:2                   (Label 16)
*>  20.2.0.0/24        0.0.0.0          0          100        0       ?

Route Distinguisher: 1:1                   (Label 16)
*>i 101.0.0.0/24       1.1.1.1          0          100        0       100 i

Route Distinguisher: 1:2                   (Label 16)
*>  201.0.0.0/24       0.0.0.0          0          100        0       100 i
Total number of entries 4
```

# 10. Check VRF tables at egress PE for EBGP routes

– Verify CE LAN routes appear here

```
PE2# show ip rib vrf vrf_1

Displaying ipv4 routes in RIB

Origin Codes: R - RIP, O - OSPFv2, B - BGP
              C - connected, S - static, L - local
Type Codes:   E - External BGP, I - Internal BGP
              IA - OSPF inter area, ia - OSPF intra area
              E1 - OSPF external type 1, E2 - OSPF external type 2
              EV - BGP EVPN, V - BGP VPN
* indicates selected for forwarding

VRF: vrf_1

Prefix               Nexthop          Interface     VRF(egress)   Origin/  Distance/    Age
                                                                   Type     Metric
-------------------------------------------------------------------------------------------
*20.1.0.0/24         1.1.1.1          -             -             B/V      [200/0]      00h:23m:48s
*20.2.0.0/24         -                1/1/26        -             C        [0/0]        -
*20.2.0.1/32         -                1/1/26        -             L        [0/0]        -
*101.0.0.0/24        1.1.1.1          -             -             B/V      [200/0]      00h:14m:14s
*201.0.0.0/24        20.2.0.2         -             -             B/E      [20/0]       00h:14m:12s

Total Route Count : 5
```

# 10. Check VRF tables at egress PE for EBGP routes

– Verify CE LAN routes appear here

```
PE2# show ip route vrf vrf_1

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: vrf_1

Prefix               Nexthop          Interface     VRF(egress)   Origin/  Distance/    Age
                                                                   Type     Metric
-------------------------------------------------------------------------------------------
20.1.0.0/24          1.1.1.1          1/1/25        default         B/V      [200/0]      00h:21m:20s
20.2.0.0/24          -                1/1/26        -               C        [0/0]        -
20.2.0.1/32          -                1/1/26        -               L        [0/0]        -
101.0.0.0/24         1.1.1.1          1/1/25        default         B/V      [200/0]      00h:11m:46s
201.0.0.0/24         20.2.0.2         1/1/26        -               B/E      [20/0]       00h:11m:44s

Total Route Count : 5
```

# 11. Check CE routing table

- Check routes on CE
- Verify remote CE routes are learnt

```
CE1# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix                Nexthop           Interface     VRF(egress)   Origin/  Distance/    Age
                                                                    Type     Metric
-------------------------------------------------------------------------------------------
20.0.0.0/24           -                 1/1/43        -               C        [0/0]        -
20.0.0.2/32           -                 1/1/43        -               L        [0/0]        -
20.1.0.0/24           -                 1/1/44        -               C        [0/0]        -
20.1.0.1/32           -                 1/1/44        -               L        [0/0]        -
20.2.0.0/24           20.1.0.2          1/1/44        -               B/E      [20/0]       00h:13m:15s
101.0.0.0/24          20.0.0.1          1/1/43        -               B/I      [200/0]      00h:01m:10s
101.0.1.0/24          20.0.0.1          1/1/43        -               B/I      [200/0]      00h:01m:10s
101.0.2.0/24          20.0.0.1          1/1/43        -               B/I      [200/0]      00h:01m:10s
101.0.3.0/24          20.0.0.1          1/1/43        -               B/I      [200/0]      00h:01m:10s
101.0.4.0/24          20.0.0.1          1/1/43        -               B/I      [200/0]      00h:01m:12s
201.0.0.0/24          20.1.0.2          1/1/44        -               B/E      [20/0]       00h:01m:12s
201.0.1.0/24          20.1.0.2          1/1/44        -               B/E      [20/0]       00h:01m:12s
201.0.2.0/24          20.1.0.2          1/1/44        -               B/E      [20/0]       00h:01m:12s
201.0.3.0/24          20.1.0.2          1/1/44        -               B/E      [20/0]       00h:01m:12s
201.0.4.0/24          20.1.0.2          1/1/44        -               B/E      [20/0]       00h:01m:12s

Total Route Count : 15
```

Remote CE WAN subnet

Remote CE LAN subnet

# Demo

# AOS-CX 10.09 MPLS Demonstrations

MPLS L3 VPN
- OSPF
- MG-BGP (iBGP)
- MPLS LDP
- CE-PE eBGP

– 3 NODE LSR

– 5 NODE LSR

MPLS QoS

Multicast deployment

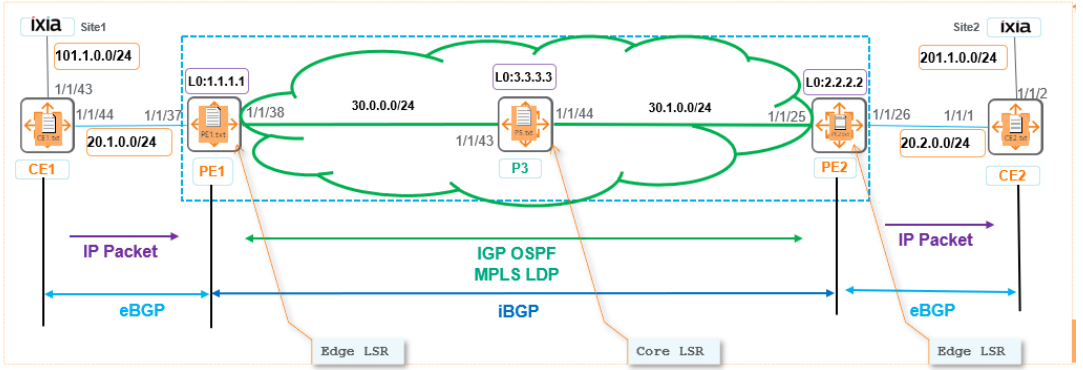


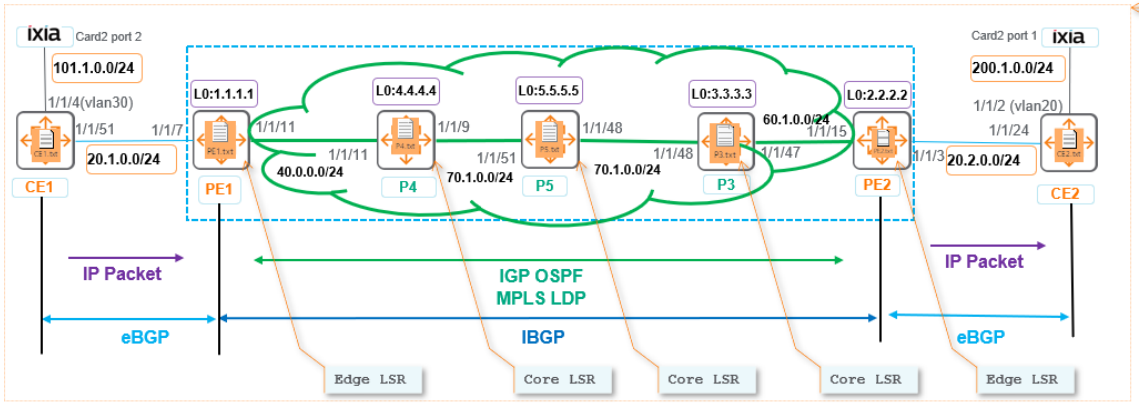

# MPLS Caveats (not supported) as part of 10.09

Only supported on 8360

VSX PEs are not supported

Use dual standalone PEs for redundancy

No multicast between MPLS PEs

- PIM multicast over GRE tunnels is also not supported
- Use multicast over VXLAN tunnels between CEs if multicast is required as a workaround

No IPv6

No Fast ReRoute (FRR)/ Graceful Restart (GR)

No entropy label support

No implicit null/PHP

No MPLS SNMP (REST is supported)

VXLAN EVPN cannot be enabled when L3 MPLS VRF is enabled (l3vpn-only is must for l3vpn vrf)

Sub-interfaces or split-ports, e.g. 1/1/1.10 or 1/1/1:1 are not supported

# AOS-CX 10.09 MPLS L3VPN (with RR)

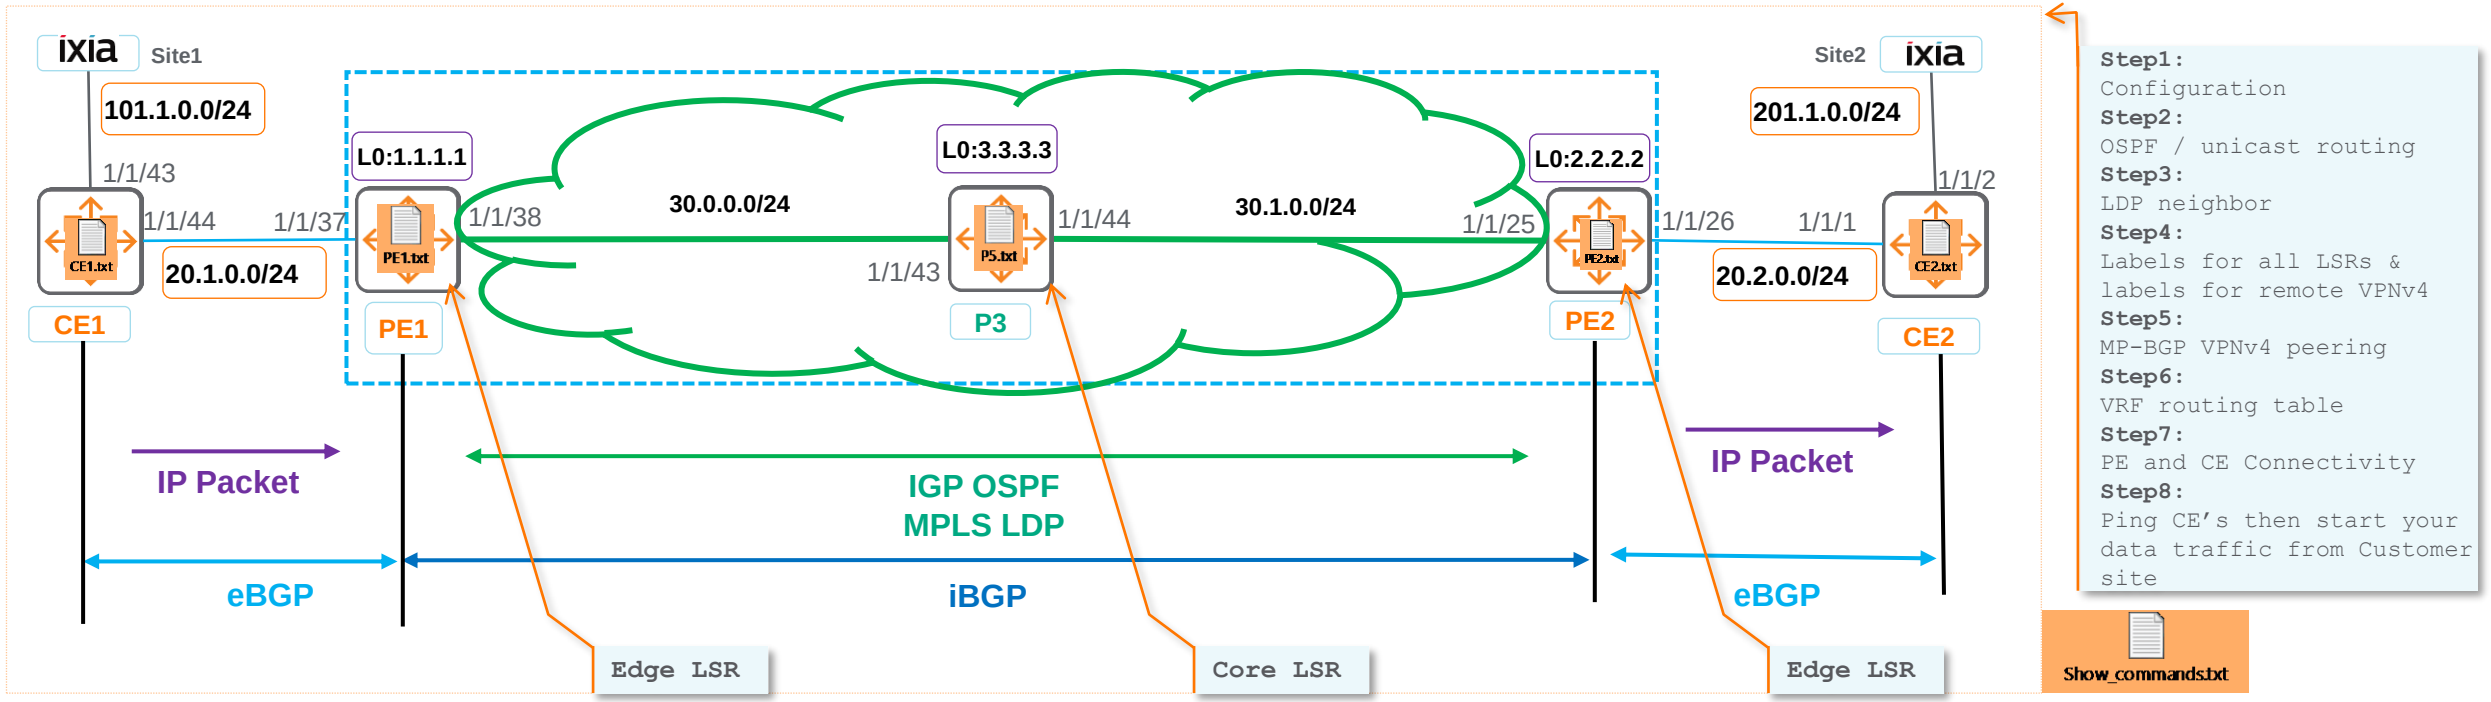

# AOS-CX 10.09 MPLS L3VPN (without RR)

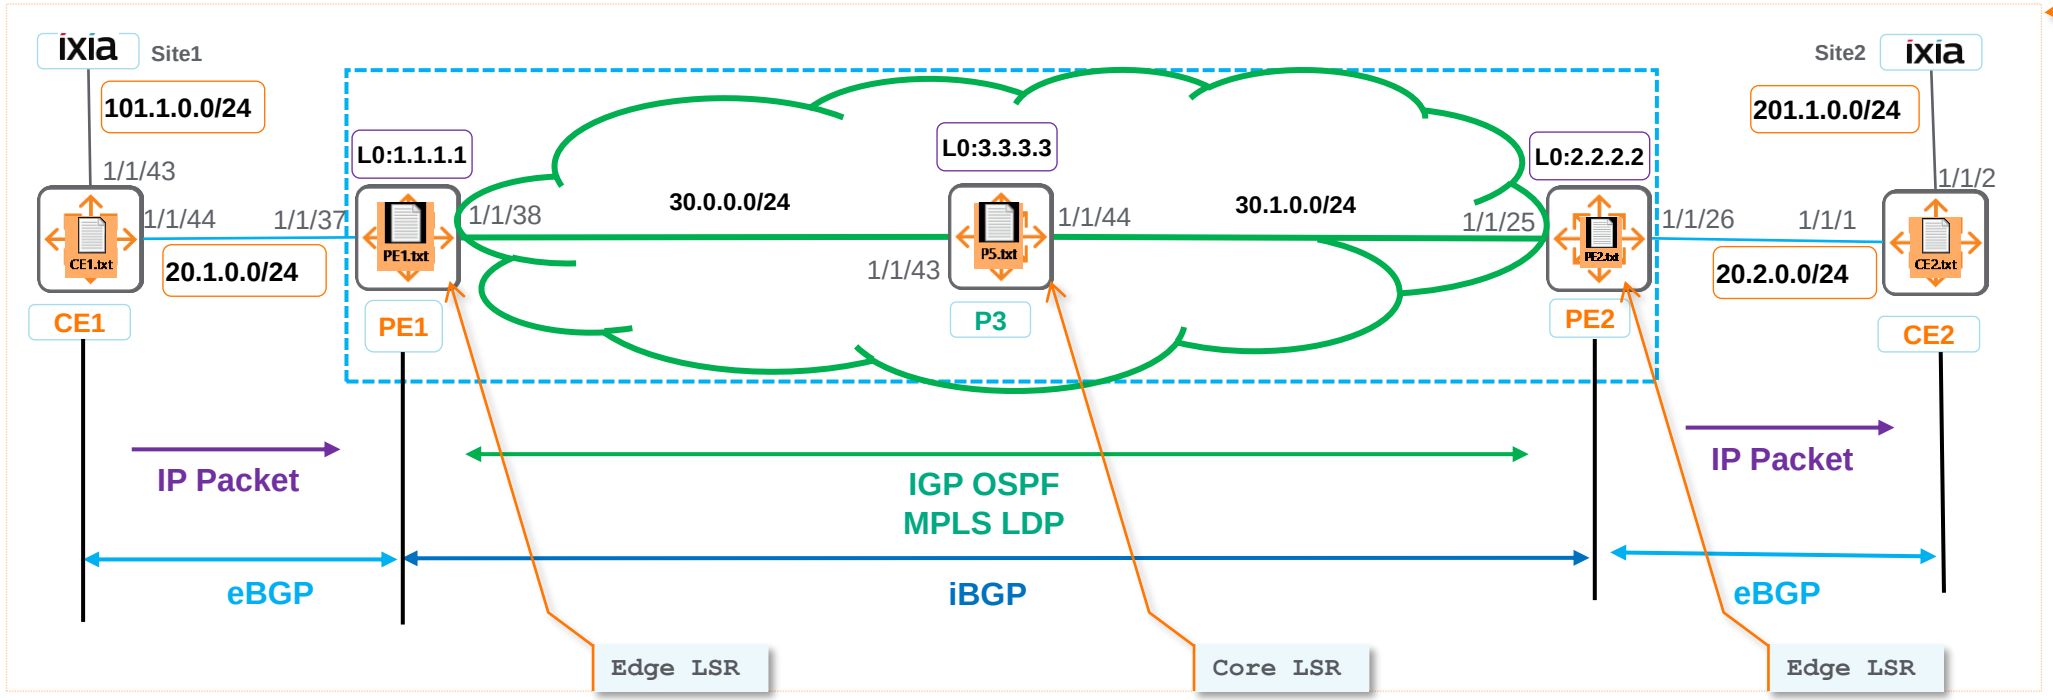


**Step1:**
Configuration
**Step2:**
OSPF / unicast routing
**Step3:**
LDP neighbor
**Step4:**
Labels for all LSRs & labels for remote VPNv4
**Step5:**
MP-BGP VPNv4 peering
**Step6:**
VRF routing table
**Step7:**
PE and CE Connectivity
**Step8:**
Ping CE's then start your data traffic from Customer site

# 10.09 AOS-CX MPLS L3VPN (without RR)

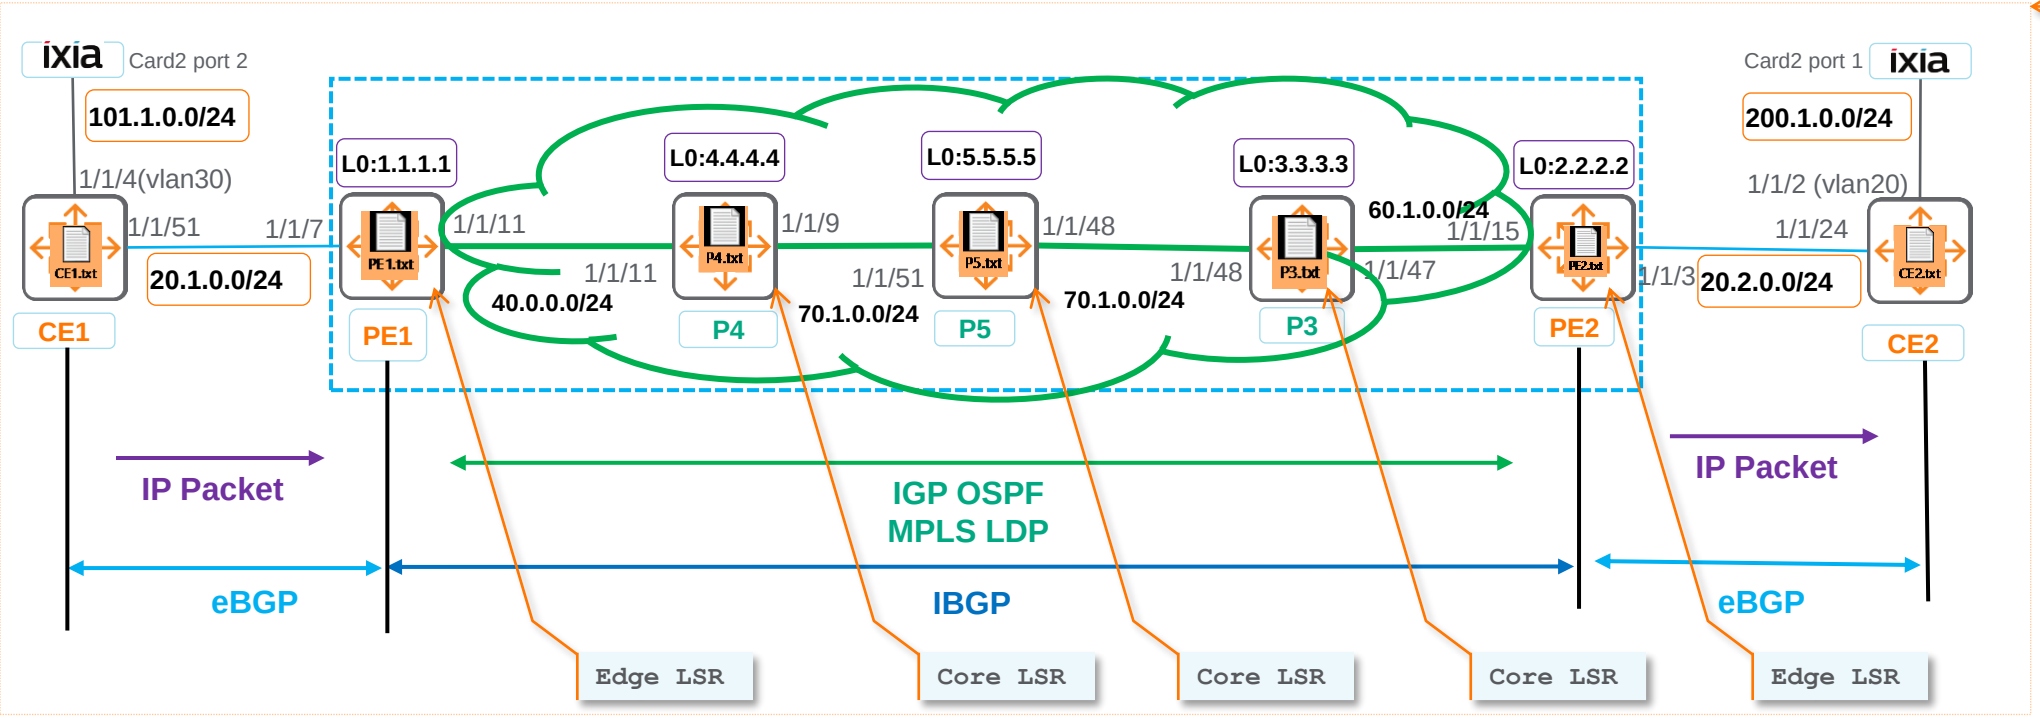


**Step1:**
Configuration
**Step2:**
OSPF / unicast routing
**Step3:**
LDP neighbor
**Step4:**
Labels for all LSRs & labels for remote VPNv4
**Step5:**
MP-BGP VPNv4 peering
**Step6:**
VRF routing table
**Step7:**
PE and CE Connectivity
**Step8:**
Ping CE's

Show_commands_5_node_LSR.txt

# 10.09 AOS-CX MPLS L3VPN (with RR)

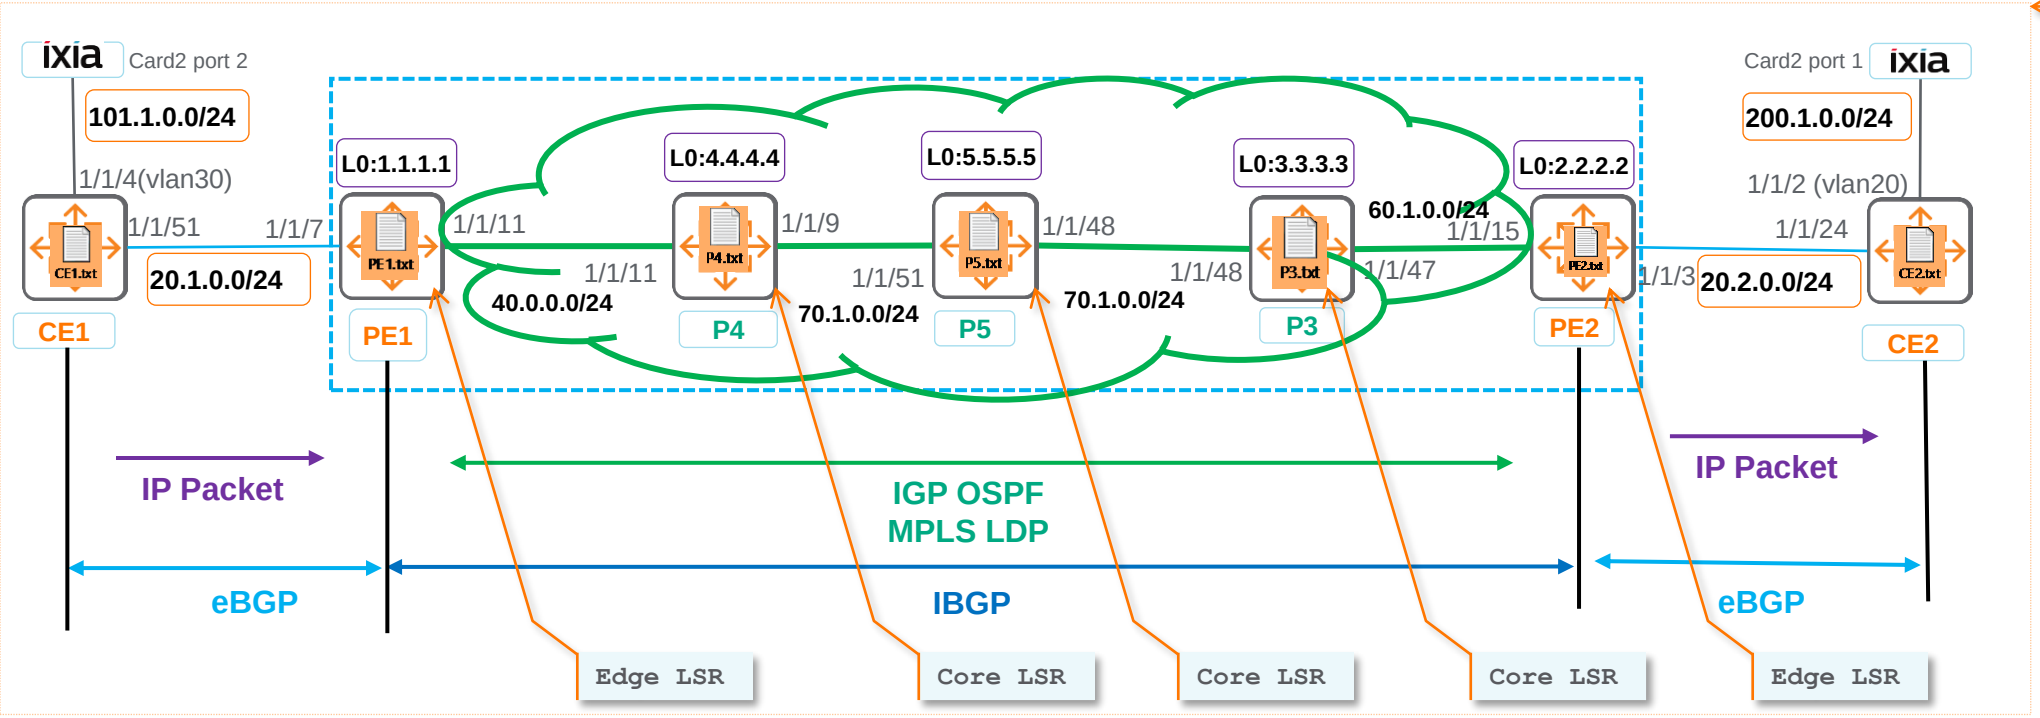


**Step1:**
Configuration
**Step2:**
OSPF / unicast routing
**Step3:**
LDP neighbor
**Step4:**
Labels for all LSRs & labels for remote VPNv4
**Step5:**
MP-BGP VPNv4 peering
**Step6:**
VRF routing table
**Step7:**
PE and CE Connectivity
**Step8:**
Ping CE's

Show_commands_5_node_LSR.txt

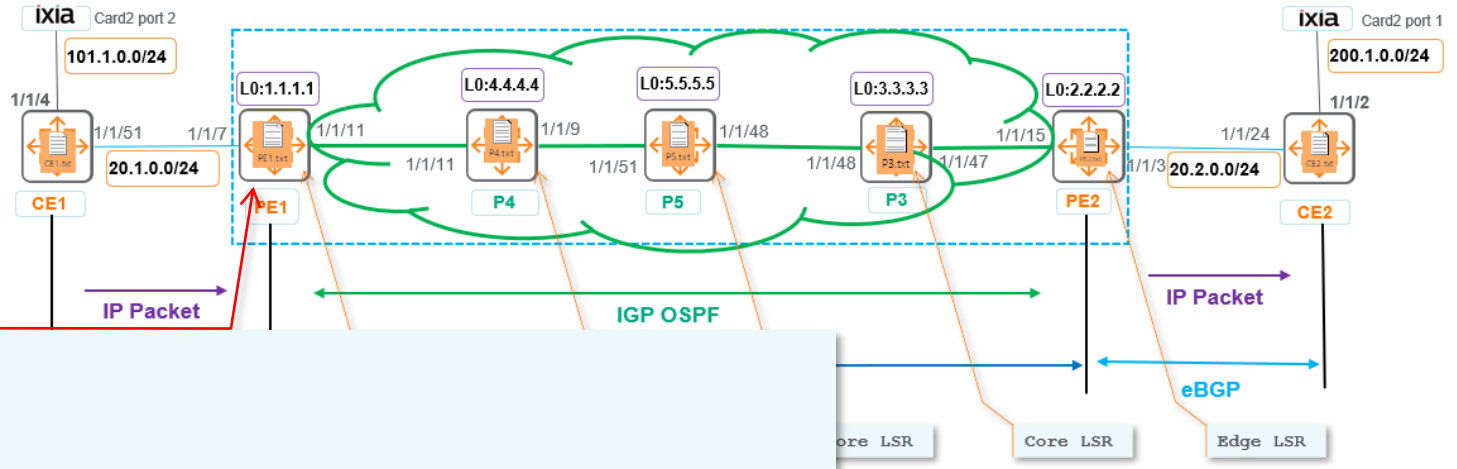


```
PE1# show mpls forwarding
MPLS Bindings
Entry Bindings    : 11
Exit Bindings     : 3
Transit Bindings  : 4
PHP Mode          : Explicit-Null
QoS Mode          : Uniform
TTL Propagation   : Uniform
Entry Bindings:
Origin   Prefix              Ingress          Nexthop            Outgoing  Egress       Egress           Status
                             VRF              Address            Label     Interface    VRF
-------------------------------------------------------------------------------------------------------------------
LDP      2.2.2.2/32          default          40.0.0.2           22        1/1/11       default          operational
LDP      3.3.3.3/32          default          40.0.0.2           23        1/1/11       default          operational
LDP      4.4.4.4/32          default          40.0.0.2           exp-null  1/1/11       default          operational
LDP      5.5.5.5/32          default          40.0.0.2           24        1/1/11       default          operational
BGP      10.1.1.0/24         vrf_1            2.2.2.2            16        1/1/11       default          operational
BGP      20.2.0.0/24         vrf_1            2.2.2.2            16        1/1/11       default          operational
BGP      200.1.0.0/24        vrf_1            2.2.2.2            16        1/1/11       default          operational
BGP      200.1.1.0/24        vrf_1            2.2.2.2            16        1/1/11       default          operational
BGP      200.1.2.0/24        vrf_1            2.2.2.2            16        1/1/11       default          operational
BGP      200.1.3.0/24        vrf_1            2.2.2.2            16        1/1/11       default          operational
BGP      200.1.4.0/24        vrf_1            2.2.2.2            16        1/1/11       default          operational
Exit Bindings:
Origin   Prefix              Incoming  Service   Egress           Status
                             Label     Label     VRF
-------------------------------------------------------------------------
static   n/a                 exp-null  -         default          operational
BGP      n/a                 imp-null  47        vrf_1            operational
static   n/a                 7         -         default          operational
Transit Bindings:
Origin   Prefix              Incoming  Egress       Egress          Nexthop            Outgoing  Status
                             Label     Interface    VRF             Address            Label
-------------------------------------------------------------------------------------------------------------
LDP      4.4.4.4/32          48        1/1/11       default         40.0.0.2           exp-null  operational
LDP      2.2.2.2/32          49        1/1/11       default         40.0.0.2           22        operational
LDP      3.3.3.3/32          50        1/1/11       default         40.0.0.2           23        operational
LDP      5.5.5.5/32          51        1/1/11       default         40.0.0.2           24        operational
PE1#
```

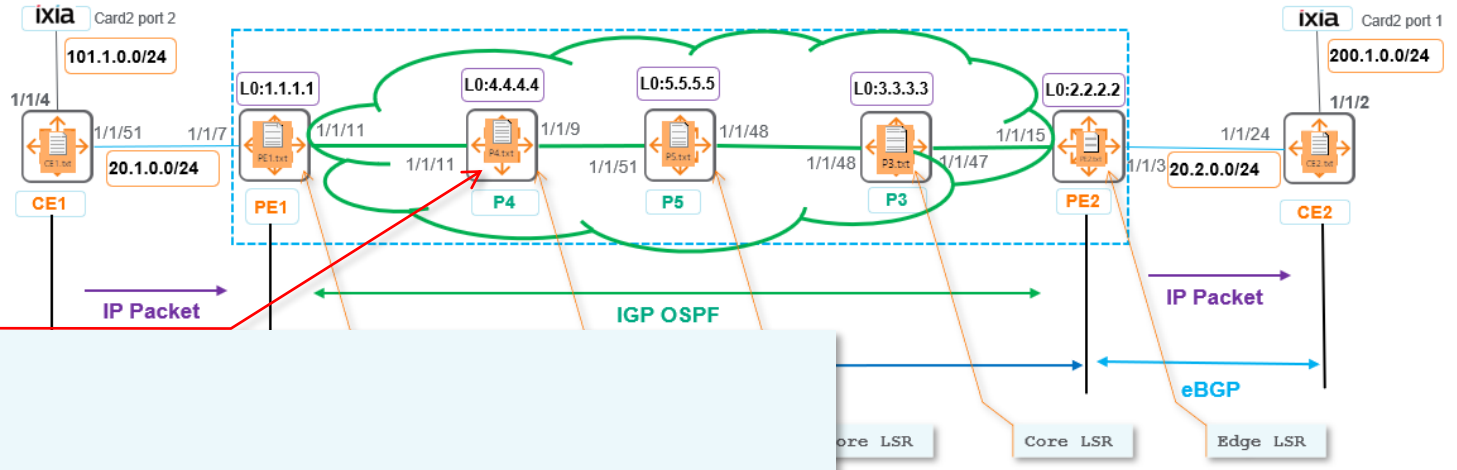


```
P4# show mpls forwarding
MPLS Bindings
Entry Bindings    : 4
Exit Bindings     : 2
Transit Bindings : 4
PHP Mode          : Explicit-Null
QoS Mode          : Uniform
TTL Propagation   : Uniform

Entry Bindings:

Origin   Prefix             Ingress          Nexthop            Outgoing  Egress       Egress           Status
                             VRF              Address            Label     Interface    VRF
-------------------------------------------------------------------------------------------------------------
LDP      1.1.1.1/32          default          40.0.0.1           exp-null  1/1/11       default          operational
LDP      2.2.2.2/32          default          70.1.0.2           20        1/1/9        default          operational
LDP      3.3.3.3/32          default          70.1.0.2           21        1/1/9        default          operational
LDP      5.5.5.5/32          default          70.1.0.2           exp-null  1/1/9        default          operational

Exit Bindings:

Origin   Prefix             Incoming  Service   Egress           Status
                            Label     Label     VRF
-------------------------------------------------------------------------
static   n/a                exp-null  -         default          operational
static   n/a                7         -         default          operational

Transit Bindings:

Origin   Prefix             Incoming  Egress       Egress           Nexthop            Outgoing  Status
                            Label     Interface    VRF              Address            Label
-------------------------------------------------------------------------------------------------------
LDP      2.2.2.2/32         22        1/1/9        default          70.1.0.2           20        operational
LDP      3.3.3.3/32         23        1/1/9        default          70.1.0.2           21        operational
LDP      5.5.5.5/32         24        1/1/9        default          70.1.0.2           exp-null  operational
LDP      1.1.1.1/32         30        1/1/11       default          40.0.0.1           exp-null  operational

P4#
```

Show_P4.txt

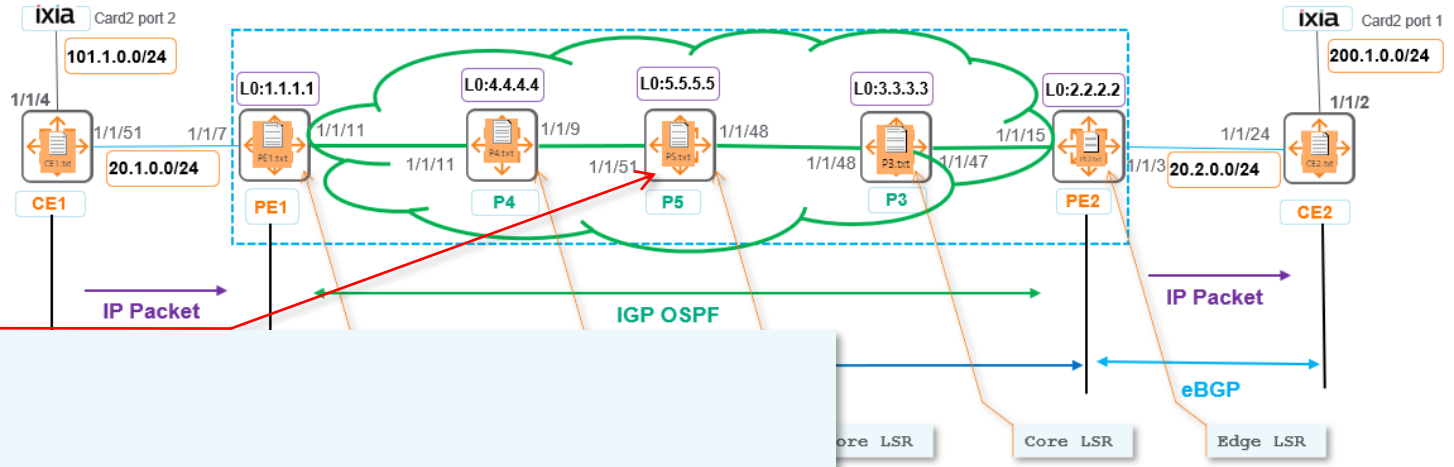


```
P5# show mpls forwarding

MPLS Bindings

Entry Bindings    : 4
Exit Bindings     : 2
Transit Bindings  : 4
PHP Mode          : Explicit-Null
QoS Mode          : Uniform
TTL Propagation   : Uniform

Entry Bindings:

Origin   Prefix             Ingress          Nexthop          Outgoing  Egress       Egress           Status
                             VRF              Address          Label     Interface    VRF
-------------------------------------------------------------------------------------------------------------
LDP      1.1.1.1/32          default          70.1.0.1         30        1/1/51       default          operational
LDP      2.2.2.2/32          default          60.1.0.1         19        1/1/48       default          operational
LDP      3.3.3.3/32          default          60.1.0.1         exp-null  1/1/48       default          operational
LDP      4.4.4.4/32          default          70.1.0.1         exp-null  1/1/51       default          operational

Exit Bindings:

Origin   Prefix             Incoming  Service   Egress           Status
                            Label     Label     VRF
-----------------------------------------------------------------------
static   n/a                exp-null  -         default          operational
static   n/a                7         -         default          operational

Transit Bindings:

Origin   Prefix             Incoming  Egress       Egress           Nexthop          Outgoing  Status
                            Label     Interface    VRF              Address          Label
-------------------------------------------------------------------------------------------------------
LDP      2.2.2.2/32         20        1/1/48       default          60.1.0.1         19        operational
LDP      3.3.3.3/32         21        1/1/48       default          60.1.0.1         exp-null  operational
LDP      4.4.4.4/32         23        1/1/51       default          70.1.0.1         exp-null  operational
LDP      1.1.1.1/32         29        1/1/51       default          70.1.0.1         30        operational

P5#
```

Show_P5.txt

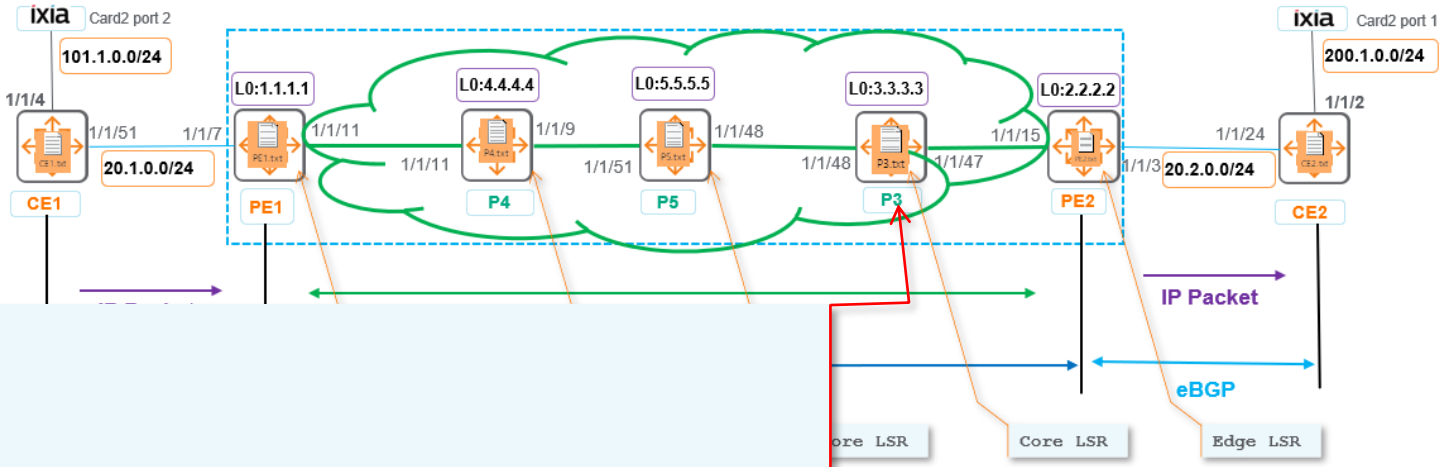


```
P3# sh mpls forwarding

MPLS Bindings

Entry Bindings    : 4
Exit Bindings     : 2
Transit Bindings  : 4
PHP Mode          : Explicit-Null
QoS Mode          : Uniform
TTL Propagation   : Uniform

Entry Bindings:

Origin   Prefix             Ingress          Nexthop            Outgoing  Egress       Egress           Status
                             VRF              Address            Label     Interface    VRF
-------------------------------------------------------------------------------------------------------------
LDP      1.1.1.1/32          default          60.1.0.2           29        1/1/48       default          operational
LDP      2.2.2.2/32          default          30.1.0.2           exp-null  1/1/47       default          operational
LDP      4.4.4.4/32          default          60.1.0.2           23        1/1/48       default          operational
LDP      5.5.5.5/32          default          60.1.0.2           exp-null  1/1/48       default          operational

Exit Bindings:

Origin   Prefix             Incoming  Service   Egress           Status
                            Label     Label     VRF
-----------------------------------------------------------------------
static   n/a                exp-null  -         default          operational
static   n/a                7         -         default          operational

Transit Bindings:

Origin   Prefix             Incoming  Egress       Egress           Nexthop            Outgoing  Status
                            Label     Interface    VRF              Address            Label
-------------------------------------------------------------------------------------------------------
LDP      2.2.2.2/32         19        1/1/47       default          30.1.0.2           exp-null  operational
LDP      5.5.5.5/32         22        1/1/48       default          60.1.0.2           exp-null  operational
LDP      4.4.4.4/32         24        1/1/48       default          60.1.0.2           23        operational
LDP      1.1.1.1/32         30        1/1/48       default          60.1.0.2           29        operational

P3#
```

Show_P3.txt

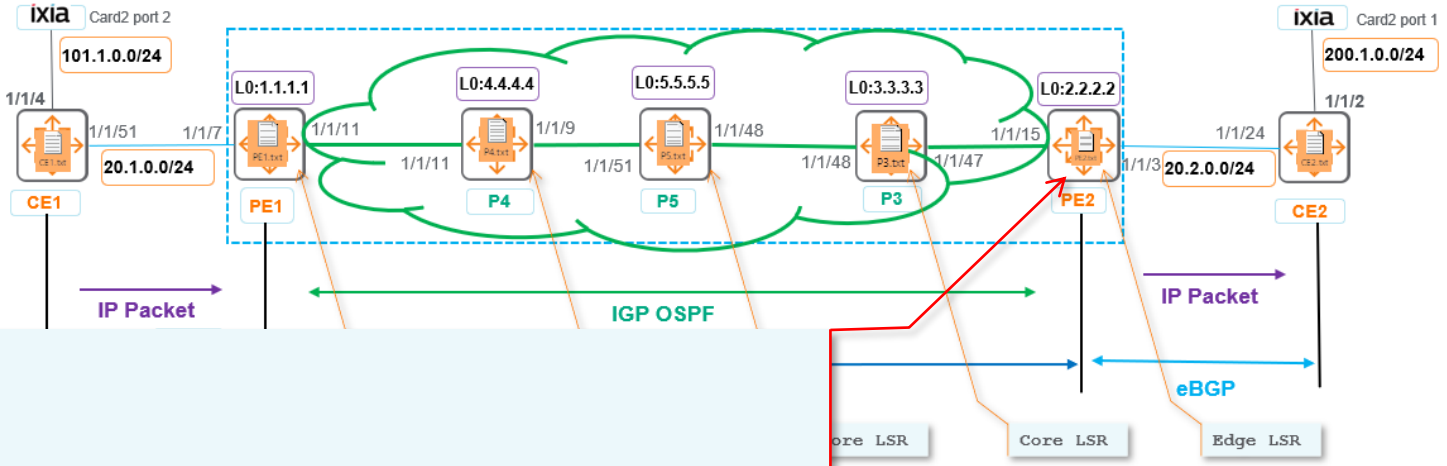


```
PE2# show mpl forwarding
MPLS Bindings
Entry Bindings    : 11
Exit Bindings     : 3
Transit Bindings : 4
PHP Mode          : Explicit-Null
QoS Mode          : Uniform
TTL Propagation   : Uniform
Entry Bindings:
Origin    Prefix              Ingress           Nexthop           Outgoing  Egress       Egress           Status
                              VRF               Address           Label     Interface    VRF
-------------------------------------------------------------------------------------------------------------------
LDP       1.1.1.1/32          default           30.1.0.1          30        1/1/15       default          operational
LDP       3.3.3.3/32          default           30.1.0.1          exp-null  1/1/15       default          operational
LDP       4.4.4.4/32          default           30.1.0.1          24        1/1/15       default          operational
LDP       5.5.5.5/32          default           30.1.0.1          22        1/1/15       default          operational
BGP       20.1.0.0/24         vrf_1             1.1.1.1           47        1/1/15       default          operational
BGP       30.30.30.0/24       vrf_1             1.1.1.1           47        1/1/15       default          operational
BGP       101.1.0.0/24        vrf_1             1.1.1.1           47        1/1/15       default          operational
BGP       101.1.1.0/24        vrf_1             1.1.1.1           47        1/1/15       default          operational
BGP       101.1.2.0/24        vrf_1             1.1.1.1           47        1/1/15       default          operational
BGP       101.1.3.0/24        vrf_1             1.1.1.1           47        1/1/15       default          operational
BGP       101.1.4.0/24        vrf_1             1.1.1.1           47        1/1/15       default          operational
Exit Bindings:
Origin    Prefix              Incoming  Service   Egress          Status
                              Label     Label     VRF
-------------------------------------------------------------------------
static    n/a                 exp-null  -         default         operational
BGP       n/a                 imp-null  16        vrf_1           operational
static    n/a                 7         -         default         operational
Transit Bindings:
Origin    Prefix              Incoming  Egress        Egress          Nexthop           Outgoing  Status
                              Label     Interface     VRF             Address           Label
-------------------------------------------------------------------------------------------------------------
LDP       3.3.3.3/32          18        1/1/15        default         30.1.0.1          exp-null  operational
LDP       5.5.5.5/32          23        1/1/15        default         30.1.0.1          22        operational
LDP       4.4.4.4/32          25        1/1/15        default         30.1.0.1          24        operational
LDP       1.1.1.1/32          31        1/1/15        default         30.1.0.1          30        operational
PE2#
```

Show_PE2.txt

# MPLS L3VPN

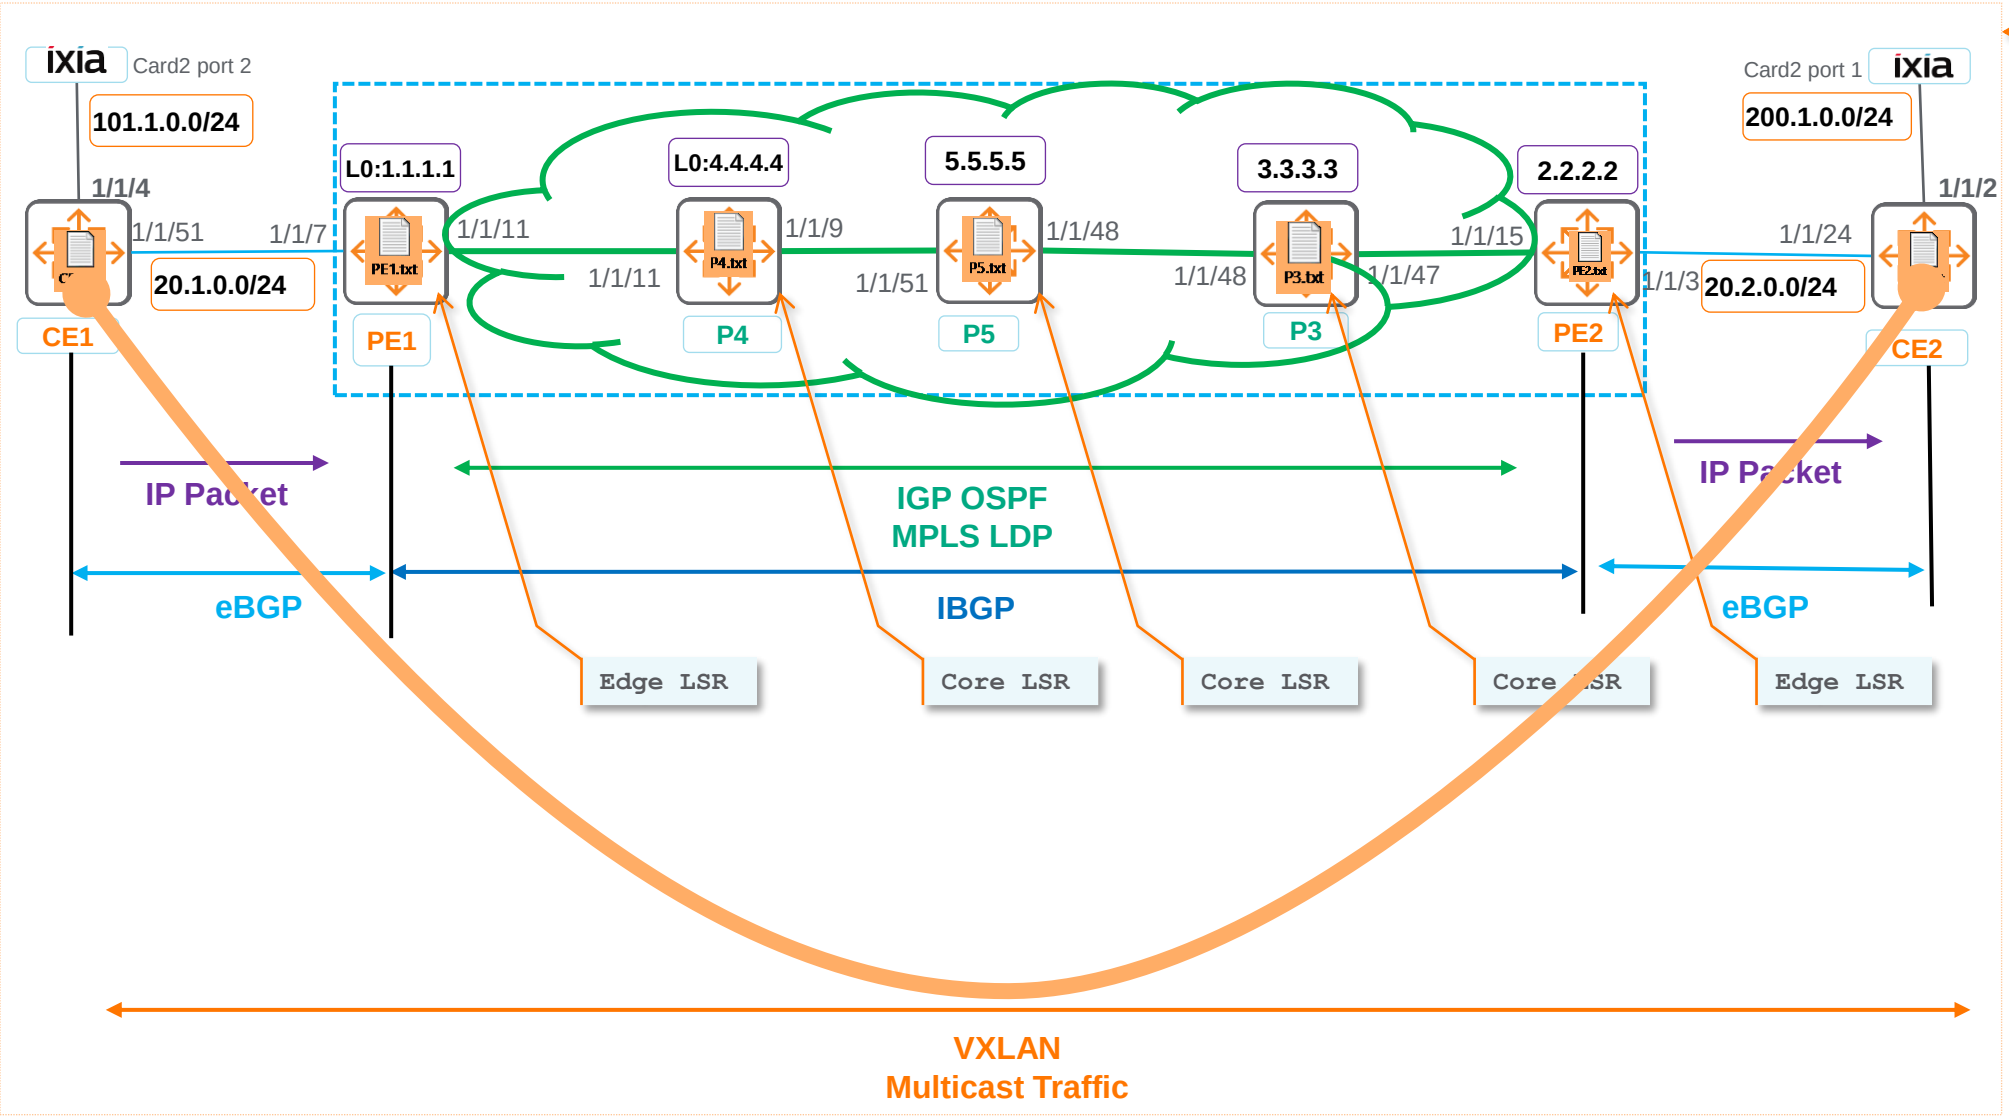


**Step1:**
Configuration
**Step2:**
OSPF / unicast routing
**Step3:**
LDP neighbor
**Step4:**
Labels for all LSRs & labels for remote VPNv4
**Step5:**
MP-BGP VPNv4 peering
**Step6:**
VRF routing table
**Step7:**
PE and CE Connectivity
**Step8:**
Ping CE's
**Step9:**
vxlan CEs

aruba
a Hewlett Packard Enterprise company

# Thank you

daryl.wan@hpe.com
yashavantha.n.n@hpe.com