Deploy MPLS L3 VPN

Issue Date: [201609]
Revision: [01]

APNIC

Acknowledgement

- Cisco Systems

Course Outline

- MPLS L3 VPN Models
- L3 VPN Terminologies
- MPLS VPN Operation
  - Control Panel
  - Data Plane
  - Forwarding function
- Function of RD and RT
- Configuration Examples
- MPLS L3 VPN Service Deployment
  - Multi-homed VPN Sites
  - Hub and Spoke
  - Extranet VPN
  - Internet Access Services

MPLS L3VPN Principle

APNIC

MPLS VPN Models

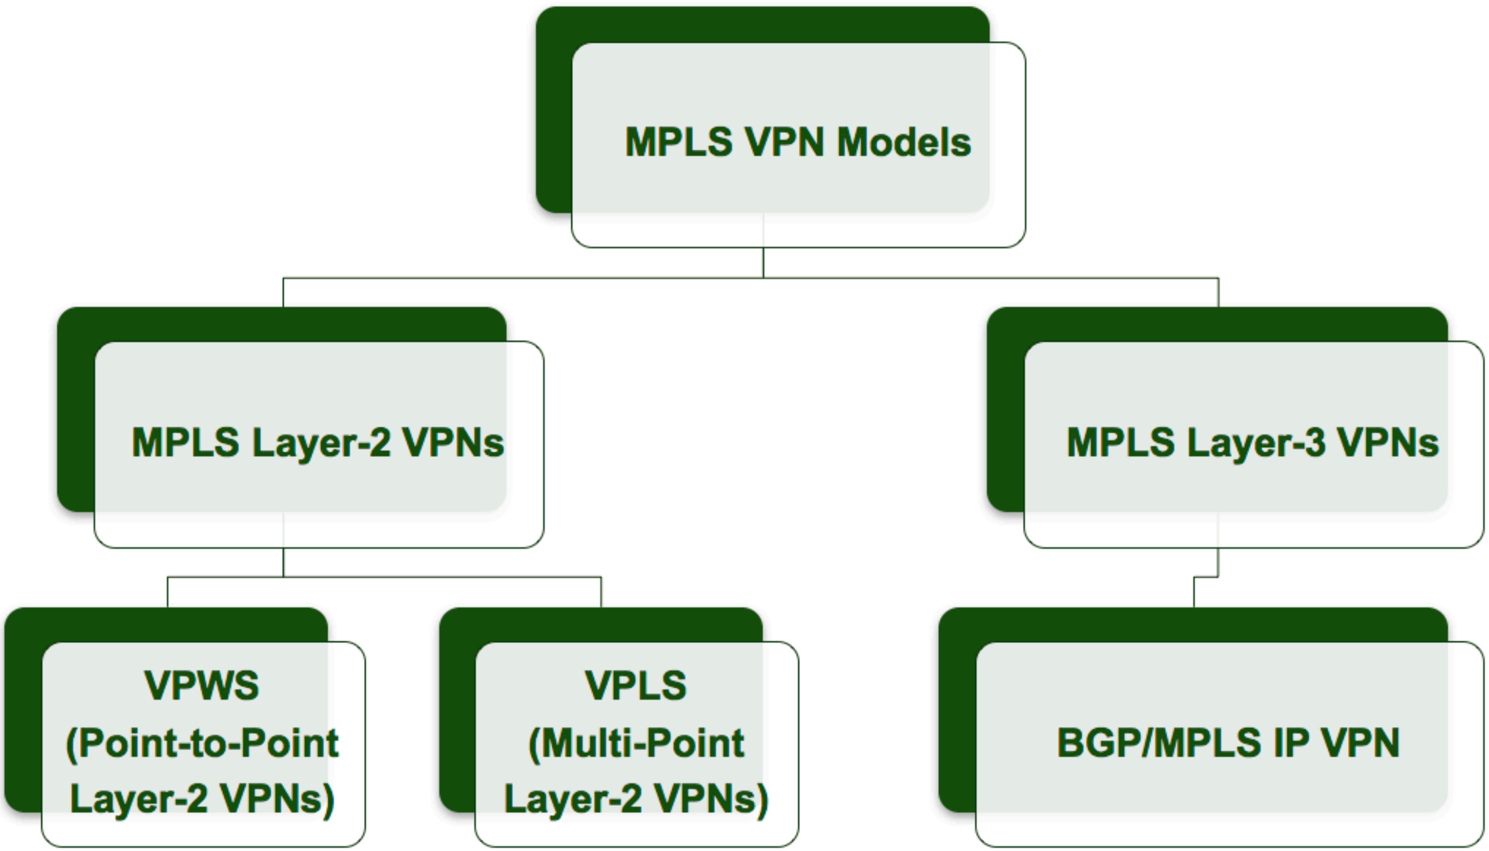

Advantages of MPLS Layer-3 VPN

- Scalability
- Security
- Easy to Create
- Flexible Addressing
- Integrated Quality of Service (QoS) Support
- Straightforward Migration

MPLS L3VPN Topology

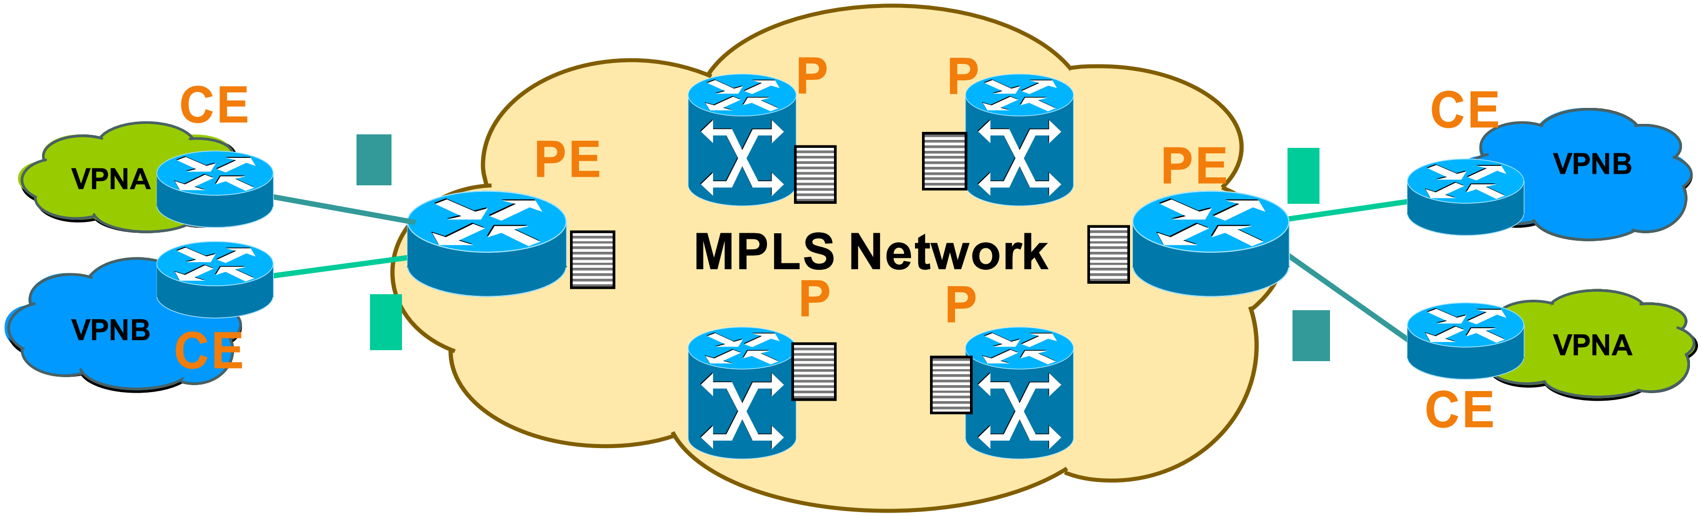

- PE: Provider Edge Router
- P : Provider Router
- CE: Customer Edge Router

Virtual Routing and Forwarding Instance

- Virtual routing and forwarding table
  - On PE router
  - Separate instance of routing (RIB) and forwarding table
- A VRF defines the VPN membership of a customer site attached to a PE device.
- VRF associated with one or more customer interfaces

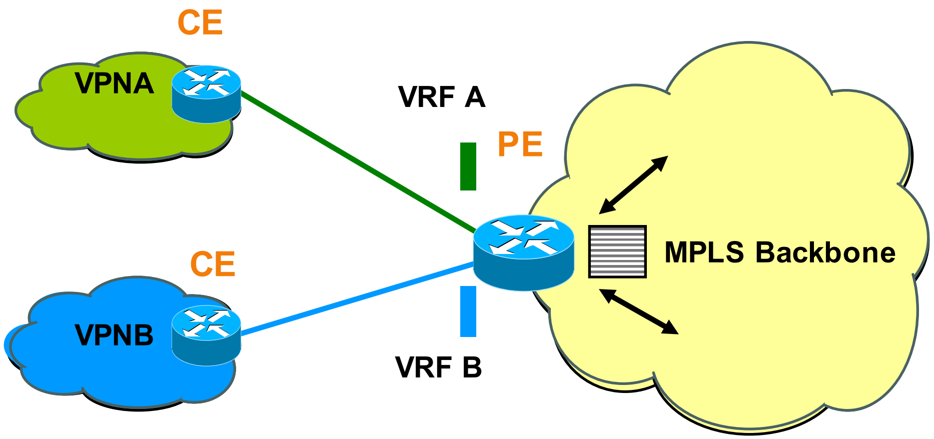

Routes Transfer between CE and PE

- PE installs the internal routes (IGP) in global routing table
- PE installs the VPN customer routes in VRF routing tables
  - VPN routes are learned from CE routers or remote PE routers
  - VRF-aware routing protocol (static, RIP, BGP, OSPF, IS-IS) on each PE

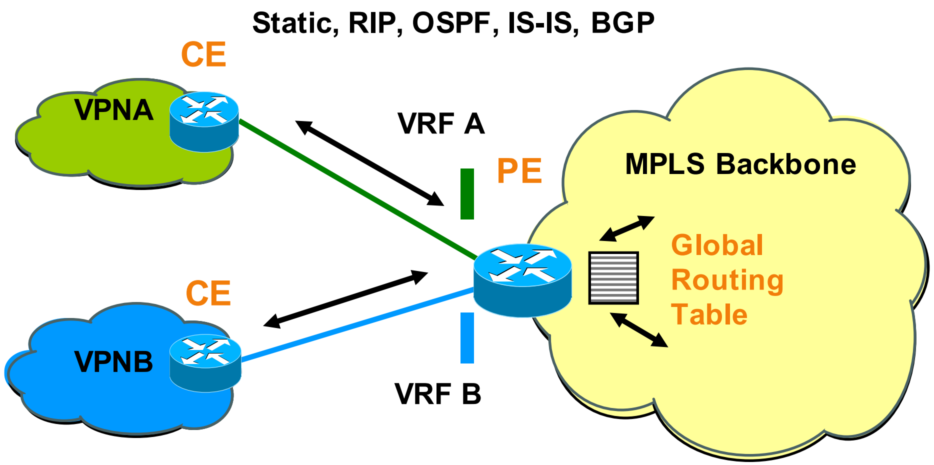

APNIC

Control Plane: Multi-Protocol BGP

- PE routers distribute VPN routes to each other via MP-BGP.
- MP-BGP customizes the VPN Customer Routing Information as per the Locally Configured VRF Information at the PE using:
  - Route Distinguisher (RD)
  - Route Target (RT)
  - VPN Label

What is RD

- Route distinguisher is an 8-octet field prefixed to the customer's IPv4 address. RD makes the customer's IPv4 address unique inside the SP MPLS network.
- RD is configured in the VRF at PE

VPNv4 Address:

Route Distinguisher (8 bytes)		IPv4 Address (4 bytes)

Example:

Type 0	100:1	10.1.1.1
Type 1	192.168.19.1:1	10.1.1.1
Type 2	65538:10	10.1.1.1

Route Advertisement: RD

- VPN customer IPv4 prefix is converted into a VPNv4 prefix by appending the RD to the IPv4 address
- PE devices use MP-BGP to advertise the VPNv4 address

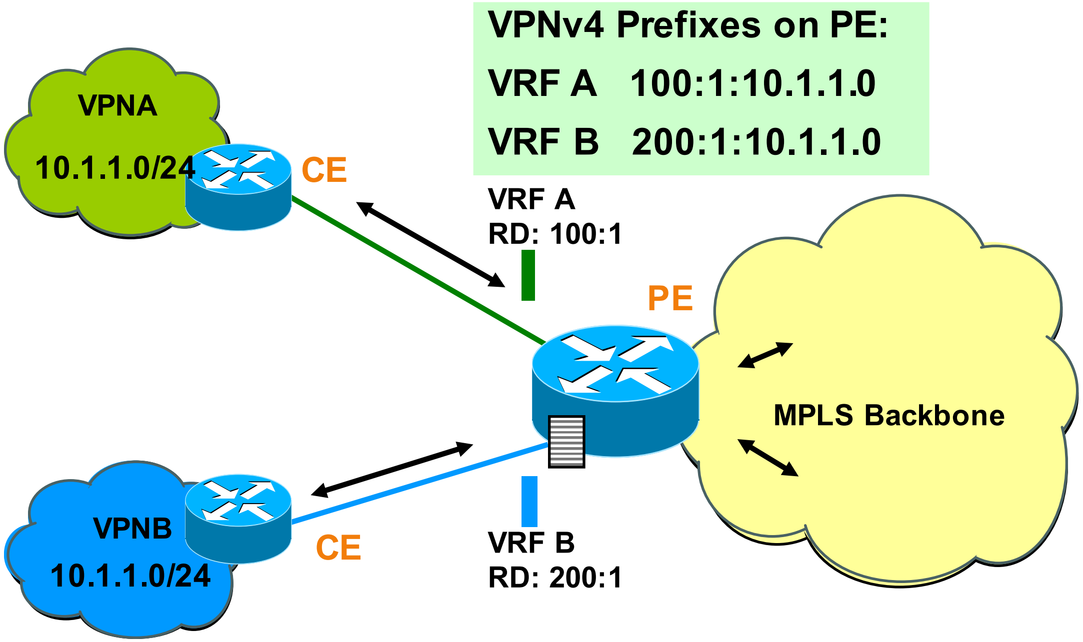

What is RT

- Route Target is a BGP extended community attribute, is used to control VPN routes advertisement.

Route Target (8 bytes)

Example:

Type 0	100:1
Type 1	192.168.1.1:1
Type 2	65538:10

- Two types of RT:
  - Export RT
  - Import RT

Route Advertisement: RT

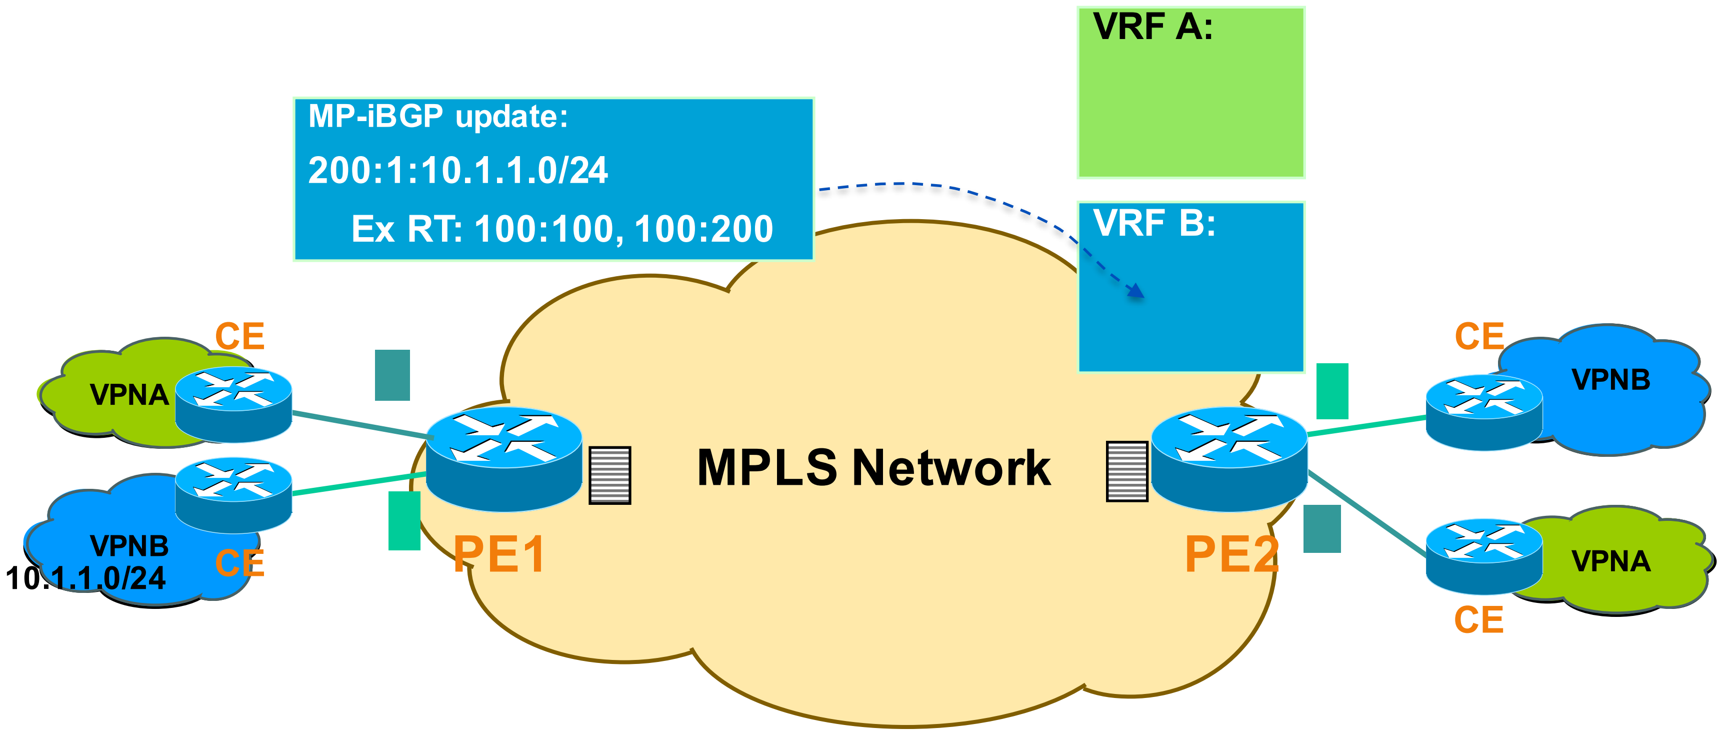

	Import RT	Export RT
VRF A	100:1	100:1
VRF B	100:100 100:200	**100:100** 100:200

	Import RT	Export RT
VRF A	100:1 100:2 100:3	100:1 100:2
VRF B	**100:100**	100:100

Using RT to Build VPN Topologies

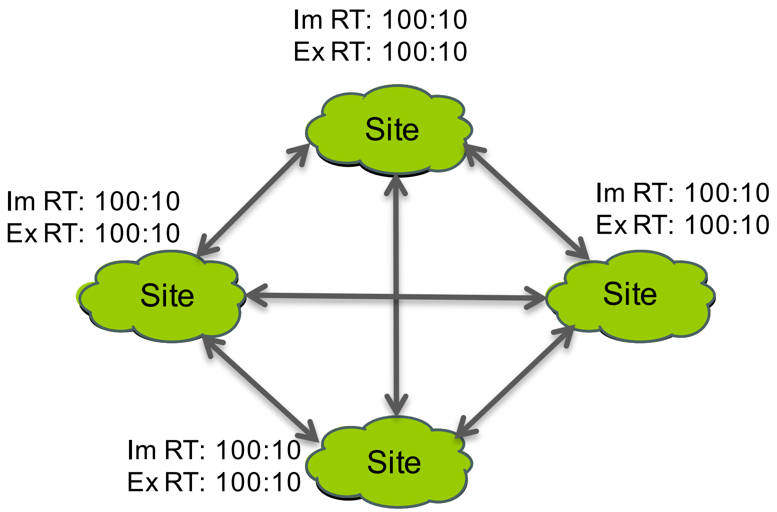

Full Mesh

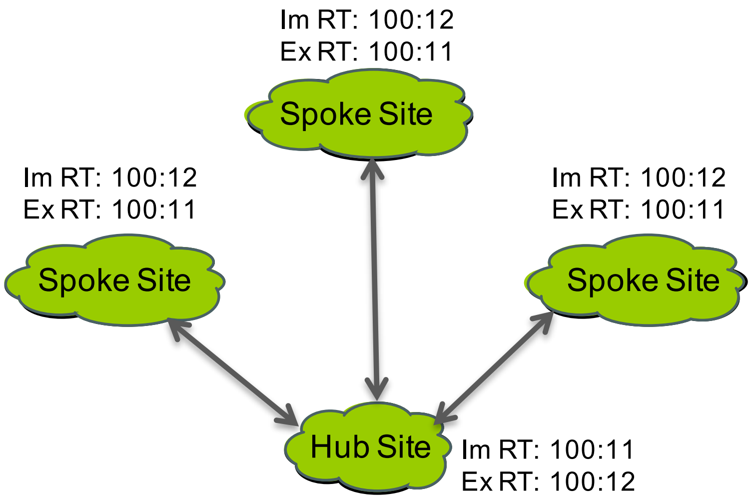

Hub Spoke

In a full-mesh VPN, each site in the VPN can communicate with every other site in that same VPN.

In a hub-and-spoke VPN, the spoke sites in the VPN can communicate only with the hub sites; they cannot communicate with other spoke sites.

VPN Label

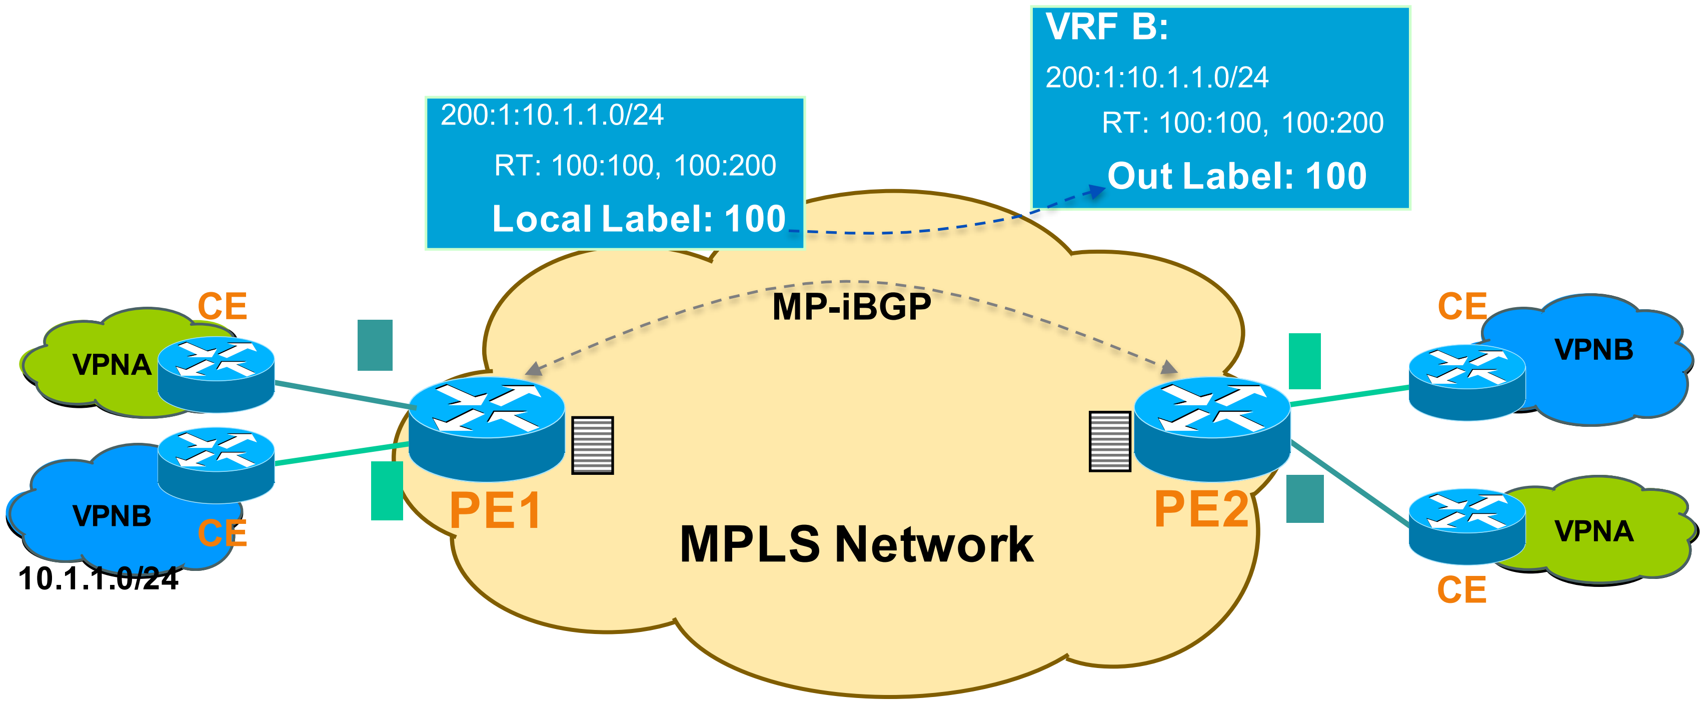

- PE adds the label to the NLRI field.

Control Plane Walkthrough(1/2)

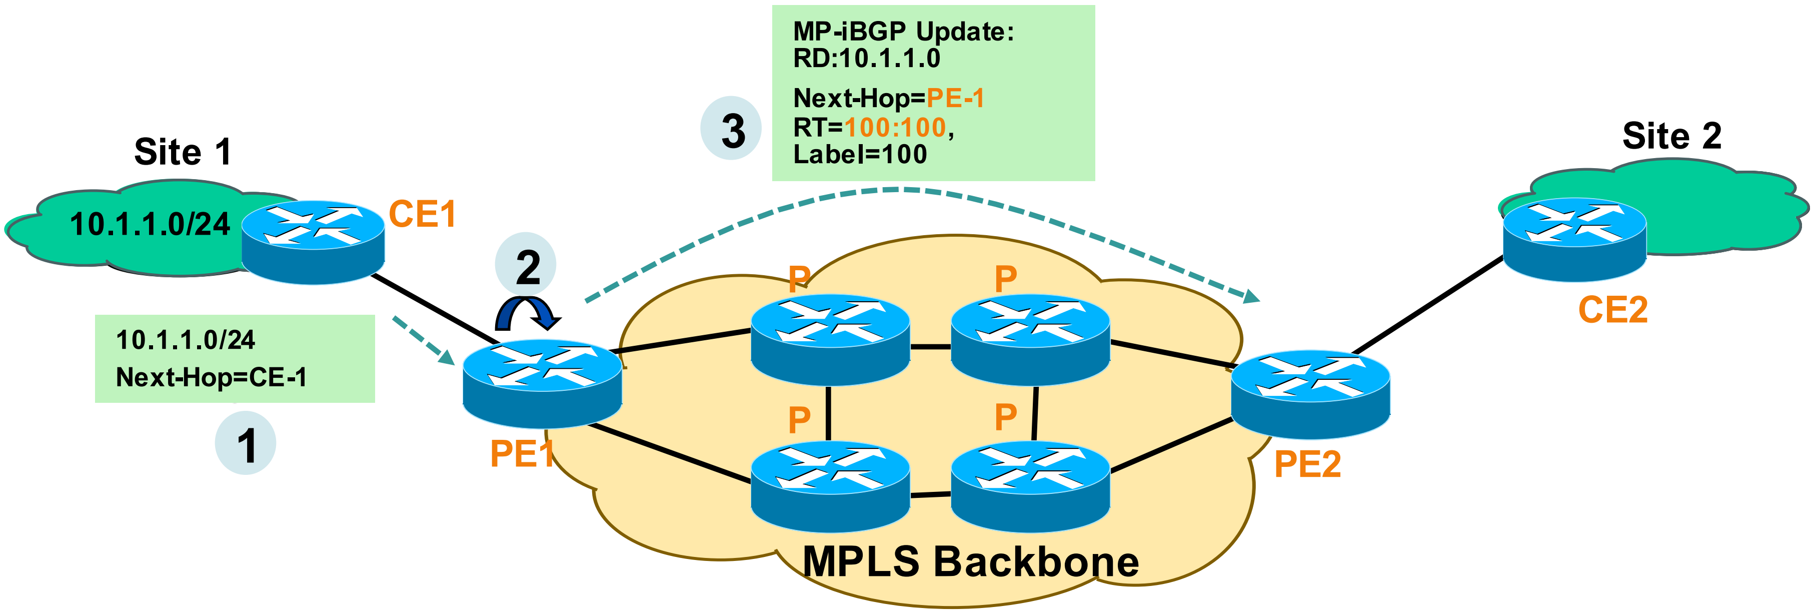

1. PE1 receives an IPv4 update (eBGP/OSPF/ISIS/RIP)
2. PE1 converts it into VPNv4 address and constructs the MP-iBGP UPDATE message
   - Associates the RT values (export RT =100:100) per VRF configuration
   - Rewrites next-hop attribute to itself
   - Assigns a label (100); Installs it in the MPLS forwarding table.
3. PE1 sends MP-iBGP update to other PE routers

Control Plane Walkthrough(2/2)

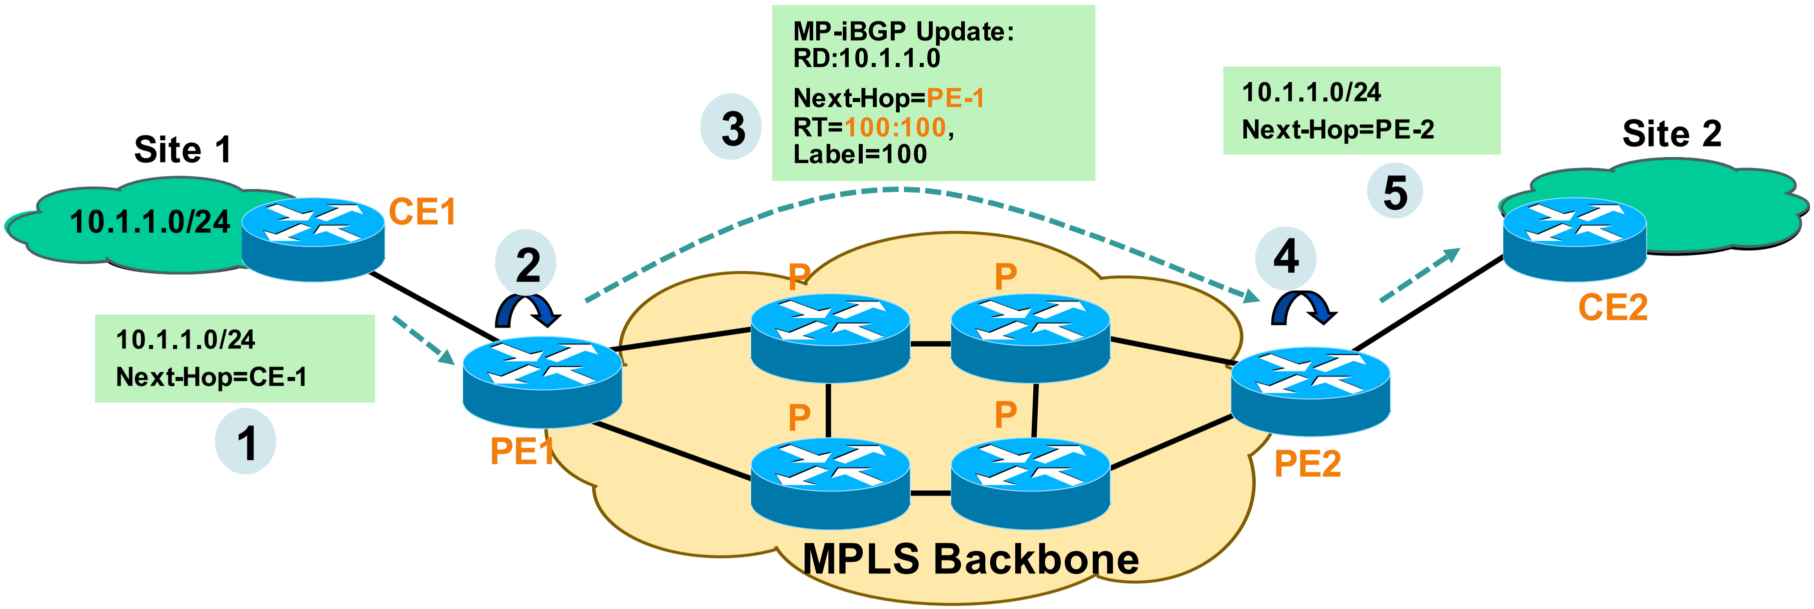

4. PE2 receives and checks whether the RT=200:1 is locally configured as import RT within any VRF, if yes, then
   - PE2 translates VPNv4 prefix back to IPv4 prefix
   - Updates the VRF CEF table for 10.1.1.0/24 with label=100
5. PE2 advertises this IPv4 prefix to CE2

Control Plane: Tunnel Label

Local Label	Prefix	Out Interface	Out Label
Pop-Label	1.1.1.1/32	-	-

Local Label	Prefix	Out Interface	Out Label
50	1.1.1.1/32	Eth0/1	Pop-Label

Local Label	Prefix	Out Interface	Out Label
25	1.1.1.1/32	Eth0/0	50

Local Label	Prefix	Out Interface	Out Label
-	1.1.1.1/32	Eth0/1	25

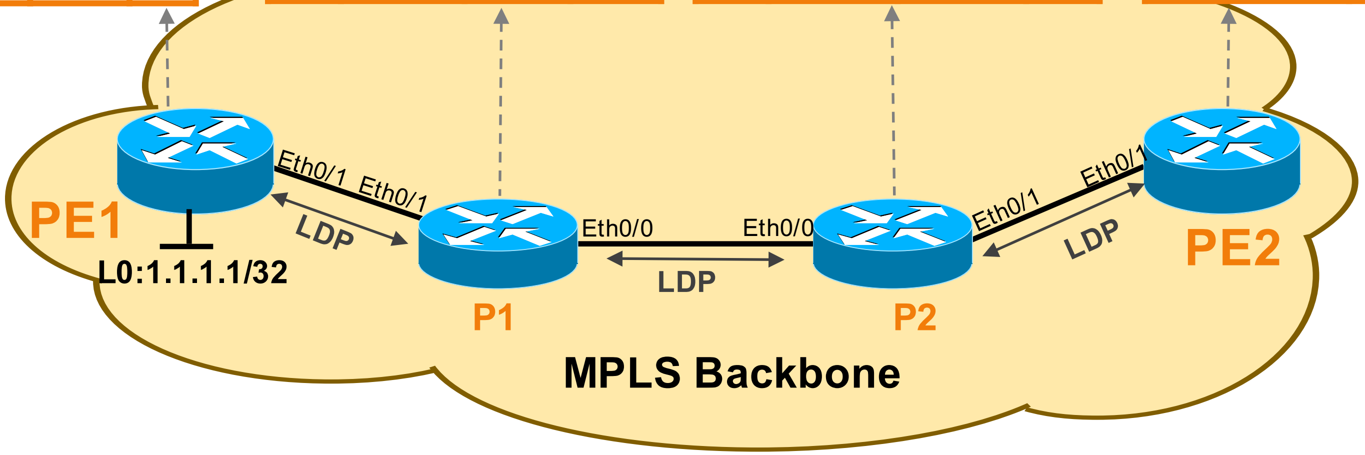

- LDP runs on the MPLS backbone network to build the public LSP. The tunnel label is also called transport label or public label.
- Local label mapping are sent to connected nodes. Receiving nodes update forwarding table.

Data Plane

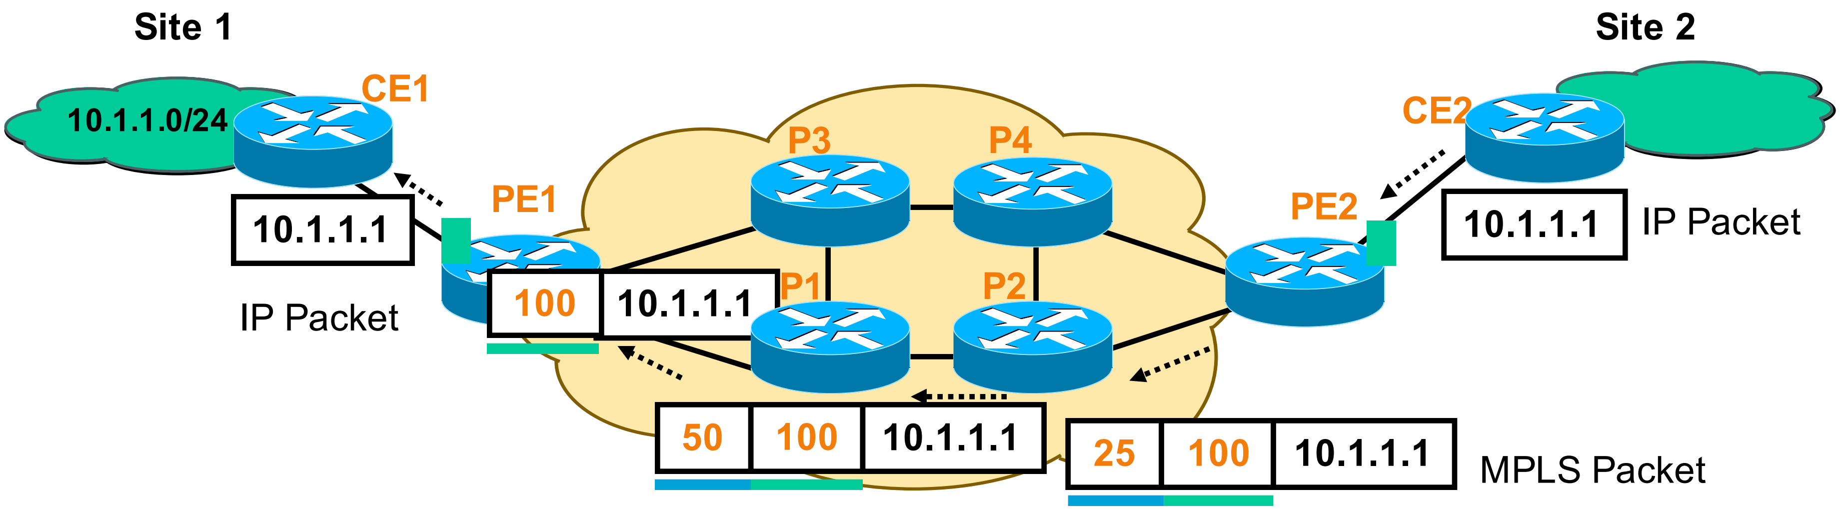

- PE2 imposes two labels for each IP packet going to site2
  - Tunnel label is learned via LDP; corresponds to PE1 address
  - VPN label is learned via BGP; corresponds to the VPN address
- P1 does the Penultimate Hop Popping (PHP)
- PE1 retrieves IP packet (from received MPLS packet) and forwards it to CE1.

Configuration Example

- Task: Configure MPLS L3VPN on Cisco **IOS** (Version 15.2) to make the following CEs communication with each other.
- Prerequisite configuration:
  - 1. IP address configuration on PE & P routers
  - 2. IGP configuration on PE & P routers
    - Make sure all the routers in public network can reach each other.

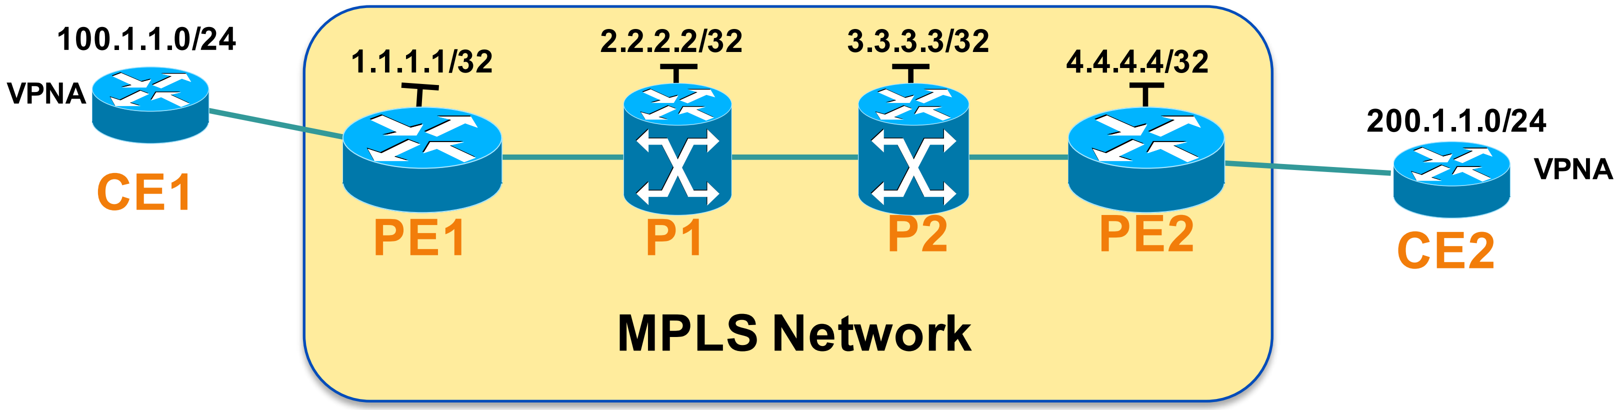

Configure MPLS & LDP

- Configuration steps:
  - 1. Configure MPLS and LDP on PE & P routers

```
ip cef
mpls ldp router-id loopback 0

interface ethernet1/0
mpls ip
mpls label protocol ldp

interface ethernet1/1
mpls ip
mpls label protocol ldp
```

Configure VRF

- Configuration steps:
  - 2. Configure VRF instance on PE routers

```
vrf definition VPNA
 rd 100:10
 route-target export 100:100
 route-target import 100:100
 !
 address-family ipv4
 exit-address-family
!
```

  - Bind PE-CE interface under VRF

```
interface FastEthernet0/0
 vrf forwarding VPNA
 ip address 10.1.1.1 255.255.255.252
```

Configure MP-iBGP

- Configuration steps:
  - 3. Enable MP-iBGP neighbors in vpnv4 address-family on PE routers

```
router bgp 100
 neighbor 4.4.4.4 remote-as 100
 neighbor 4.4.4.4 update-source loopback 0
 !
 address-family vpnv4
  neighbor 4.4.4.4 activate
  neighbor 4.4.4.4 send-community both
 exit-address-family
!
```

Configure PE-CE eBGP Neighbour

- Configuration steps:
  - 4. Adding PE-CE eBGP neighbour in VRF context of BGP on PE

```
router bgp 100
 address-family ipv4 vrf VPNA
  neighbor 10.1.1.2 remote-as 65001
  neighbor 10.1.1.2 activate
 exit-address-family
!
```

Adding PE-CE eBGP neighbour in BGP on CE

```
router bgp 65001
 neighbor 10.1.1.1 remote-as 100
 !
 address-family ipv4
  network 100.1.1.0 mask 255.255.255.0
  neighbor 10.1.1.1 activate
 exit-address-family
!
ip route 100.1.1.0 255.255.255.0 null 0
```

Verify Results – VRF Routing Table

- Check the routes of VRF VPNA on PE.

```
PE1#show bgp vpnv4 unicast vrf VPNA
BGP table version is 4, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:10 (default for vrf VPNA)
 *>  100.1.1.0/24     10.1.1.2                 0             0 65001 i
 *>i 200.1.1.0        4.4.4.4                  0    100      0 65002 i
```

Verify Results – VPN Reachability

- CE can learn the routes from each other:

```
CE2#show ip route
....
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.1.2.0/30 is directly connected, FastEthernet0/1
L        10.1.2.2/32 is directly connected, FastEthernet0/1
      100.0.0.0/24 is subnetted, 1 subnets
B        100.1.1.0 [20/0] via 10.1.2.1, 00:38:26
      200.1.1.0/24 is variably subnetted, 2 subnets, 2 masks
S        200.1.1.0/24 is directly connected, Null0
C        200.1.1.1/32 is directly connected, Loopback1
```

Configuration Example

- Task: Configure MPLS L3VPN on **Juniper Junos** (Version 12.1) to make the following CEs communication with each other.
- Prerequisite configuration:
  - 1. IP address configuration on PE & P routers
  - 2. IGP configuration on PE & P routers
    - Make sure all the routers in public network can reach each other.

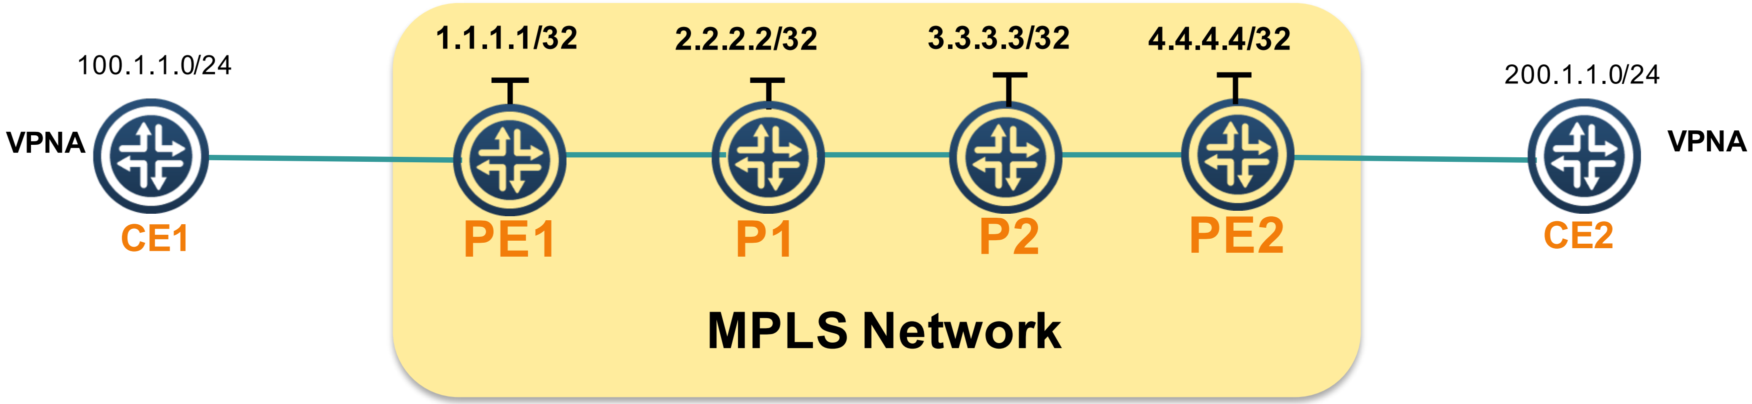

Configure MPLS & LDP

- Configuration steps:
  - 1. Configure MPLS and LDP on PE & P routers
  - This is the example on PE1.

```
interfaces {
    em0 {
        unit 0 {
            family inet {
                address 10.0.12.1/30;
            }
            family mpls;
        }
    }
}
```

```
protocols {
    mpls {
        interface em0.0;
    }
     ldp {
        interface em0.0;
}}
```

Configure VRF

- Configuration steps:
  - 2. Configure VRF instance on PE routers.

```
routing-instances {
    VPNA {
        instance-type vrf;
        interface em1.0;
        route-distinguisher 100:10;
        vrf-target target:100:100;
    }
}
```

VPN instance and parameters, Interface em1.0 has been added in the VPNA

```
em1 {
        unit 0 {
            family inet {
                address 10.0.1.2/30;
            }
        }
    }
```

This is the interface configuration from PE to CE, as a normal interface

Configure MP-iBGP

- Configuration steps:
  - 3. Enable MP-iBGP neighbors in vpnv4 address-family on PE routers

```
routing-options {
    router-id 1.1.1.1;
    autonomous-system 100;
}
```

```
protocols {
    bgp {
        local-address 1.1.1.1;
        family inet-vpn {
            unicast;
        }
        group PE1-PE2 {
            type internal;
            neighbor 4.4.4.4;
} }
```

Configure PE-CE eBGP Neighbour

- Configuration steps:
  - 4. Adding PE-CE eBGP neighbour in VPN on PE

```
routing-instances {
    VPNA {
        instance-type vrf;
        interface em1.0;
        route-distinguisher 100:10;
        vrf-target target:100:100;
        protocols {
            bgp {
                group PE1-CE1 {
                    type external;
                    peer-as 65001;
                    neighbor 10.0.1.1;
}}}}}
```

Configure PE-CE eBGP Neighbour

- Configuration steps:
  - 4. Adding CE-PE eBGP neighbour in BGP on CE

```
routing-options {
    autonomous-system 65001;
}
protocols {
    bgp {
        group CE1-PE1 {
            type external;
            peer-as 100;
            neighbor 10.0.1.2;
        }
    }
}
```

CE1 is in AS 65001, sets up the neighbor with AS100.

Advertise Static Route on CE

- Configuration steps:
  - 5. Advertise routes on CE routers, CE1 advertises 100.1.1.0/24, CE2 advertises 200.1.1.0/24

```
routing-options {
    generate {
        route 100.1.1.0/24 passive;
}}
```

Generate a static route.

```
policy-options {
    policy-statement ADVERTISE-PREFIX {
        from {
            route-filter 100.1.1.0/24 exact;
        }
        then accept;
}}
```

Define the route policy

```
protocols {
    bgp {
        group CE1-PE1 {
            export ADVERTISE-PREFIX;
}}}
```

Apply the policy in eBGP neighbor, only advertise 100.1.1.0/24

Verify Results – VRF Routing Table

- Check the routes of VRF VPNA on PE.

```
root@PE1> show route receive-protocol  bgp 4.4.4.4

inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

VPNA.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
  Prefix    Nexthop        MED     Lclpref    AS path
* 10.0.2.0/30              4.4.4.4                      100         I
* 200.1.1.0/24             4.4.4.4                      100         65002 I


mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)


bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
  Prefix    Nexthop        MED     Lclpref    AS path
  100:20:10.0.2.0/30
*                          4.4.4.4                      100         I
  100:20:200.1.1.0/24
*                          4.4.4.4                      100         65002 I
```

RD on PE2 is 100:20

Check VPN Routes in BGP

- Check the detailed route of VRF VPNA on PE received from remote PE.

```
root@PE1> show route receive-protocol bgp 4.4.4.4 detail
......(Omitted)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

......(Omitted)


* 100:20:200.1.1.0/24  (1 entry, 0 announced)
     Import Accepted
     Route Distinguisher: 100:20
     VPN Label: 300016
     Nexthop: 4.4.4.4
     Localpref: 100
     AS path: 65002 I
     Communities: target:100:100
```

Verify Results – VPN Reachability

- CE can learn the routes from each other:

```
root@CE1> show route


inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


10.0.1.0/30         *[Direct/0] 00:44:34
                     > via em1.0
10.0.1.1/32         *[Local/0] 00:44:34
                       Local via em1.0
10.0.2.0/30         *[BGP/170] 00:04:23, localpref 100
                       AS path: 100 I
                     > to 10.0.1.2 via em1.0
100.1.1.0/24        *[Aggregate/130] 00:16:30
                       Reject
200.1.1.0/24        *[BGP/170] 00:04:24, localpref 100
                       AS path: 100 65002 I
                     > to 10.0.1.2 via em1.0
```

Configuration Example

- Task: Configure MPLS L3VPN on **Huawei VRP** (Version 5.1) to make the following CEs communication with each other.
- Prerequisite configuration:
  - 1. IP address configuration on PE & P routers
  - 2. IGP configuration on PE & P routers
    - Make sure all the routers in public network can reach each other.

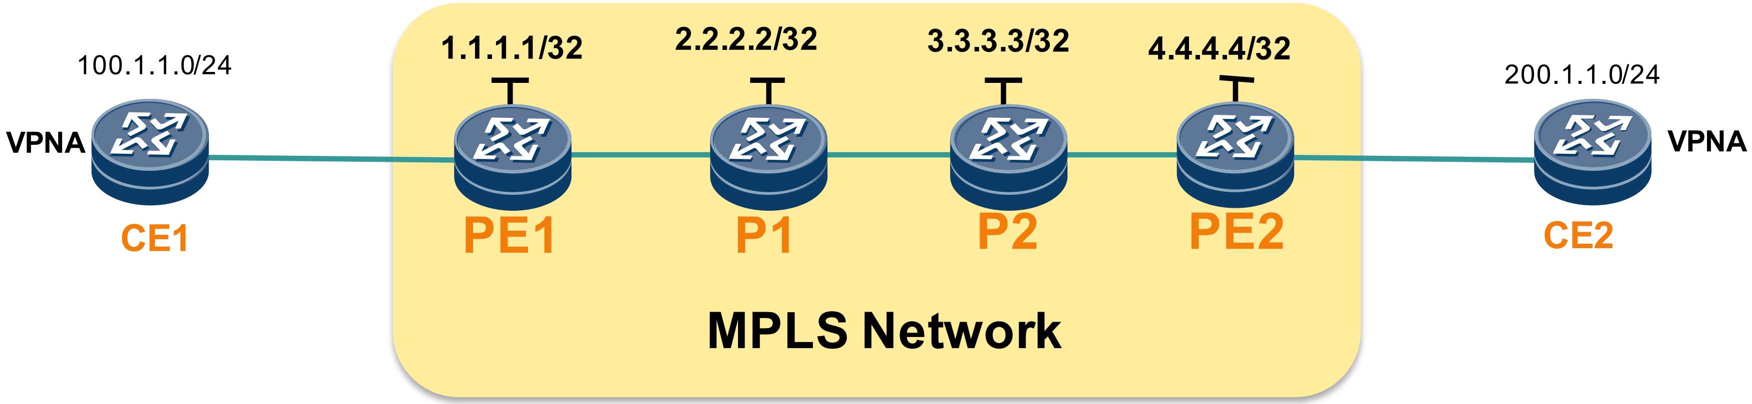

Configure MPLS & LDP

- Configuration steps:
  - 1. Configure MPLS and LDP on PE & P routers

```
[PE1] mpls lsr-id 1.1.1.1
[PE1] mpls
Info: Mpls starting, please wait... OK!
[PE1-mpls] quit
[PE1] mpls ldp
[PE1-mpls-ldp] quit
[PE1] interface gigabitethernet 0/0/0
[PE1-GigabitEthernet0/0/0] mpls
[PE1-GigabitEthernet0/0/0] mpls ldp
[PE1-GigabitEthernet0/0/0] quit
```

Configure VRF

- Configuration steps:
  - 2. Configure VRF instance on PE routers

```
[PE1] ip vpn-instance VPNA
[PE1-vpn-instance-VPNA] ipv4-family
[PE1-vpn-instance-VPNA-af-ipv4] route-distinguisher 100:10
[PE1-vpn-instance-VPNA-af-ipv4] vpn-target 100:100 both
 IVT Assignment result:
Info: VPN-Target assignment is successful.
 EVT Assignment result:
Info: VPN-Target assignment is successful.
[PE1-vpn-instance-VPNA-af-ipv4] quit
```

  - Bind PE-CE interface under VRF

```
[PE1] interface gigabitethernet 0/0/1
[PE1-GigabitEthernet0/0/1] ip binding vpn-instance vpna
Info: All IPv4 related configurations on this interface are removed!
Info: All IPv6 related configurations on this interface are removed!
[PE1-GigabitEthernet0/0/1] ip address 10.1.1.1 30
[PE1-GigabitEthernet0/0/1] quit
```

Configure MP-iBGP

- Configuration steps:
  - 3. Enable MP-iBGP neighbors in vpnv4 address-family on PE routers

```
[PE1] bgp 100
[PE1-bgp] peer 4.4.4.4 as-number 100
[PE1-bgp] peer 4.4.4.4 connect-interface loopback 0
[PE1-bgp] ipv4-family vpnv4
[PE1-bgp-af-vpnv4] peer 4.4.4.4 enable
[PE1-bgp-af-vpnv4] quit
[PE1-bgp] quit
```

Configure PE-CE eBGP Neighbour

- Configuration steps:
  - 4. Adding PE-CE eBGP neighbour in VRF context of BGP on PE

```
[PE1] bgp 100
[PE1-bgp] ipv4-family vpn-instance VPNA
[PE1-bgp-vpna] peer 10.1.1.2 as-number 65001
[PE1-bgp-vpna] quit
```

Adding CE-PE eBGP neighbour in BGP on CE

```
[CE1] ip route-static 100.1.1.0 24 null 0
[CE1] bgp 65001
[CE1-bgp] peer 10.1.1.2 as-number 100
[CE1-bgp] network 100.1.1.0 24
[CE1-bgp] quit
```

Verify Results – VRF Routing Table

- Check the routes of VRF VPNA on PE.

```
<PE1> display bgp vpnv4 vpn-instance VPNA routing-table

 BGP Local router ID is 10.0.0.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 VPN-Instance VPNA, Router ID 10.0.0.1:

 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   100.1.1.0/24       10.1.1.2        0                      0      65001i
 *>i  200.1.1.0          4.4.4.4         0          100         0      65002i
```

Check VPN Routes in BGP

- Check the detailed route of VRF VPNA on PE.

```
<PE1> display bgp vpnv4 vpn-instance VPNA routing-table 200.1.1.0

 BGP local router ID : 1.1.1.1
 Local AS number : 100

 VPN-Instance  VPNA, Router ID 1.1.1.1:
 Paths:    1 available, 1 best, 1 select
 BGP routing table entry information of 200.1.1.0/24:
 Label information (Received/Applied):  1028/NULL
 From: 4.4.4.4 (4.4.4.4)
 Route Duration: 00h00m04s
 Relay Tunnel Out-Interface: GigabitEthernet0/0/0
 Relay token: 0x18
 Original nexthop: 4.4.4.4
 Qos information : 0x0
 Ext-Community:RT  <100 : 100>
 AS-path 65002, origin igp, MED 0, localpref 100, pref-val 0, valid, internal, b
est, select, active, pre 255, IGP cost 3
 Advertised to such 1 peers:
    10.1.1.2
```

Verify Results – VPN Reachability

- CE can learn the routes from each other:

```
[CE2]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 7        Routes : 7

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.1.2.0/30  Direct  0    0           D   10.1.2.2        GigabitEthernet
0/0/1
       10.1.2.2/32  Direct  0    0           D   127.0.0.1       GigabitEthernet
0/0/1
      100.1.1.0/24  EBGP    255  0           D   10.1.2.1        GigabitEthernet
0/0/1
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
      200.1.1.0/24  Static  60   0           D   0.0.0.0         NULL0
      200.1.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
```

Questions?

Issue Date:

Revision:

APNIC

MPLS L3VPN Services

APNIC

MPLS L3VPN Services

- **Multi-homed VPN Sites**
- Hub and Spoke Service
- Extranet Service
- Internet Access Service

VPN Multihoming Scenarios

- In an MPLS VPN Layer 3 environment, it is common for customers to multihome their networks to provide link redundancy.

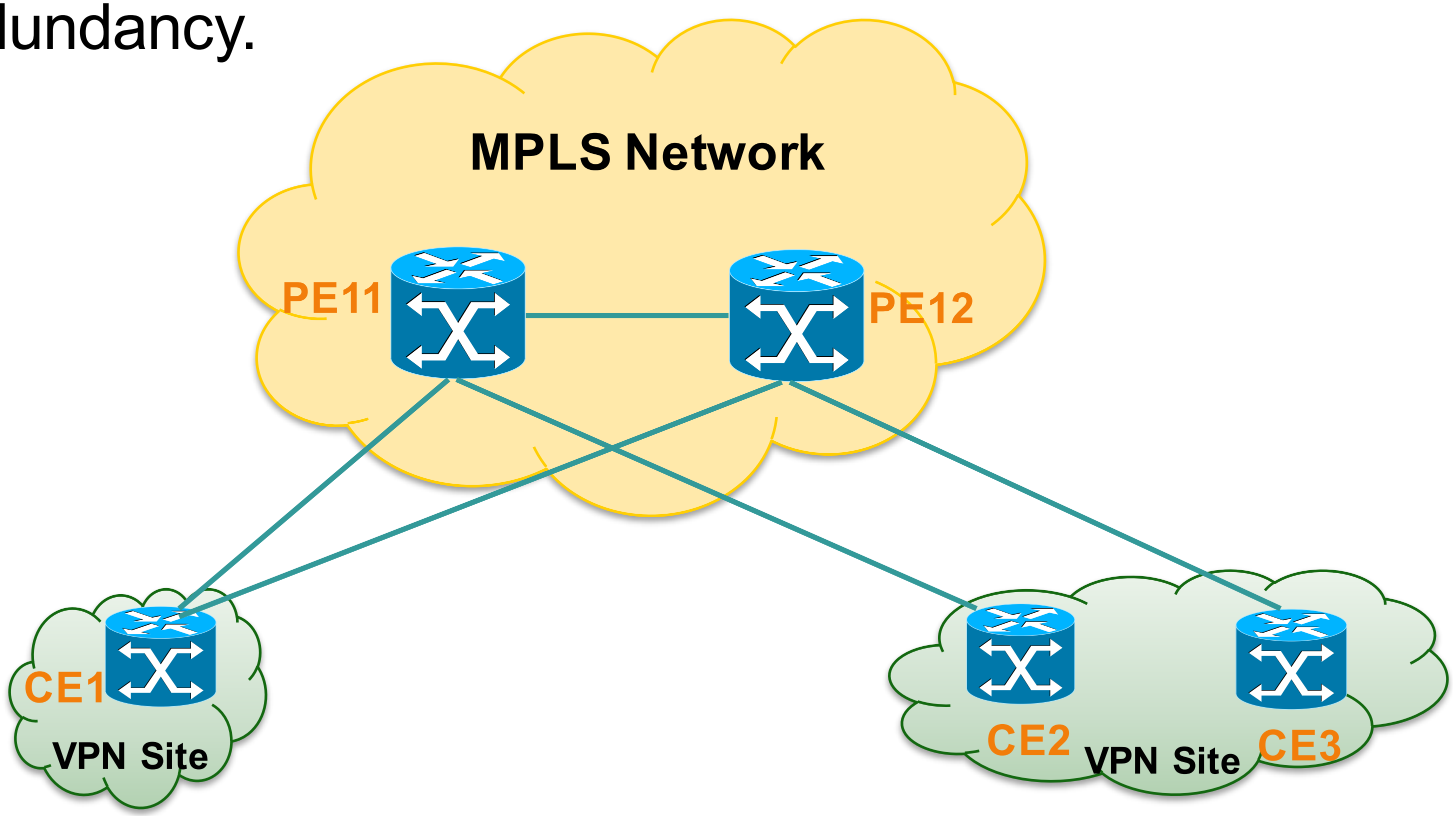

VPN Route Advertisement

- VPN route advertisement from multihomed VPN site.

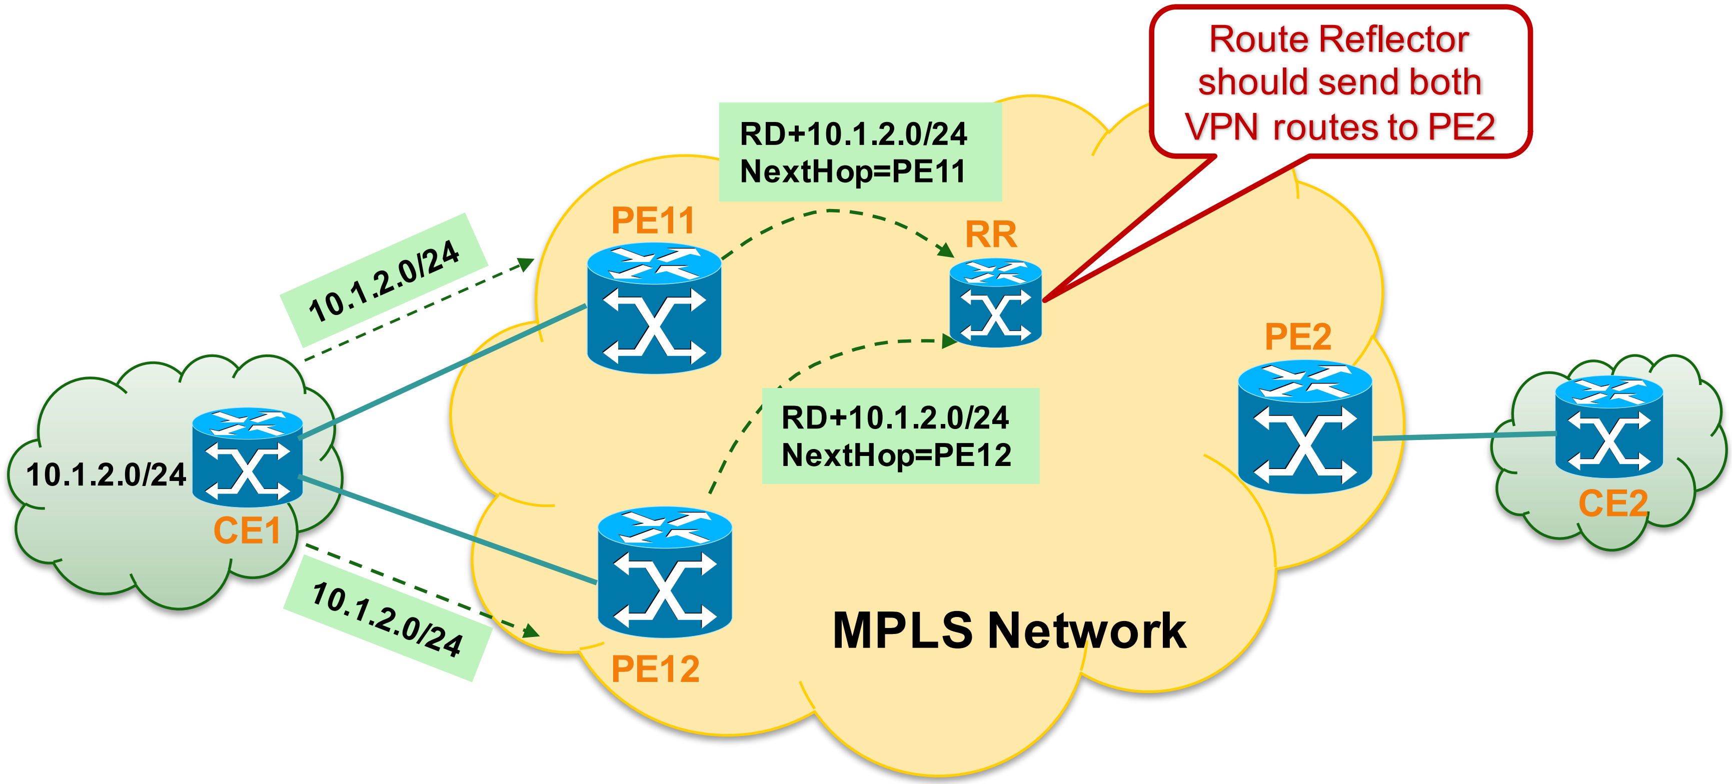

VPN Route Advertisement– Unique RD

- Configure unique RD per VRF per PE for multihomed site/interfaces

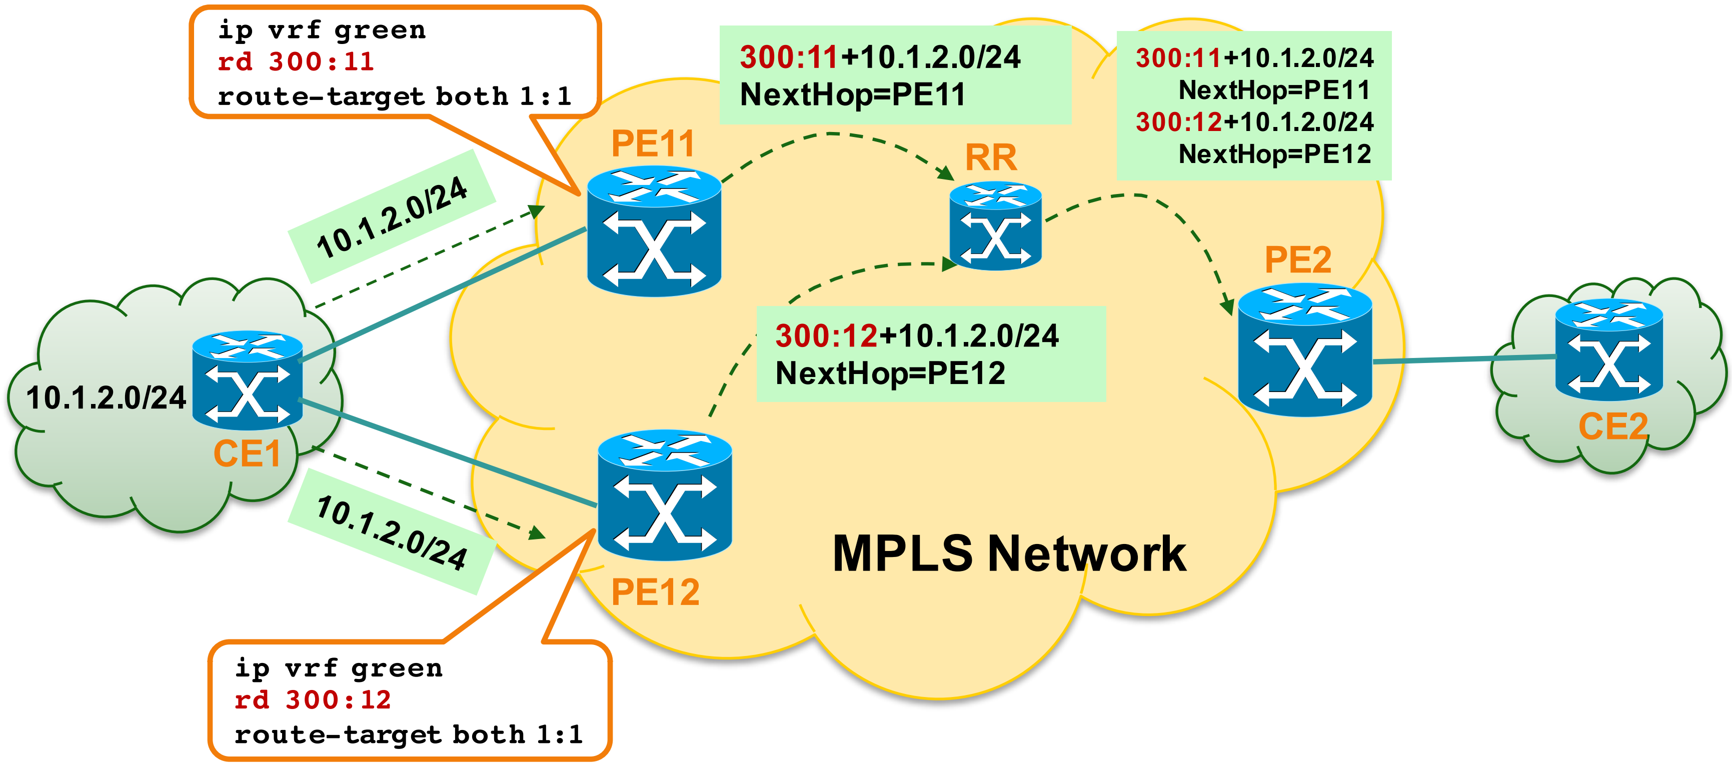

Load Sharing Configuration

- To implement load sharing between PE11 and PE12, enable BGP multipath at remote PE routers such as PE2.

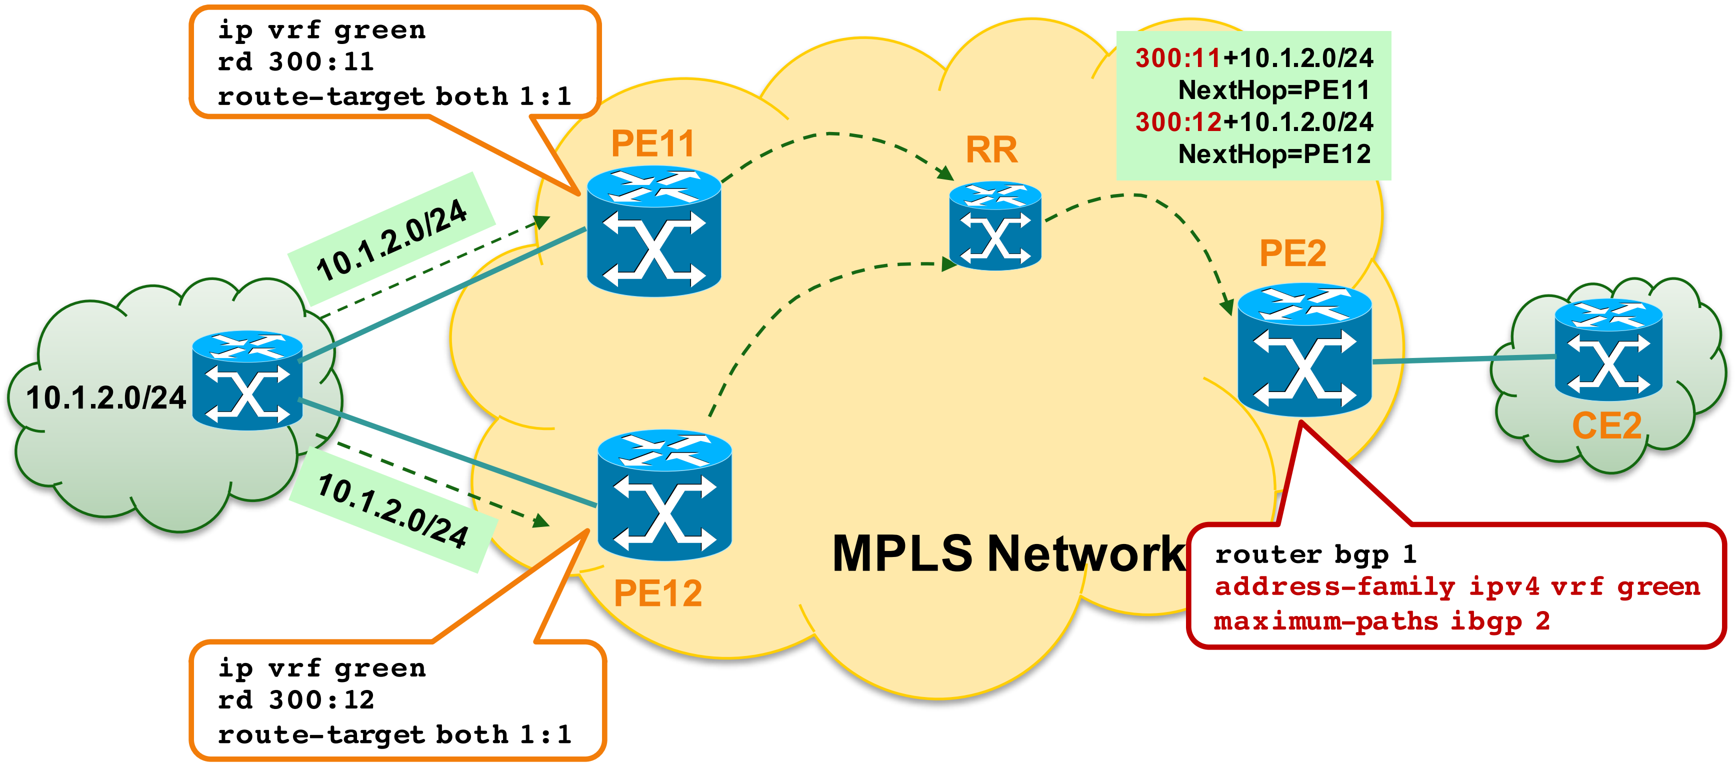

PE-CE Link Failure

- After detecting the PE-CE link failure, PE11 sends BGP message to withdraw the VPN routes, traffic will be dropped on PE11 before PE2 completes BGP route convergence.

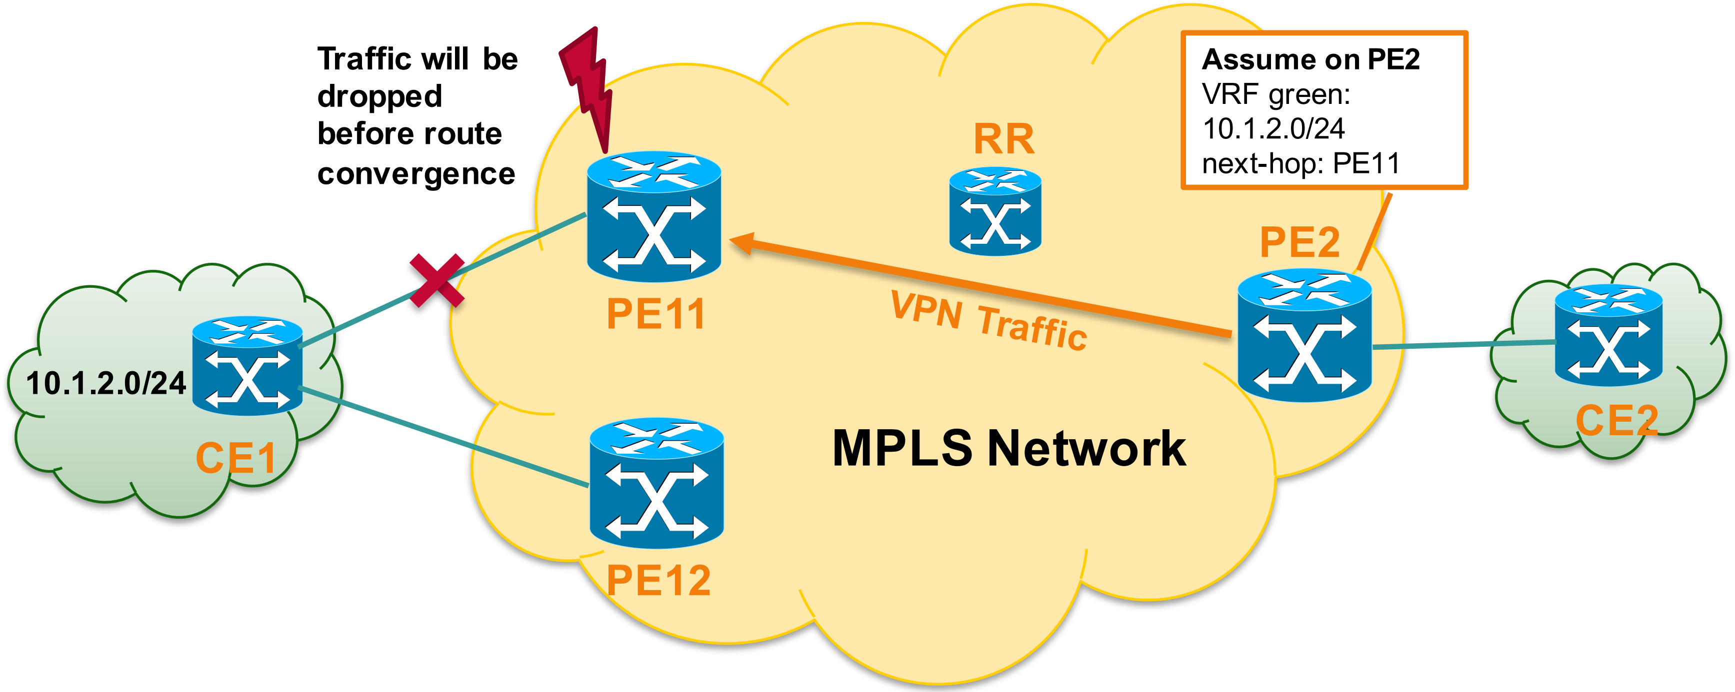

VPN Fast Convergence – PIC Edge

- Use **PIC Edge** feature to minimize the loss due to the PE-CE link failure from sec to msec.
- **P**refix **I**ndependent **C**onvergence is a method for speeding up convergence of the FIB under failover conditions.

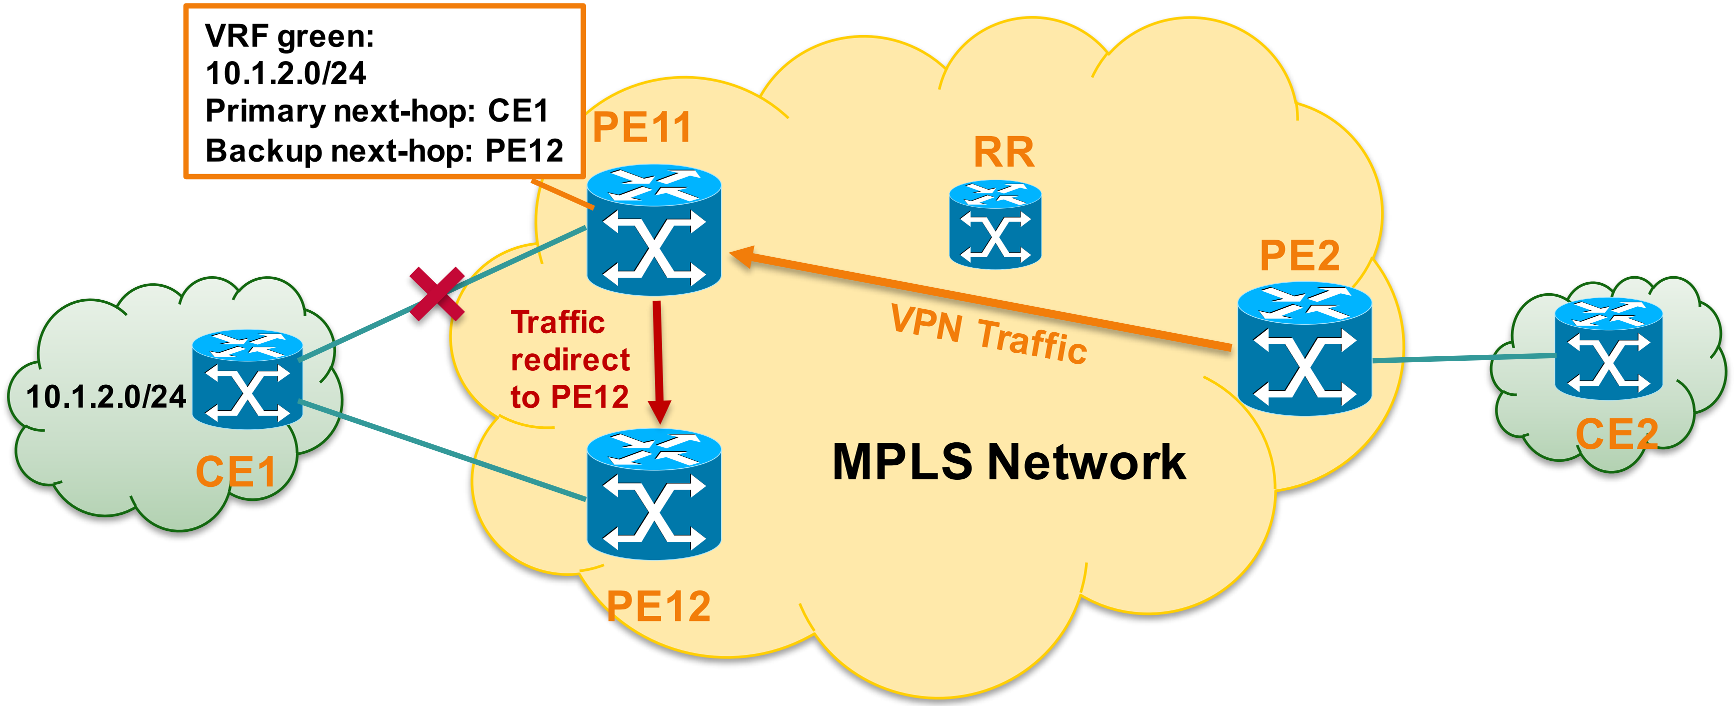

PE Node Failure

- When PE11 router fails, traffic will be lost before PE2 completes BGP route convergence.

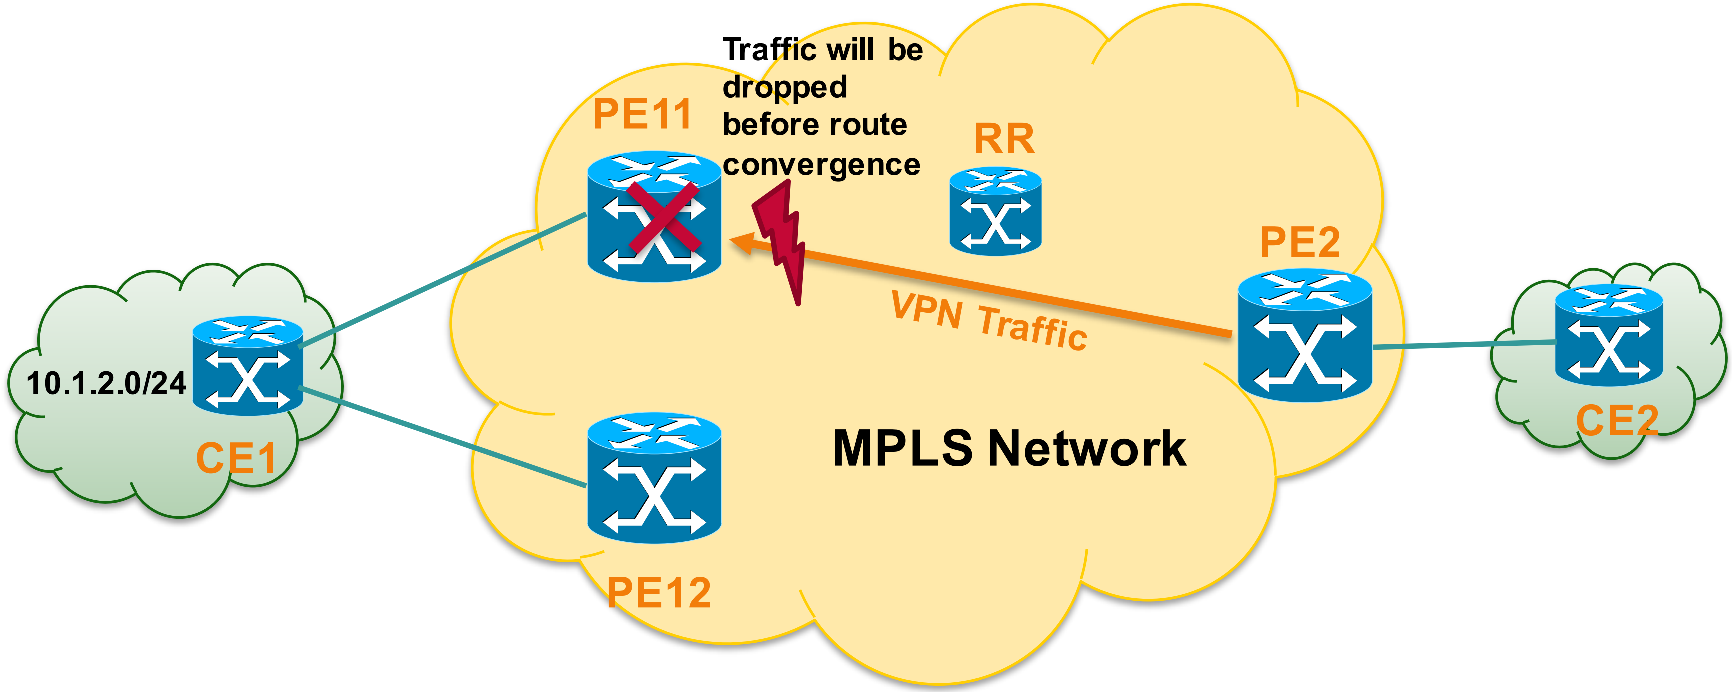

VPN Fast Convergence – PIC Edge

- PE2 uses the alternative VPN route for forwarding until global convergence is complete, this reduces traffic loss.

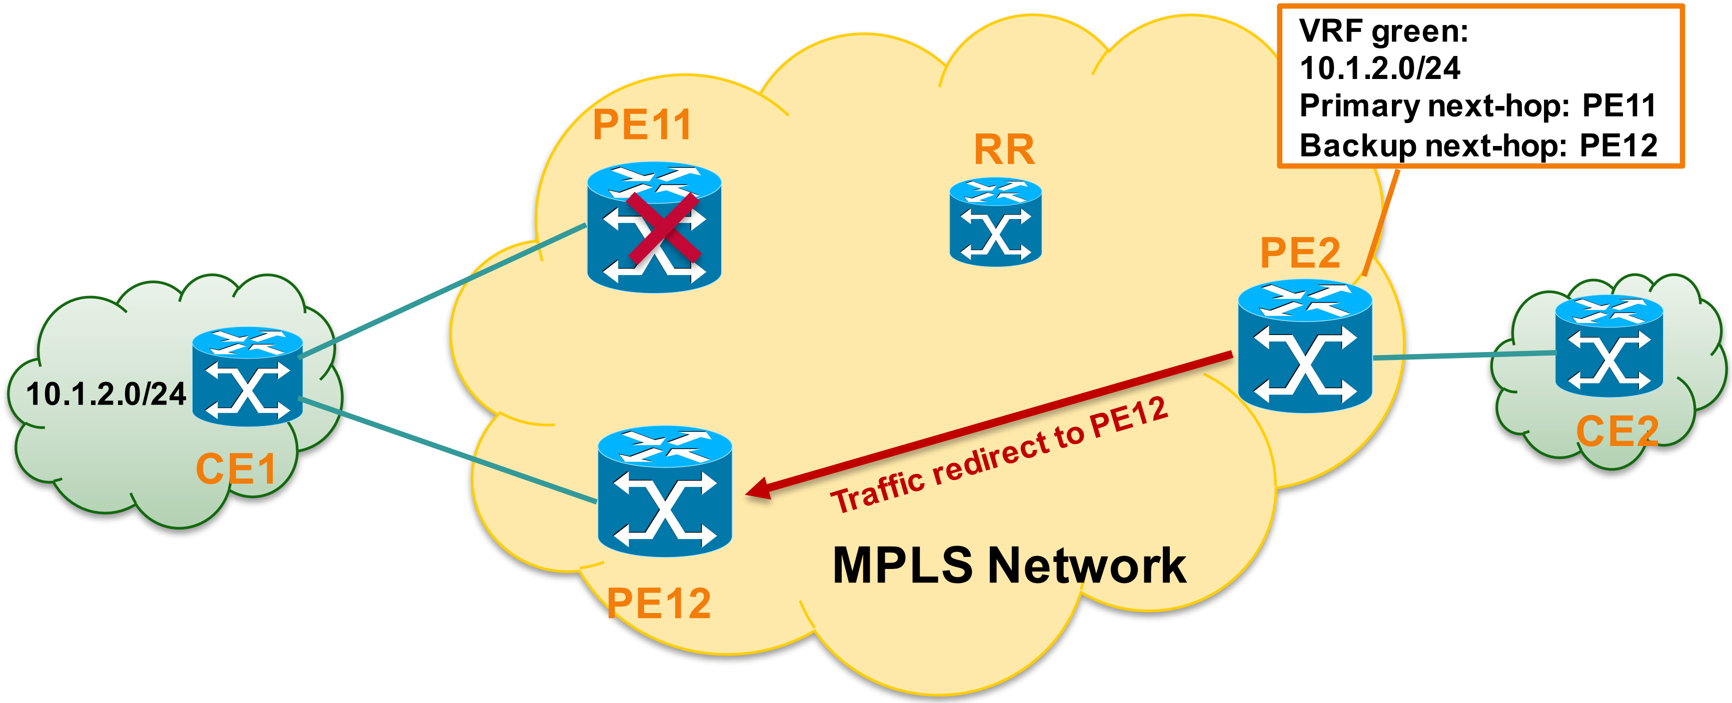

MPLS L3VPN Services

- Multi-homed VPN Sites
- **Hub and Spoke Service**
- Extranet Service
- Internet Access Service

Hub and Spoke Service

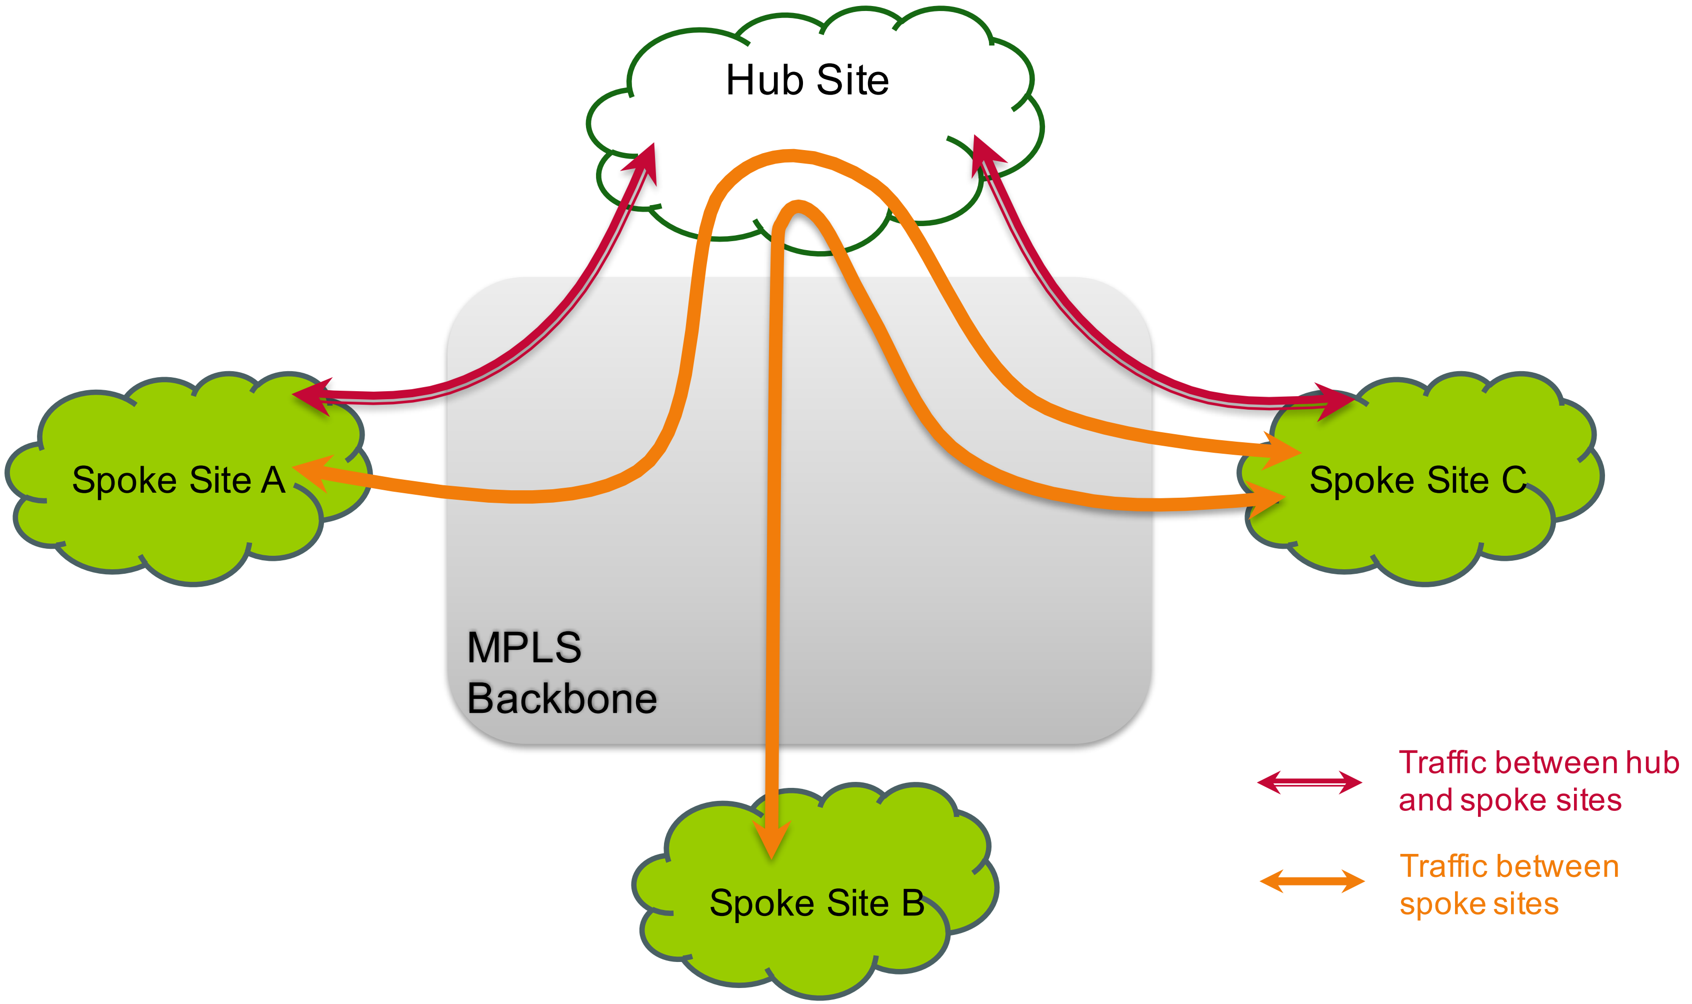

Option 1 - Single Interface

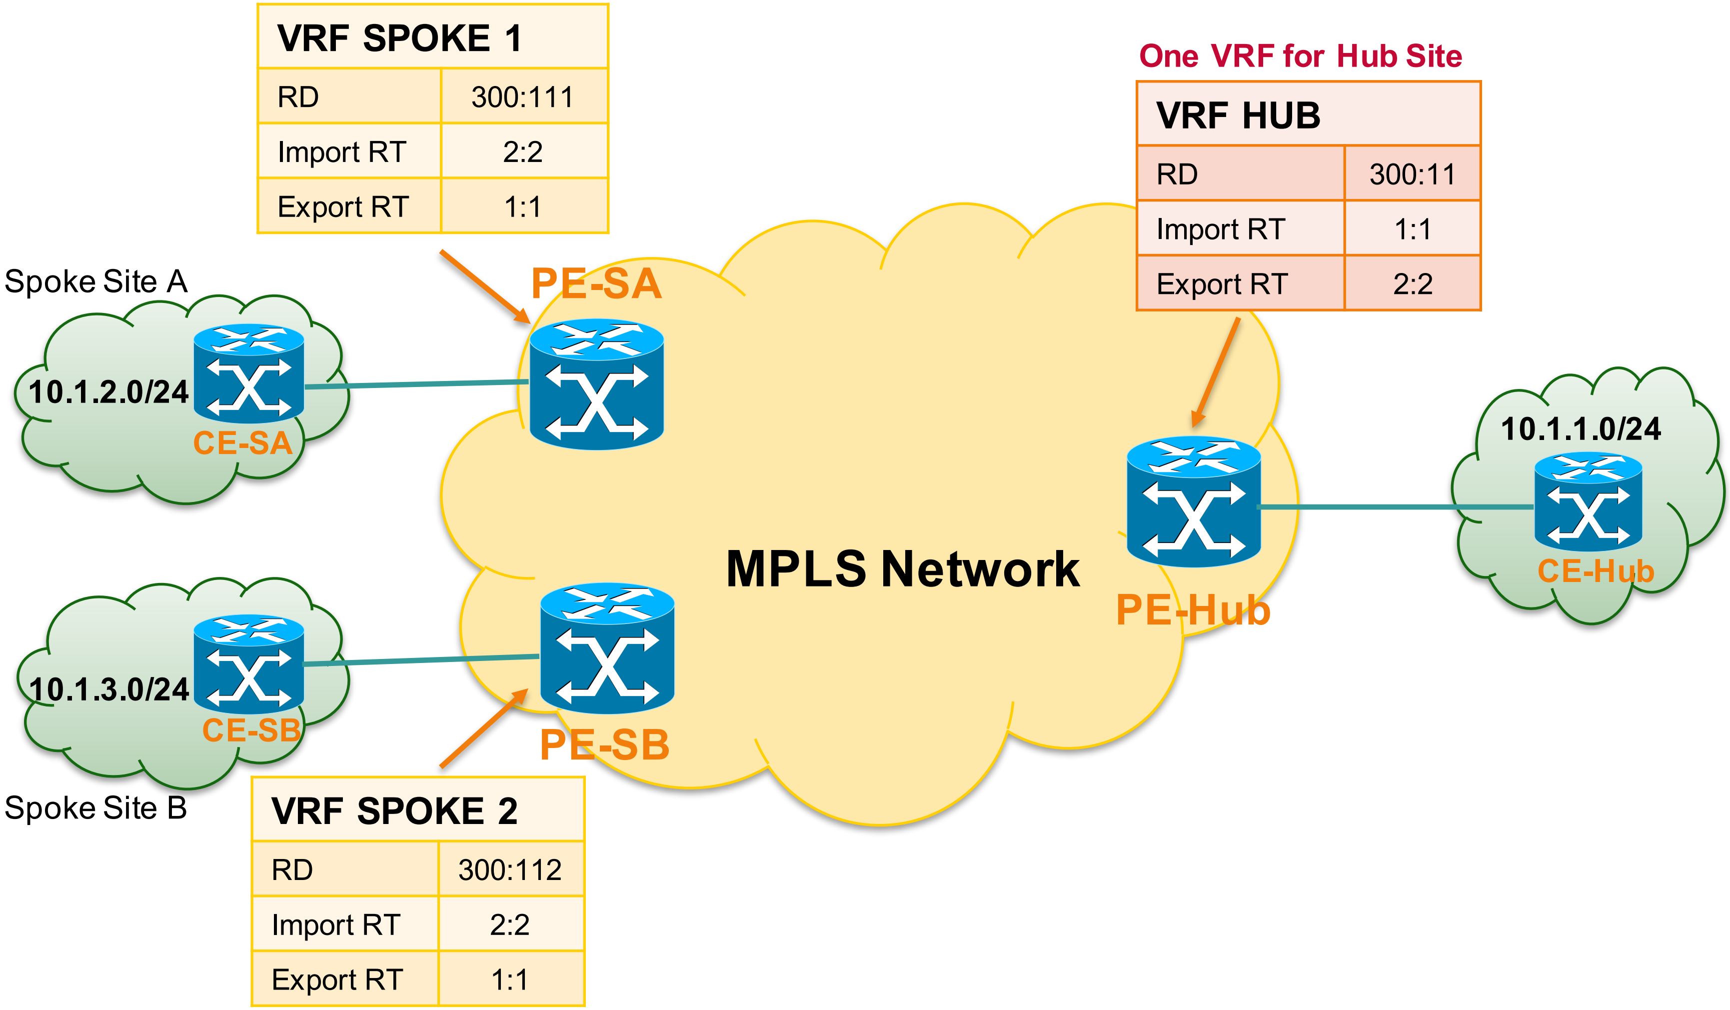

VRF SPOKE 1	
RD	300:111
Import RT	2:2
Export RT	1:1

VRF HUB	
RD	300:11
Import RT	1:1
Export RT	2:2

VRF SPOKE 2	
RD	300:112
Import RT	2:2
Export RT	1:1

Control Plane – from Spoke to Hub

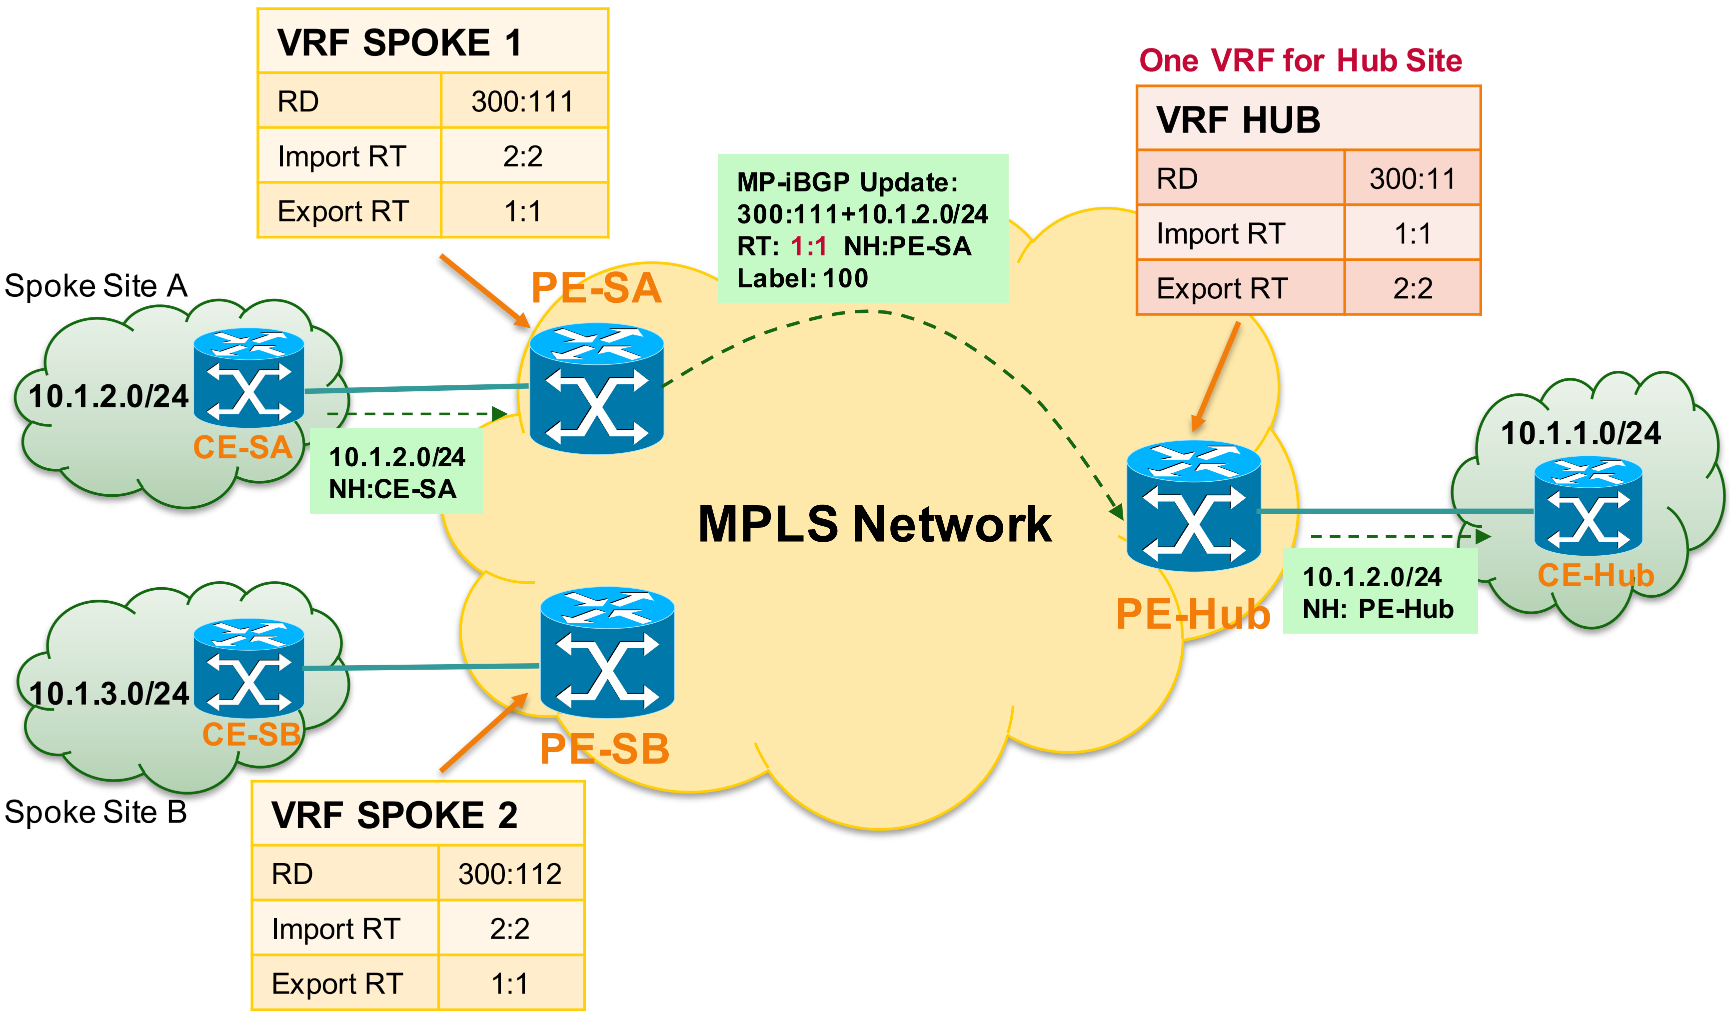

Control Plane – from Hub to Spoke

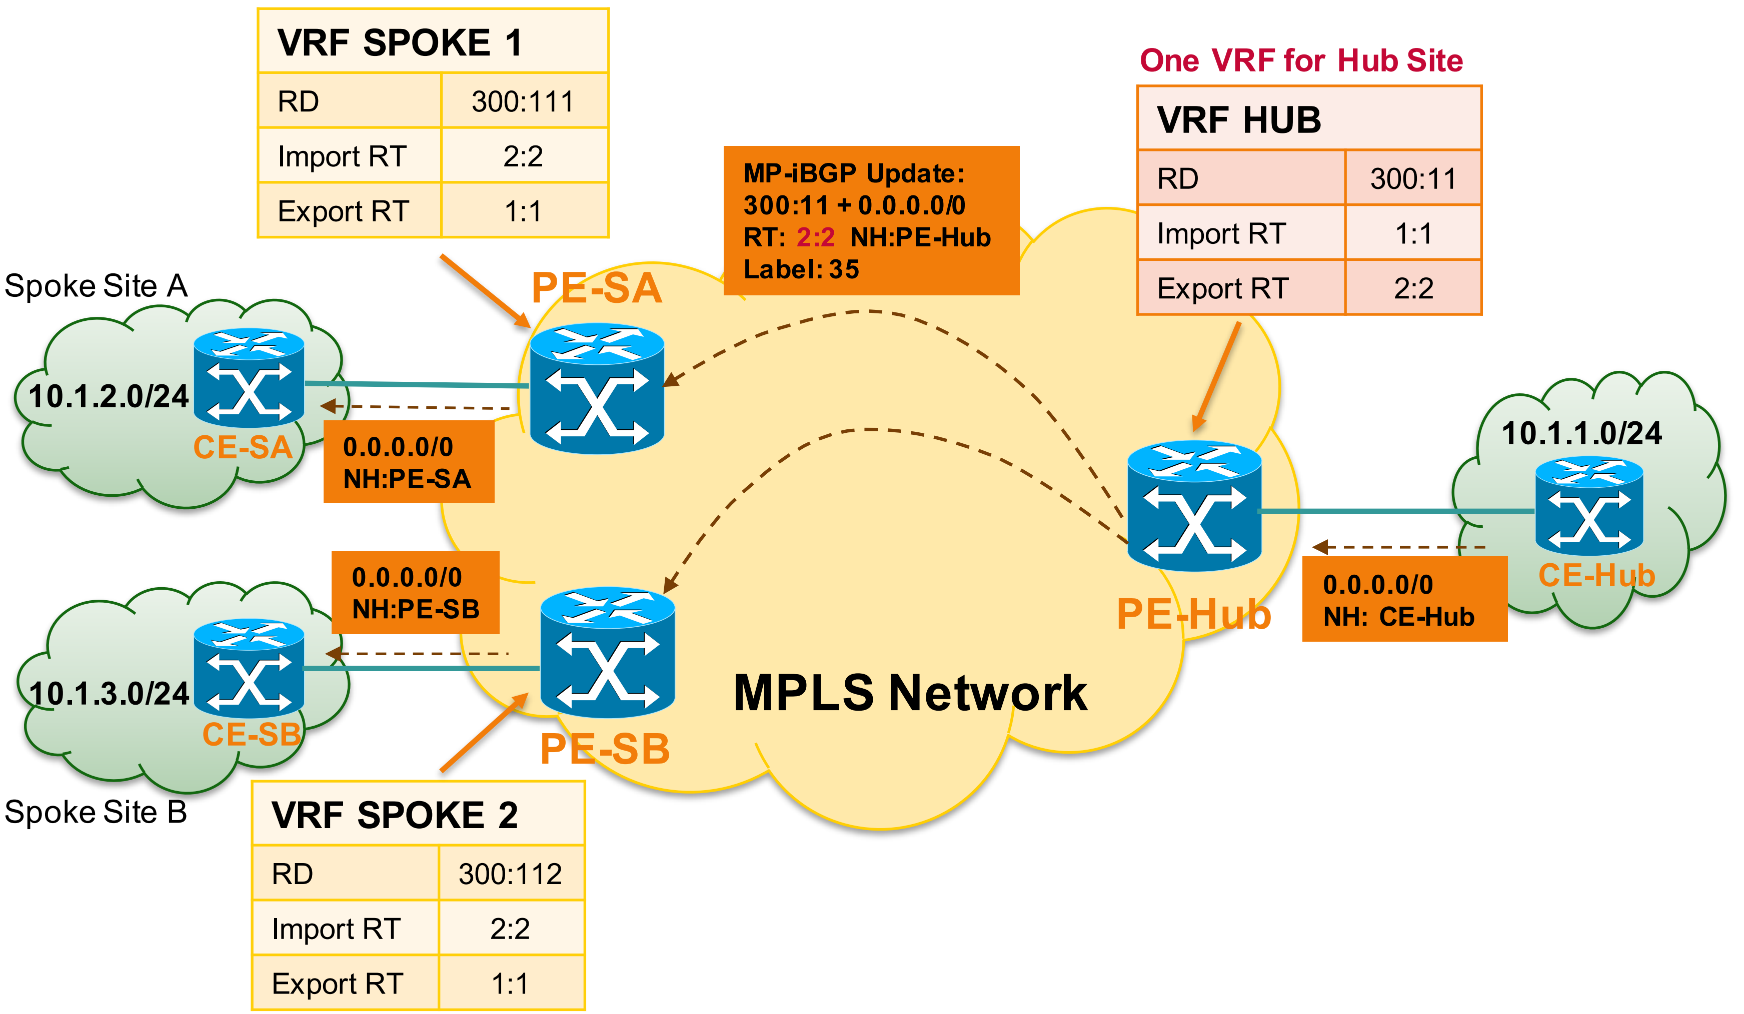

Data Plane – Traffic between Spoke Sites

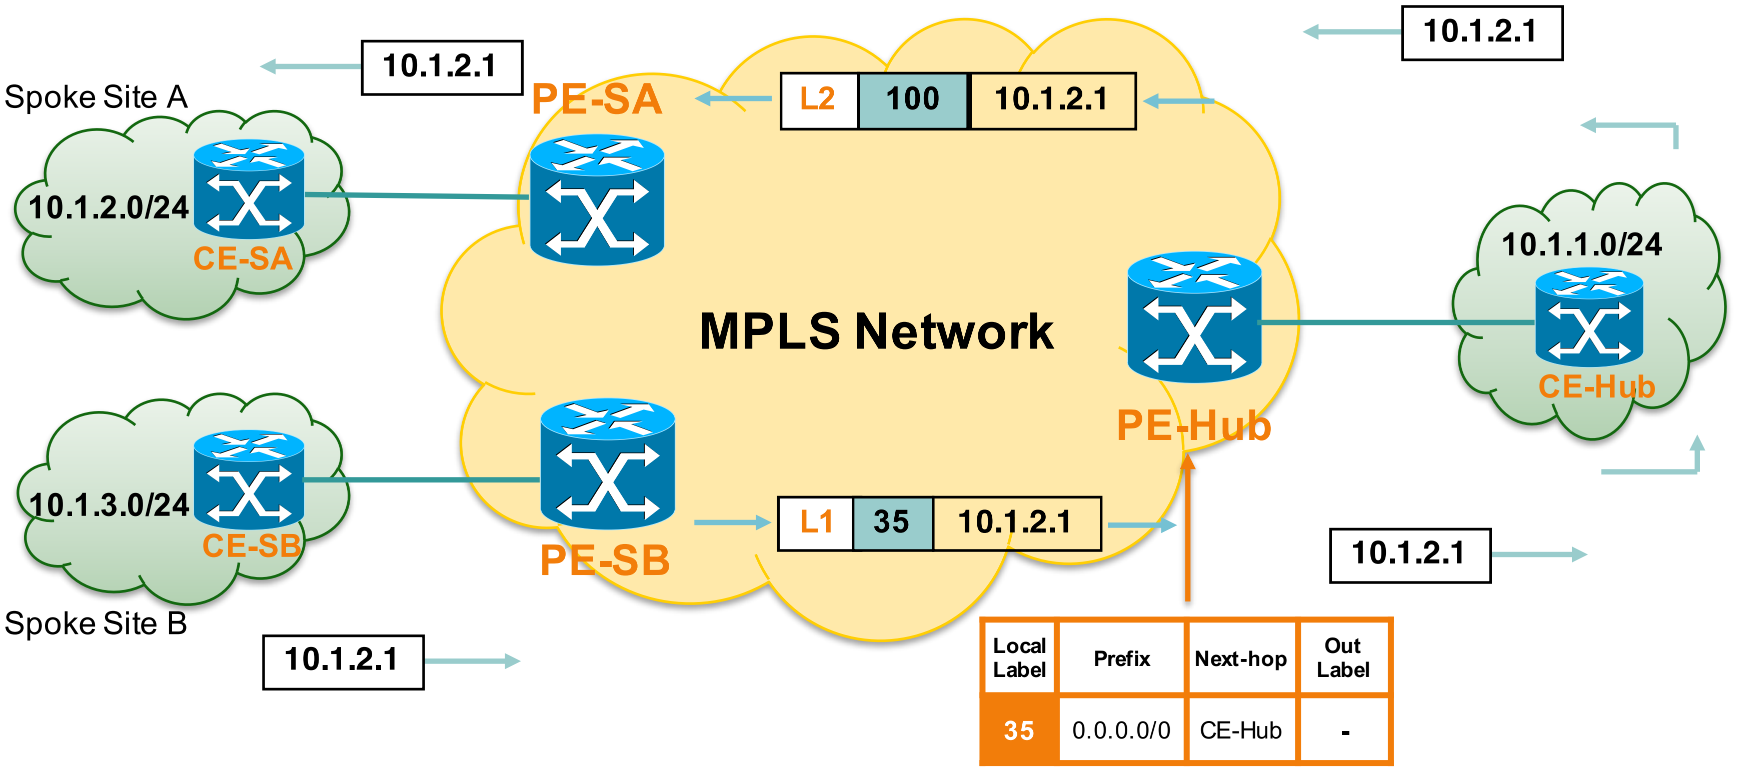

Local Label	Prefix	Next-hop	Out Label
35	0.0.0.0/0	CE-Hub	-

Option 2 – Two Interfaces

If more specific spoke CE routes need to be exchanged between spoke CE routers, option 2 can be selected.

VRF SPOKE 1	
RD	300:111
Import RT	2:2
Export RT	1:1

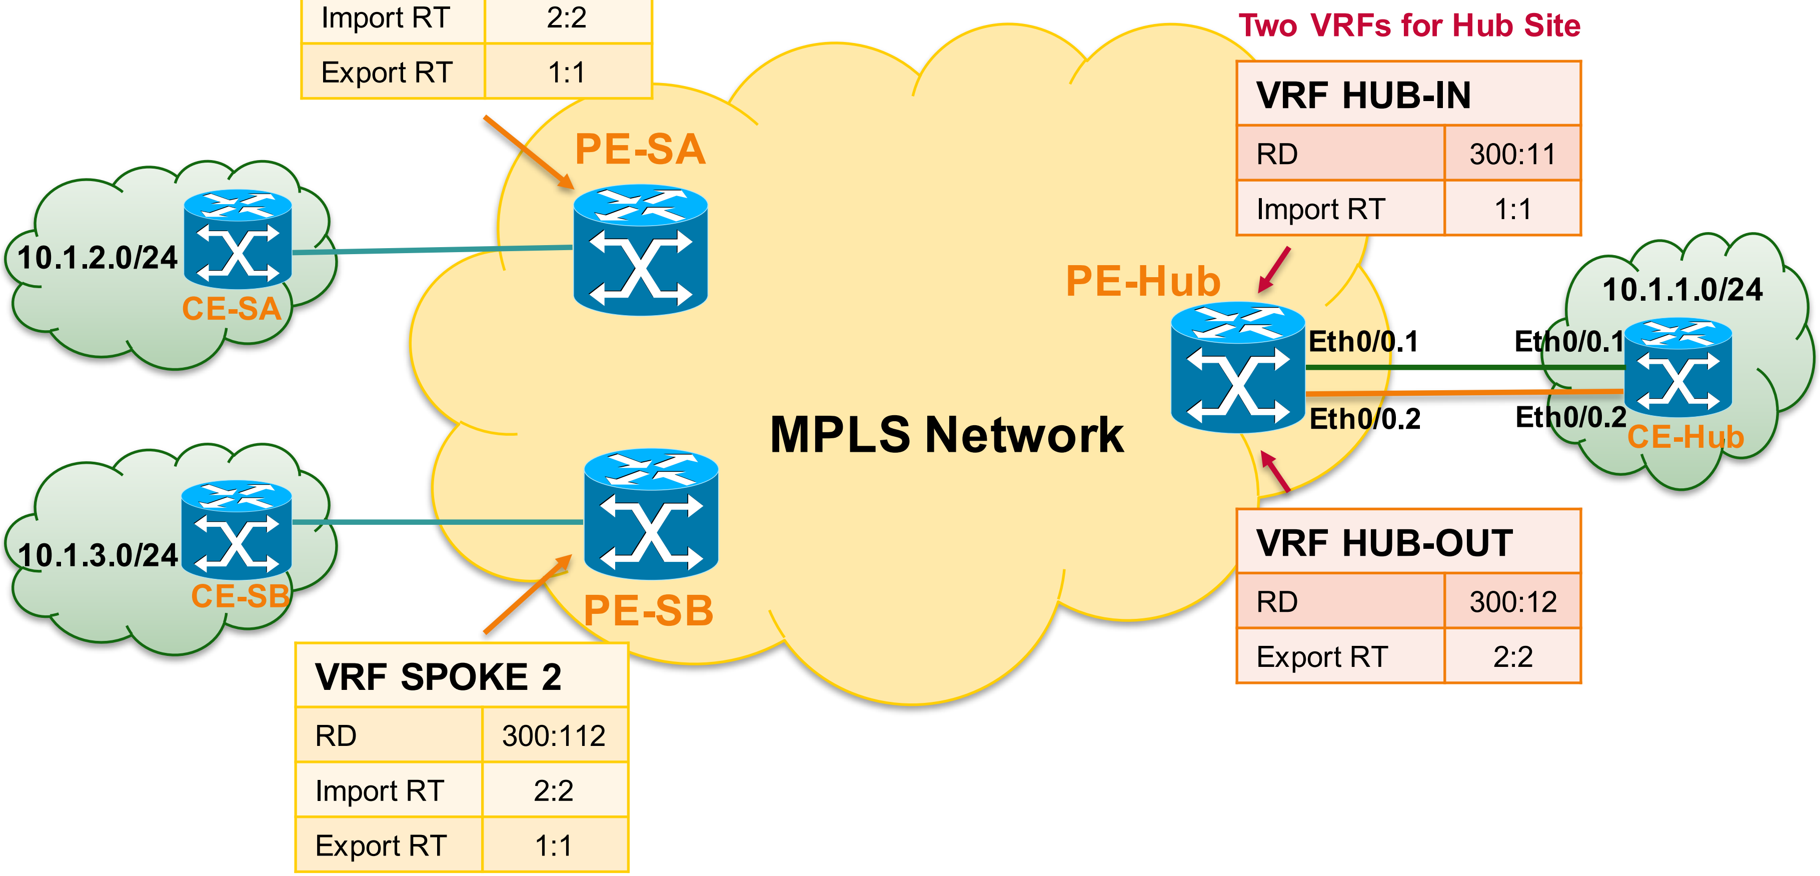

VRF HUB-IN	
RD	300:11
Import RT	1:1

VRF HUB-OUT	
RD	300:12
Export RT	2:2

VRF SPOKE 2	
RD	300:112
Import RT	2:2
Export RT	1:1

Option 2 – Control Plane (Hub in)

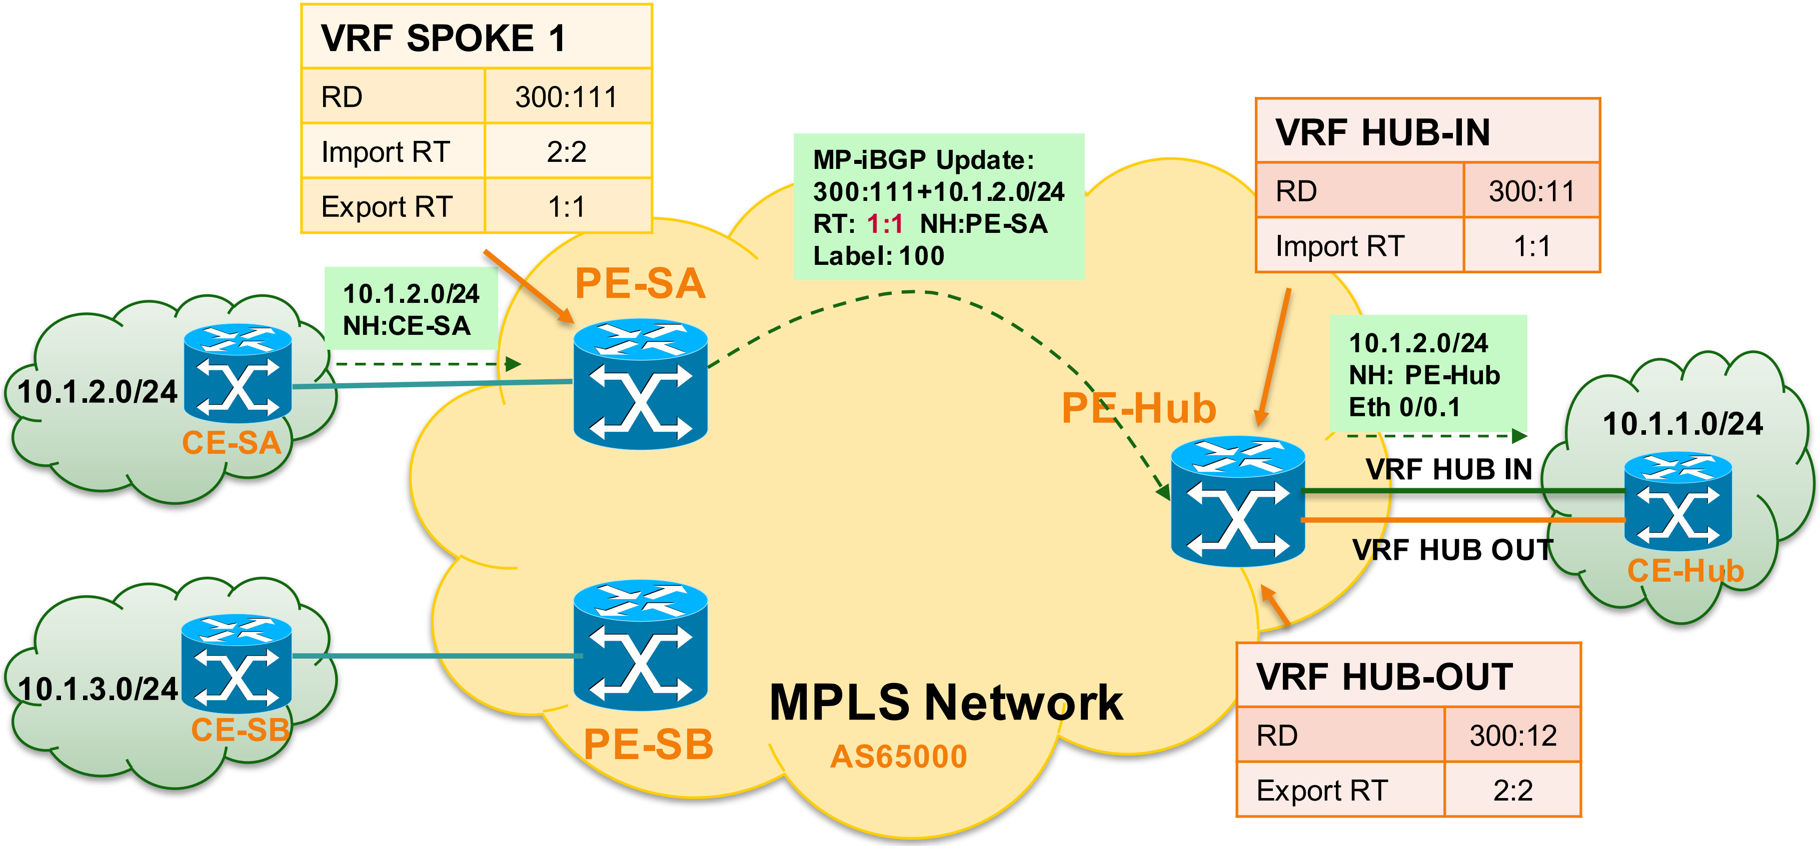

Option 2 – Control Plane (Hub out)

- Deployment of allowas-in feature

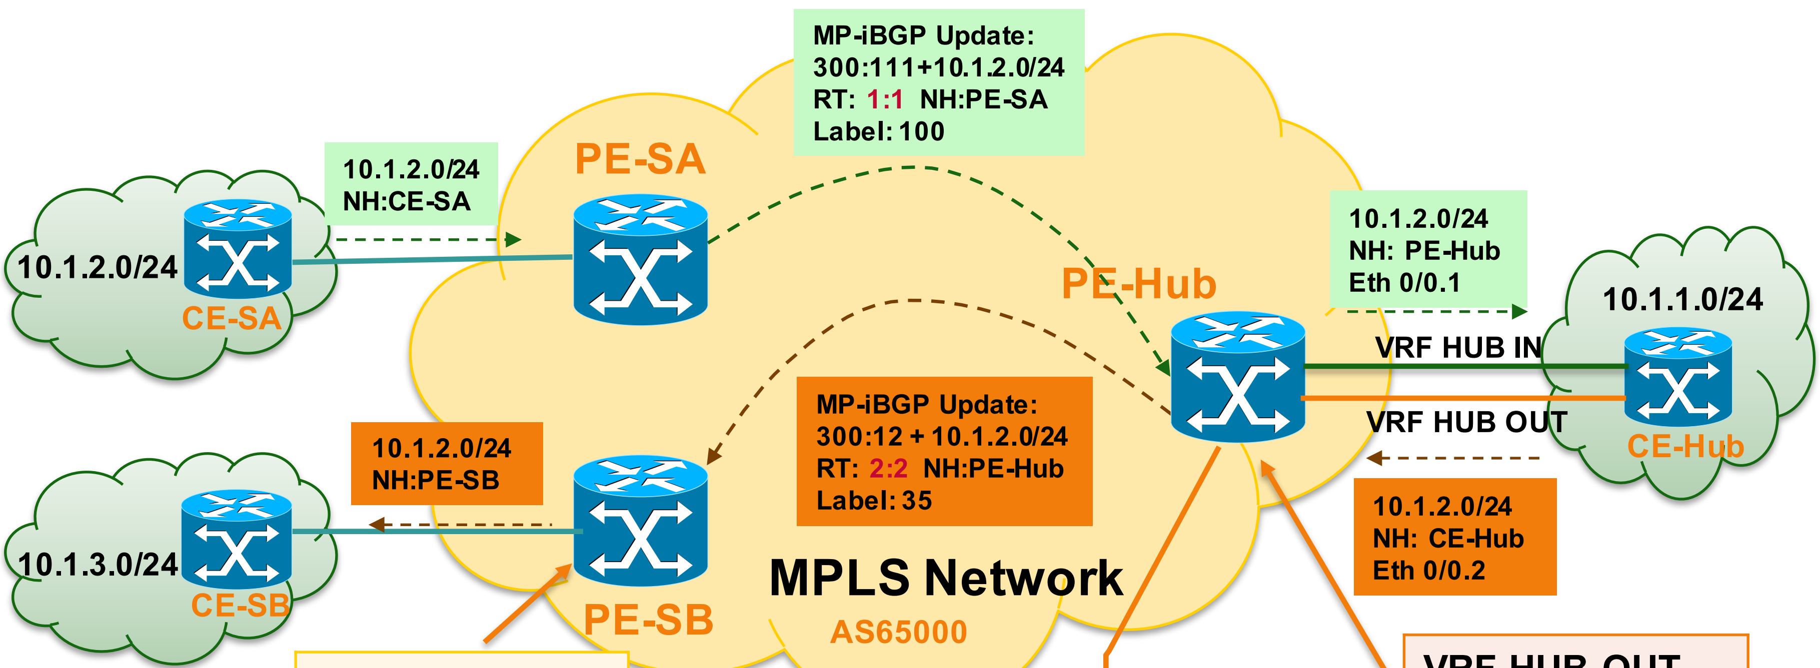

VRF SPOKE 2	
RD	300:112
Import RT	2:2
Export RT	1:1

```
router bgp 65000
 address-family ipv4 vrf HUB-OUT
 neighbor <CE> allowas-in 2
```

VRF HUB-OUT	
RD	300:12
Export RT	2:2

Option 2 – Data Plane

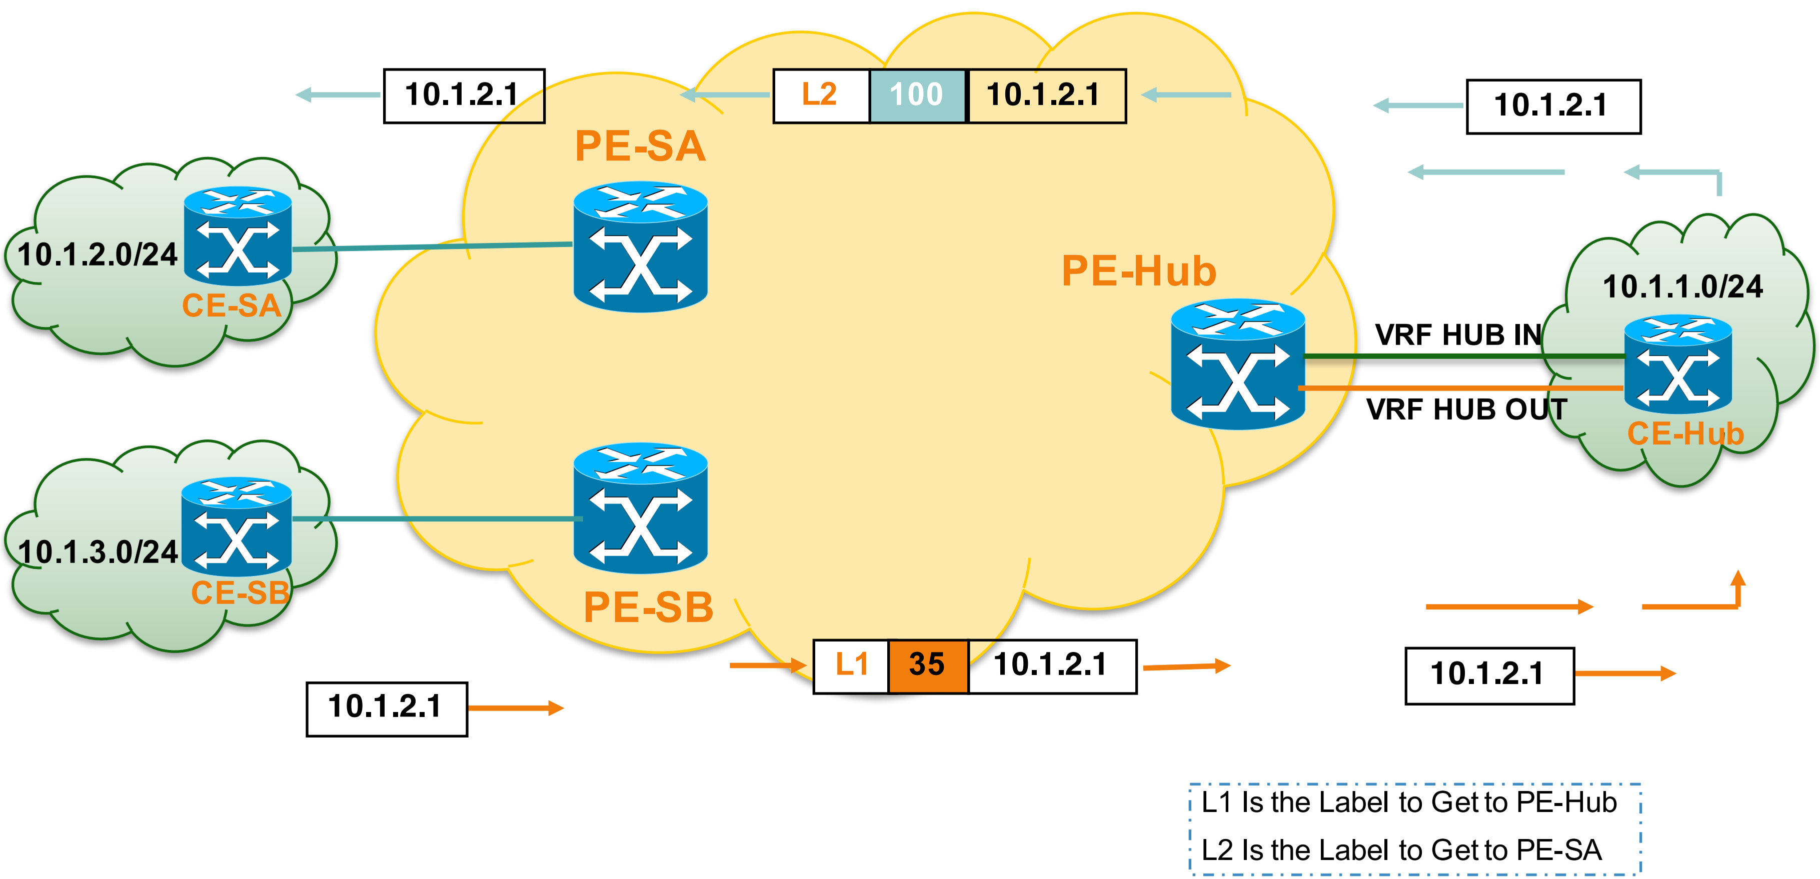

MPLS L3VPN Services

- Multi-homed VPN Sites
- Hub and Spoke Service
- **Extranet Service**
- Internet Access Service

Extranet Service

- Communication between VPNs may be required i.e., External intercompany communication (dealers with manufacturer, retailer with wholesale provider, etc.)

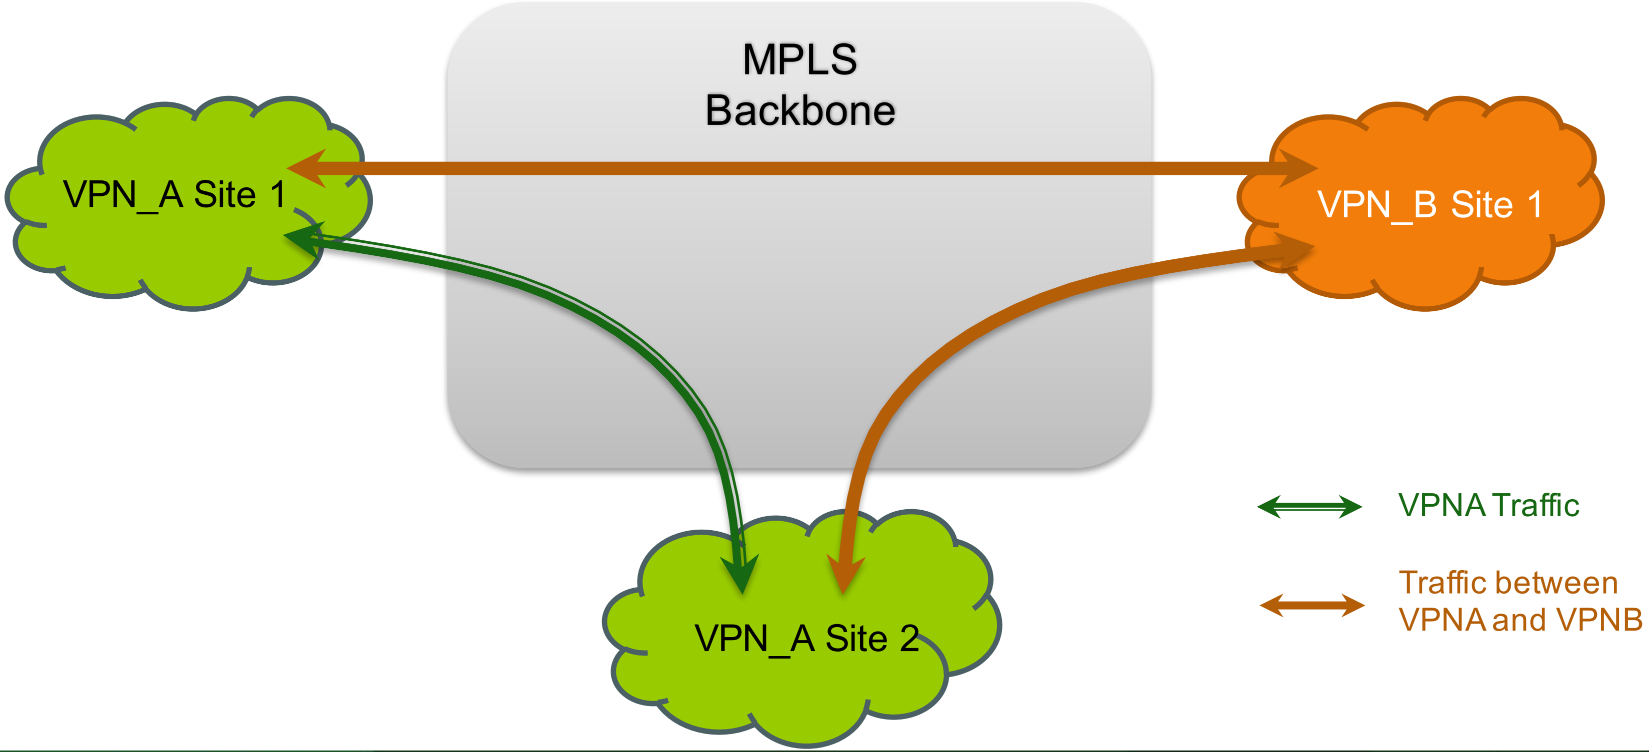

Extranet VPN – Simple Extranet

- Designing RT to implement the communication.

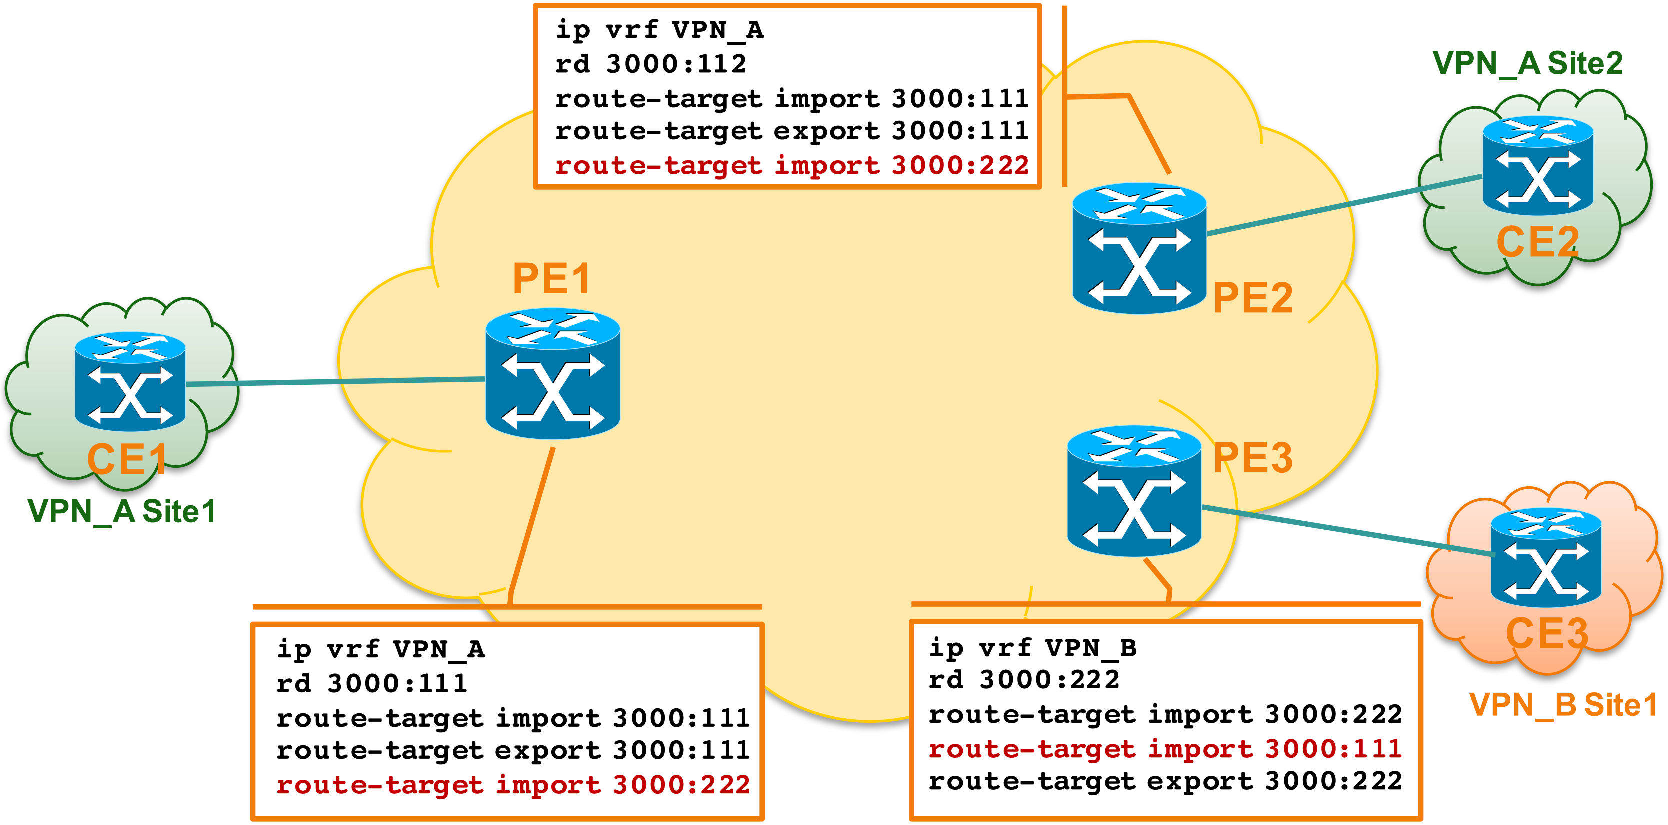

More Complex Scenario

- If only allow VPNB Site1 to communicate with the servers in VPNA Site1.

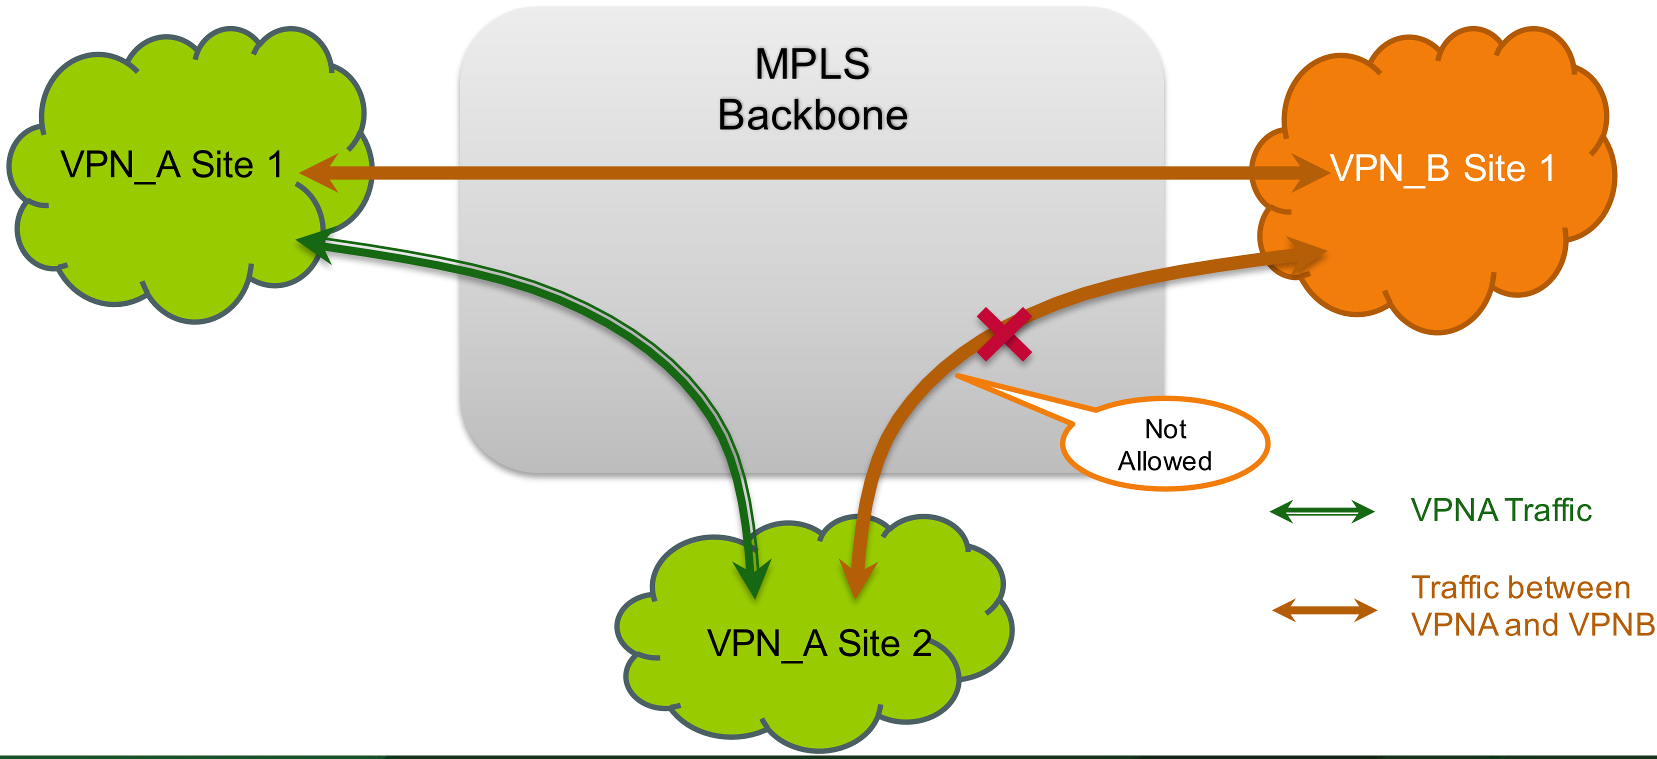

Extranet VPN – Advanced Extranet

```
ip vrf VPN_A
rd 3000:112
route-target import 3000:111
route-target export 3000:111
```

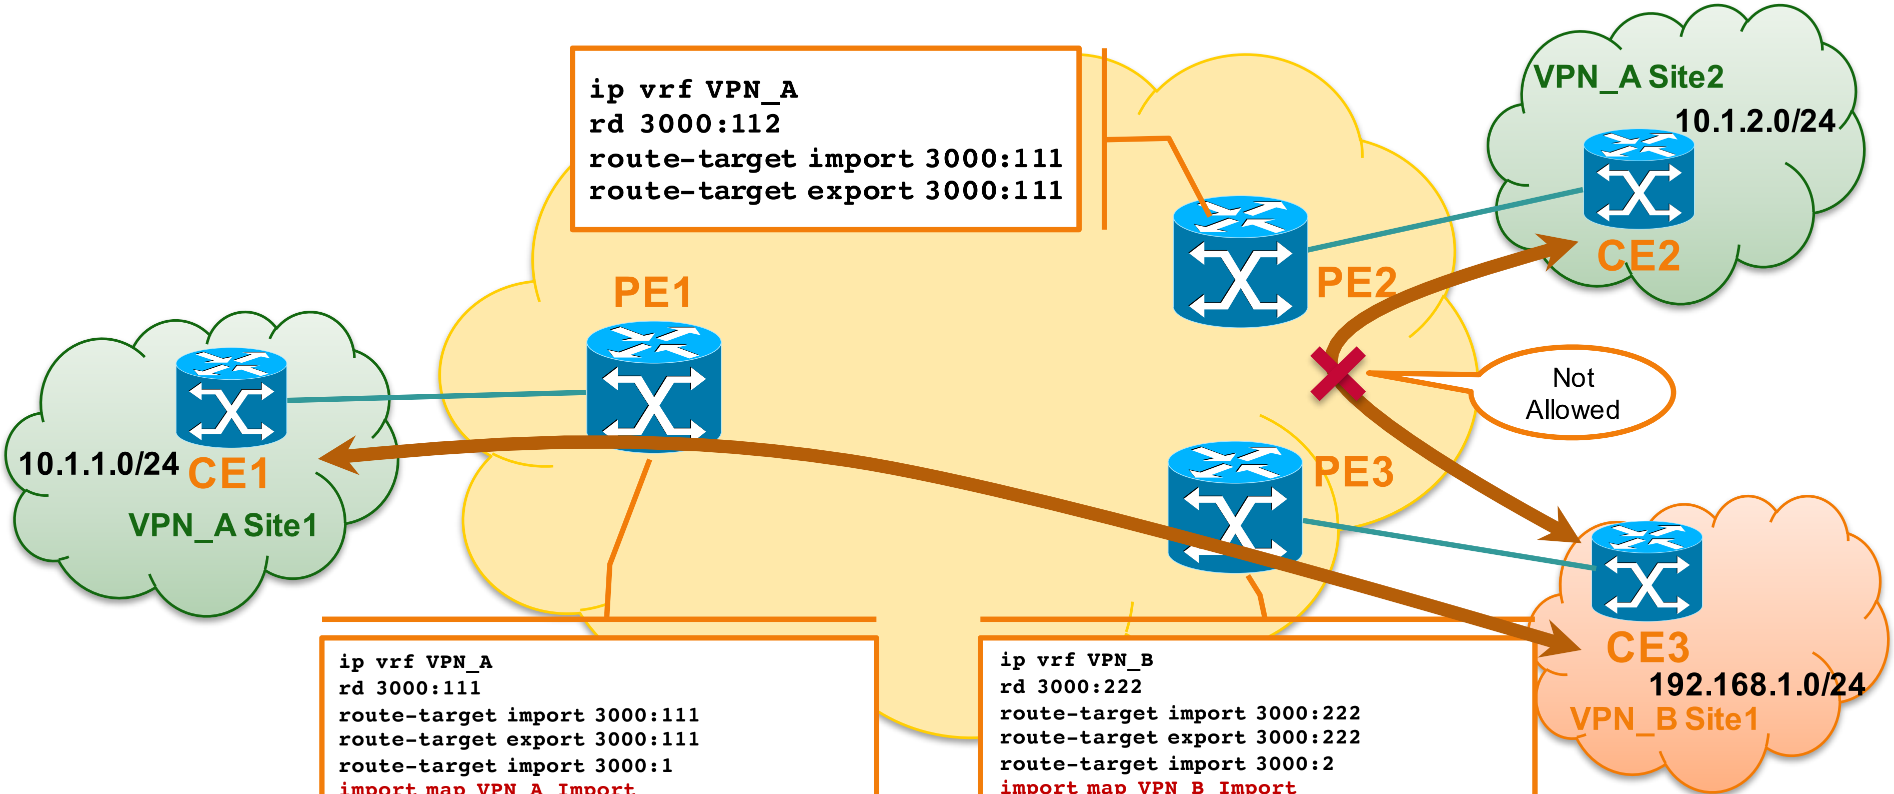


```
ip vrf VPN_A
rd 3000:111
route-target import 3000:111
route-target export 3000:111
route-target import 3000:1
import map VPN_A_Import
export map VPN_A_Export
!
route-map VPN_A_Export permit 10
 match ip address 1
 set extcommunity rt  3000:2 additive
!
route-map VPN_A_Import permit 10
 match ip address 2
!
access-list 1 permit 10.1.1.0.0 0.0.0.255
access-list 2 permit 192.168.1.0 0.0.0.255
```

```
ip vrf VPN_B
rd 3000:222
route-target import 3000:222
route-target export 3000:222
route-target import 3000:2
import map VPN_B_Import
export map VPN_B_Export
!
route-map VPN_B_Export permit 10
 match ip address 2
 set extcommunity rt  3000:1 additive
!
route-map VPN_B_Import permit 10
 match ip address 1
!
access-list 1 permit 10.1.1.0.0 0.0.0.255
access-list 2 permit 192.168.1.0 0.0.0.255
```

Lack of 'Additive' Would Result in 3000:222 Being Replaced with 3000:1. We Don't Want That.

MPLS L3VPN Services

- Multi-homed VPN Sites
- Hub and Spoke Service
- Extranet Service
- **Internet Access Service**

Internet Access Service to VPN Customers

- Internet access service could be provided as another value-added service to VPN customers
- Security mechanism must be in place at both provider network and customer network
  - To protect from the Internet vulnerabilities

Internet Access: Design Options

1. VRF Specific Default Route
2. Separate PE-CE Sub-interfaces
3. Extranet with Internet-VRF

Option 1: VRF Specific Default Route

```
ip route 71.8.0.0 255.255.0.0 Serial0

Router bgp 65000
 no bgp default ipv4-unicast
 network 71.8.0.0 mask 255.255.0.0
 neighbor 192.168.1.1 remote 65000
 neighbor 192.168.1.1 activate
 neighbor 192.168.1.1 next-hop-self
 neighbor 192.168.1.1 update-source loopback0
```

② Add the static route pointing to VRF interface;
Announce it to neighbors.
For traffic (Internet→VPN)

PE2: Routing Table

Destination	Label/Interface
71.8.0.0/16	192.168.1.2

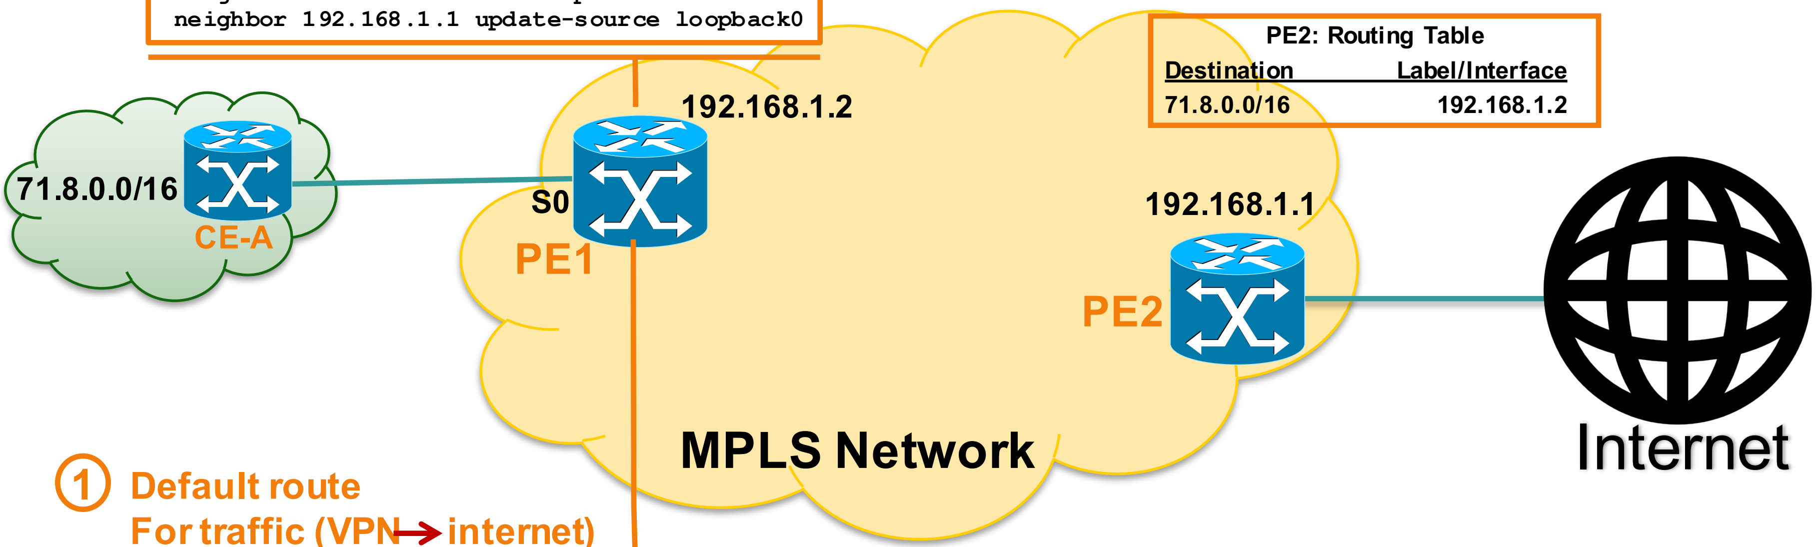

① **Default route**
For traffic (VPN→internet)

```
ip route vrf VPN-A 0.0.0.0 0.0.0.0 192.168.1.1 global
```

Option 1: Data Plane

PE1: Global Routing/FIB Table

Destination	Label/Interface
192.168.1.1/32	Label=30
71.8.0.0/16	Serial 0

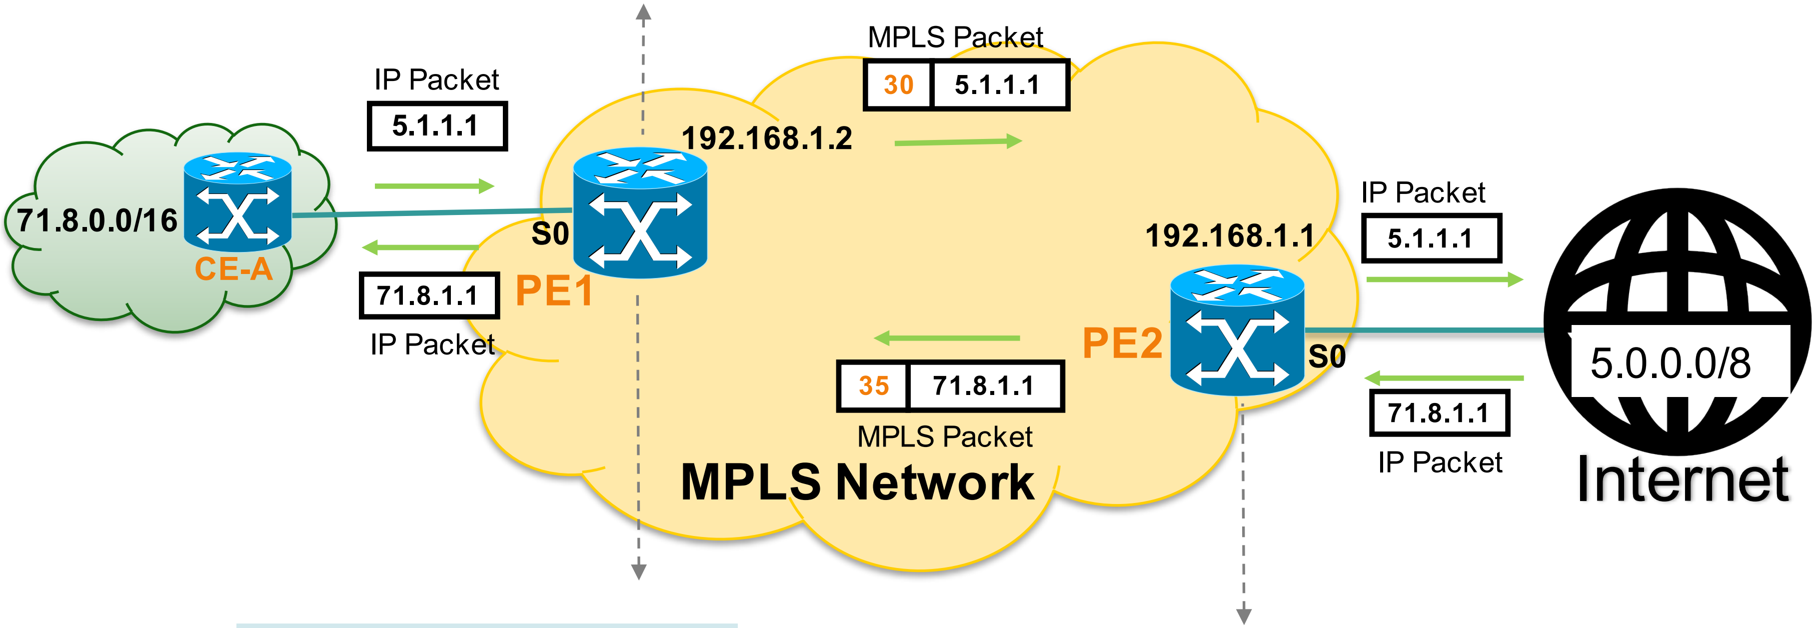

PE1: VRF Routing/FIB Table

Destination	Label/Interface
0.0.0.0/0	192.168.1.1 (Global)
Site-1	Serial 0

PE2: Global Table and LFIB

Destination	Label/Interface
192.168.1.2/32	Label=35
71.8.0.0/16	192.168.1.2
5.1.0.0/16	Serial 0

Option 2: Separate PE-CE Sub-interfaces

One sub-interface associated to VRF

```
Interface Ethernet0/0.1
 encapsulation dot1q 100
 ip vrf forwarding VPN-A
 ip address 192.168.20.1 255.255.255.0
```

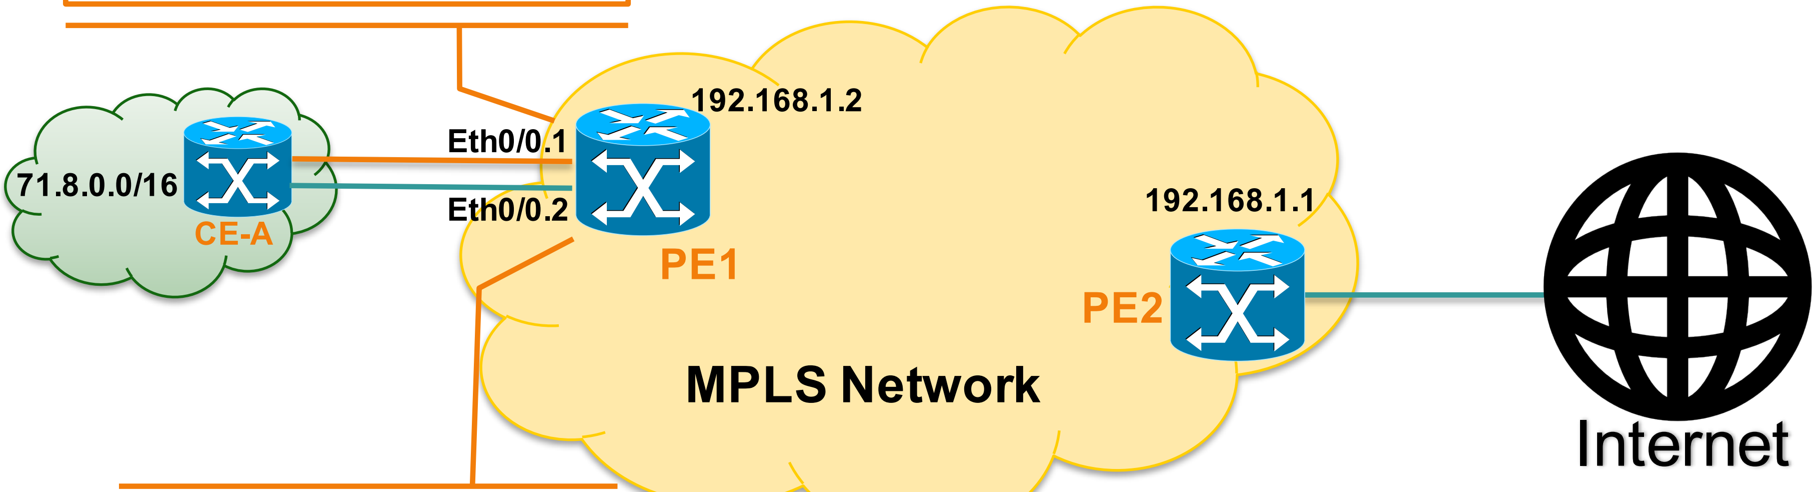


```
Interface Ethernet0/0.2
 encapsulation dot1q 200
 ip address 71.8.10.1 255.255.0.0
```

One sub-interface (global) for Internet routing

Option 2: Data Plane

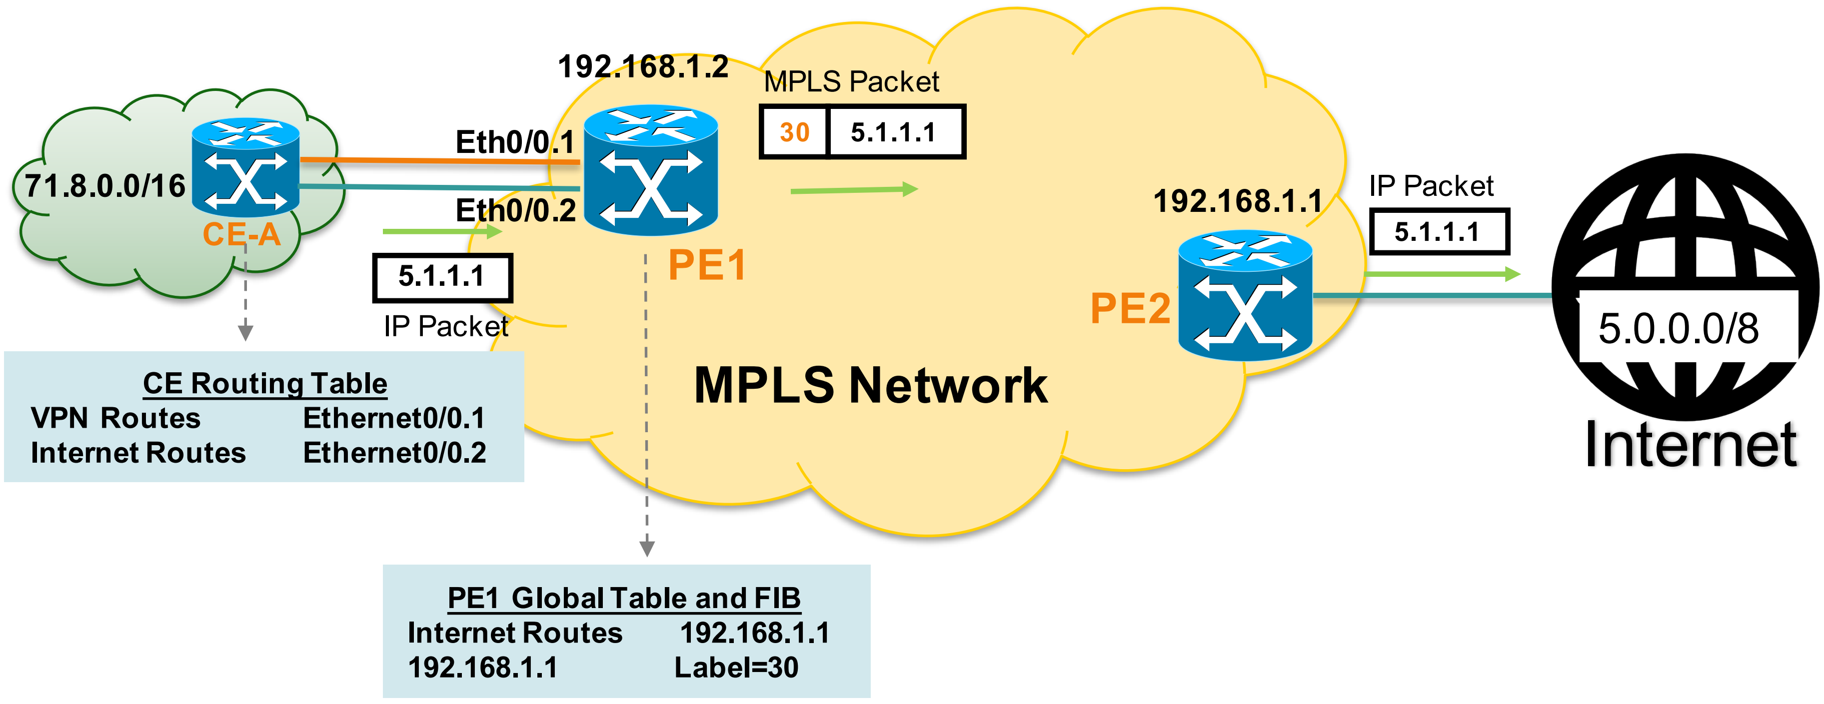

Option 3: Extranet with Internet-VRF

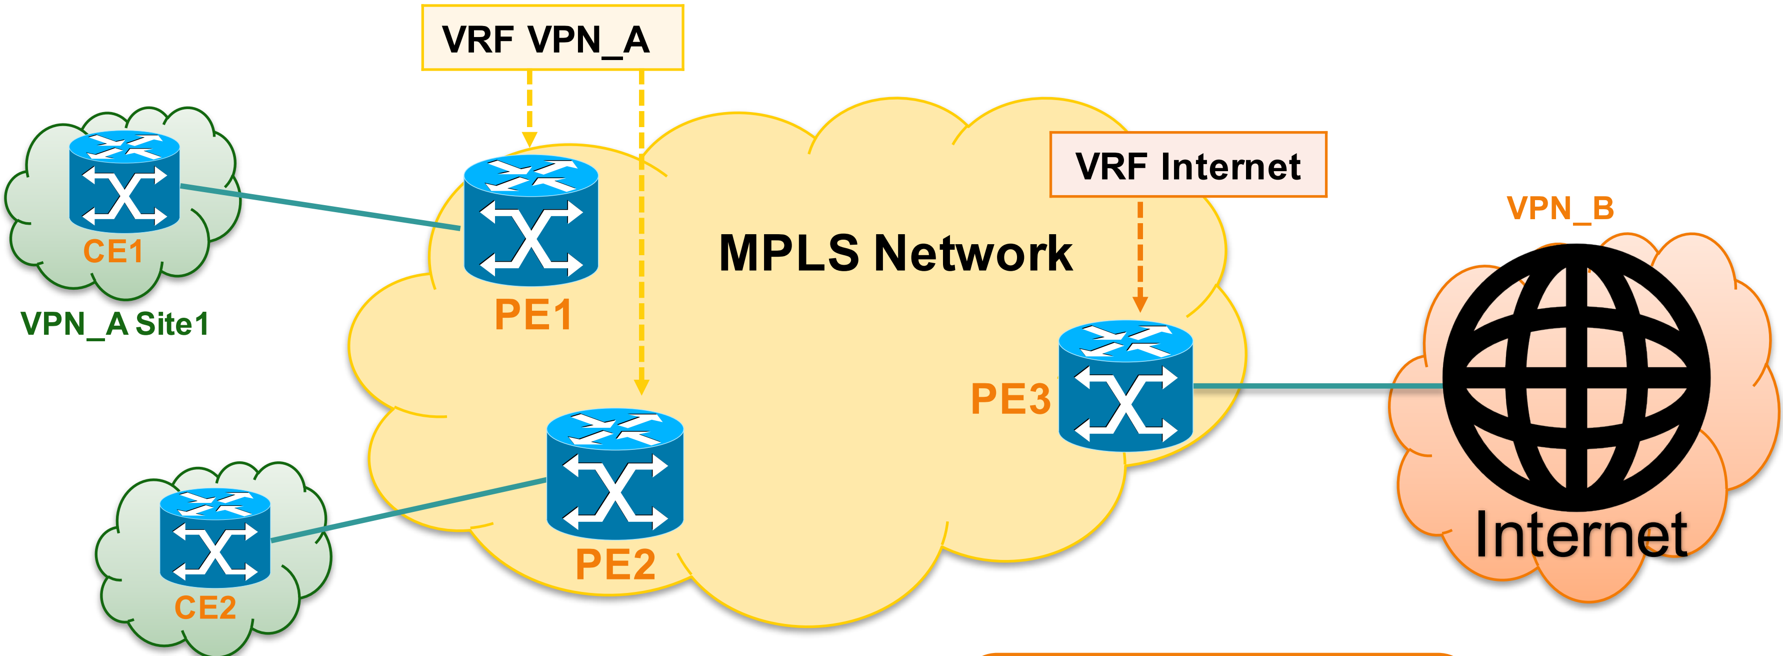

Design RT to implement the VRF communication

Option 3: Extranet with Internet-VRF

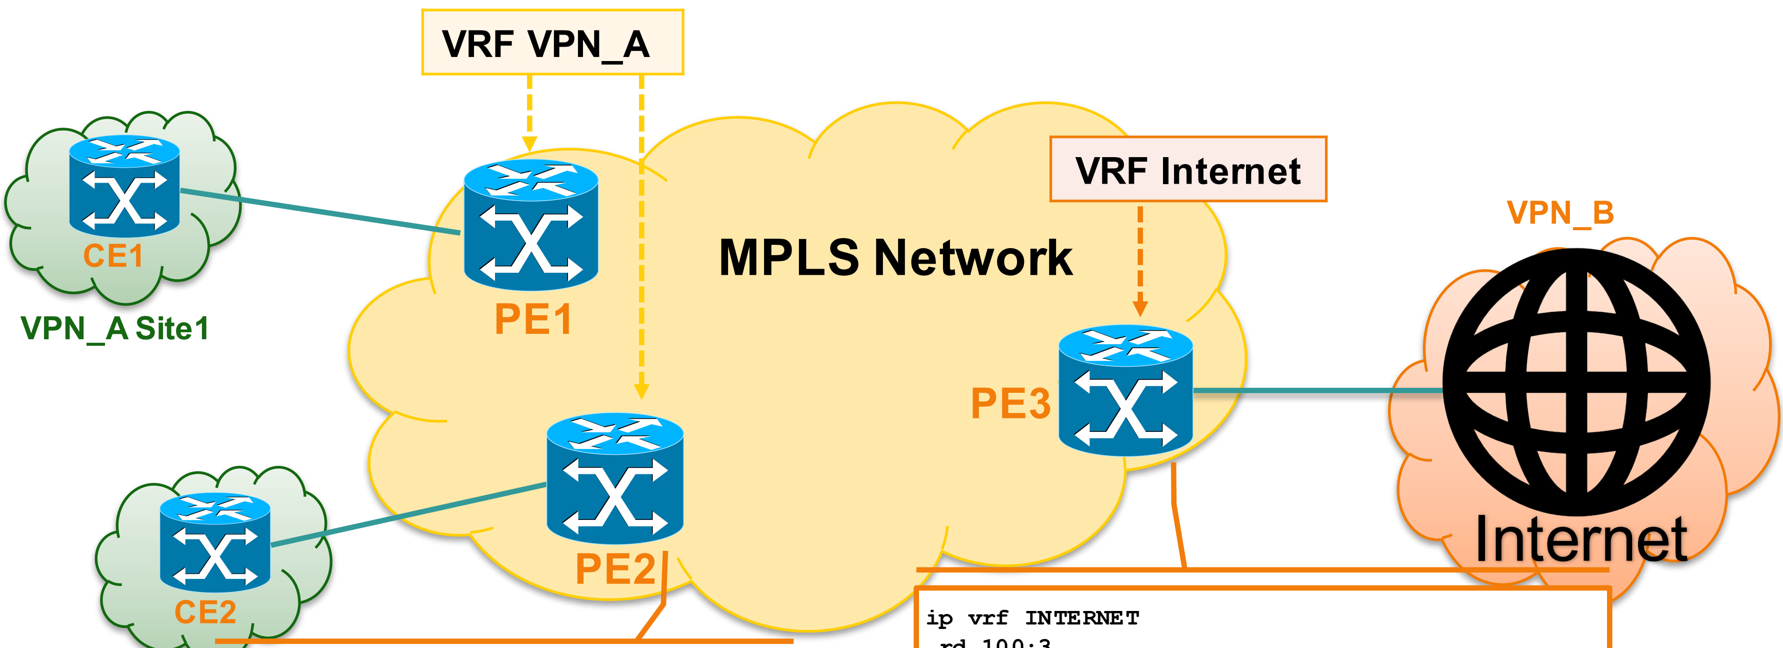


```
ip vrf VPNA
 rd 100:2
 route-target export 1:10
 route-target import 1:10
 route-target import 100:10
```

```
ip vrf INTERNET
 rd 100:3
 route-target export 100:10
 route-target import 1:10

router bgp 100
 address-family ipv4 vrf INTERNET
 network 0.0.0.0 0.0.0.0

ip route vrf INTERNET 0.0.0.0 0.0.0.0 200.1.1.2
```

Best Practice (1)

1. Use RR to scale BGP; deploy RRs in pair for the redundancy
   - Keep RRs out of the forwarding paths and disable CEF (saves memory)
2. Consider unique RD per VRF per PE,
   - Helpful for many scenarios such as multi-homing, hub&spoke etc.
3. Utilize SP's public address space for PE-CE IP addressing
   - Helps to avoid overlapping; Use /31 subnetting on PE-CE interfaces

Best Practice (2)

4. Limit number of prefixes per-VRF and/or per-neighbor on PE
   - Max-prefix within VRF configuration; Suppress the inactive routes
   - Max-prefix per neighbor (PE-CE) within OSPF/RIP/BGP VRF af
5. Leverage BGP Prefix Independent Convergence (PIC) for fast convergence <100ms (IPv4 and IPv6):
   - PIC Core
   - PIC Edge
   - Best-external advertisement
   - Next-hop tracking (ON by default)
6. Consider RT-constraint for Route-reflector scalability
7. Consider 'BGP slow peer' for PE or RR – faster BGP convergence

Conclusion

- MPLS based IP/VPN is the most optimal L3VPN technology
  - Any-to-any IPv4 or IPv6 VPN topology
  - Partial-mesh, Hub and Spoke topologies also possible
- Various IP/VPN services for additional value/revenue
- IP/VPN paves the way for virtualization & Cloud Services
  - Benefits whether SP or Enterprise.

Questions?

Issue Date:

Revision:

APNIC