Deploying MPLS L3VPN

BUILT FOR THE HUMAN NETWORK

Abstract

- **This session** describes the implementation of IP Virtual Private Networks (IP VPNs) using MPLS. It is the most common Layer 3 VPN technology, as standardized by IETF RFC2547/4364, realizing IP connectivity between VPN site and MPLS network.
- Service Providers have been using IP VPN to provide scalable site-to-site/WAN connectivity to Enterprises/SMBs' for more than a decade. Enterprises have been using it to address network segmentation (virtualization and traffic separation) inside the site e.g. Campus, Data Center. This technology realizes IP connectivity between VPN site and MPLS network.
- The session will cover:
  - IP VPN Technology Overview (RFC2547/RFC4364)
  - IP VPN Configuration Overview
  - IP VPN-based services (multihoming, Hub&Spoke, extranet, Internet, NAT, VRF-lite, etc.)
  - Best Practices

Terminology

- LSR: label switch router
- LSP: label switched path
  - The chain of labels that are swapped at each hop to get from one LSR to another
- VRF: VPN routing and forwarding
  - Mechanism in Cisco IOS® used to build per-customer RIB and FIB
- MP-BGP: multiprotocol BGP
- PE: provider edge router interfaces with CE routers
- P: provider (core) router, without knowledge of VPN
- VPNv4: address family used in BGP to carry MPLS-VPN routes
- RD: route distinguisher
  - Distinguish same network/mask prefix in different VRFs
- RT: route target
  - Extended community attribute used to control import and export policies of VPN routes
- LFIB: label forwarding information base
- FIB: forwarding information base

Agenda

- IP/VPN Overview
- IP/VPN Services
- Best Practices
- Conclusion

Agenda

- IP/VPN Overview
  - Technology Overview (How It Works)
  - Configuration Overview
- IP/VPN Services
- Best Practices
- Conclusion

IP/VPN Technology Overview

- More than one routing and forwarding tables
- Control plane—VPN route propagation
- Data or forwarding plane—VPN packet forwarding

IP/VPN Technology

MPLS IP/VPN Topology / Connection Model

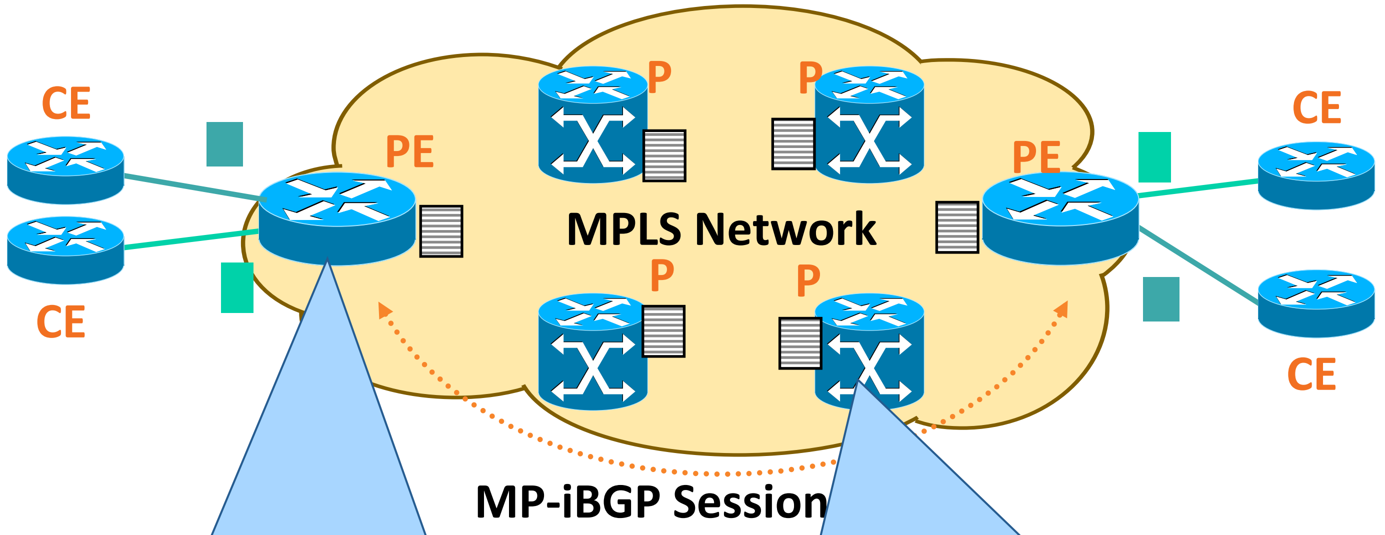

PE Routers

- Sit at the Edge
- Use MPLS with P routers
- Uses IP with CE routers
- Distributes VPN information through MP-BGP to other PE routers

P Routers

- Sit inside the network
- Forward packets by looking at labels
- P and PE routers share a common IGP

IP/VPN Technology Overview

Separate Routing Tables at PE

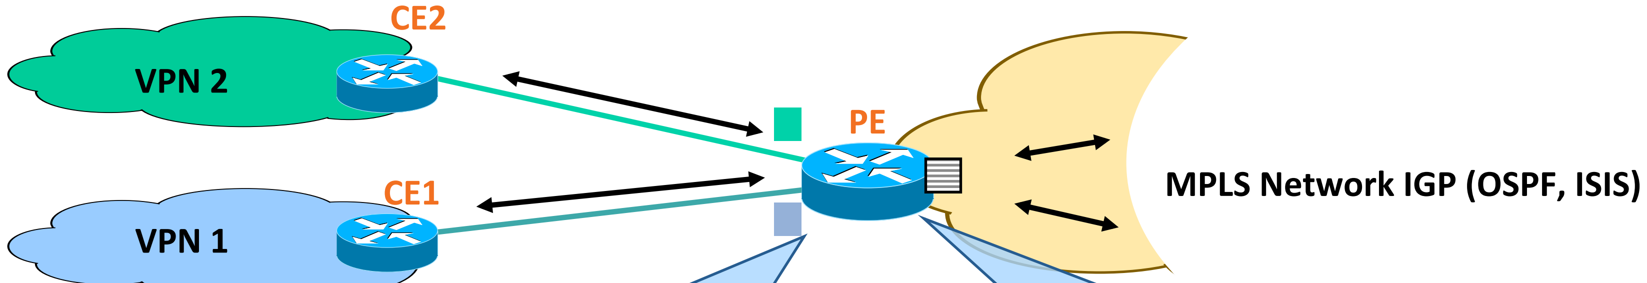

Customer Specific Routing Table

- Routing (RIB) and forwarding table (CEF) dedicated to VPN customer
  - VPN1 routing table
  - VPN2 routing table
- Referred to as VRF table for <named VPN>

```
IOS: "show ip route vrf <name>"
IOS-XR:"sh route vrf <name> ipv4
NX-OS: "sh ip route vrf <name>"
```

Global Routing Table

- Created when IP routing is enabled on PE.
- Populated by OSPF, ISIS, etc. running inside the MPLS network

```
IOS: "show ip route"
IOS-XR:"sh route ipv4 unicast"
NX-OS: "sh ip route"
```

IP/VPN Technology Overview

Virtual Routing and Forwarding Instance

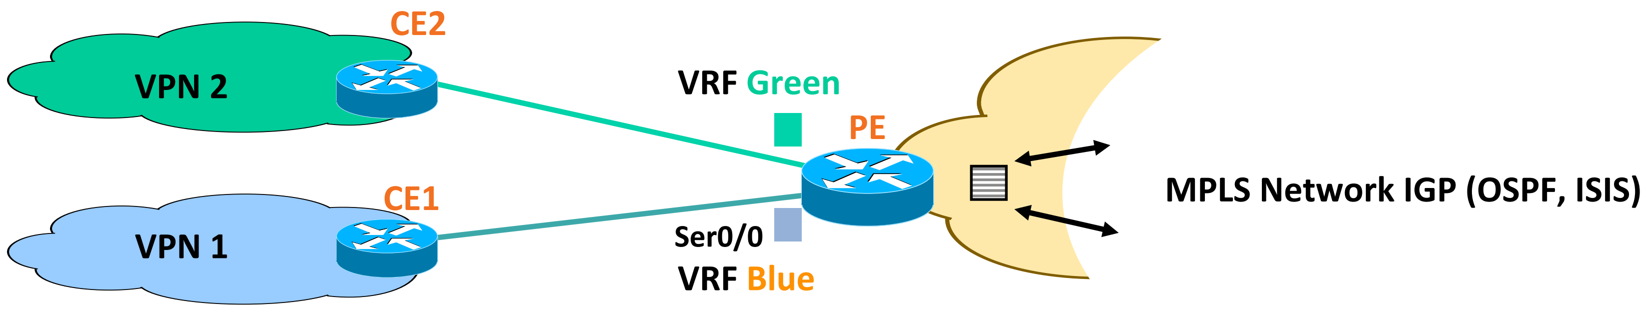

- What's a Virtual Routing and Forwarding (VRF) ?
  - Representation of VPN customer inside the MPLS network
  - Each VPN is associated with at least one VRF
- VRF configured on each PE and associated with PE-CE interface(s)
  - Privatize an interface, i.e., coloring of the interface
- No changes needed at CE

IOS_PE(conf)#ip vrf blue
IOS_PE(conf)#interface Ser0/0
IOS_PE(conf)#ip vrf forwarding blue

IP/VPN Technology Overview

Virtual Routing and Forwarding Instance

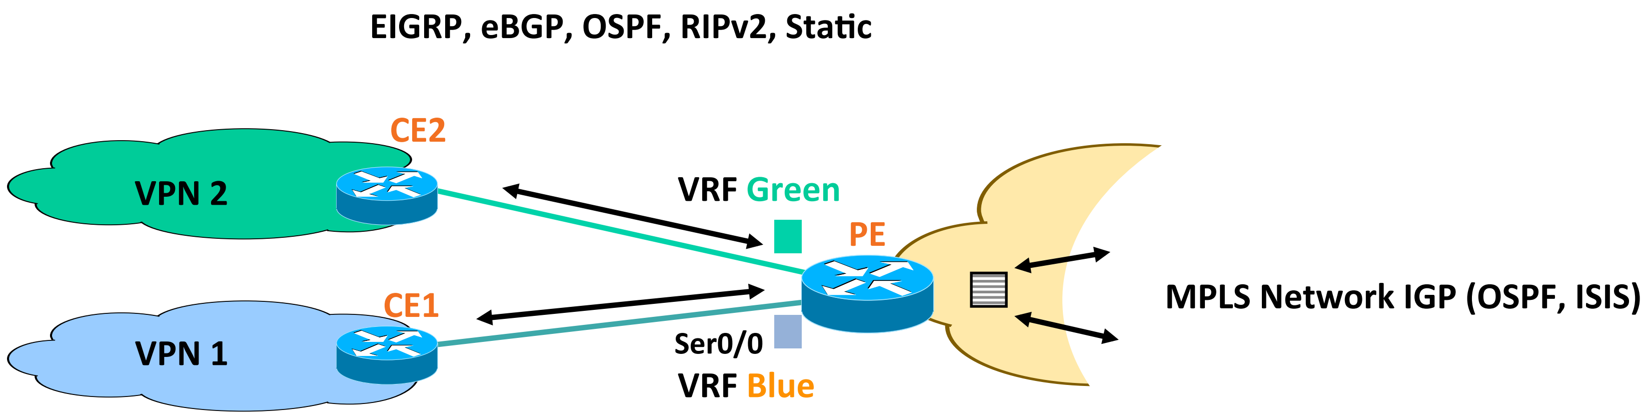

- PE installs the internal routes (IGP) in global routing table
- PE installs the VPN customer routes in VRF routing table(s)
  - VPN routes are learned from CE routers or remote PE routers
  - VRF-aware routing protocol (static, RIP, BGP, EIGRP, OSPF) on each PE
- VPN customers can use overlapping IP addresses
  - BGP plays a key role. Let’s understand few BGP specific details.….

IP/VPN Technology Overview

Control Plane = Multi-Protocol BGP (MP-BGP)

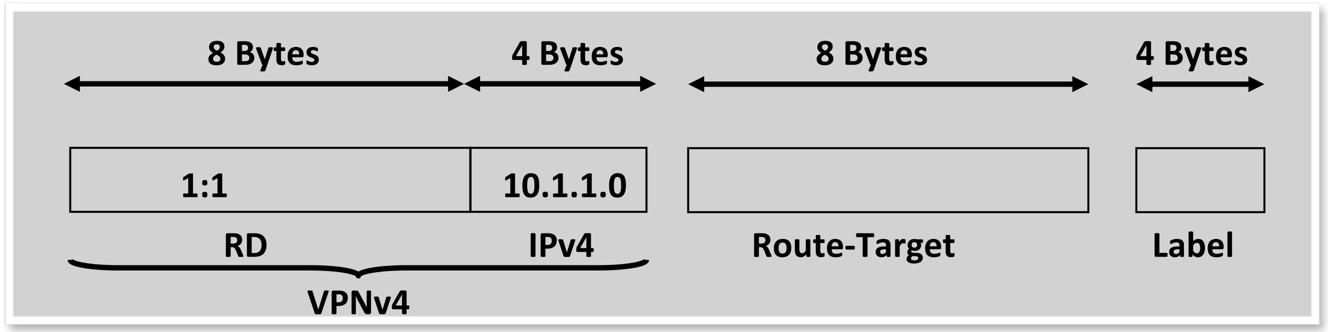

MP-BGP UPDATE Message Showing VPNv4 Address, RT, Label only

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)
- Label

IP/VPN Technology Overview: Control Plane

MP-BGP UPDATE Message Capture

- Visualize how the BGP UPDATE message advertising VPNv4 routes looks like.
- Notice the Path Attributes.

Route Target = 3:3

VPNv4 Prefix 1:1:200.1.62.4/30 ; Label = 23

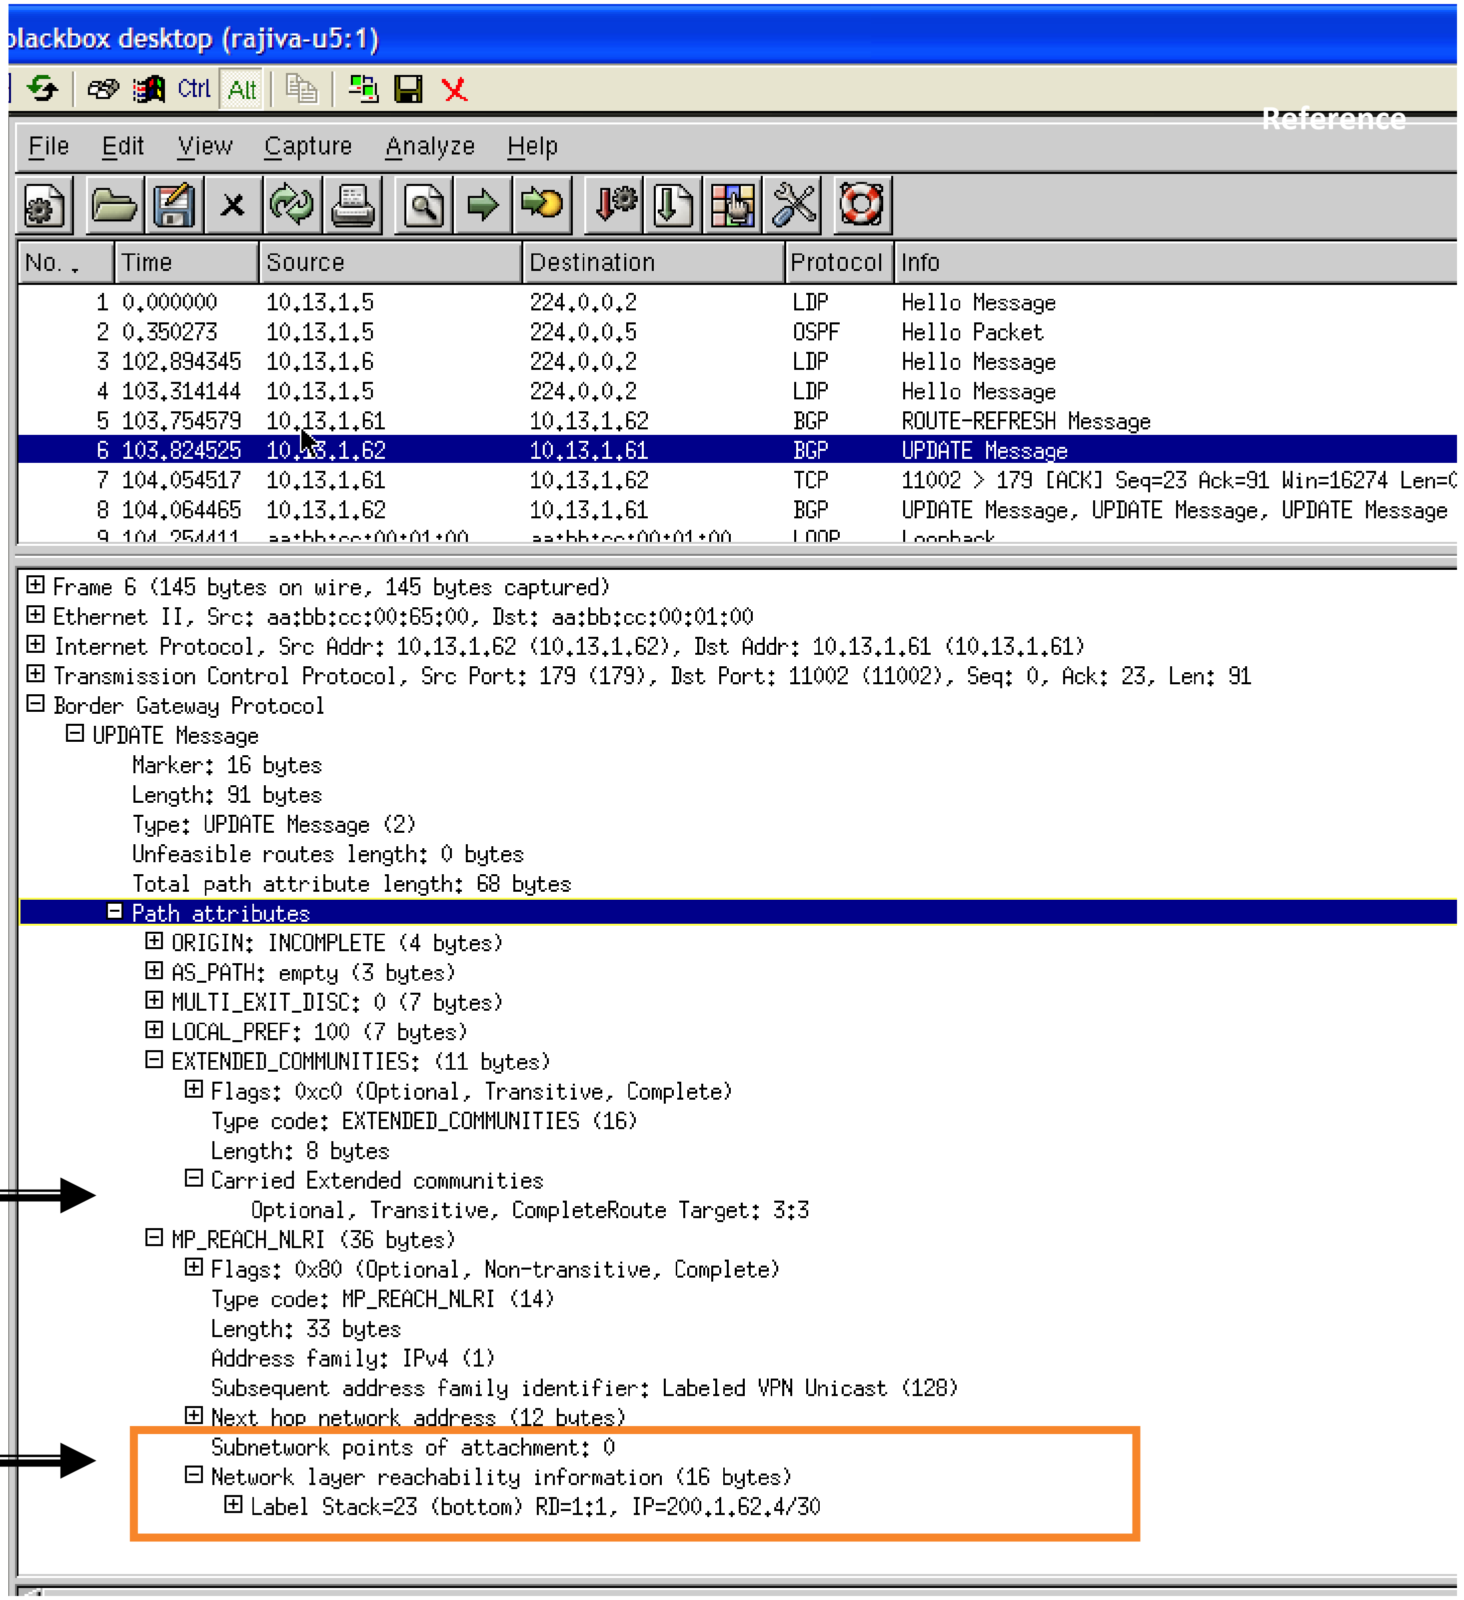

IP/VPN Technology Overview: Control Plane

Route-Distinguisher (rd)

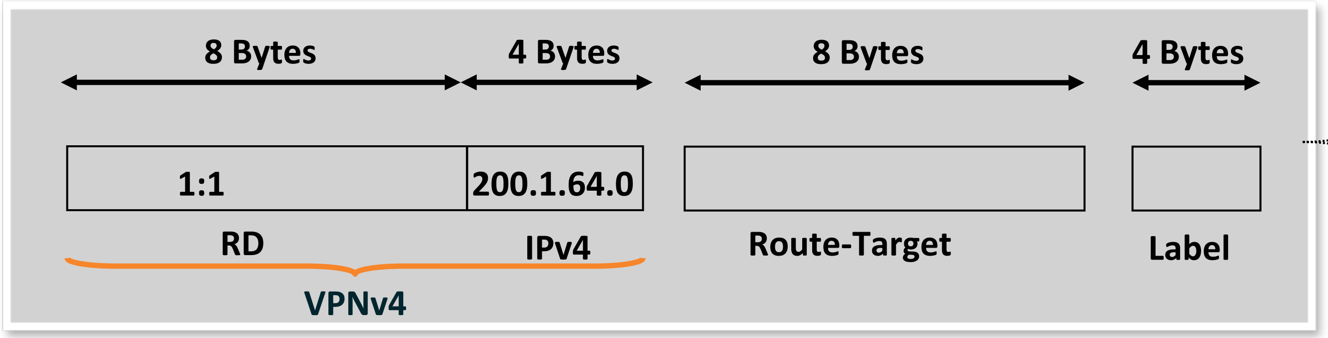

MP-BGP UPDATE Message Showing VPNv4 Address, RT, Label only

- VPN customer IPv4 prefix is converted into a VPNv4 prefix by appending the RD (1:1, say) to the IPv4 address (200.1.64.0, say) => 1:1:200.1.64.0
  – Makes the customer's IPv4 address unique inside the SP MPLS network.
- Route Distinguisher (rd) is configured in the VRF at PE
  – RD is not a BGP attribute, just a field.

```
IOS_PE#
!
ip vrf green
 rd 1:1
!
```

* After 12.4(3)T, 12.4(3) 12.2(32)S, 12.0(32)S etc., RD Configuration within VRF Has Become **Optional**. Prior to That, It Was Mandatory.

IP/VPN Technology Overview: Control Plane

Route-Target (rt)

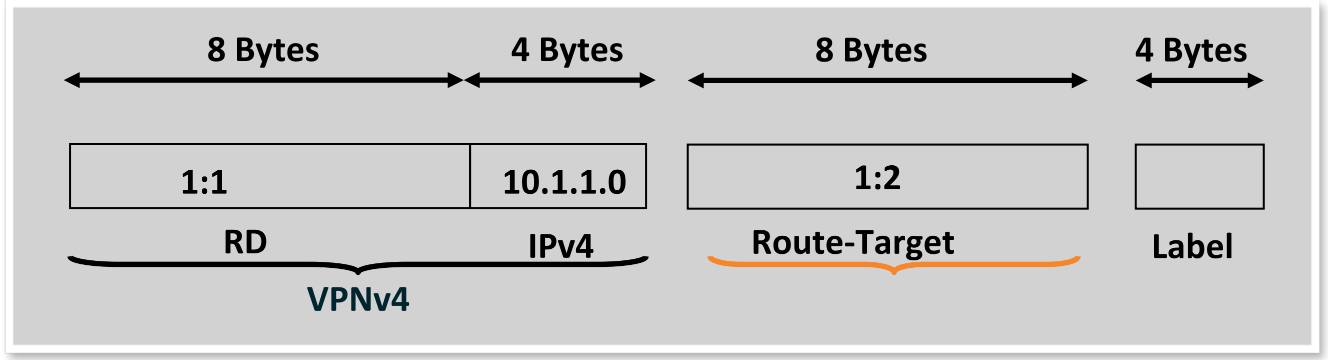

- Route-target (rt) identifies which VRF(s) keep which VPN prefixes
  - – rt is an 8-byte extended community attribute.
- Each VRF is configured with a set of route-targets at PE
  - – Export and Import route-targets must be the same for any-to-any IP/VPN
- Export route-target values are attached to VPN routes in PE->PE MP-iBGP advertisements

```
IOS_PE#
!
ip vrf green
 route-target import 3:3
 route-target export 3:3
 route-target export 10:3
!
```

IP/VPN Technology Overview: Control Plane

Label

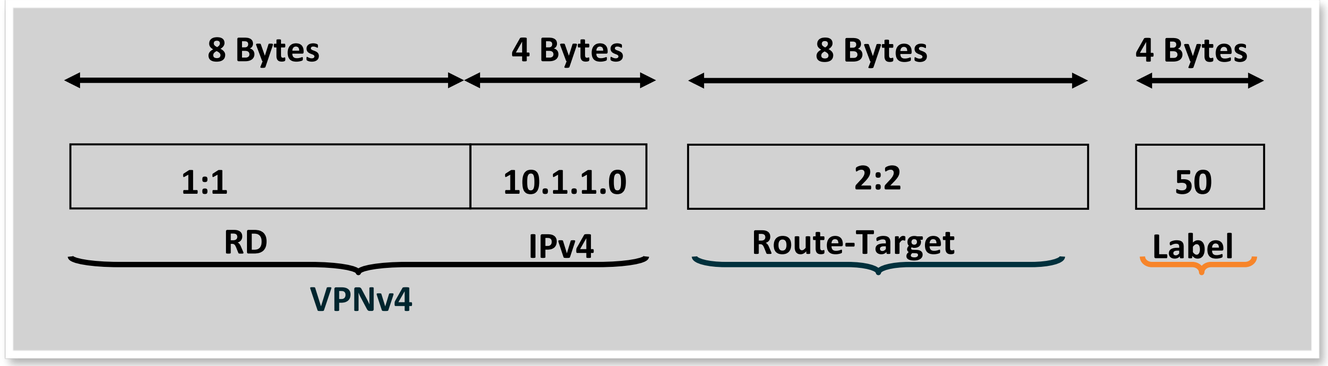

- PE assigns a label for the VPNv4 prefix;
  - Next-hop-self towards MP-iBGP neighbors by default i.e. PE sets the NEXT-HOP attribute to its own address (loopback)
  - Label is not an attribute.
- PE addresses used as BGP next-hop must be uniquely known in IGP
  - Do not summarize the PE loopback addresses in the core

IP/VPN Technology Overview: Control Plane

Putting it all together

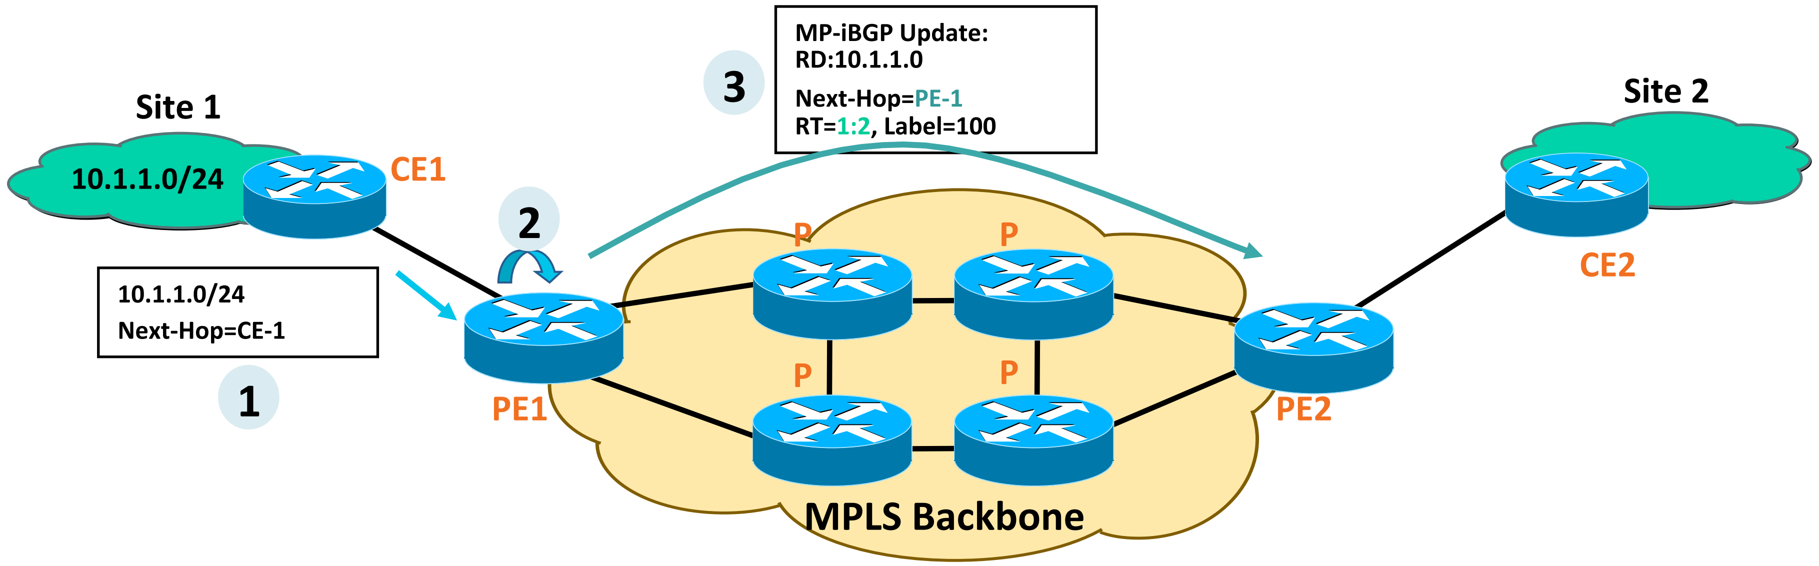

- PE1 receives an IPv4 update (eBGP/OSPF/ISIS/RIP/EIGRP)
- PE1 translates it into VPNv4 address and constructs the MP-iBGP UPDATE message
  - Associates the RT values (export RT =1:2, say) per VRF configuration
  - Rewrites next-hop attribute to itself
  - Assigns a label (100, say); Installs it in the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

IP/VPN Technology Overview: Control Plane

Putting it all together

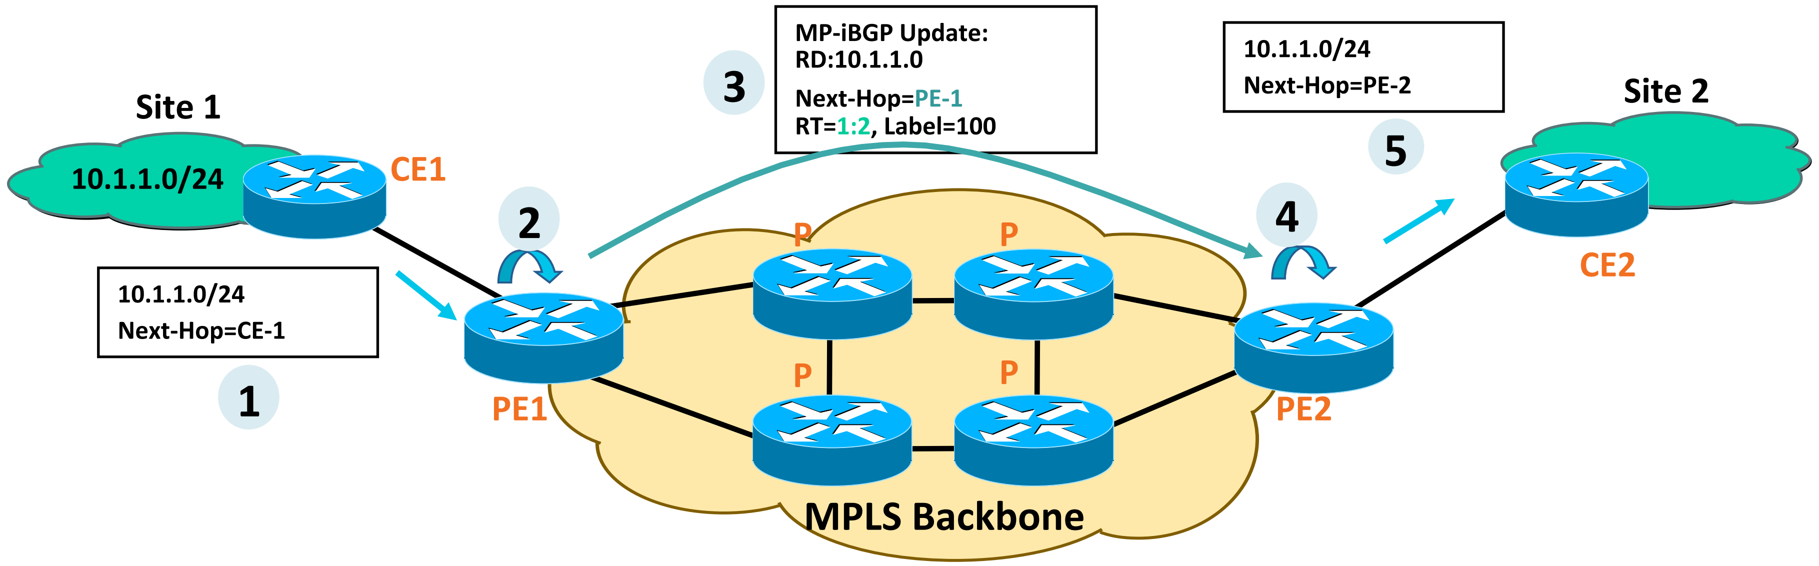

- PE2 receives and checks whether the RT=1:2 is locally configured as 'import RT' within any VRF, if yes, then
  - PE2 translates VPNv4 prefix back to IPv4 prefix
  - Updates the VRF CEF Table for 10.1.1.0/24 with label=100
- PE2 advertises this IPv4 prefix to CE2 (using whatever routing protocol)

Control Plane is now ready

IP/VPN Technology Overview

Forwarding Plane

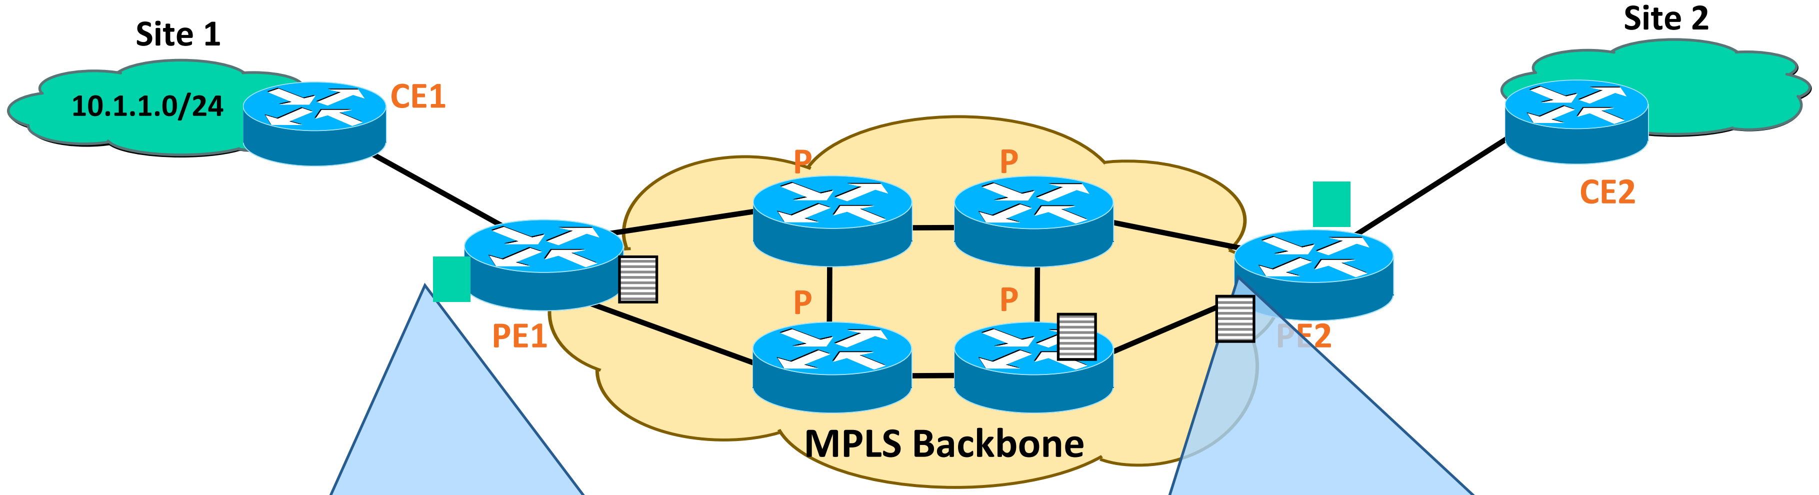

Customer Specific Forwarding Table

- Stores VPN routes with associated labels
- VPN routes learned via BGP
- Labels learned via BGP

```
IOS:show ip cef vrf <name>
NX-OS: show forwarding vrf <name>
IOS-XR: show cef vrf <name> ipv4
```

Global Forwarding Table

- Stores next-hop i.e. PE routes with associated labels
- Next-hop i.e. PE routes learned through IGP
- Label learned through LDP or RSVP

```
IOS:show ip cef
NX-OS: show forwarding ipv4
IOS-XR: show cef ipv4
```

IP/VPN Technology Overview: Forwarding Plane

Packet Forwarding

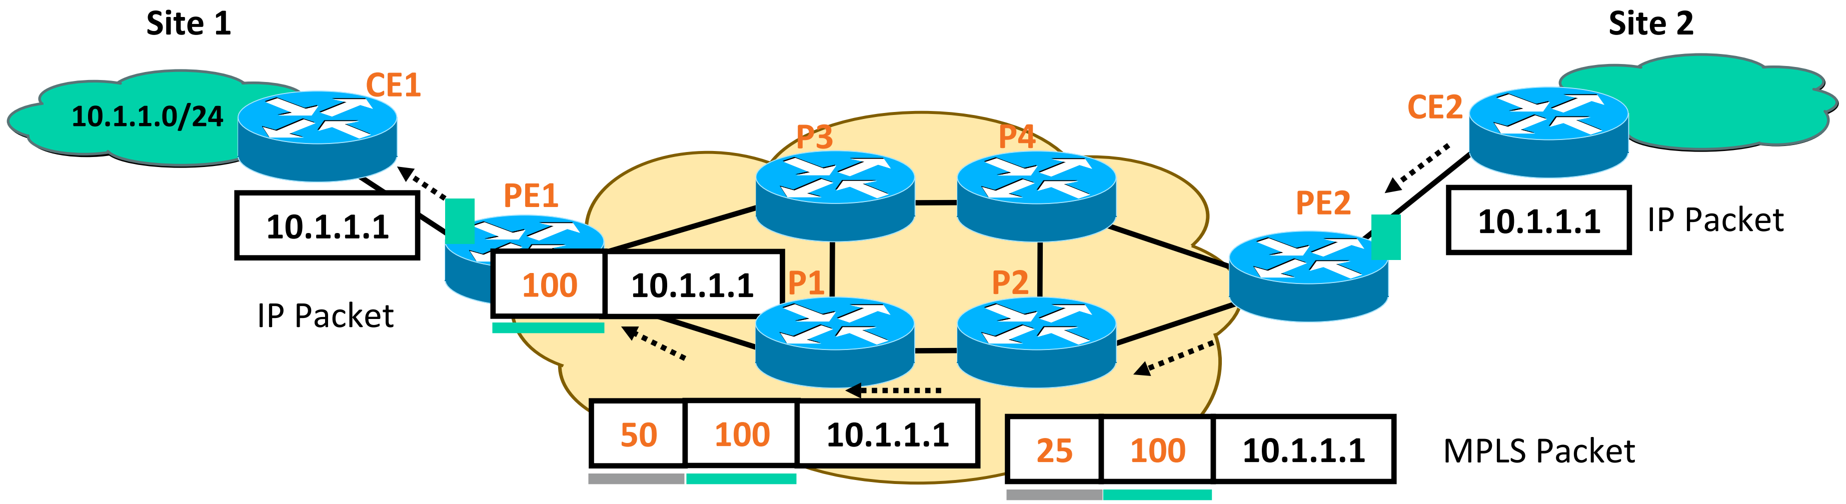

- PE2 imposes two labels (MPLS headers) for each IP packet going to site2
  - Outer label is learned via LDP; Corresponds to PE1 address (e.g. IGP route)
  - Inner label is learned via BGP; corresponds to the VPN address (BGP route)
- P1 does the Penultimate Hop Popping (PHP)
- PE1 retrieves IP packet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

IP/VPN Technology: Forwarding Plane

MPLS IP/VPN Packet Capture

- This capture might be helpful if you never captured an MPLS packet before.

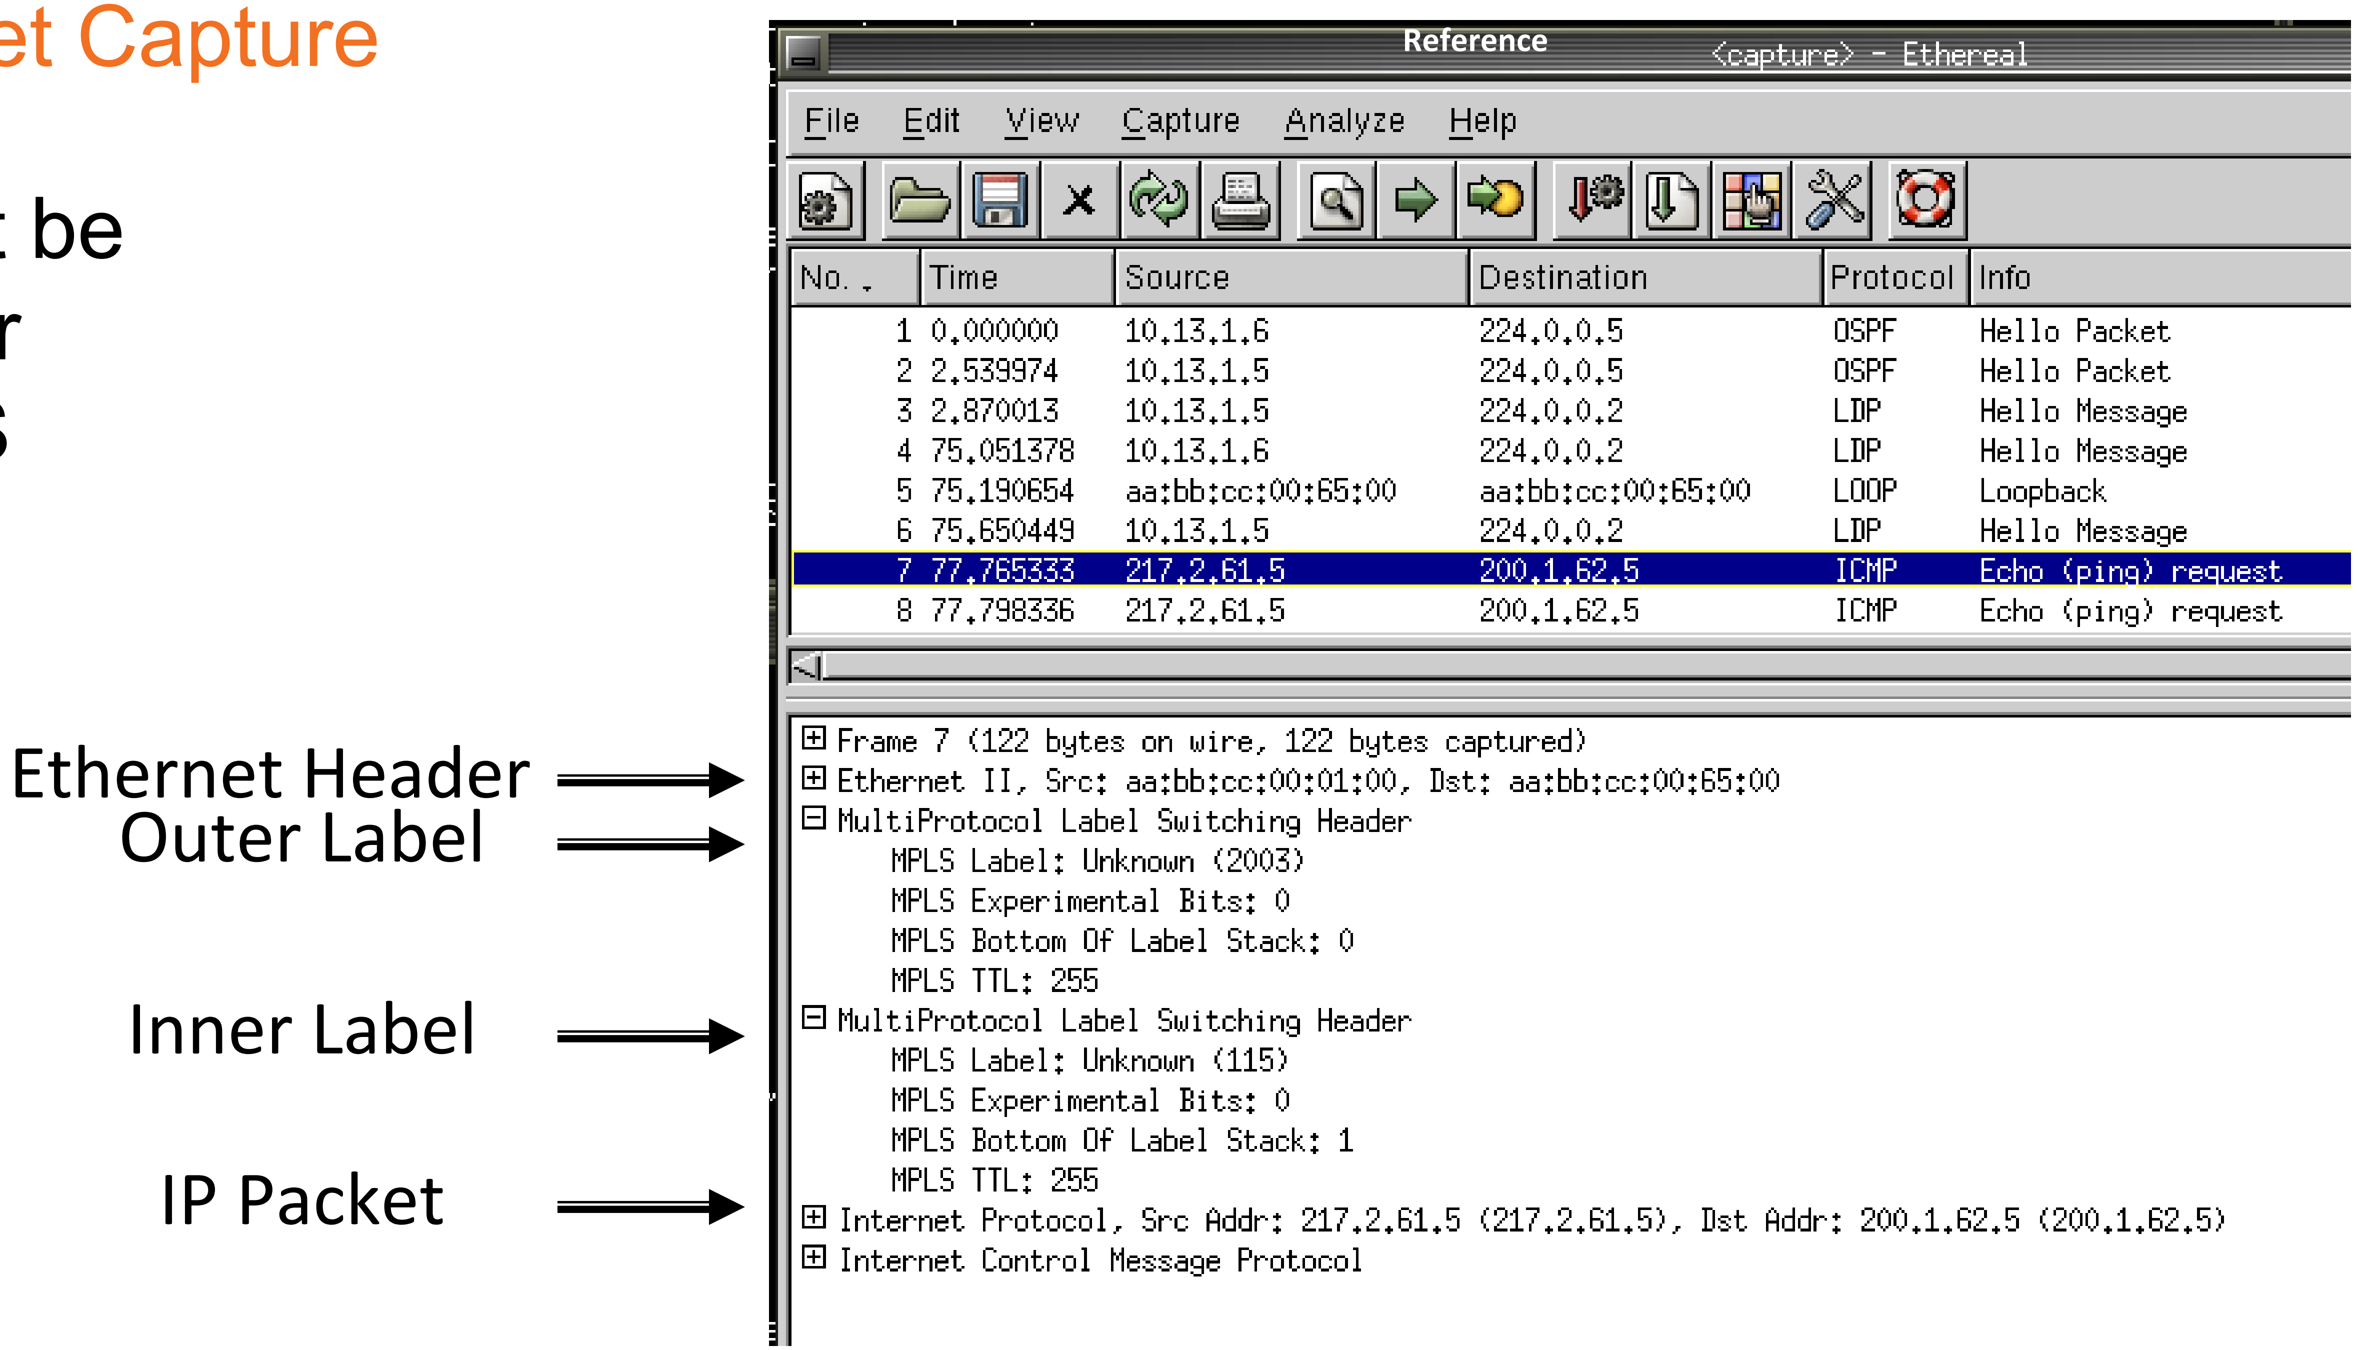

Agenda

- IP/VPN Overview
  - Technology Overview
  - Configuration Overview (IOS, IOS-XR and NX-OS)
- IP/VPN Services
- Best Practices
- Conclusion

MPLS based IP/VPN Sample Configuration (IOS)

Reference

VRF Definition

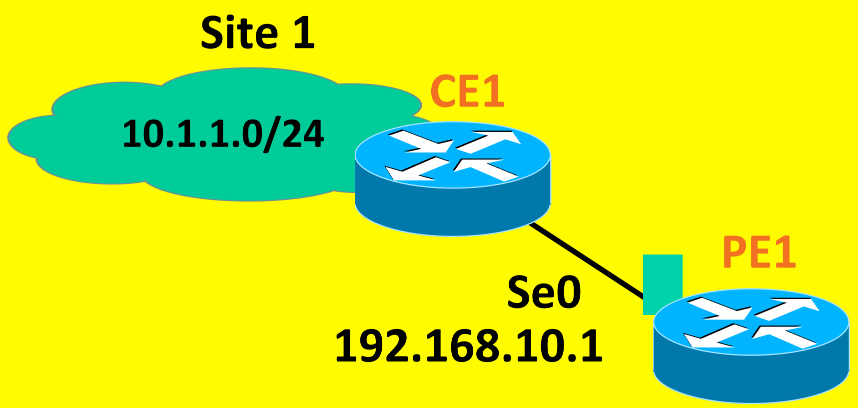

PE1

```
ip vrf VPN-A
 rd 1:1
 route-target export  100:1
 route-target import  100:1
```

```
interface Serial0
 ip address 192.168.10.1/24
 ip vrf forwarding VPN-A
```

```
vrf definition VPN-A
 rd 1:1
 address-family ipv4
  route-target export  100:1
  route-target import  100:1
```

```
interface Serial0
 ip address 192.168.10.1/24
 vrf forwarding VPN-A
```

PE-P Configuration

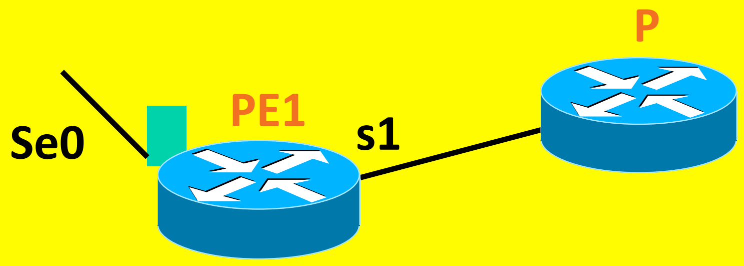

PE1

```
Interface Serial1
 ip address 130.130.1.1 255.255.255.252
 mpls ip
```

```
router ospf 1
 network 130.130.1.0 0.0.0.3 area 0
```

MPLS based IP/VPN Sample Configuration (IOS)

Reference

PE: MP-IBGP Config

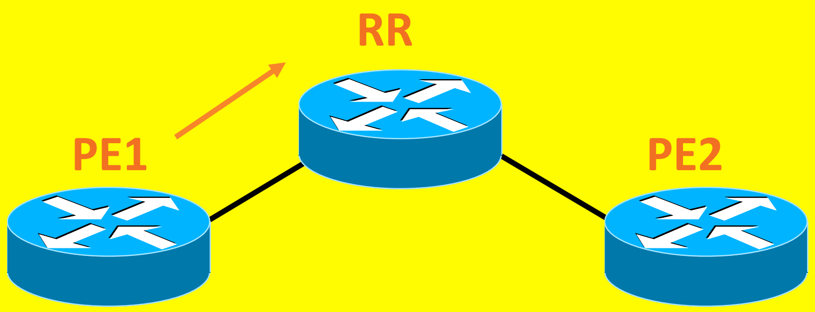

PE1

```
router bgp 1
 neighbor 1.2.3.4 remote-as 1
 neighbor 1.2.3.4 update-source loopback0
!
 address-family vpnv4
  neighbor 1.2.3.4 activate
  neighbor 1.2.3.4 send-community both
!
```

RR: MP-IBGP Config

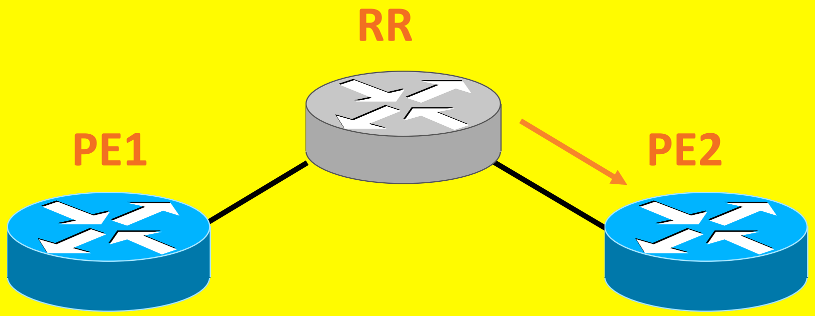

RR

```
router bgp 1
 no bgp default route-target filter
 neighbor 1.2.3.6 remote-as 1
 neighbor 1.2.3.6 update-source loopback0
!
 address-family vpnv4
 neighbor 1.2.3.6 route-reflector- client
 neighbor 1.2.3.6 activate
!
```

MPLS based IP/VPN Sample Configuration (IOS)

PE-CE Routing: BGP

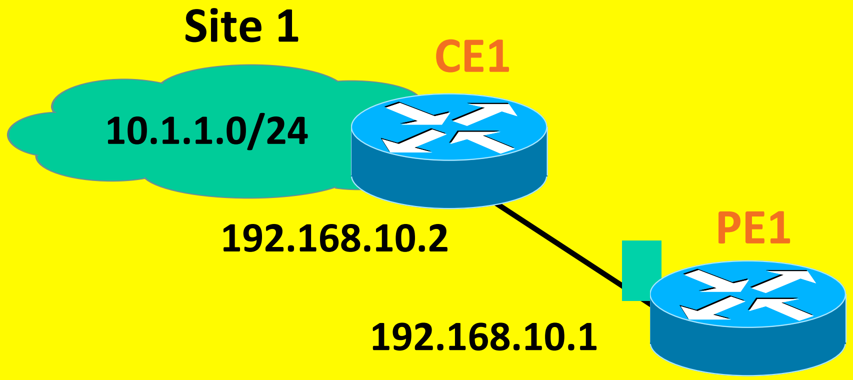

PE1

```
router bgp 1
!
address-family ipv4 vrf VPN-A
  neighbor 192.168.10.2 remote-as 2
  neighbor 192.168.10.2 activate
  exit-address-family
!
```

PE-CE Routing: OSPF

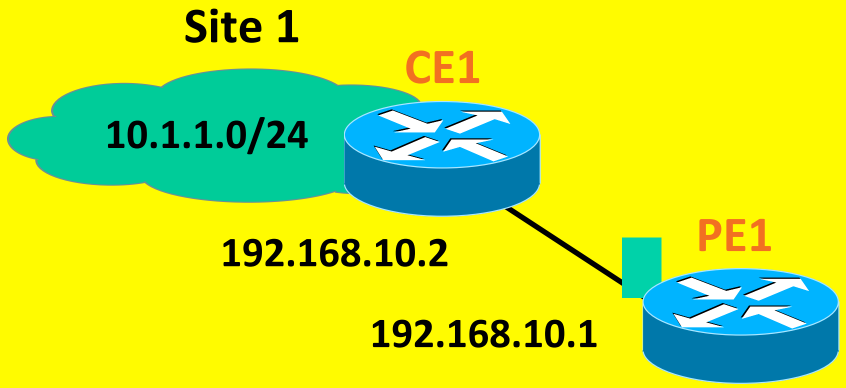

PE1

```
router ospf 1
!
router ospf 2 vrf VPN-A
 network 192.168.10.0 0.0.0.255 area 0
 redistribute bgp 1 subnets
!
```

MPLS based IP/VPN Sample Configuration (IOS)

Reference

PE-CE Routing: RIP

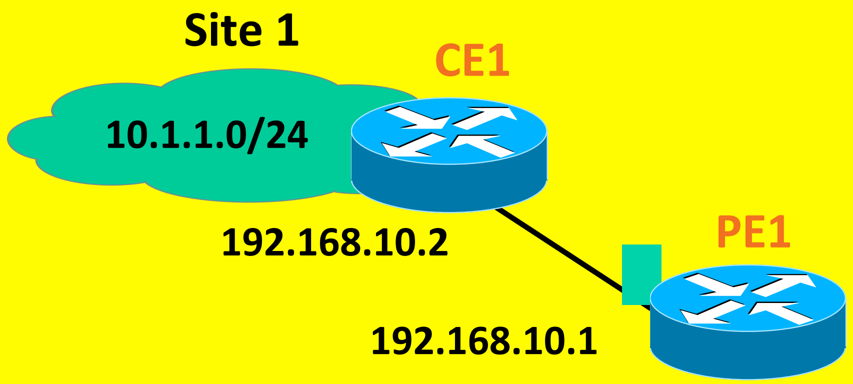

PE1

```
router rip
!
 address-family ipv4 vrf VPN-A
  version 2
  no auto-summary
  network 192.168.10.0
  redistribute bgp 1 metric transparent
!
```

PE-CE Routing: EIGRP

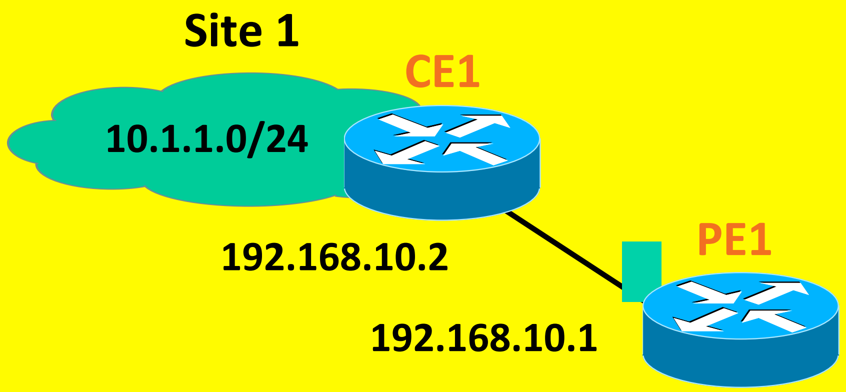

PE1

```
router eigrp 1
!
 address-family ipv4 vrf VPN-A
  no auto-summary
  network 192.168.10.0 0.0.0.255
  autonomous-system 10
  redistribute bgp 1 metric 100000 100
255 1 1500
!
```

MPLS based IP/VPN Sample Configuration (IOS)

PE-CE Routing: Static

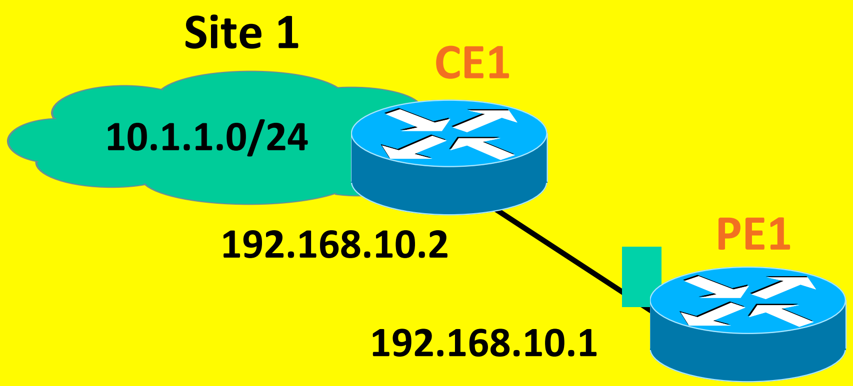

PE1

```
ip route vrf VPN-A 10.1.1.0 255.255.255.0
192.168.10.2
```

MPLS based IP/VPN Sample Configuration (IOS)

Reference

If PE-CE Protocol Is Non-BGP, then Redistribution of Local VPN Routes into MP-IBGP Is Required (Shown Below)

PE-RR (VPN Routes to VPNv4)

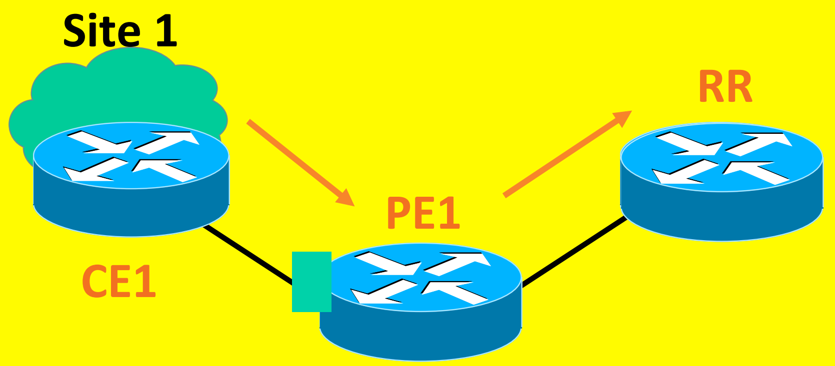

PE1

```
router bgp 1
 neighbor 1.2.3.4 remote-as 1
 neighbor 1.2.3.4 update-source loopback 0

 address-family ipv4 vrf VPN-A
  redistribute {rip|connected|static|eigrp|ospf}
```

- Having familiarized with IOS based config, let’s peek through IOS-XR and NX-OS config for VPNs

MPLS based IP/VPN Sample Config (IOS-XR)

Reference

VRF Definition

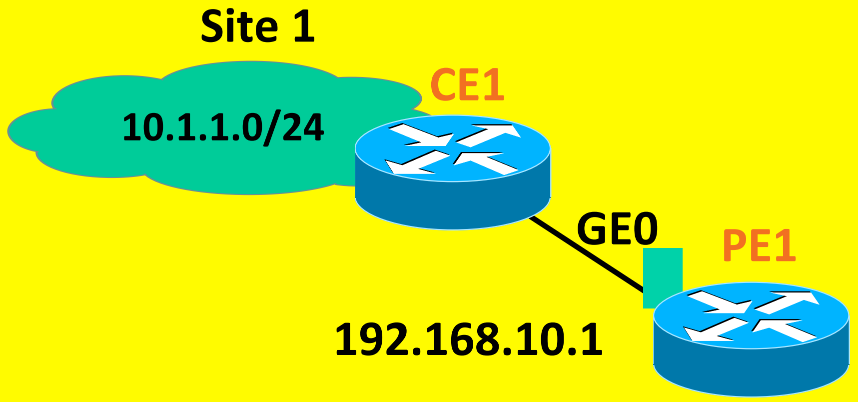

PE1

```
vrf VPN-A
 address-family ipv4 unicast
   import route-target 100:1
   export route-target 100:1
!
router bgp 1
 vrf VPN-A
  rd 1:1
```

```
Interface GigaEthernet0/0/0/0
 ipv4 address 192.168.10.1 255.255.255.0
 vrf VPN-A
```

PE-P Configuration

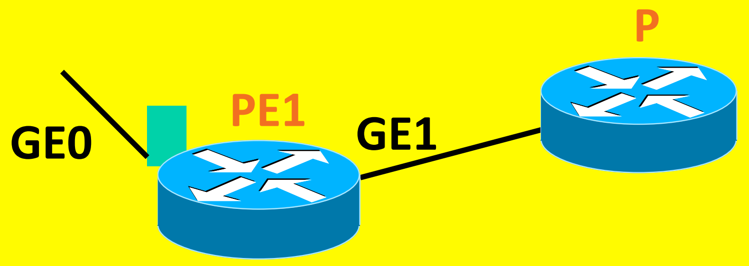

PE1

```
mpls ldp
 route-id 1.2.3.1
 interface GigabitEthernet0/0/0/1
 !
!
mpls oam
```

```
router ospf 1
```


MPLS based IP/VPN Sample Config (IOS-XR)

Reference

PE: MP-IBGP Config

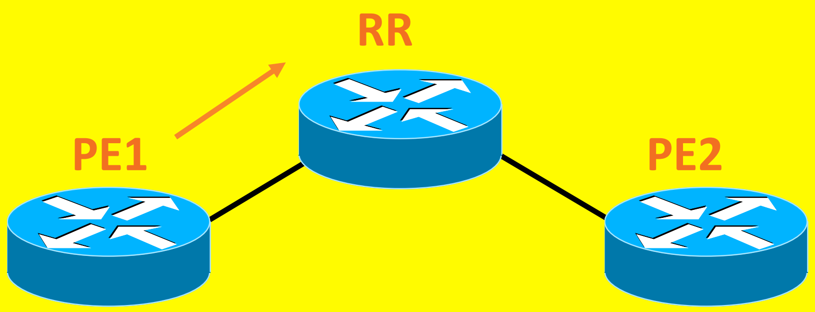

PE1

```
router bgp 1
 router-id 1.2.3.1
 address-family vpnv4 unicast
 !
 neighbor 1.2.3.4
  remote-as 1
  update-source loopback0
  address-family vpnv4 unicast
 !
!
```

RR: MP-IBGP Config

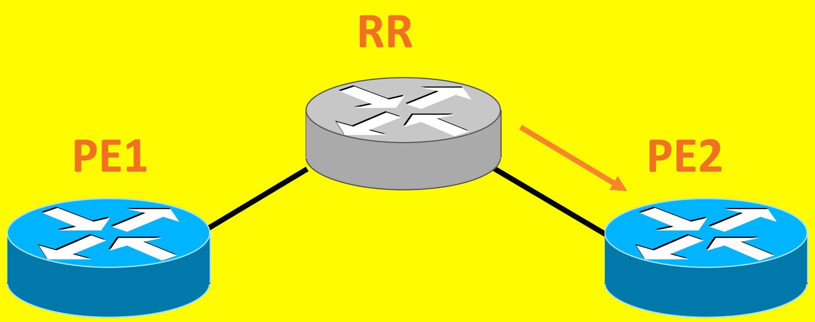

RR

```
router bgp 1
 router-id 1.2.3.4
 address-family vpnv4 unicast
 !
 neighbor 1.2.3.1
  remote-as 1
  update-source loopback0
  address-family vpnv4 unicast
  route-reflector-client
 !
!
```

MPLS based IP/VPN Sample Config (IOS-XR)

Reference

PE-CE Routing: BGP

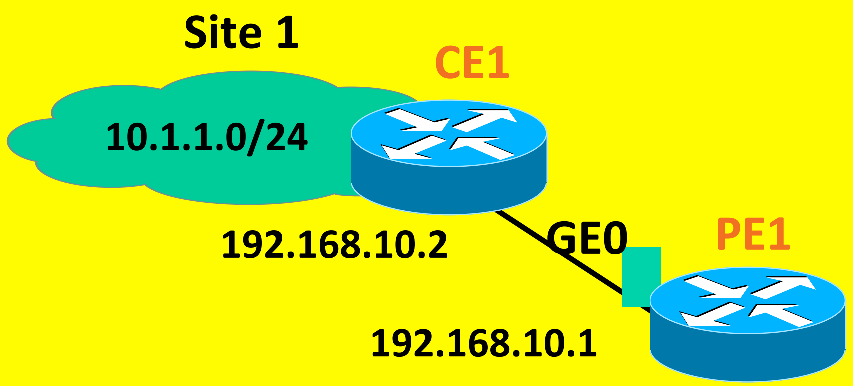

PE1

```
router bgp 1
 vrf VPN-A
  address-family ipv4 unicast
  neighbor 192.168.10.2
   remote-as 2
   address-family ipv4 unicast
    route-policy pass-all in|out
   !
  !
 !
!
```

PE-CE Routing: OSPF

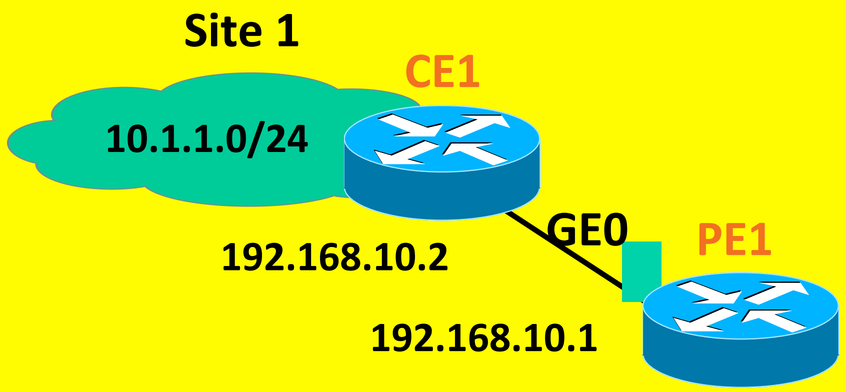

PE1

```
router ospf 2
 vrf VPN-A
  redistribute bgp 1
  area 0
   interface GigabitEthernet0/0/0/1
  !
 !
```

Reference

MPLS based IP/VPN Sample Config (IOS-XR)

PE-CE Routing: RIP

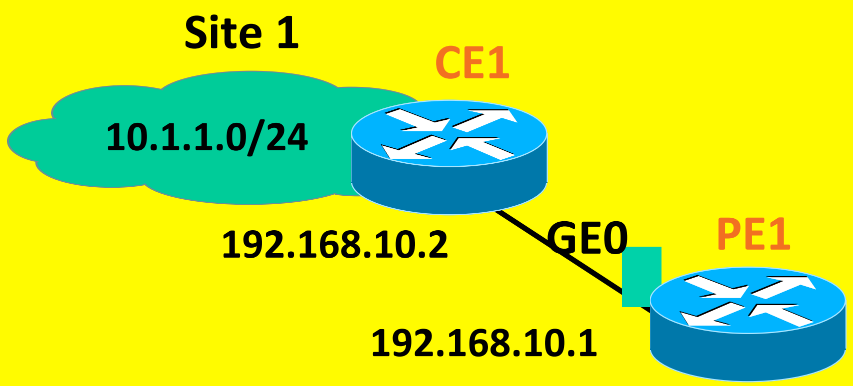

PE1

```
router rip
 vrf VPN-A
  interface GigabitEthernet0/0/0/0
  redistribute bgp 1
 !
```

PE-CE Routing: EIGRP

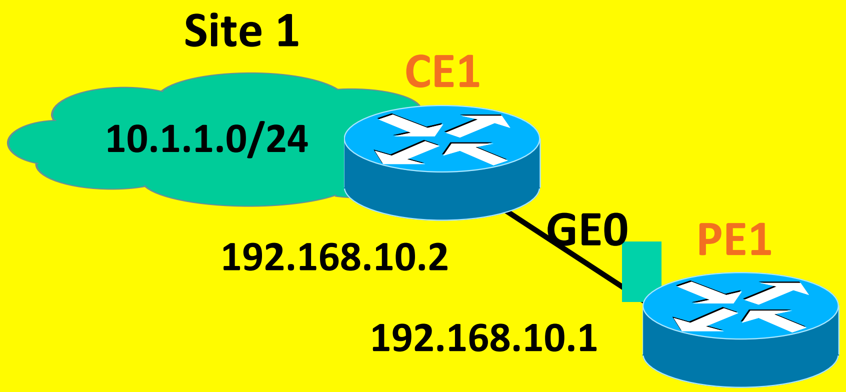

PE1

```
router eigrp 1
 vrf VPN-A
 address-family ipv4
  as 10
  default-metric 100000 100 255 1 1500
  interface GigabitEthernet0/0/0/0
  redistribute bgp 1
```

MPLS based IP/VPN Sample Config (IOS-XR)

Reference

PE-CE Routing: Static

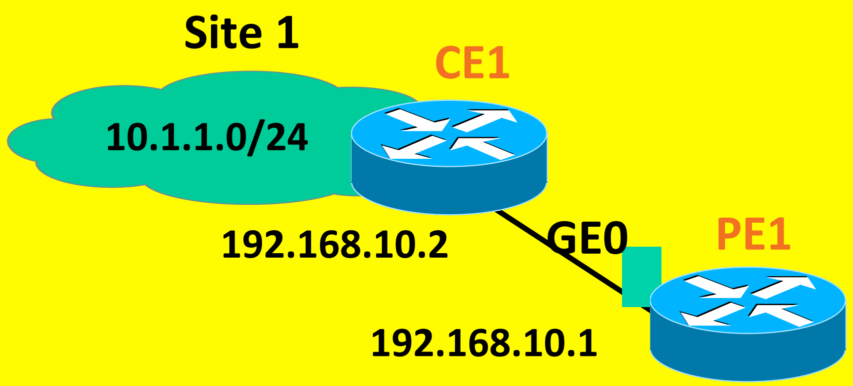

PE1

```
router static
  vrf VPN-A
    address-family ipv4 unicast
      ip route 10.1.1.0/8 192.168.10.2
```

MPLS based IP/VPN Sample Config (IOS-XR)

Reference

If PE-CE Protocol Is Non-BGP, then Redistribution of Local VPN Routes into MP-IBGP Is Required (Shown Below)

PE-RR (VPN Routes to VPNv4)

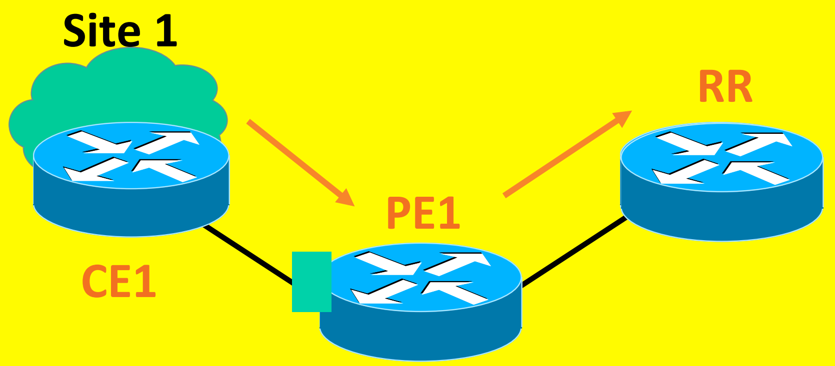

PE1

```
router bgp 1
 vrf VPN-A
  address-family ipv4 unicast
   redistribute {rip|connected|static|eigrp|ospf}
```

Agenda

- IP/VPN Overview
- IP/VPN Services
    1. Load-Sharing for Multihomed VPN Sites
    2. Hub and Spoke Service
    3. Extranet Service
    4. Internet Access Service
    5. IPv6 VPN Service
- Best Practices
- Conclusion

IP/VPN Services:

1. Loadsharing of VPN Traffic

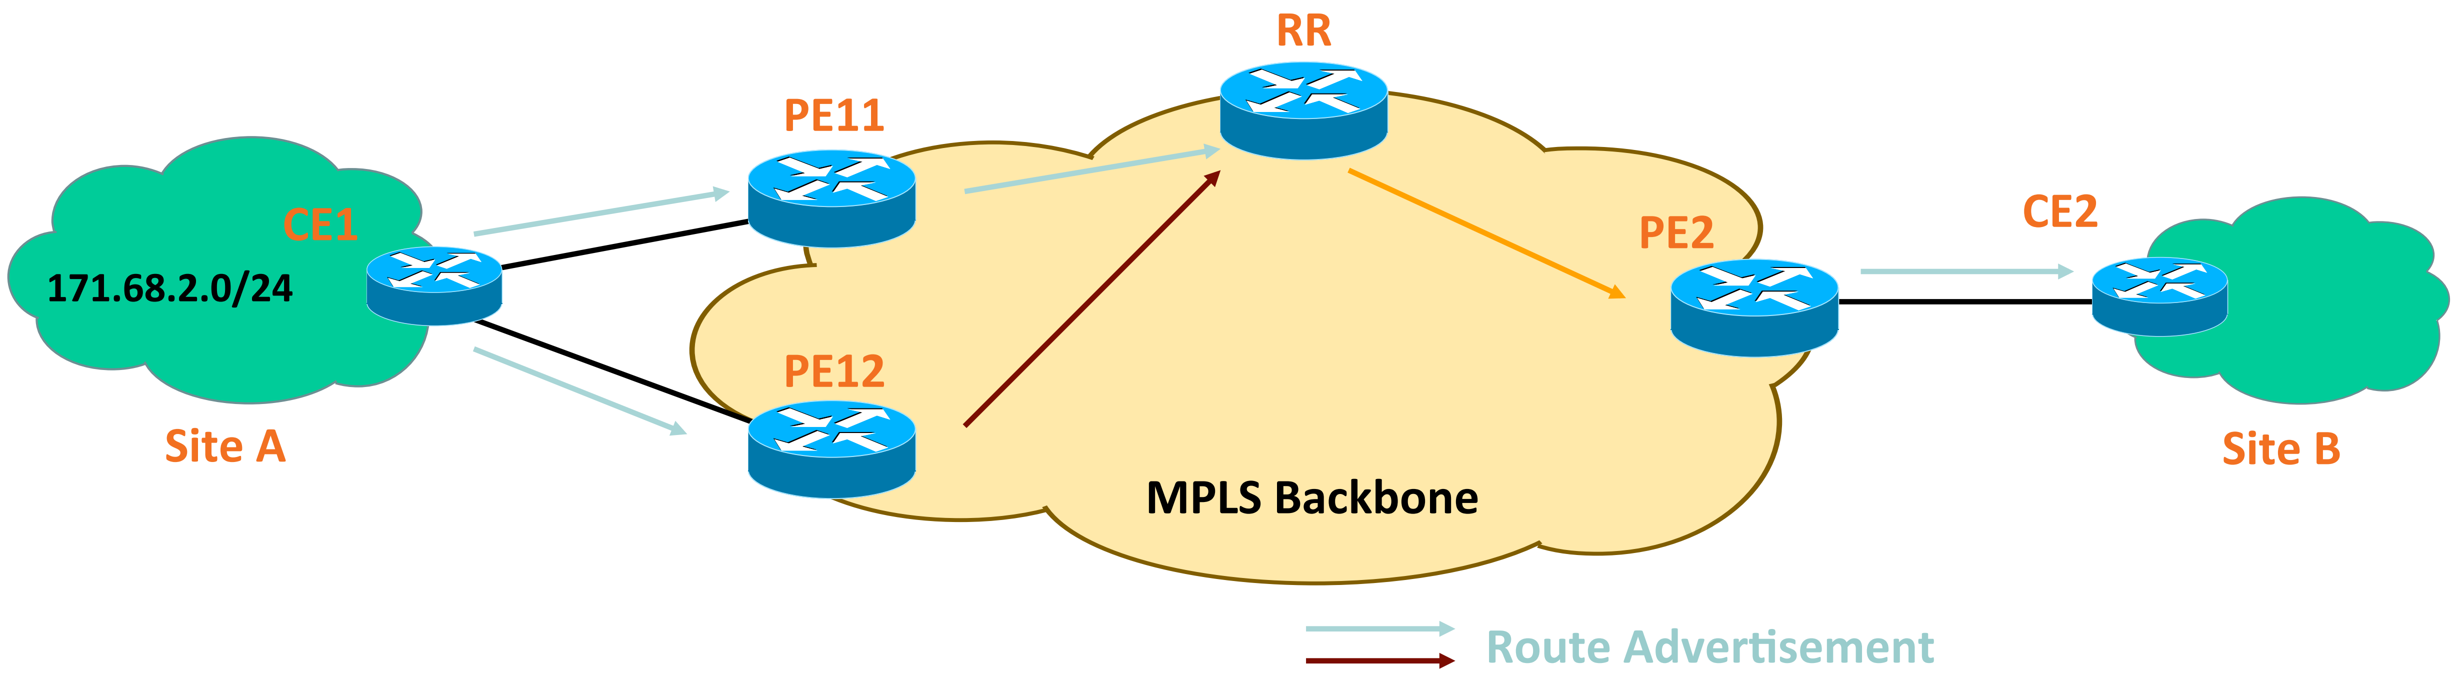

- VPN sites (such as Site A) could be multihomed
- VPN customer may demand the traffic (to the multihomed site) be loadshared

IP/VPN Services:

1. Loadsharing of VPN Traffic: Two Scenarios

1 CE → 2 PEs

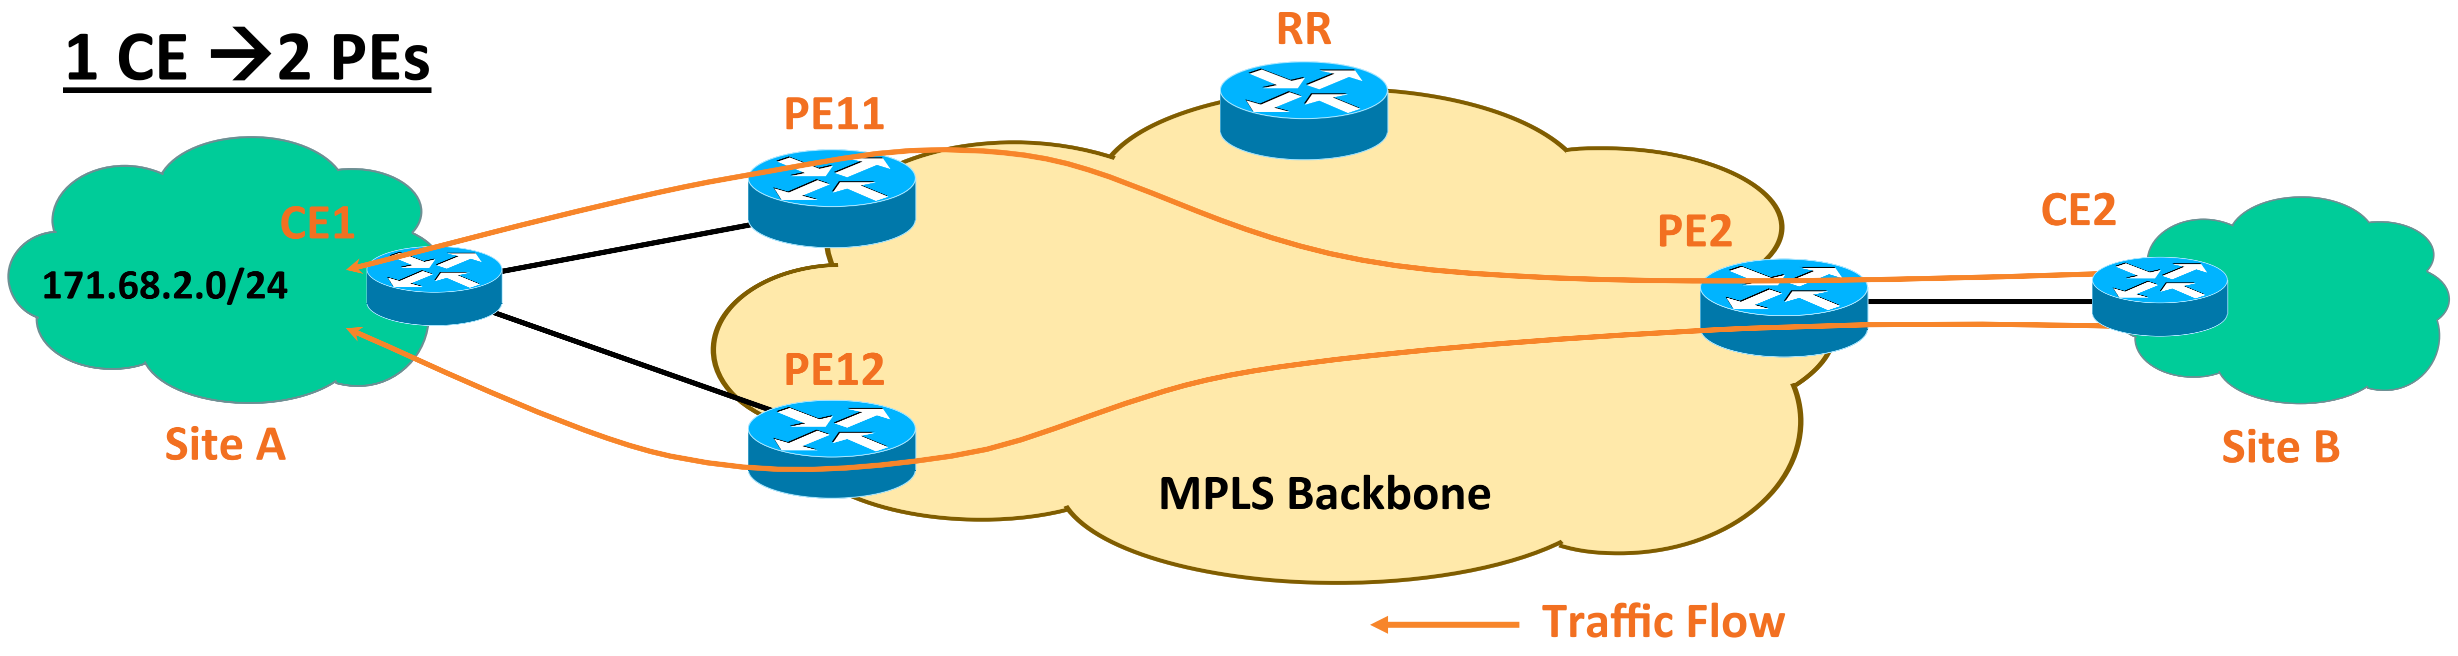

2 CEs → 2 PEs

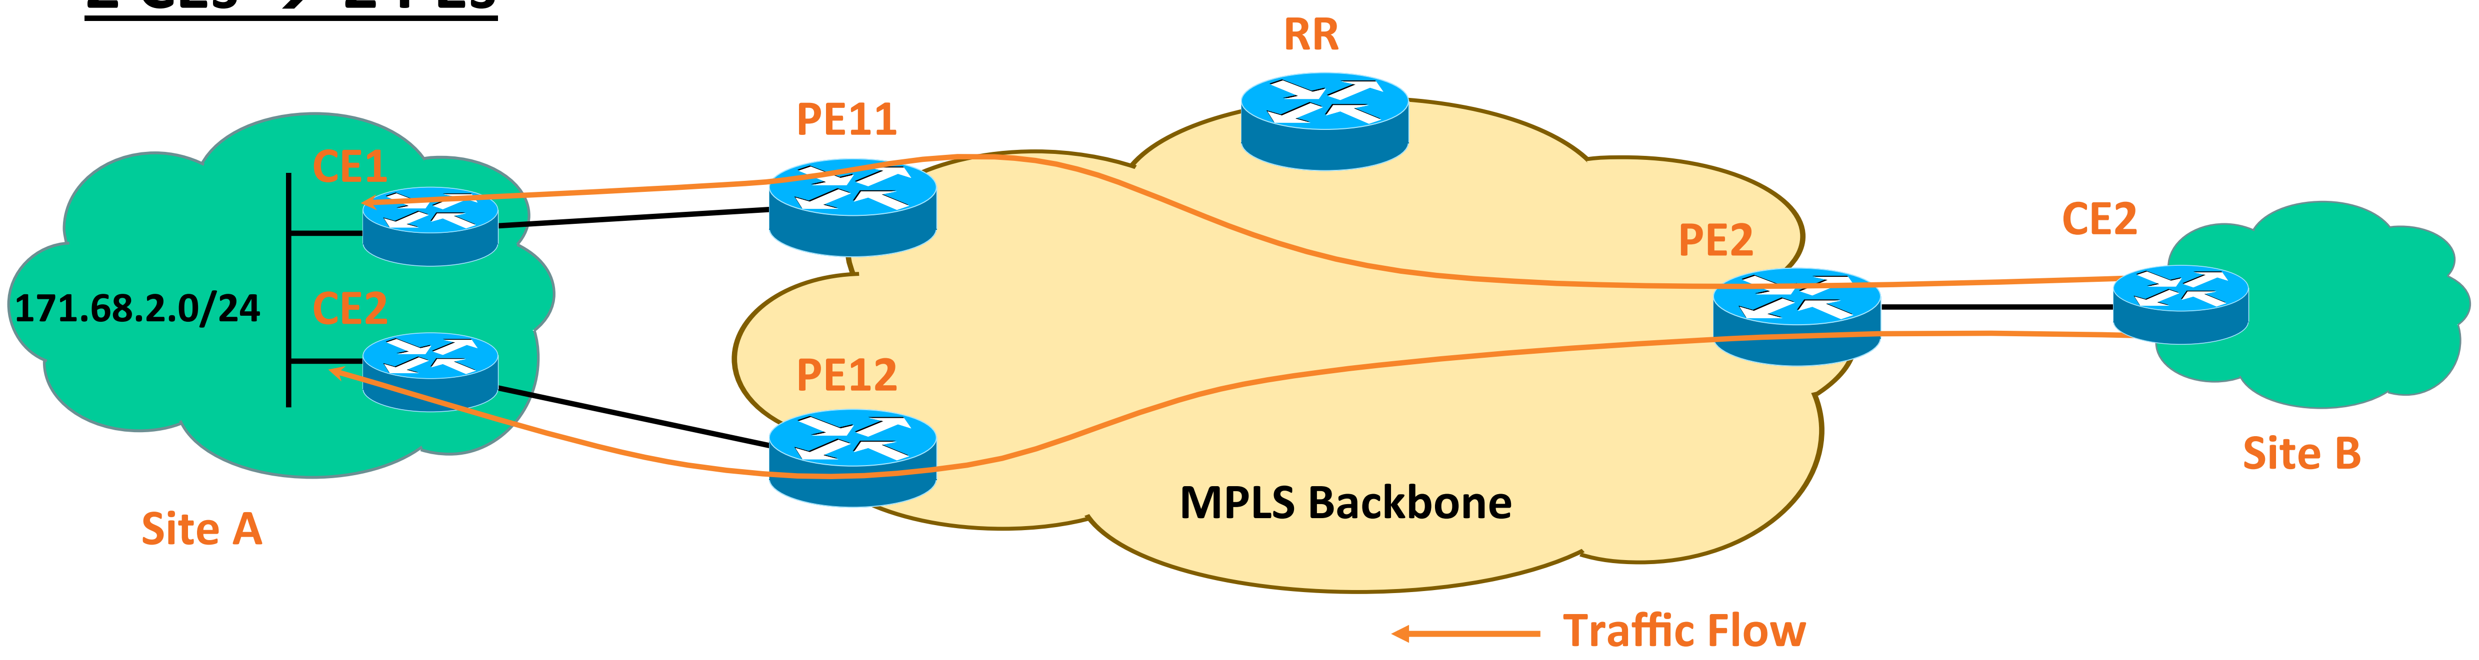

Supported in IOS, and IOS-XR.

IP/VPN Services:

1. Loadsharing of VPN Traffic: IOS Configuration

- Configure unique RD per VRF per PE for multihomed site/interfaces
  - –Assuming RR exists
- Enable BGP multipath within the relevant BGP VRF address-family at remote PE routers such as PE2 (why PE2?).

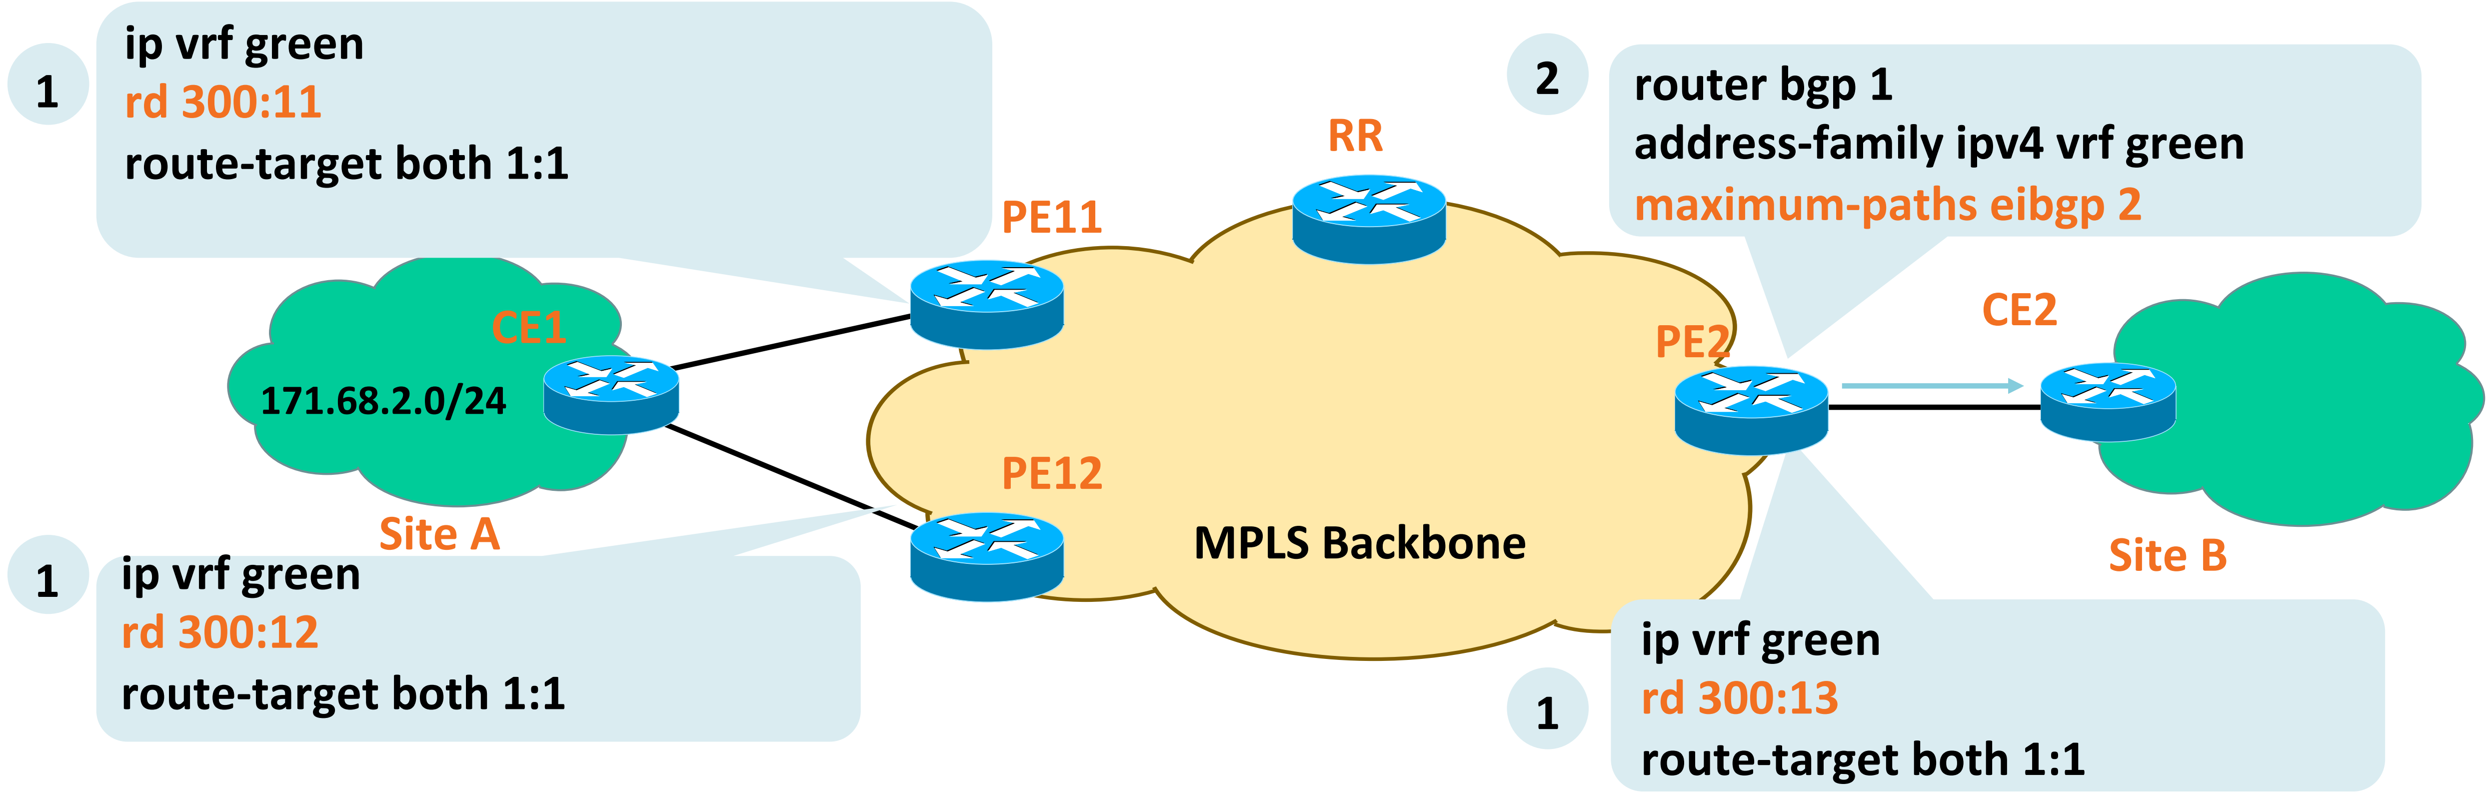

Supported in IOS, and IOS-XR

IP/VPN Services:

1. VPN Fast Convergence—PE-CE Link Failure

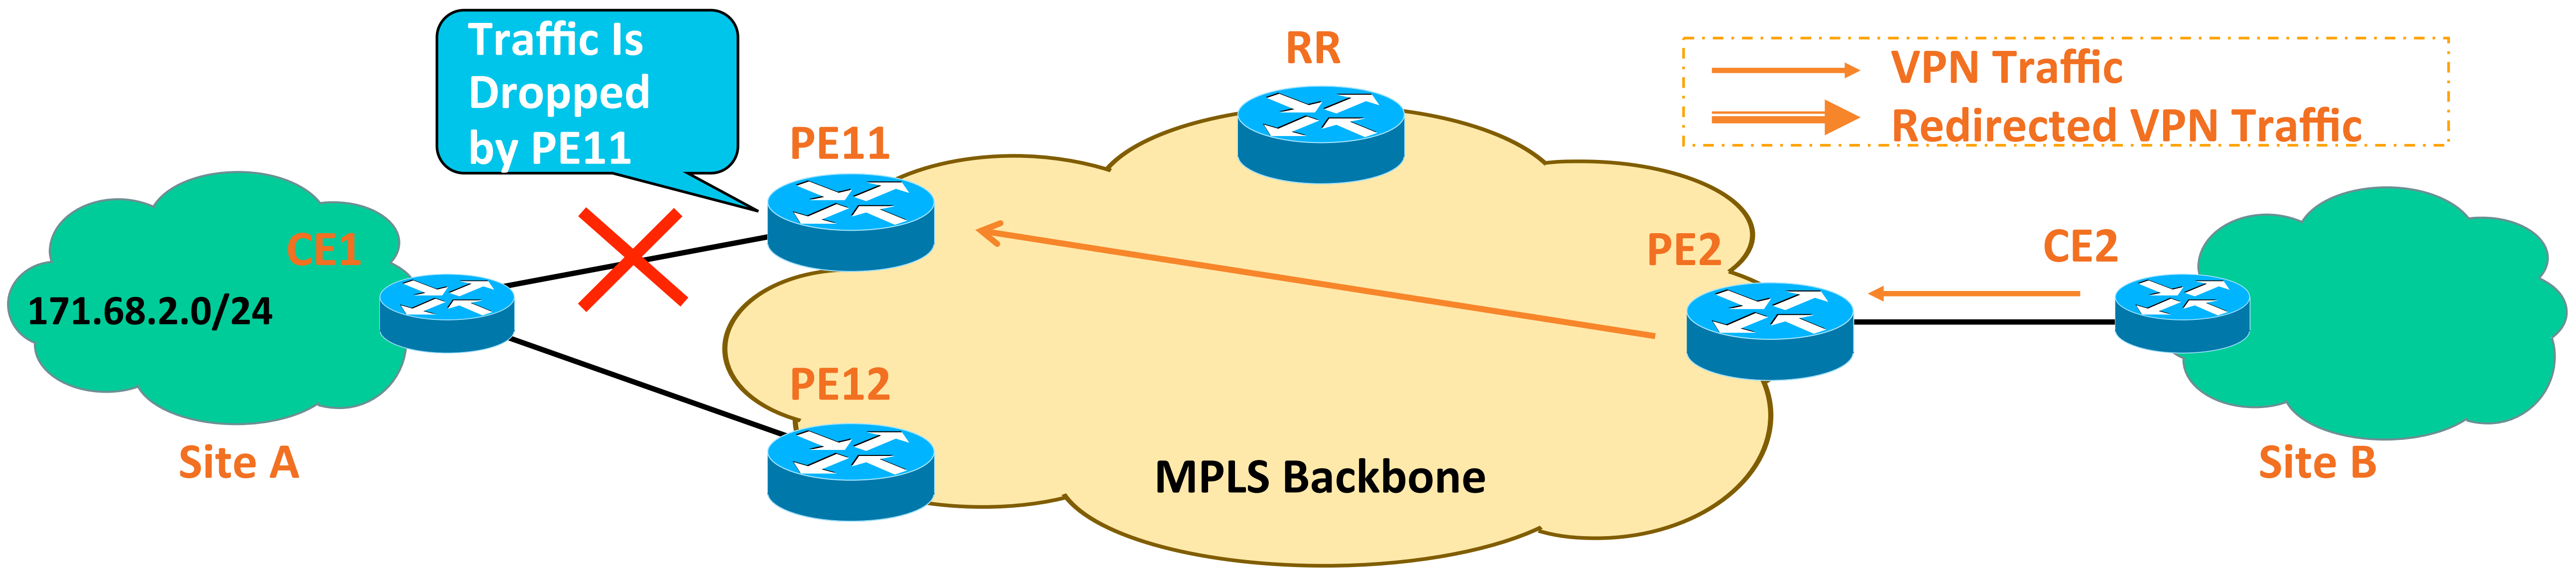

In a classic multi-homing case, PE11, upon detecting the PE-CE link failure, sends BGP message to withdraw the VPN routes towards other PE routers.

–This results in the remote PE routers selecting the alternate bestpath (if any), but until then, they keep sending the MPLS/VPN traffic to PE11, which keeps dropping the traffic.

Use **PIC Edge** feature to minimize the loss due to the PE-CE link failure **from sec to msec .**

IP/VPN Services:

1. VPN Fast Convergence—PE-CE Link Failure

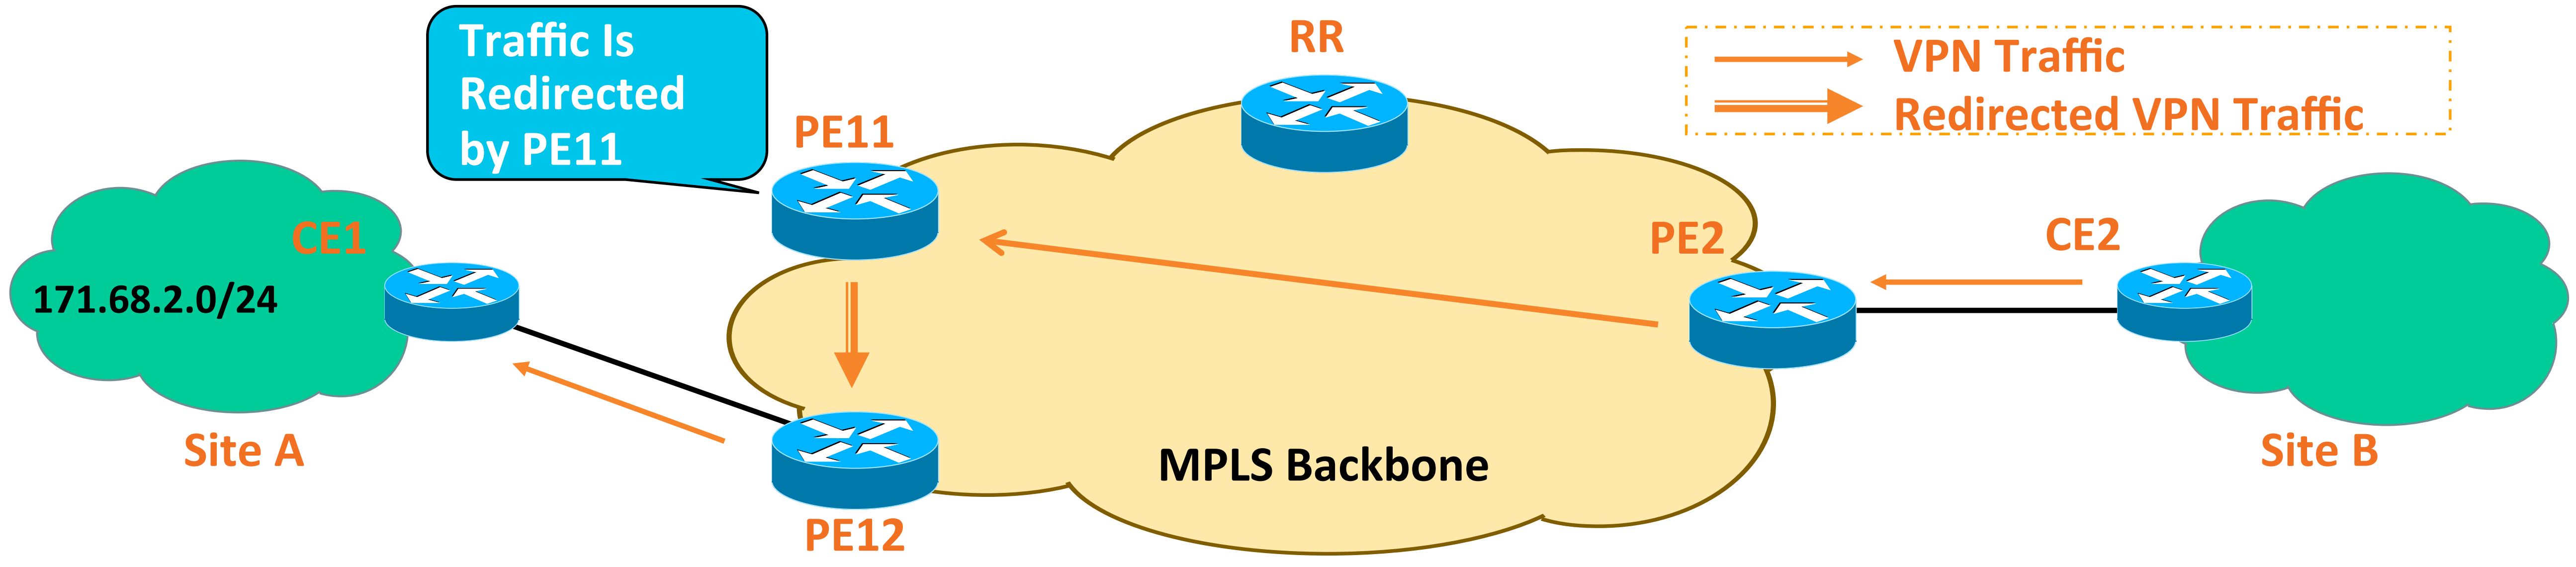

‘BGP PIC Edge’ feature helps PE11 to minimize the traffic loss from sec to msec, during local PE-CE link failure

- PE11 immediately reprograms the forwarding entry with the alternate BGP best path (which is via PE12)
- PE11 redirects the CE1 bound traffic to PE12 (with the right label)

In parallel, PE11 sends the ‘BGP withdraw message’ to RR/PE2, which will run the bestpath algorithm and removes the path learned via PE11, and then adjust their forwarding entries via PE12

This feature is independent of whether multipath is enabled on PE2 or not, however, dependent on VPN site multihoming

Agenda

- IP/VPN Overview
- IP/VPN Services
  1. Load-Sharing for Multihomed VPN Sites
  2. Hub and Spoke Service
  3. Extranet Service
  4. Internet Access Service
  5. IPv6 VPN Service
- Best Practices
- Conclusion

IP/VPN Services:

2. Hub and Spoke Service

- Many VPN deployments need to be hub and spoke
  – Spoke to spoke communication via Hub site only
- Despite MPLS based IP/VPN's implicit any-to-any, i.e., full-mesh connectivity, hub and spoke service can easily be offered
  – Done with import and export of route-target (RT) values
  – Requires unique RD per VRF per PE
- PE routers can run any routing protocol with VPN customer' hub and spoke sites independently

IP/VPN Services:

2. Hub and Spoke Service

- Two configuration Options :
    1. 1 PE-CE interface to Hub & 1 VRF;
    2. 2 PE-CE interfaces to Hub & 2 VRFs;
- Use option#1 if Hub site advertises default or summary routes towards the Spoke sites, otherwise use Option#2

- HDVRF feature* allows the option#2 to use just one PE-CE interface

* HDVRF Feature Is Discussed Later

Import and Export RT Values Must Be Different

Supported in IOS, NXOS and IOS-XR

IP/VPN Services:

2. Hub and Spoke Service: IOS Configuration – Option#1

```
ip vrf green-spoke1
 description VRF for SPOKE A
 rd 300:111
 route-target export  1:1
 route-target import  2:2
```

```
ip vrf HUB
 description VRF for HUB
 rd 300:11
 route-target import  1:1
 route-target export  2:2
```

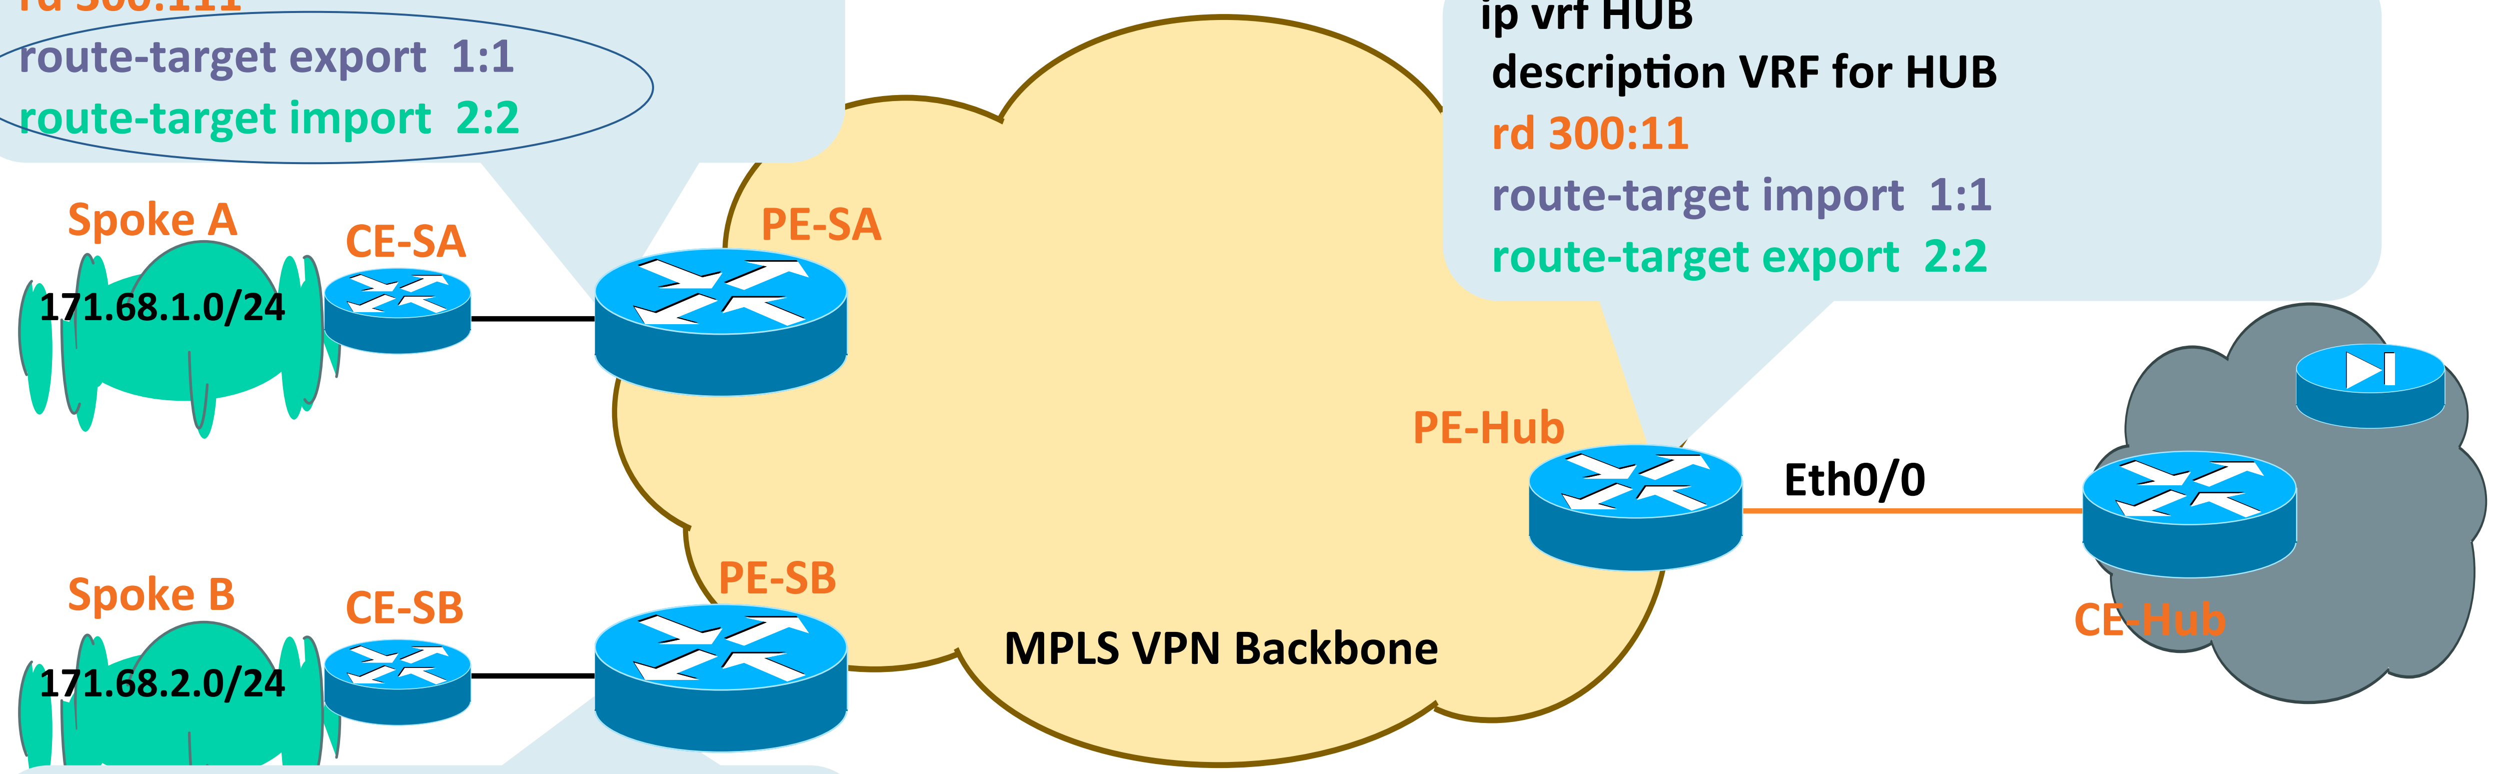


```
ip vrf green-spoke2
 description VRF for SPOKE B
 rd 300:112
 route-target export  1:1
 route-target import  2:2
```

- PE-Hub MUST advertise only default or aggregate route(s) to PE-SA/SB
- PE-Hub MUST NOT use bgp aggregation

Note: Only VRF Configuration Is Shown Here

IP/VPN Services:

2. Hub and Spoke Service: IOS Configuration – Option#2

Import and Export RT Values Must Be Different

```
ip vrf green-spoke1
 description VRF for SPOKE A
 rd 300:111
 route-target export  1:1
 route-target import  2:2
```

```
ip vrf HUB-IN
 description VRF for traffic from HUB
 rd 300:11
 route-target import  1:1
```

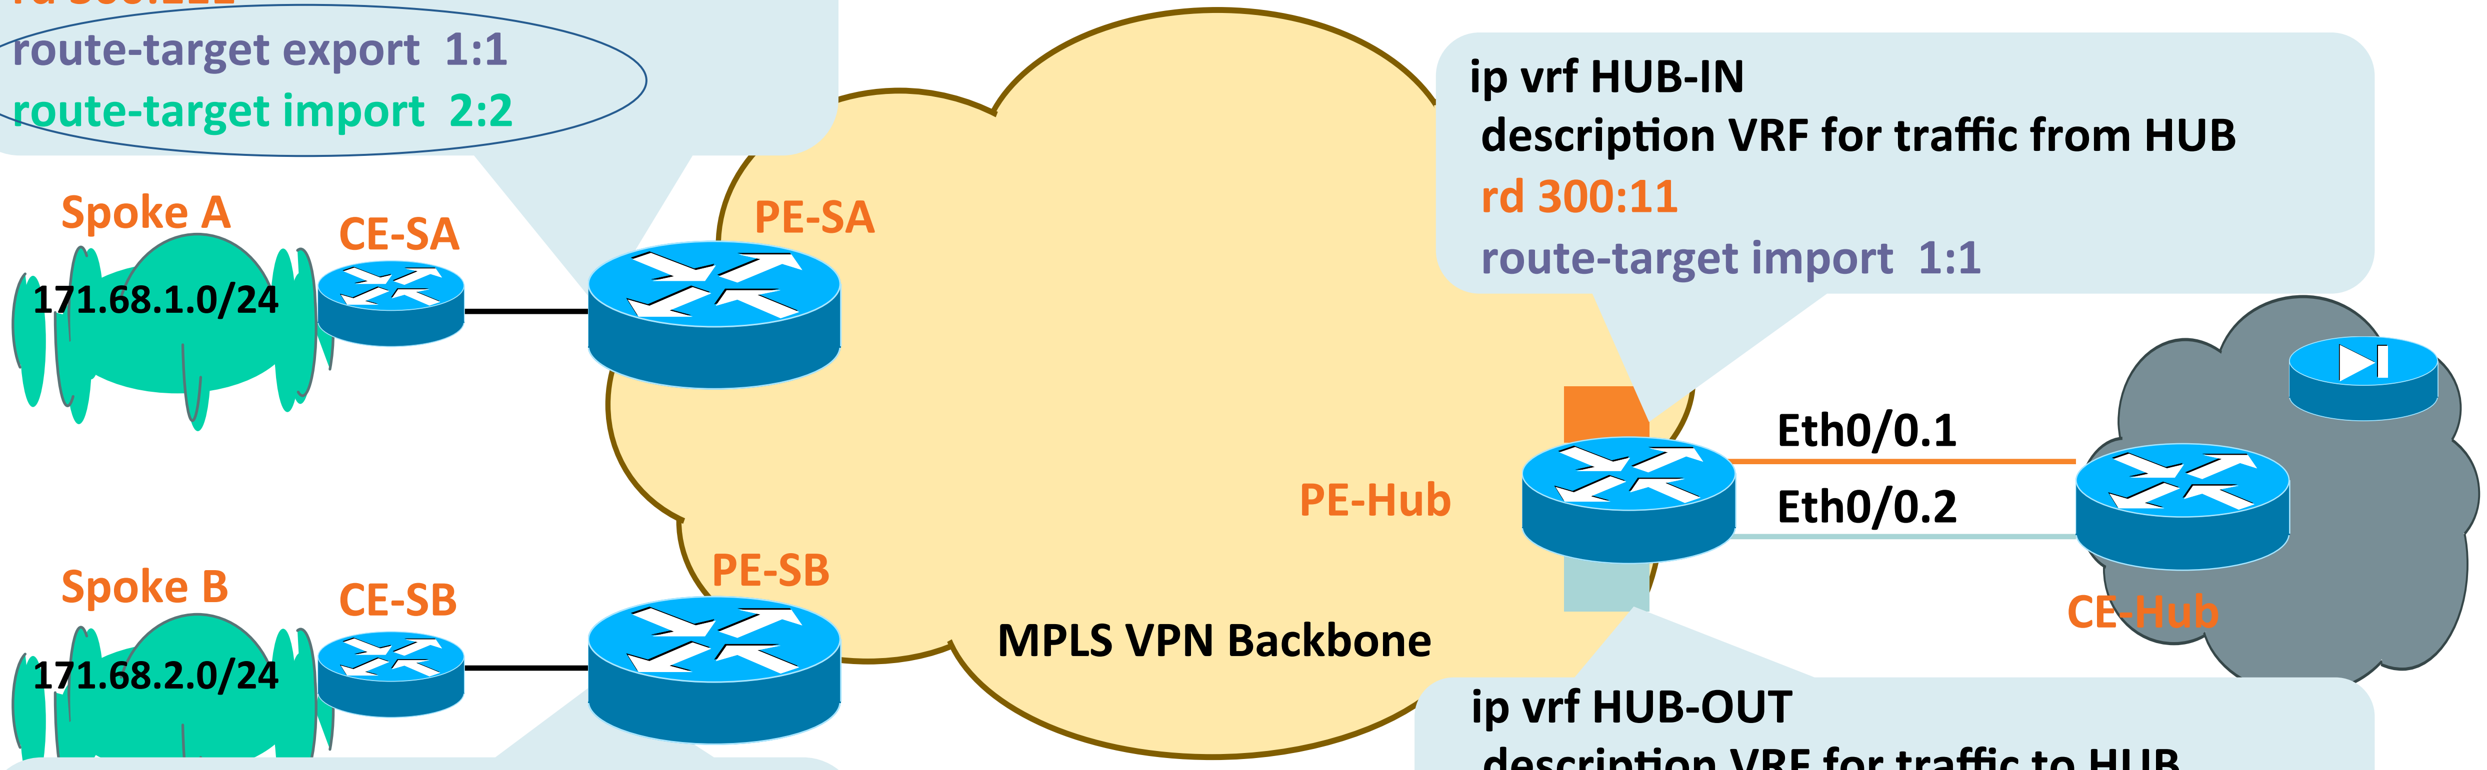


```
ip vrf green-spoke2
 description VRF for SPOKE B
 rd 300:112
 route-target export  1:1
 route-target import  2:2
```

```
ip vrf HUB-OUT
 description VRF for traffic to HUB
 rd 300:12
 route-target export  2:2
```

- **PE-Hub can advertise Spoke specific route(s) to PE-SA/SB.**
- **PE-Hub MAY use bgp aggregation.**

Note: Only VRF Configuration Is Shown Here

Supported in IOS, NXOS and IOS-XR

IP/VPN Services:

2. Hub and Spoke Service: Configuration – Option#2

- If BGP is used between every PE and CE, then allowas-in and as-override* knobs must be used at the PE_Hub**
  - Otherwise AS_PATH looping will occur

* Only If Hub and Spoke Sites Use the Same BGP ASN

** Configuration for This Is Shown on the Next Slide

Supported in IOS, NXOS and IOS-XR

IP/VPN Services:

2. Hub and Spoke Service: Configuration – Option#2

```
ip vrf green-spoke1
 description VRF for SPOKE A
 rd 300:111
 route-target export  1:1
 route-target import  2:2
```

```
router bgp <ASN>
 address-family ipv4 vrf HUB-IN
 neighbor <CE> as-override
```

```
ip vrf HUB-IN
 description VRF for traffic from HUB
 rd 300:11
 route-target import  1:1
```

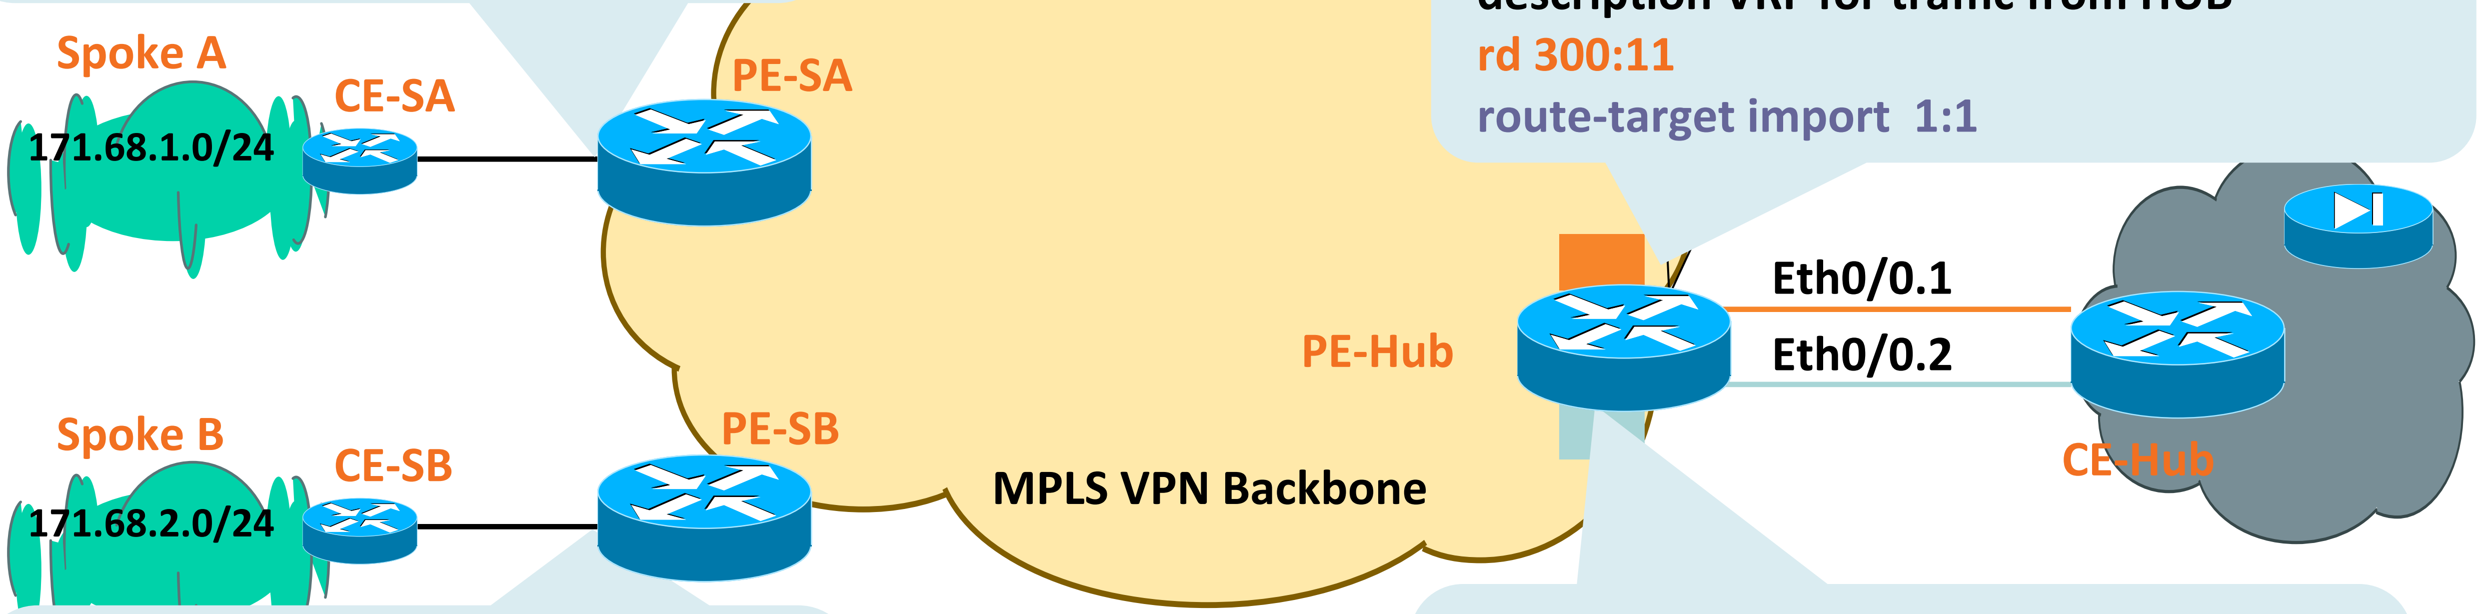


```
ip vrf green-spoke2
 description VRF for SPOKE B
 rd 300:112
 route-target export  1:1
 route-target import  2:2
```

```
ip vrf HUB-OUT
 description VRF for traffic to HUB
 rd 300:12
 route-target export  2:2
```

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

IP/VPN Services:

2. Hub and Spoke Service: Control Plane (Option#2)

Two VRFs at the PE-Hub:

- VRF HUB-IN to learn every spoke routes from remote PEs
- VRF HUB-OUT to advertise spoke routes or summary 171.68.0.0/16 routes to remote PEs

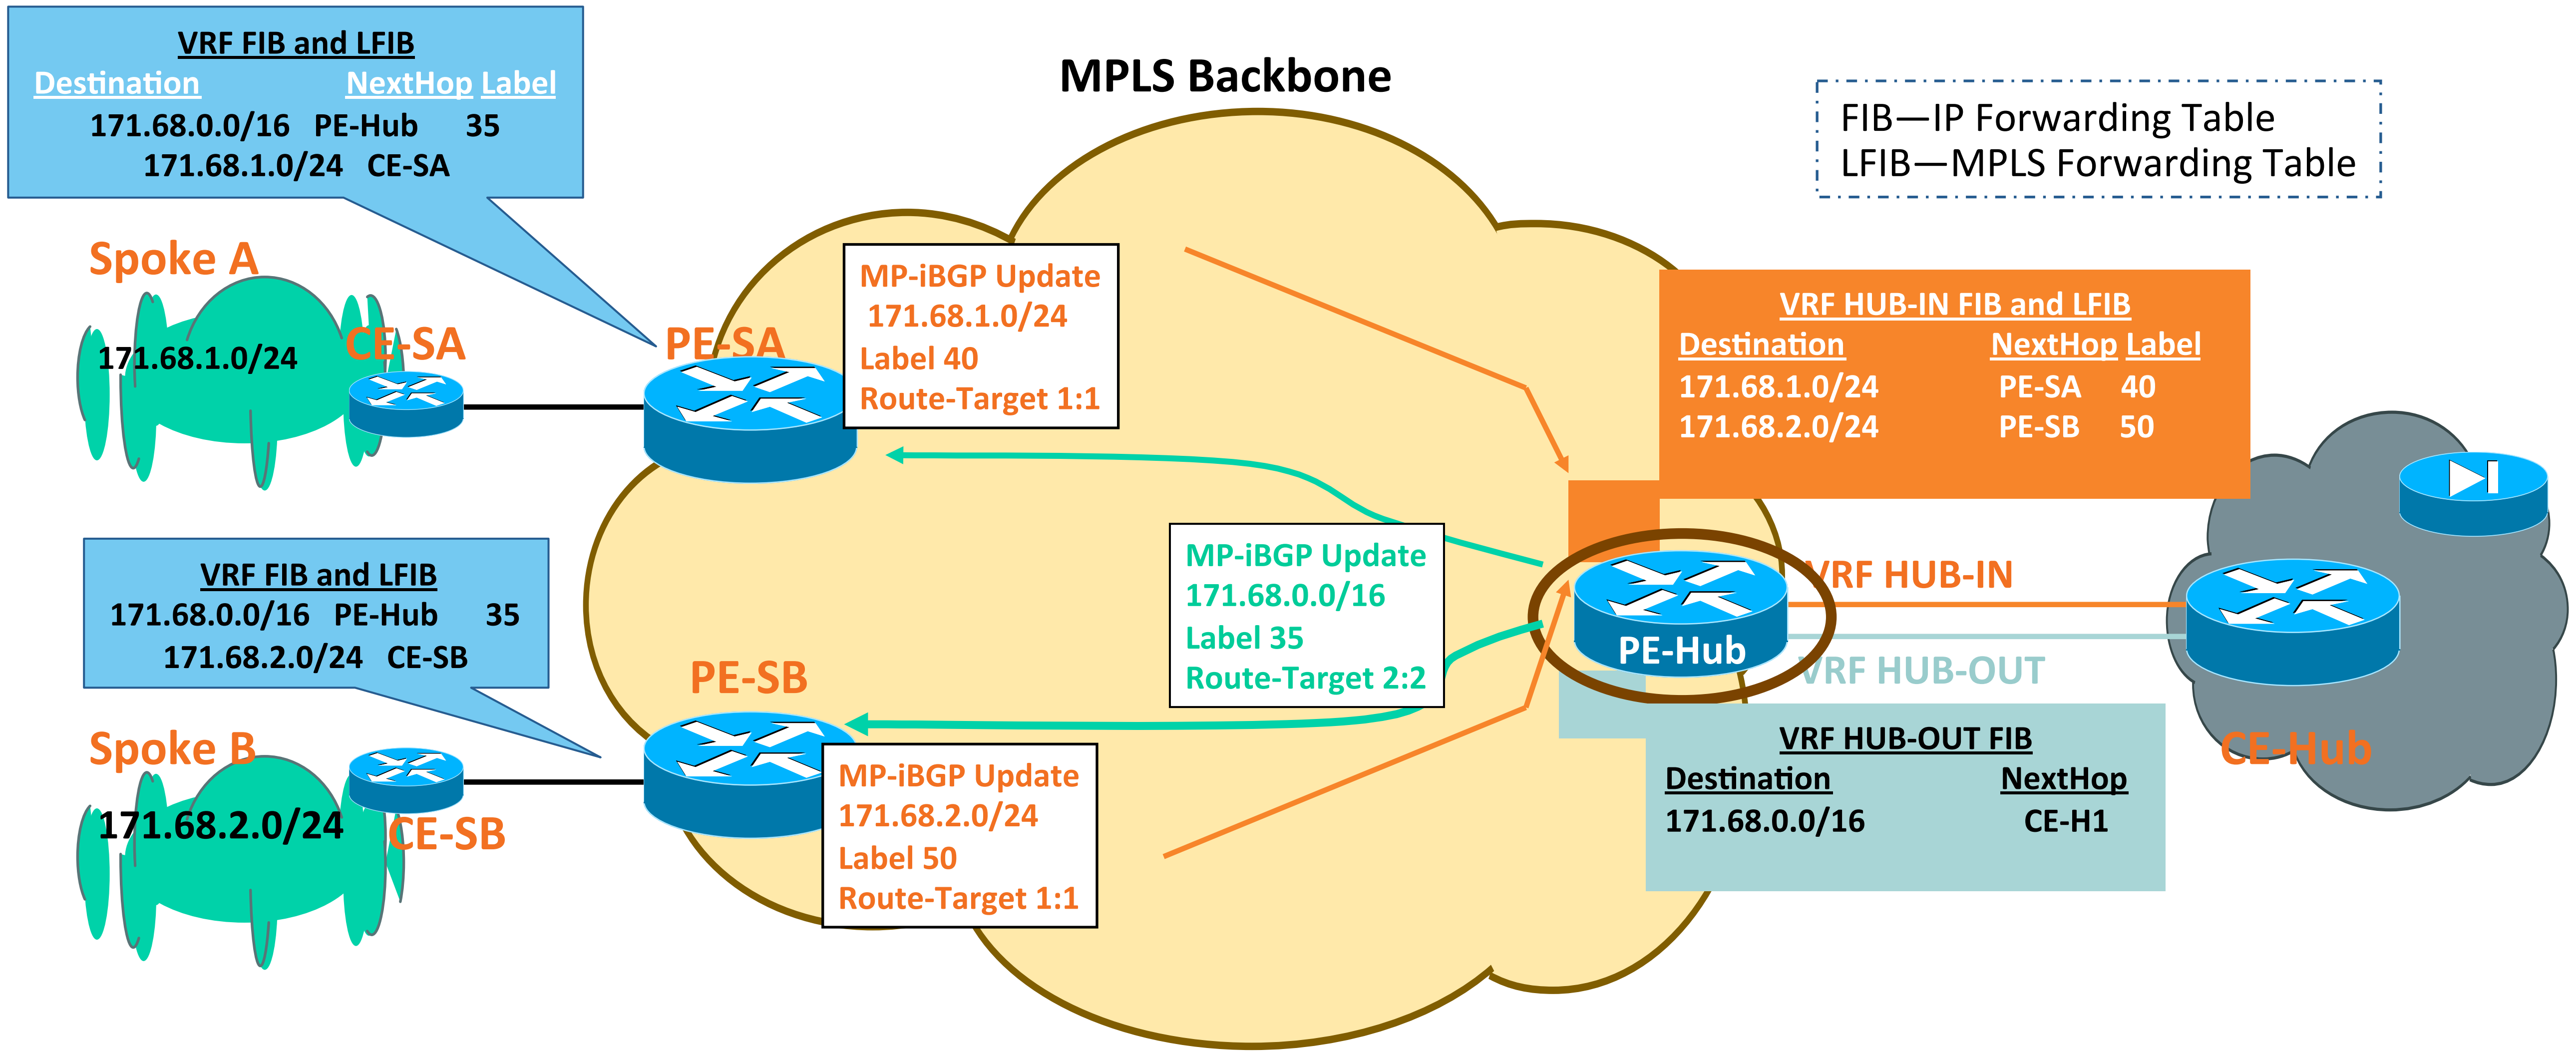

Supported in IOS, NXOS and IOS-XR

IP/VPN Services:

2. Hub and Spoke Service: Forwarding Plane (Option#2)

This Is How the Spoke-to-Spoke Traffic Flows

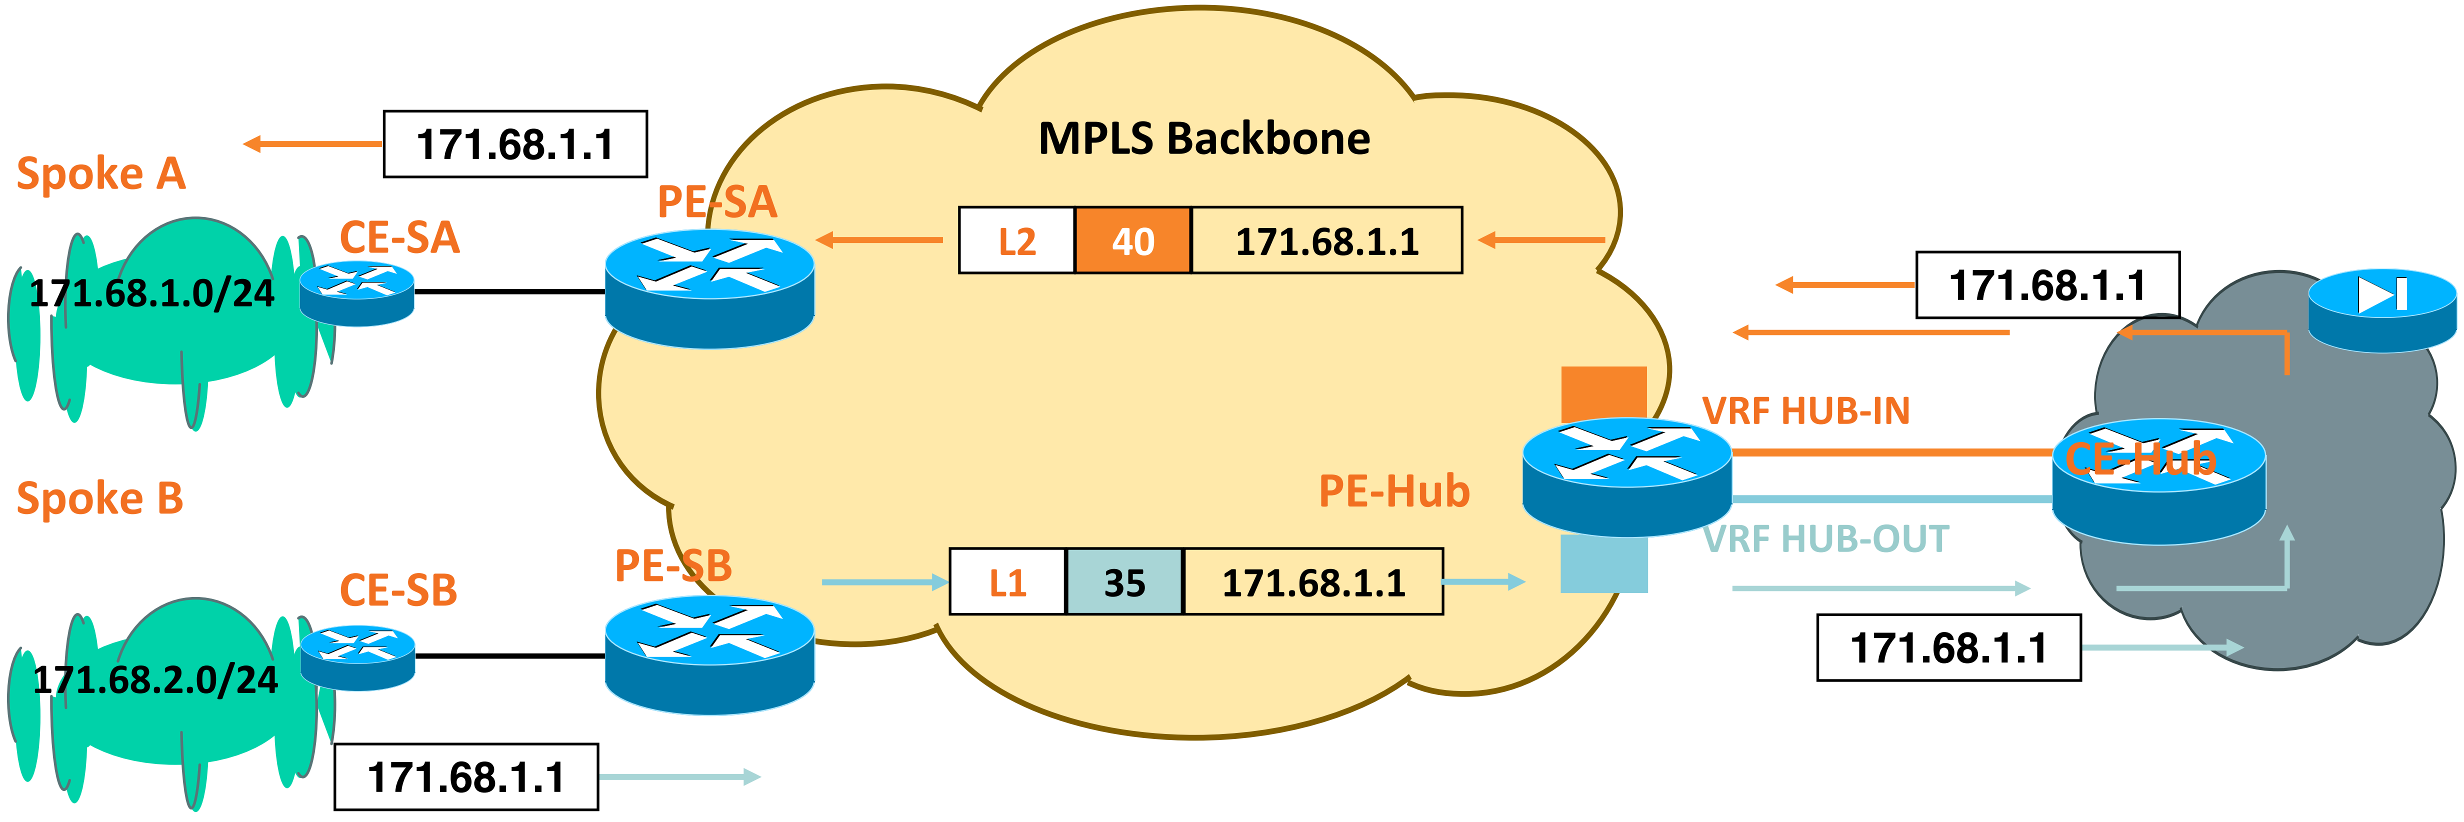

L1 Is the Label to Get to PE-Hub

L2 Is the Label to Get to PE-SA

IP/VPN Services:

2. What If Many Spoke Sites Connect to the Same PE Router?

- If more than one spoke router (CE) connects to the same PE router (within the same VRF), then such spokes can reach other without needing the hub.
  – Defeats the purpose of hub and spoke ☹

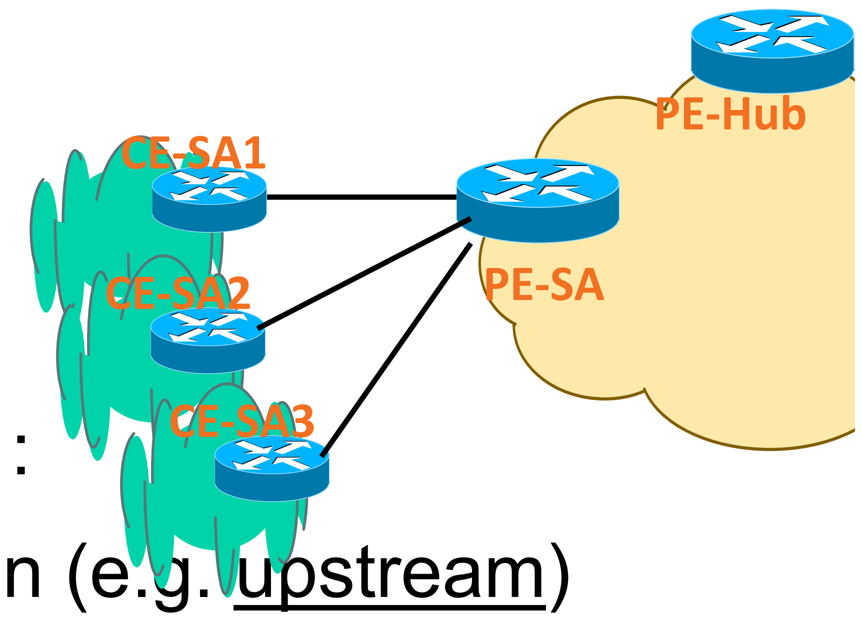

- Half-duplex VRF is the answer
  – Uses two VRFs on the PE (spoke) router :
  – A VRF for spoke->hub communication (e.g. upstream)
  – A VRF for spoke<-hub communication (e.g. downstream)

Note: 12.2(33) SRE Supports Any Interface Type (Eth, Ser, POS, Virtual-Access, etc.)

Supported in IOS

IP/VPN Services:

2. Hub and Spoke Service: Half-Duplex VRF

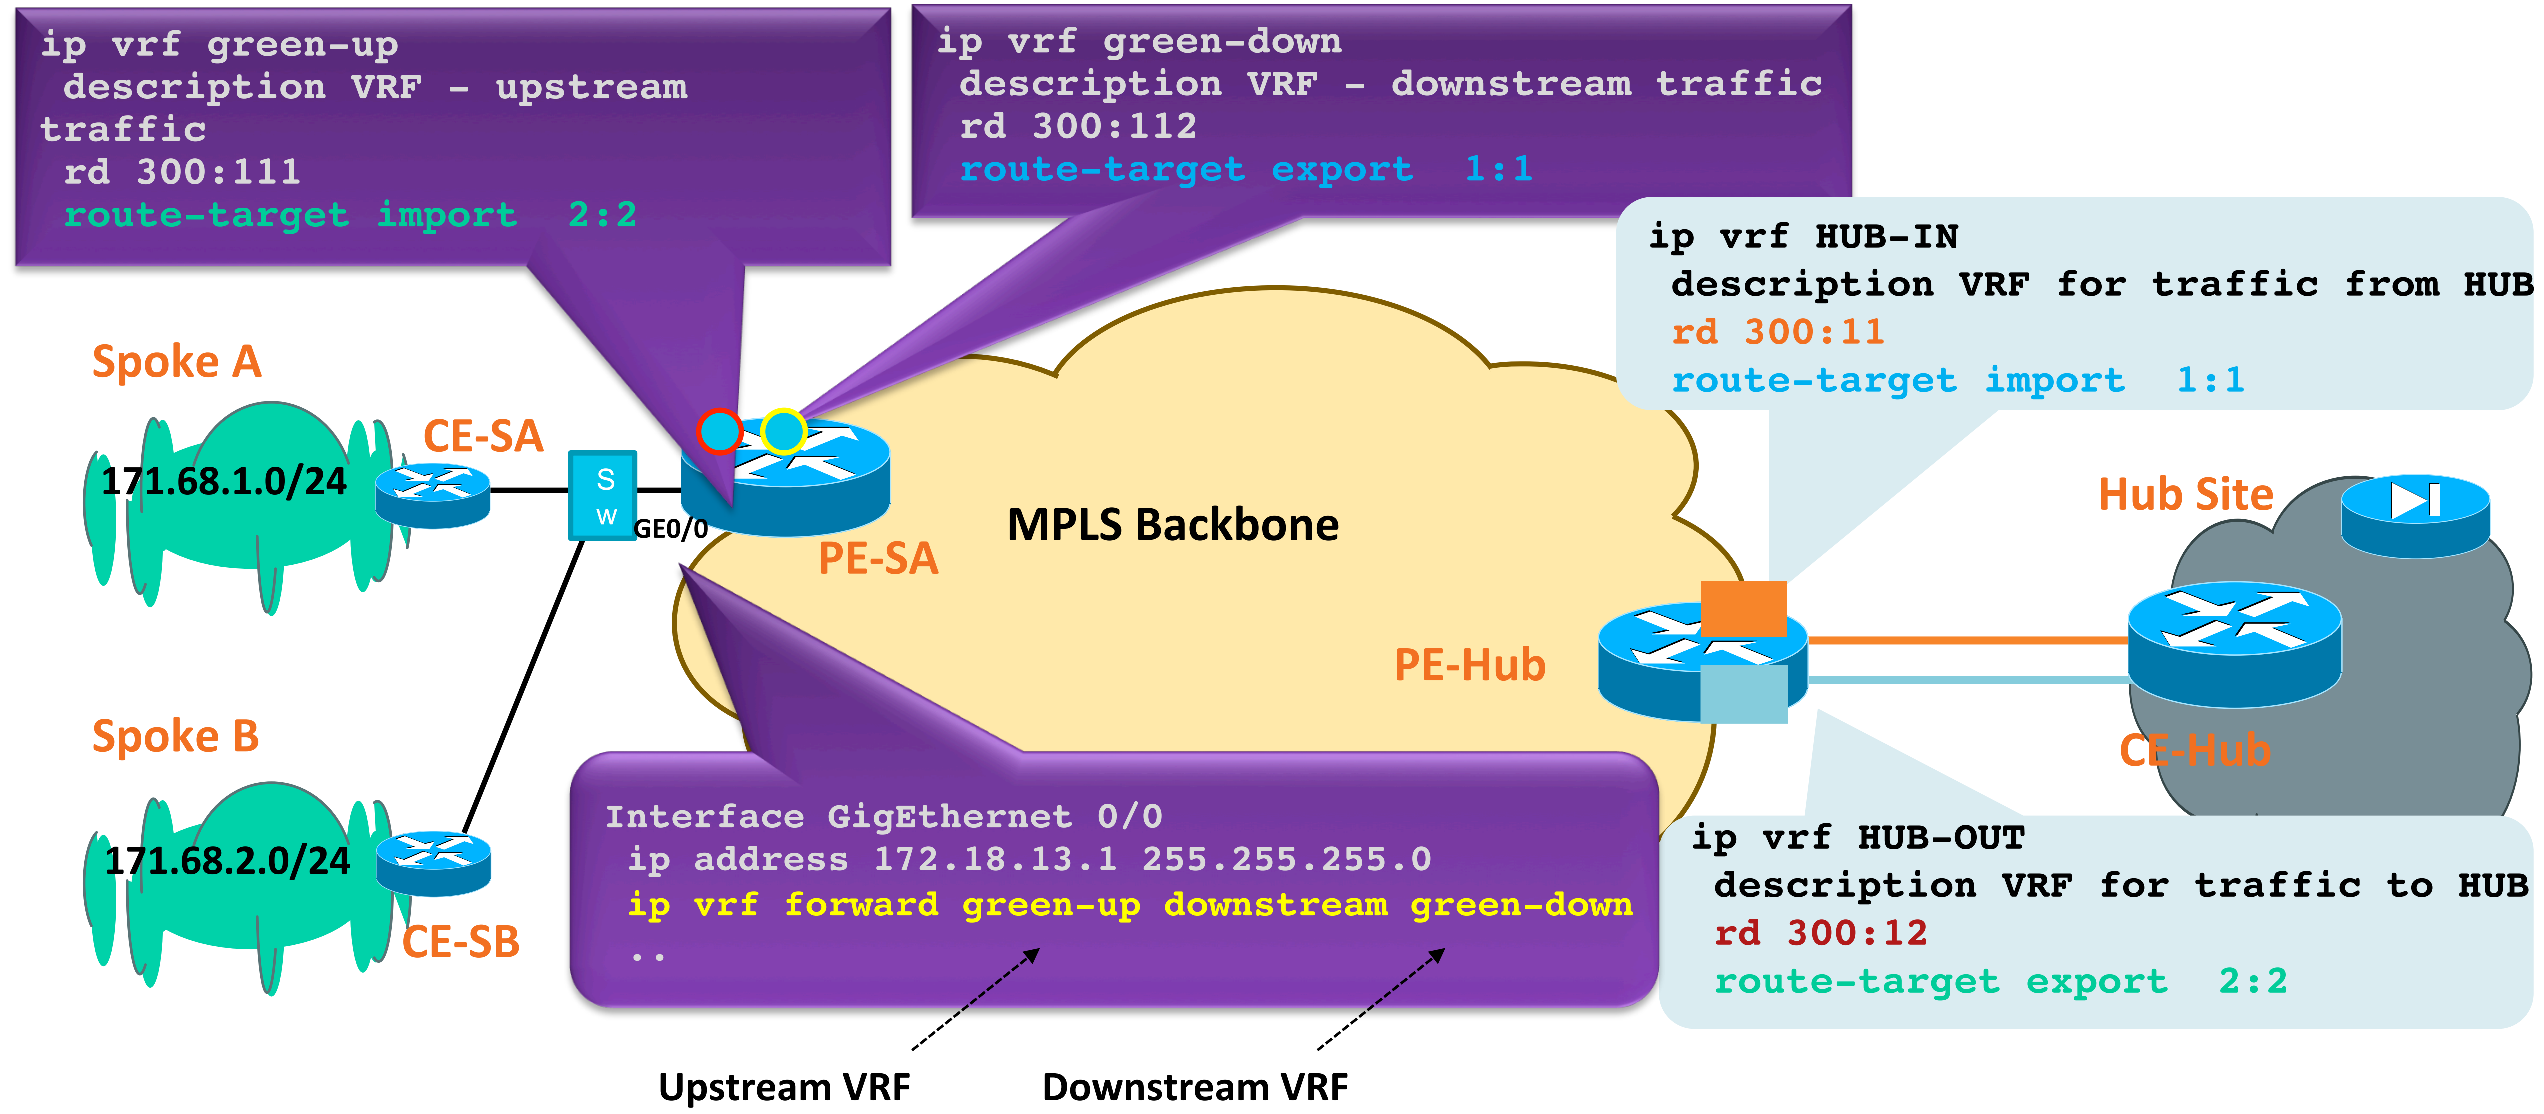

1. PE-SA installs the Spoke routes only in downstream VRF i.e. green-down
2. PE-SA installs the Hub routes only in upstream VRF i.e. green-up
3. PE-SA forwards the incoming IP traffic (from Spokes) using upstream VRF i.e. green-up routing table.
4. PE-SA forwards the incoming MPLS traffic (from Hub) using downstream VRF i.e. green-down routing table

Agenda

- IP/VPN Overview
- IP/VPN Services
  1. Load-Sharing for Multihomed VPN Sites
  2. Hub and Spoke Service
  3. Extranet Service
  4. Internet Access Service
  5. IPv6 VPN Service
- Best Practices
- Conclusion

MPLS-VPN Services

3. Extranet VPN

- MPLS based IP/VPN, by default, isolates one VPN customer from another
  - Separate virtual routing table for each VPN customer
- Communication between VPNs may be required i.e., extranet
  - External intercompany communication (dealers with manufacturer, retailer with wholesale provider, etc.)
  - Management VPN, shared-service VPN, etc.
- Needs to share the import and export route-target (RT) values within the VRFs of extranets.
  - Export-map or import-map may be used for advanced extranet.

Supported in IOS, NXOS and IOS-XR

MPLS-VPN Services

3. Extranet VPN – Simple Extranet (IOS Config sample)

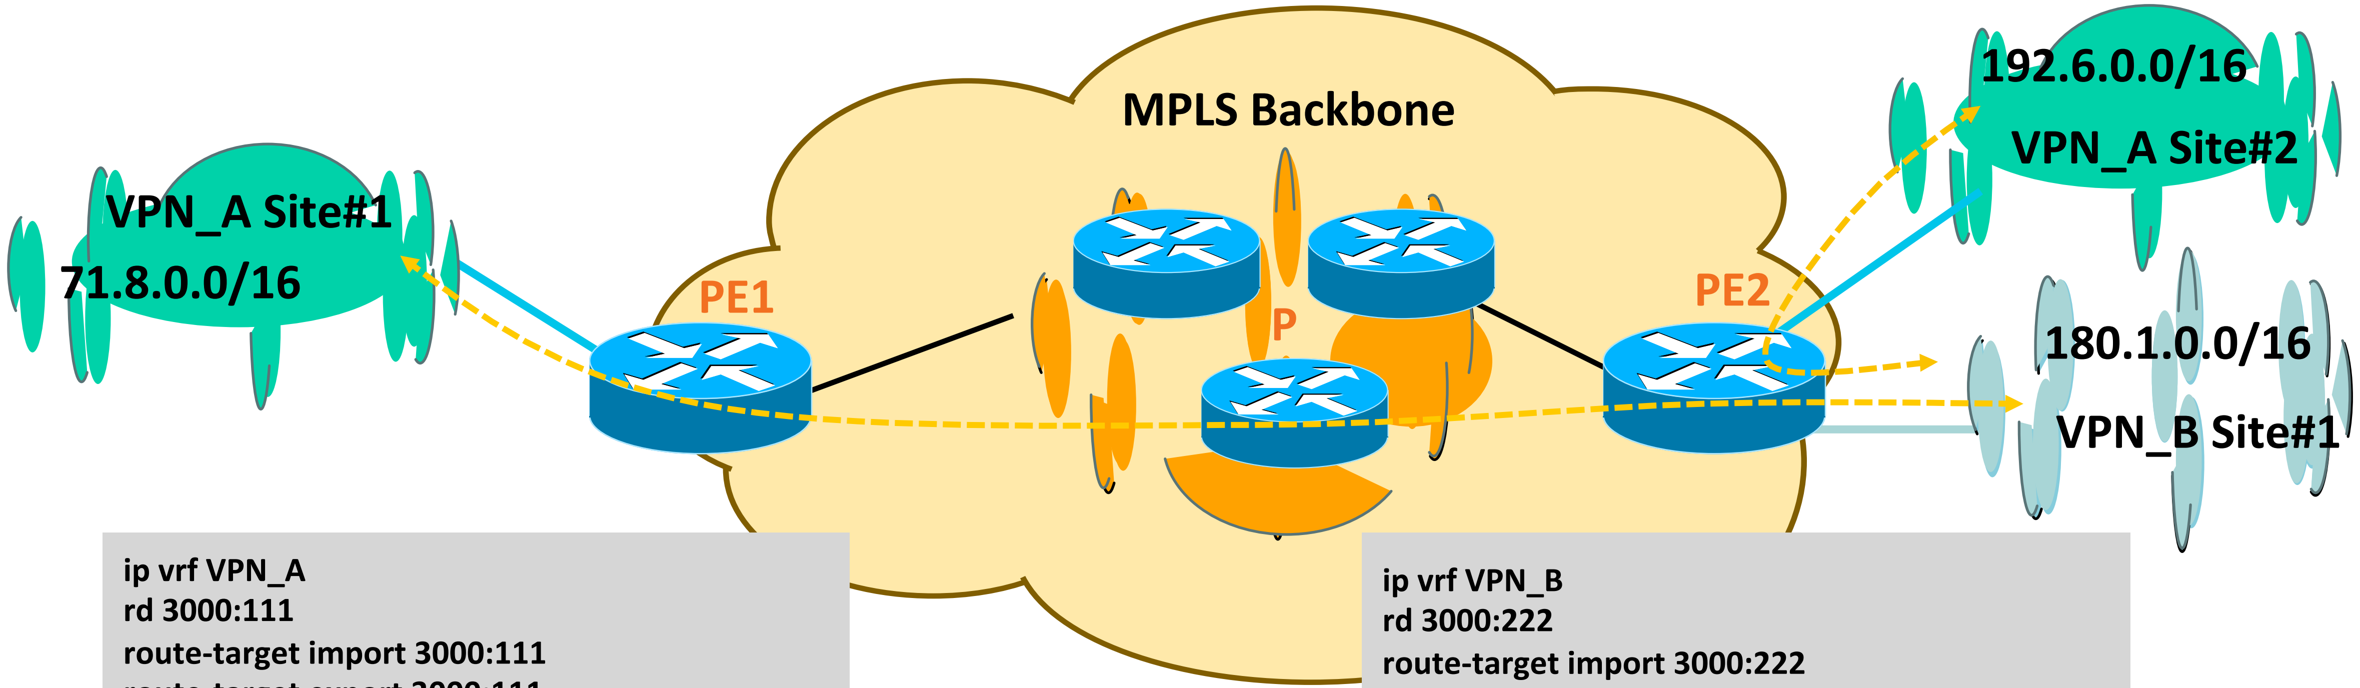


```
ip vrf VPN_A
rd 3000:111
route-target import 3000:111
route-target export 3000:111
route-target import 3000:222
```

```
ip vrf VPN_B
rd 3000:222
route-target import 3000:222
route-target export 3000:222
route-target import 3000:111
```

All Sites of Both VPN_A and VPN_B Can Communicate with Each Other

Supported in IOS, NXOS and IOS-XR

MPLS-VPN Services

3. Extranet VPN – Advanced Extranet (IOS Config sample)

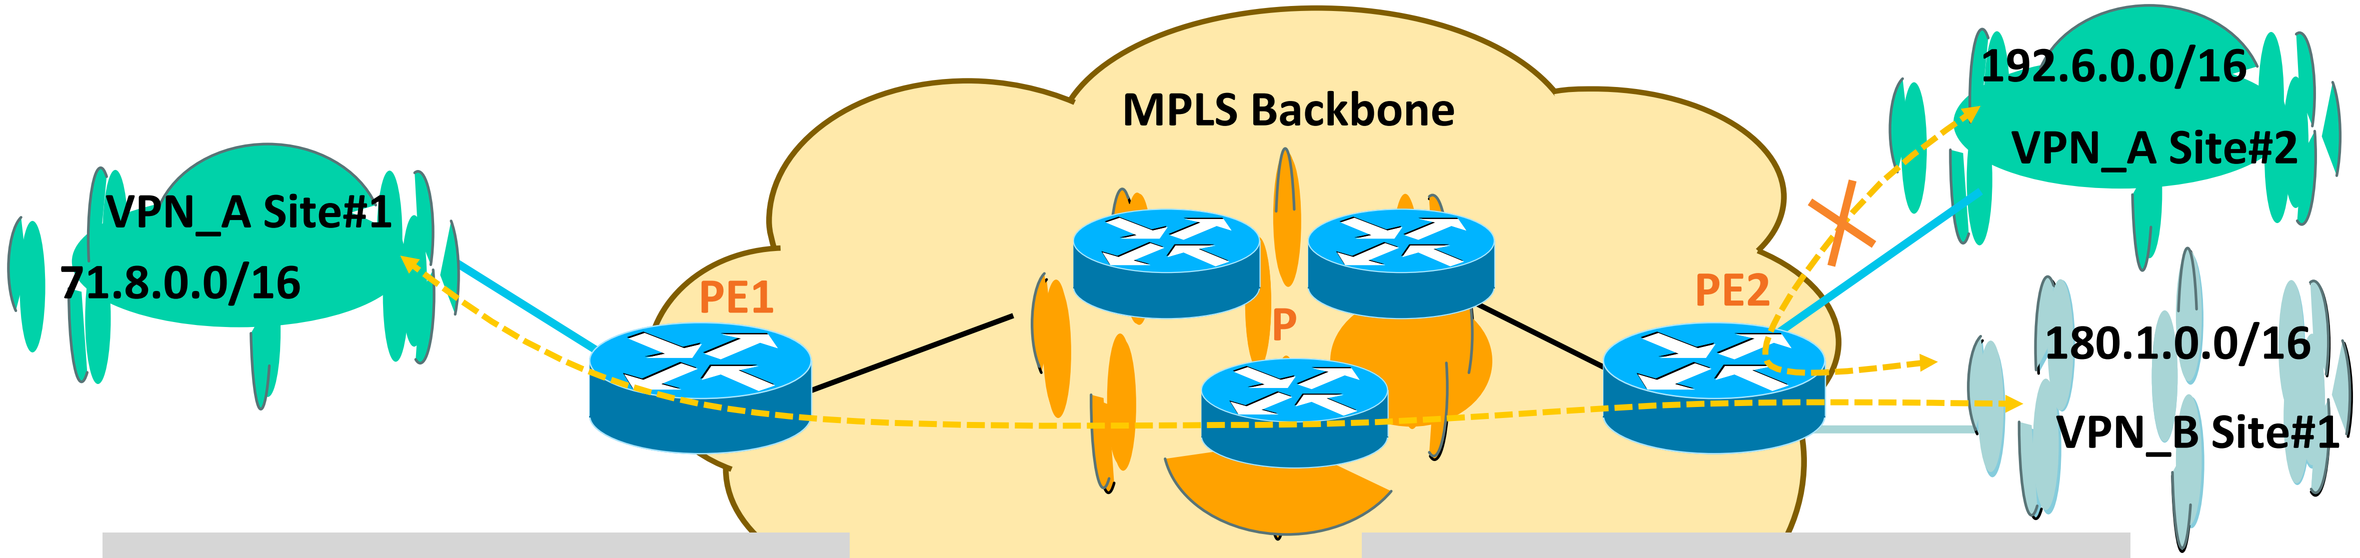


```
ip vrf VPN_A
rd 3000:111
route-target import 3000:111
route-target export 3000:111
route-target import 3000:1
import map VPN_A_Import
export map VPN_A_Export
!
route-map VPN_A_Export permit 10
 match ip address 1
 set extcommunity rt  3000:2 additive
!
route-map VPN_A_Import permit 10
 match ip address 2
!
access-list 1 permit 71.8.0.0 0.0.0.0
access-list 2 permit 180.1.0.0 0.0.0.0
```

```
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 71.8.0.0 0.0.0.0
access-list 2 permit 180.1.0.0 0.0.0.0
```

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

Only Site #1 of Both VPN_A and VPN_B Would Communicate with Each Other

Agenda

- IP/VPN Overview
- IP/VPN Services
  1. Load-Sharing for Multihomed VPN Sites
  2. Hub and Spoke Service
  3. Extranet Service
  4. Internet Access Service
  5. IPv6 VPN Service
- Best Practices
- Conclusion

MPLS-VPN Services

4. 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
- VPN customers benefit from the single point of contact for both Intranet and Internet connectivity

MPLS-VPN Services

4. Internet Access: Design Options

Four Options to Provide the Internet Service -

1. VRF specific default route with “global” keyword
2. Separate PE-CE sub-interface (non-VRF)
3. Extranet with Internet-VRF
4. VRF-aware NAT

MPLS-VPN Services

4. Internet Access: Design Options

1. VRF specific default route
   - 1.1 Static default route to move traffic from VRF to Internet (global routing table)
   - 1.2 Static routes for VPN customers to move traffic from Internet (global routing table) to VRF
2. Separate PE-CE subinterface (non-VRF)
   - May run BGP to propagate Internet routes between PE and CE
3. Extranet with Internet-VRF
   - VPN packets never leave VRF context; issue with overlapping VPN address
4. Extranet with Internet-VRF along with VRF-aware NAT
   - VPN packets never leave VRF context; works well with overlapping VPN address

IP/VPN Services: Internet Access

4.1 Option#1: VRF Specific Default Route

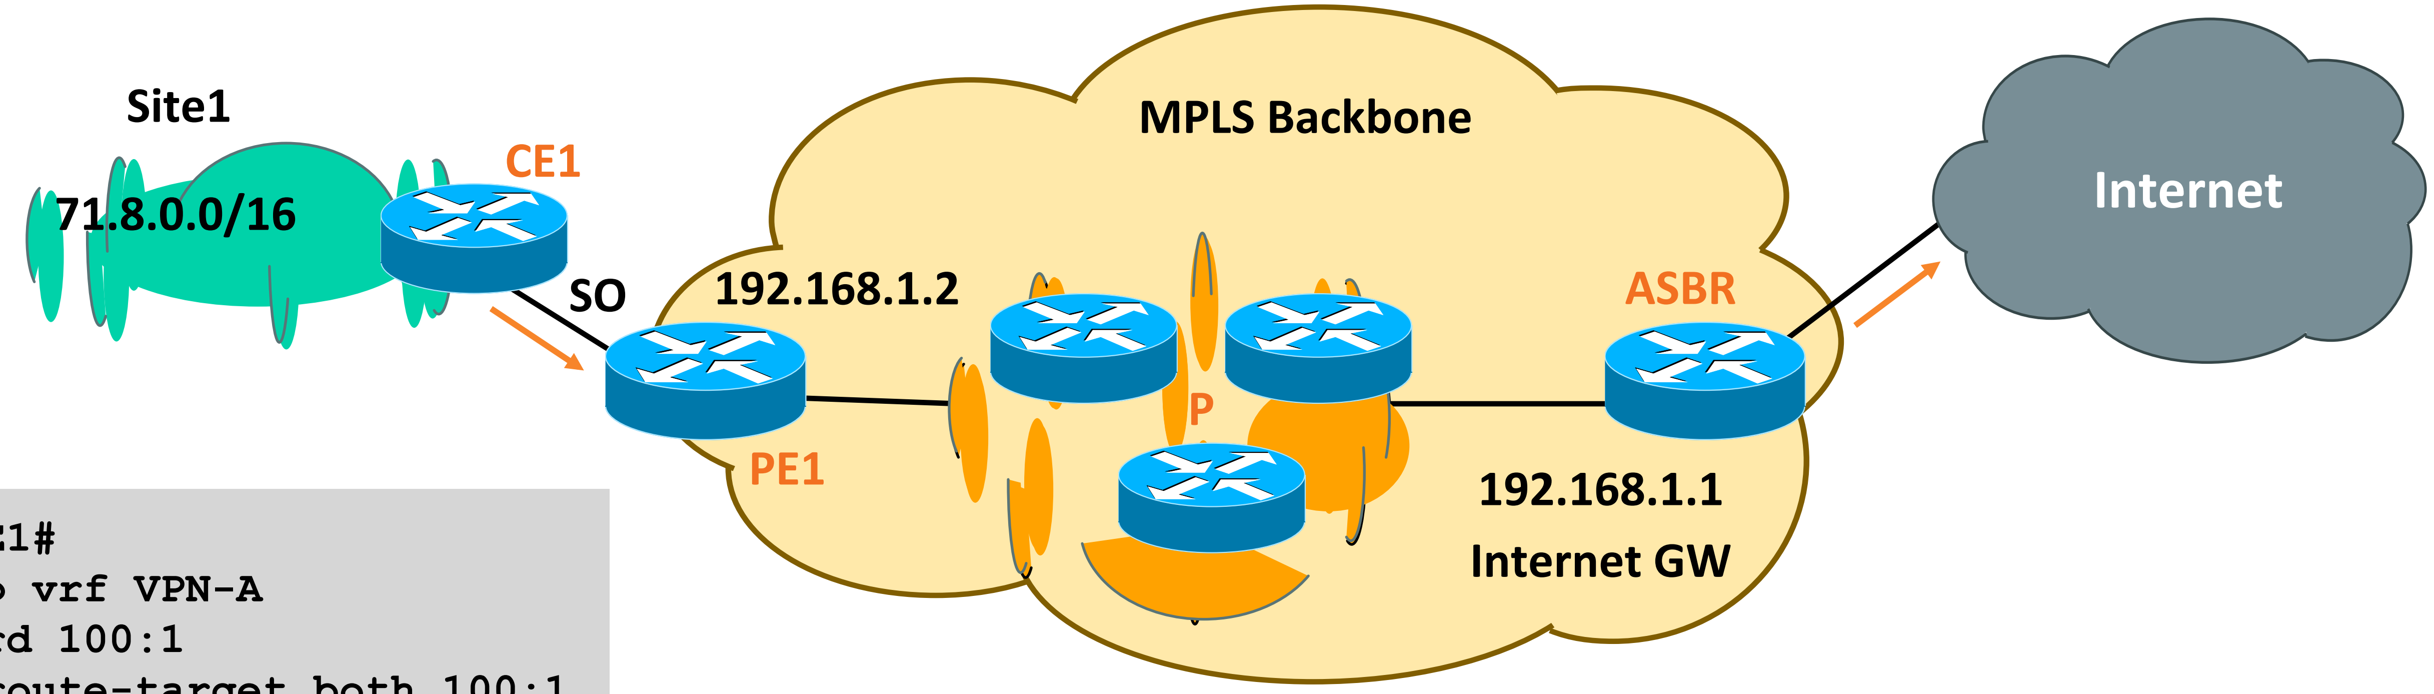


```
PE1#
ip vrf VPN-A
 rd 100:1
 route-target both 100:1
```

```
Interface Serial0
ip address 192.168.10.1 255.255.255.0
ip vrf forwarding VPN-A
```

```
Router bgp 100
 no bgp default ipv4-unicast
 redistribute static
 neighbor 192.168.1.1 remote 100
 neighbor 192.168.1.1 activate
 neighbor 192.168.1.1 next-hop-self
 neighbor 192.168.1.1 update-source loopback0
```

```
ip route vrf VPN-A 0.0.0.0 0.0.0.0 192.168.1.1 global
ip route 71.8.0.0 255.255.0.0 Serial0
```

- A default route, pointing to the ASBR, is installed into the site VRF at each PE
- The static route, pointing to the VRF interface, is installed in the global routing table and redistributed into BGP

IP/VPN Services: Internet Access

4.1 Option#1: VRF Specific Default Route (Forwarding)

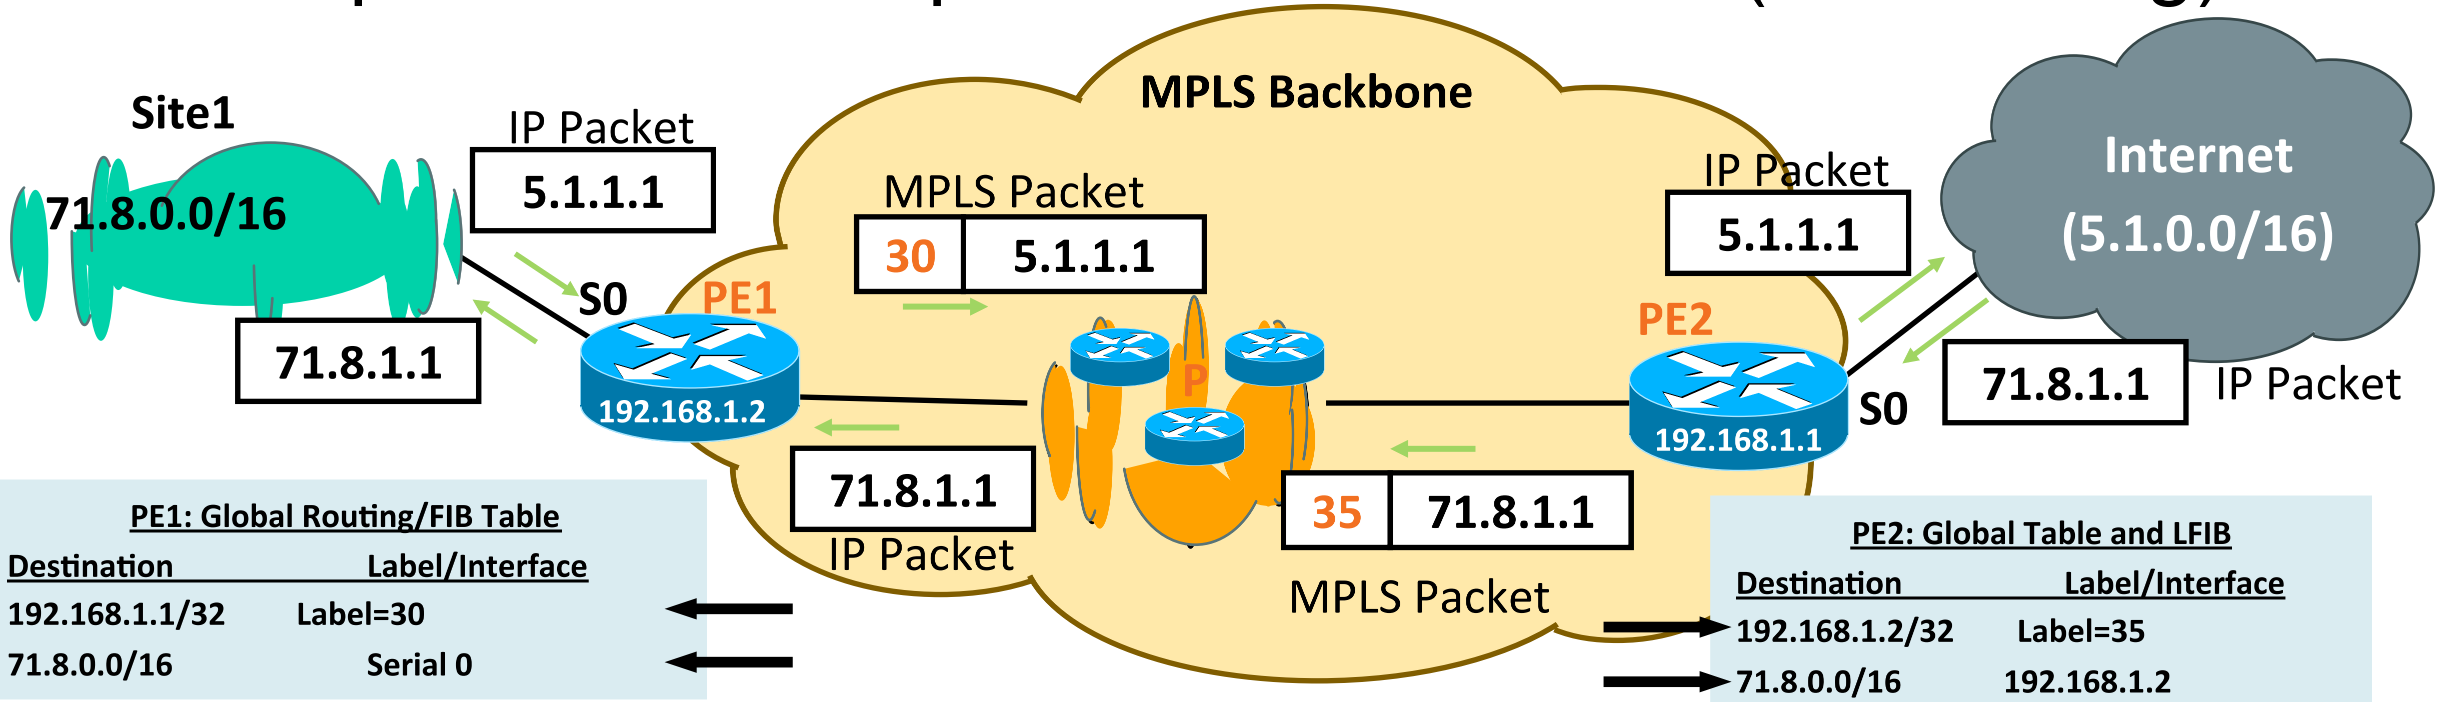

PE1: Global Routing/FIB Table

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

PE1: VRF Routing/FIB Table

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

Pros

- Different Internet gateways
- Can be used for different VRFs
- PE routers need not to hold the Internet table
- Simple configuration

Cons

- Using default route for Internet
- Routing does not allow any other default route for intra-VPN routing Increasing size of global routing table by leaking VPN routes
- Static configuration (possibility of traffic blackholing)

Supported in IOS, NXOS and IOS-XR

IP/VPN Services: Internet Access

4.2 Option#2: Separate PE-CE Subinterfaces

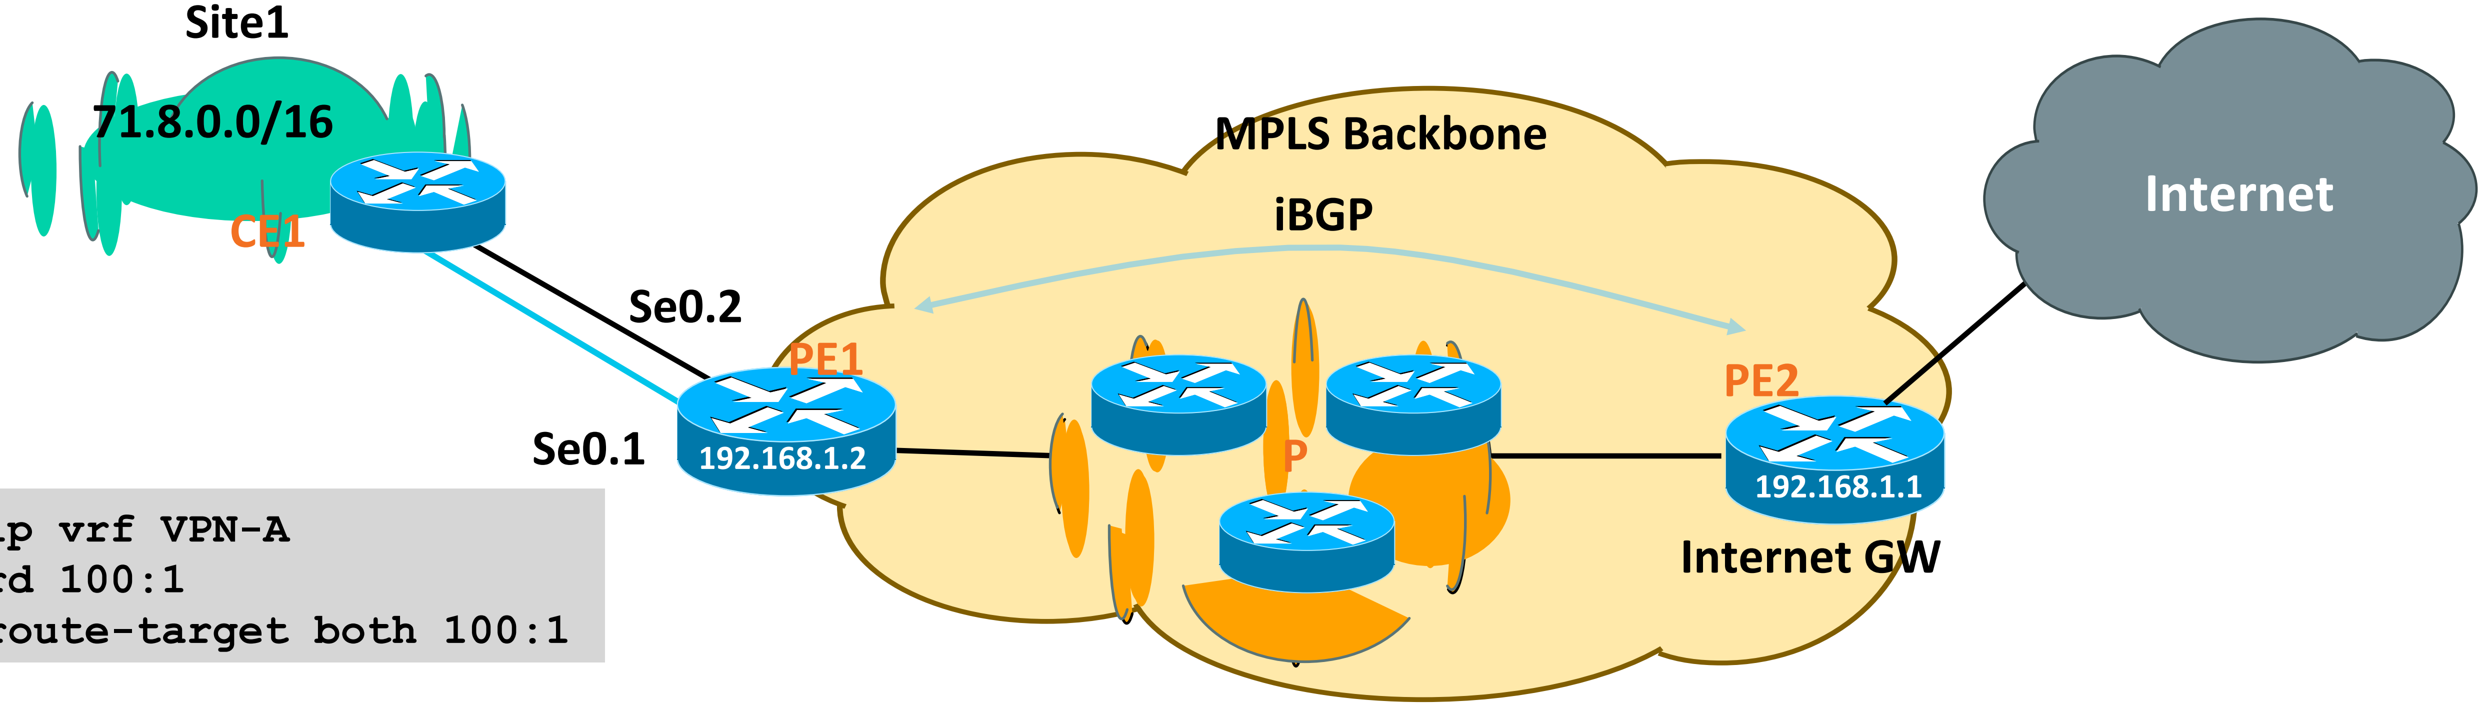


```
ip vrf VPN-A
rd 100:1
route-target both 100:1
```

```
Interface Serial0.1
 ip vrf forwarding VPN-A
 ip address 192.168.20.1 255.255.255.0
 frame-relay interface-dlci 100
!
Interface Serial0.2
 ip address 71.8.10.1 255.255.0.0
 frame-relay interface-dlci 200
!
```

```
Router bgp 100
no bgp default ipv4-unicast
neighbor 71.8.10.2 remote-as 502
```

- PE1-CE1 has one sub-interface associated to a VRF for VPN routing
- PE1-CE has another subinterface (global) for Internet routing
- PE1 may have eBGP peering with CE1 over the global interface and advertise full Internet routes or a default route to CE1
- PE2 must advertise VPN/site1 routes to the Internet.

Supported in IOS, NXOS and IOS-XR

IP/VPN Services: Internet Access

4.2 Option#2: Separate PE-CE Subinterfaces (Forwarding)

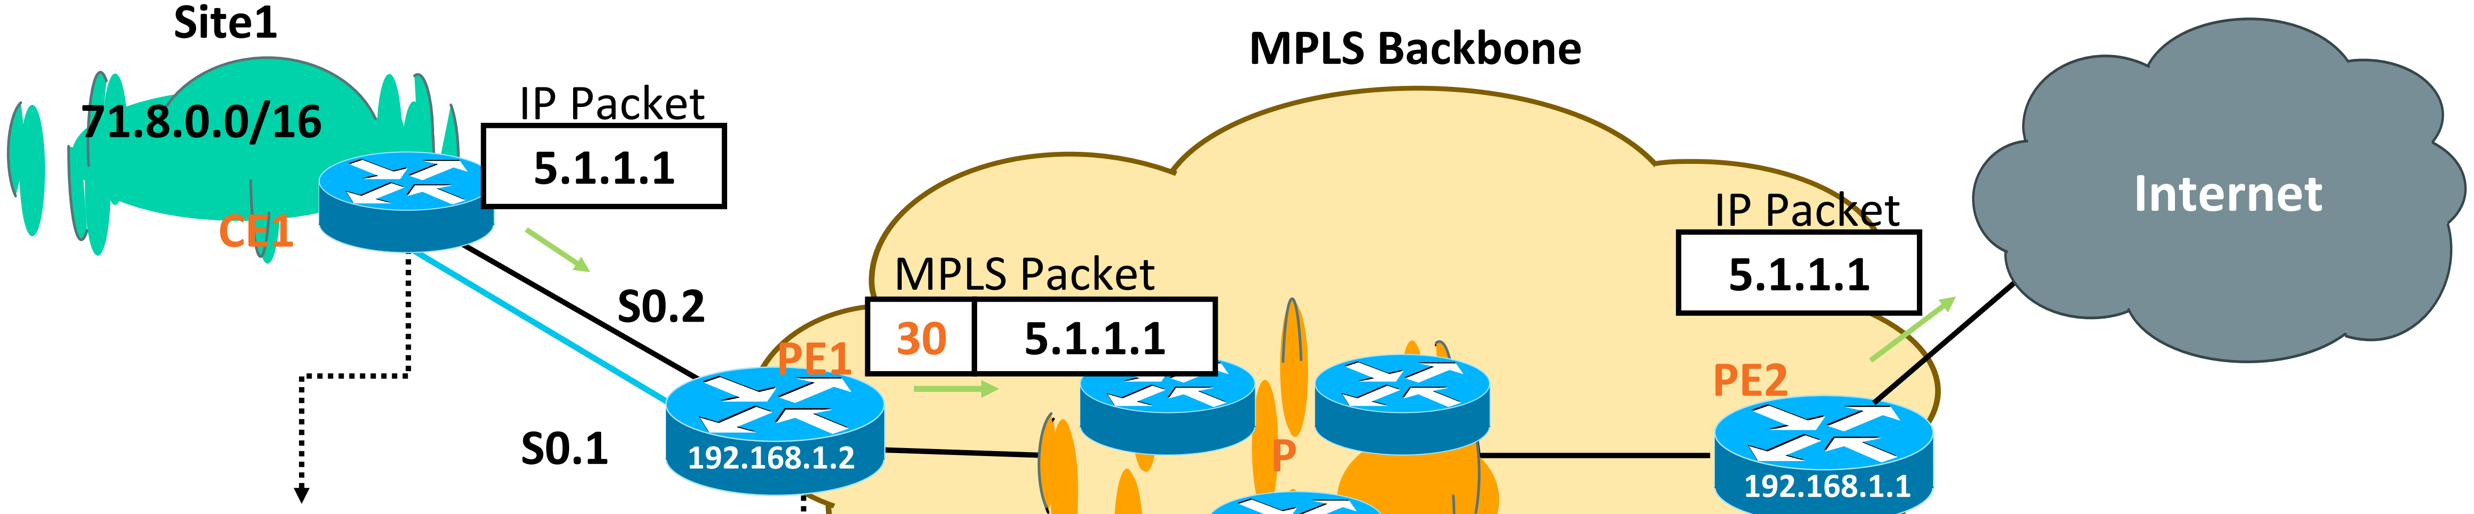

CE Routing Table

VPN Routes	Serial0.1
Internet Routes	Serial0.2

PE1 Global Table and FIB

Internet Routes	192.168.1.1
192.168.1.1	Label=30

Pros

1. CE is dual-homed and can perform Optimal Routing

2. Traffic Separation Done by CE

Cons

1. PE to Hold Full Internet Routes or default route via the Internet GW

. BGP Complexities Introduced at CE; CE1 May Need to Aggregate to Avoid AS_PATH Looping

Supported in IOS, NXOS and IOS-XR

IP/VPN Services: Internet Access

4.3 Option#3: Extranet with Internet-VRF

- The Internet routes could be placed within the VRF at the Internet-GW i.e., ASBR
- VRFs for customers could 'extranet' with the Internet VRF and receive either *default*, *partial* or full Internet routes
  - Default route is recommended
- Be careful if multiple customer VRFs, at the same PE, are importing full Internet routes
- Works well only if the VPN customers don't have overlapping addresses

IP/VPN Services: Internet Access

4.4 Option#4: Using VRF-Aware NAT

- If the VPN customers need Internet access without Internet routes, then VRF-aware NAT can be used at the Internet-GW i.e., ASBR
- The Internet GW doesn't need to have Internet routes either
- Overlapping VPN addresses is no longer a problem
- More in the "VRF-aware NAT" slides…

Agenda

- IP/VPN Overview
- IP/VPN Services
    1. Load-Sharing for Multihomed VPN Sites
    2. Hub and Spoke Service
    3. Extranet Service
    4. Internet Access Service
    5. IPv6 VPN Service
- Best Practices
- Conclusion

Supported in IOS, NXOS and IOS-XR

IP/VPN Services:

11. IPv6 VPN Service

- Similar to IPv4 VPN, IPv6 VPN can also be offered.
  – Referred to as “IPv6 VPN Provider Edge (6VPE)”.
- No modification on the MPLS core
  – Core can stay on IPv4
- PE-CE interface can be single-stack IPv6 or dual-stack
  – IPv4 and IPv6 VPNs can be offered on the same PE-CE interface
- Config and operation of IPv6 VPN are similar to IPv4 VPN

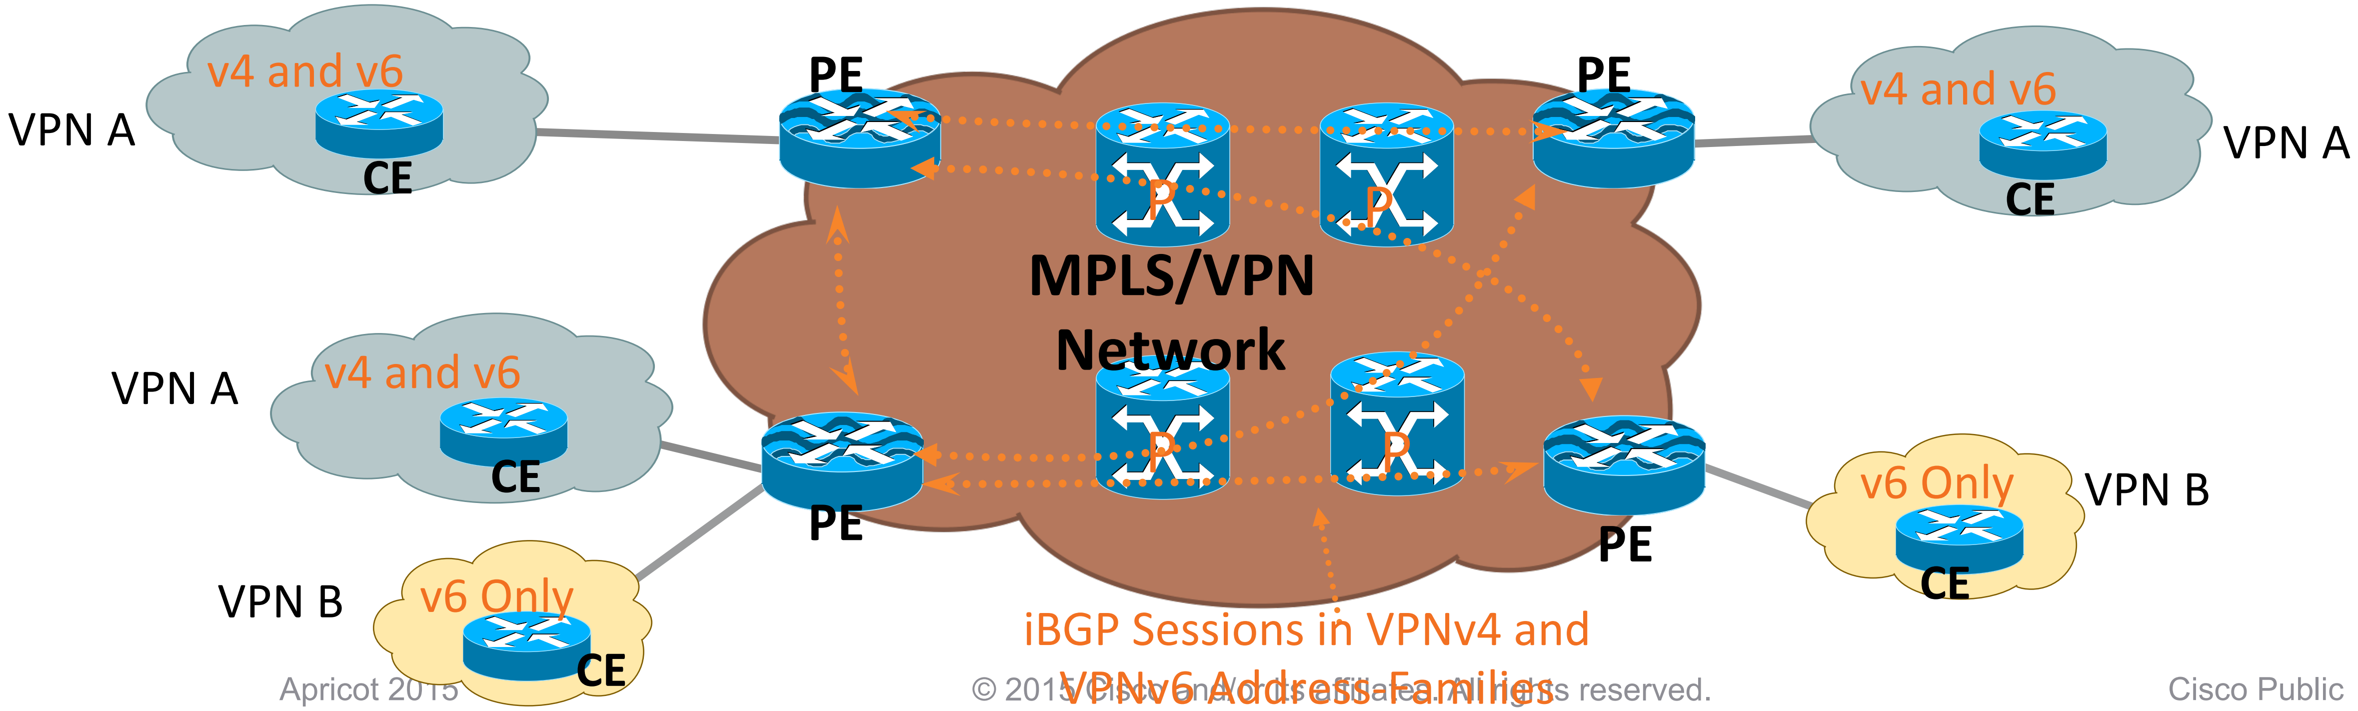

Supported in IOS, NXOS and IOS-XR

IP/VPN Services:

11. IPv6 VPN Service

```
IOS_PE#
!
vrf definition v2
 rd 2:2
 !
 address-family ipv6
  route-target export 2:2
  route-target import 2:2
!
router bgp 1
!
 address-family vpnv6
  neighbor 10.13.1.21 activate
  neighbor 10.13.1.21 send-community both
!
 address-family ipv6 vrf v2
  neighbor 200::2 remote-as 30000
  neighbor 200::2 activate
!
```

```
IOS-XR_PE#
!
vrf v2
!
 address-family ipv6 unicast
  route-target export 2:2
  route-target import 2:2
!
router bgp 1
 address-family vpnv6 unicast
 !
 neighbor 10.13.1.21
  remote-as 30000
  address-family vpnv6 unicast
!
 vrf v2
  rd 2:2
  address-family ipv6 unicast
  !
  neighbor 200::2
   remote-as 30000
   address-family ipv6 unicast
  !
```

```
NXOS_PE#
!
vrf context v2
 rd 2:2
 !
 address-family ipv6 unicast
  route-target export 2:2
  route-target import 2:2
!
router bgp 1
  neighbor 10.13.1.21
   remote-as 1
   update-source loopback0
   address-family vpnv6 unicast
      send-community extended
  !
  vrf vpn1
    neighbor 200::2
    remote-as 30000
      address-family ipv6 unicast
 !
```

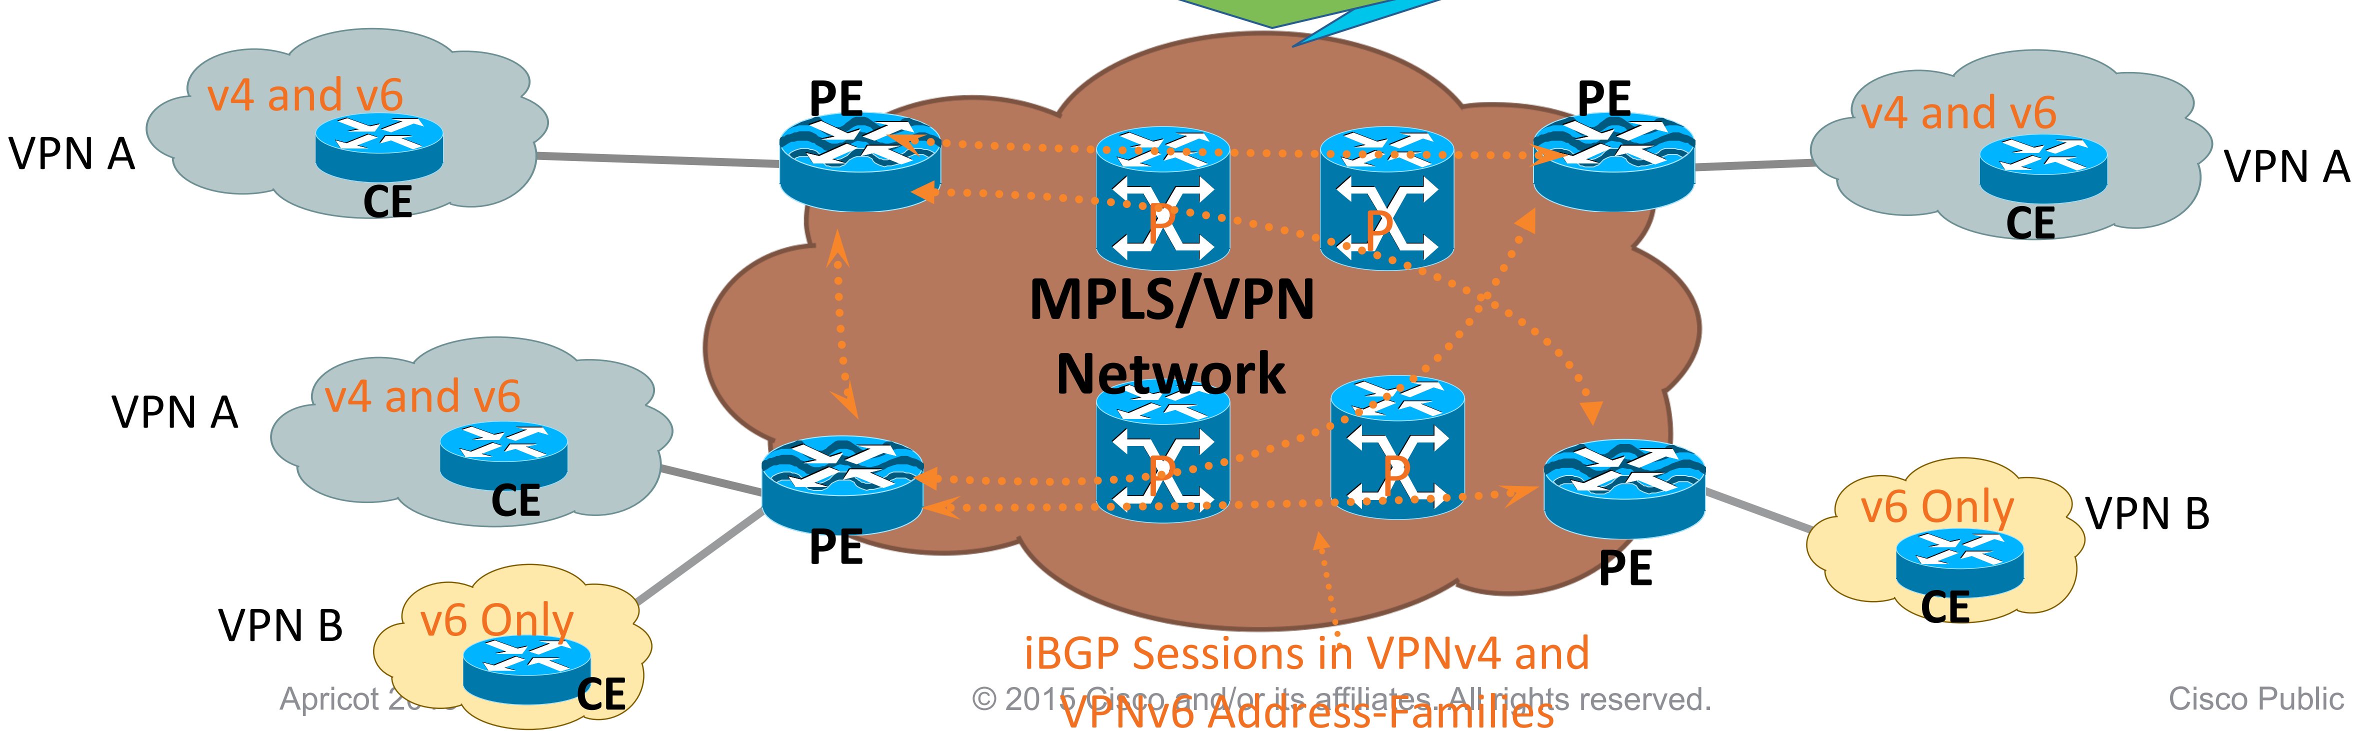

Agenda

- IP/VPN Overview
- IP/VPN Services
- Best Practices
- Conclusion

Best Practices (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. Choose AS/IP format for RT and RD i.e., ASN: X
   - Reserve first few 100s of X for the internal purposes such as filtering
3. Consider unique RD per VRF per PE,
   - Helpful for many scenarios such as multi-homing, hub&spoke etc.
4. Don't use customer names (V458:GodFatherNYC32ndSt) as the VRF names; nightmare for the NOC.
   - Consider v101, v102, v201, v202, etc. and Use VRF description for naming
5. Utilize SP's public address space for PE-CE IP addressing
   - Helps to avoid overlapping; Use /31 subnetting on PE-CE interfaces

Best Practices (2)

6. 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
7. 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)
8. Consider RT-constraint for Route-reflector scalability
9. Consider ‘BGP slow peer’ for PE or RR – faster BGP convergence

Agenda

- IP/VPN Overview
- IP/VPN Services
- Best Practices
- Conclusion

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.

BUILT FOR
THE HUMAN
NETWORK

CISCO

