cisco Live!

ALL IN

#CiscoLive

The bridge to possible

Designing IP VPNs

MPLS with/out Segment Routing

Rajiv Asati
CTO, VP/Cisco Fellow

BRKMPL-2102

Slido Poll

Abstract

- **This session** describes IP Virtual Private Networks (IP VPNs) overlays using MPLS data plane. It is the most common Layer 3 VPN technology, as standardized by IETF RFC2547/4364, enabling IPv6 (and/or IPv4) WAN connectivity among 2 or more sites, endpoints, functions etc. over IP/MPLS network(s).
- SPs have been using IP VPN to provide scalable site-to-site/WAN connectivity to Enterprises/Public Sector/SMBs for 2+ decades (and recently to create 5G slices), whereas Enterprises/Public Sectors have been using it to address network segmentation (virtualization and traffic separation) inside their sites e.g. Campus, Branch, Data Center, Cloud. The session will cover:
    - Technology Overview
    - Configuration Overview
    - Use-Cases Summary
    - Best Practices

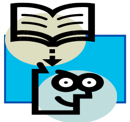

Reference

Prerequisites

- **Must** understand basic IP routing, especially BGP
- **Must** understand MPLS basics (push, pop, swap, label stacking)
- **Should** understand MPLS IP/VPN basics
- Must keep the speaker engaged…
  - …by asking bad questions ☺

Terminology

Reference

- 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 IPv4 routes
- VPNv6: address family used in BGP to carry IPv6 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)
- FIB: forwarding information base (same as CEF - Cisco Express Forwarding)
- LFIB: label forwarding information base
- 6VPE: IPv6 VPN

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

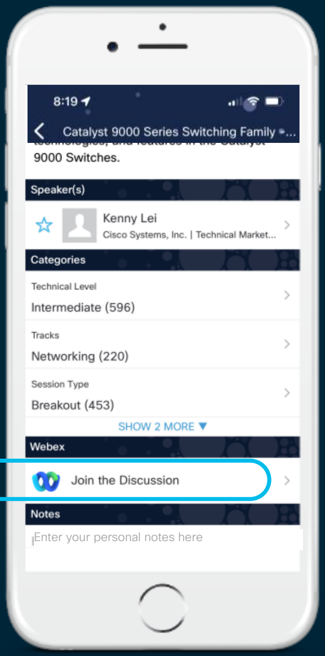

Webex spaces will be moderated by the speaker until June 17, 2022.

https://ciscolive.ciscoevents.com/ciscolivebot/#BKMPL-2102

Agenda

- IP/VPN Overview
- Use-Cases Summary
- Best Practices
- Conclusion

Agenda

- IP/VPN Overview
  - Technology Overview
  - Configuration Overview (reference only)
- Use-Cases
- Best Practices
- Conclusion

IP/VPN Technology Overview

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

IP/VPN Technology Overview

Network Topology / Connection Model

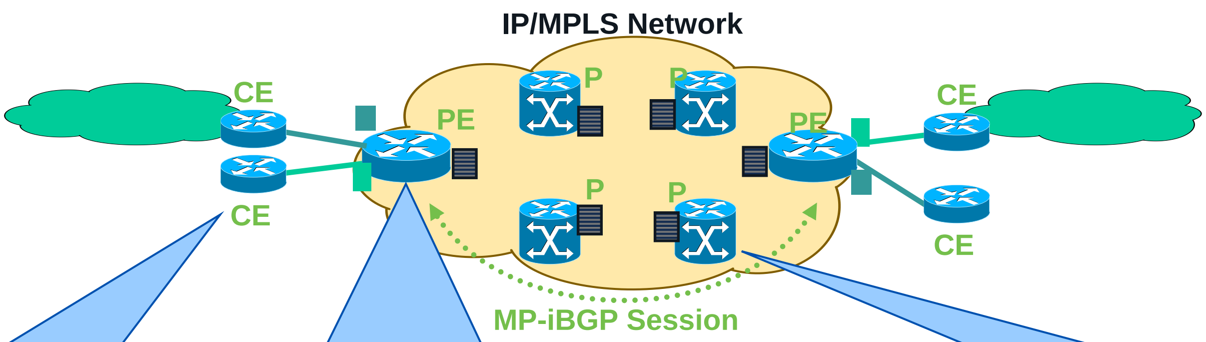

CE Routers

- Sit at the Edge
- Exchange IP traffic with PE routers (and C routers)
- Exchange IP routes with PE routers using IP routing protocol

PE Routers

- Sit at the Edge of IP/MPLS Network
- Exchange IP traffic with CE routers
- Exchange MPLS traffic with P routers
- Distributes VPN routes using MP-BGP sessions to other PE routers

P Routers

- Sit inside the network
- Exchange MPLS traffic - Forward packets by looking at MPLS labels
- Share a common IGP with PE

IP/VPN Technology Overview

Separate Routing & Forwarding Tables at PE

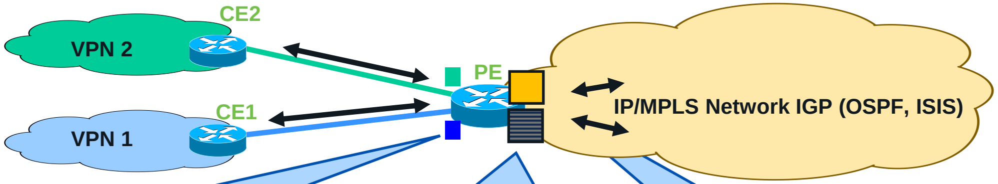

Customer Specific IP Routing Table

- Routing table (RIB) and forwarding table (FIB/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> ipv6
NX-OS: "sh ip route vrf <name>"
```

Global IP Routing Table

- Created by IP routing bestpaths.
- Populated by OSPF, ISIS, etc. running inside the MPLS network

```
IOS: "show ip route"
IOS-XR:"sh route ipv6 uni"
NX-OS: "sh ip route"
```

Global MPLS Label Database

- Created by IP routing bestpaths
- Populated by either LDP or RSVP or Routing Protocol (SR/IGP) inside the MPLS network

```
IOS: "show mpls ldp"
IOS-XR:"sh mpls ldp "
NX-OS: "sh mpls ldp"
```


IP/VPN Technology Overview

Virtual Routing and Forwarding (VRF) Instance

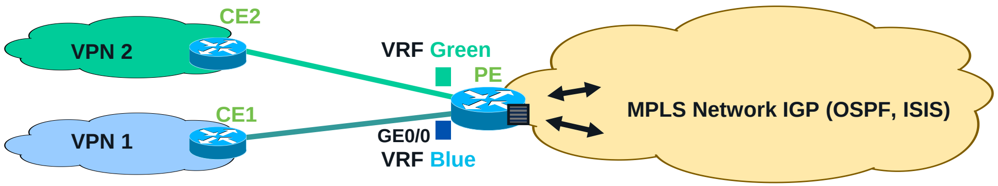

- VRF = Representation of VPN customer inside the MPLS network
  - Each customer VPN is associated with at least one VRF
- VRF configured on each PE and associated with PE-CE interface(s)
  - Privatize an interface, i.e., coloring of the interface
- No changes needed at CE

IOS_PE(conf)#ip vrf blue
IOS_PE(conf)#interface GE0/0
IOS_PE(conf)#ip vrf forwarding blue

IP/VPN Technology Overview

Virtual Routing and Forwarding Instance

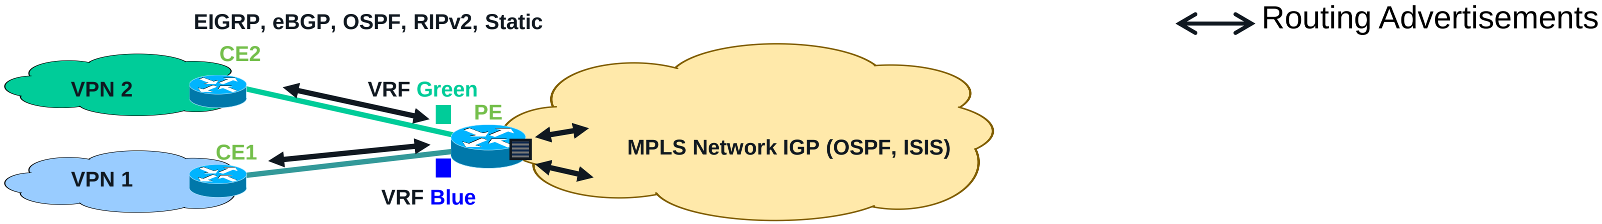

- PE installs the VPN customer' IP routes in VRF routing table(s)
  - VPN routes are learned from CE routers or remote PE routers
  - VRF-aware routing protocol (static, RIP, BGP, EIGRP, OSPF) on each PE
- PE installs the internal routes (IGP) in global routing table
- VPN customers can use overlapping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

IP/VPN Technology Overview

VPN Control Plane

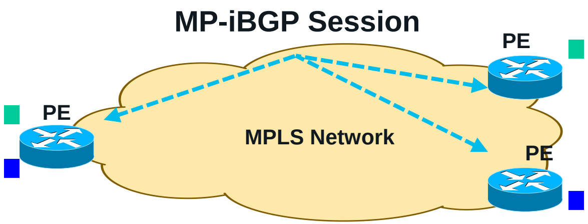

- PE routers exchange VPN routes with other PE routers using BGP
  - Multi-Protocol BGP aka MP-BGP
- PE routers advertise the IP routes to their CE routers

IP/VPN Technology Overview

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

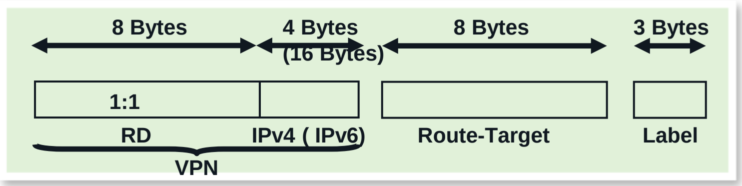

MP-BGP UPDATE Message Showing VPN route, RT, Label only

MP-BGP on PE Customizes the VPN Customer Routing Information as per the Locally Configured VRF Information using:

- Route Distinguisher (RD)
- Route Target (RT)
- Label (not configured)

IP/VPN Technology Overview: Control Plane

MP-BGP UPDATE Message Capture

Reference

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

Route Target = 3:3

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

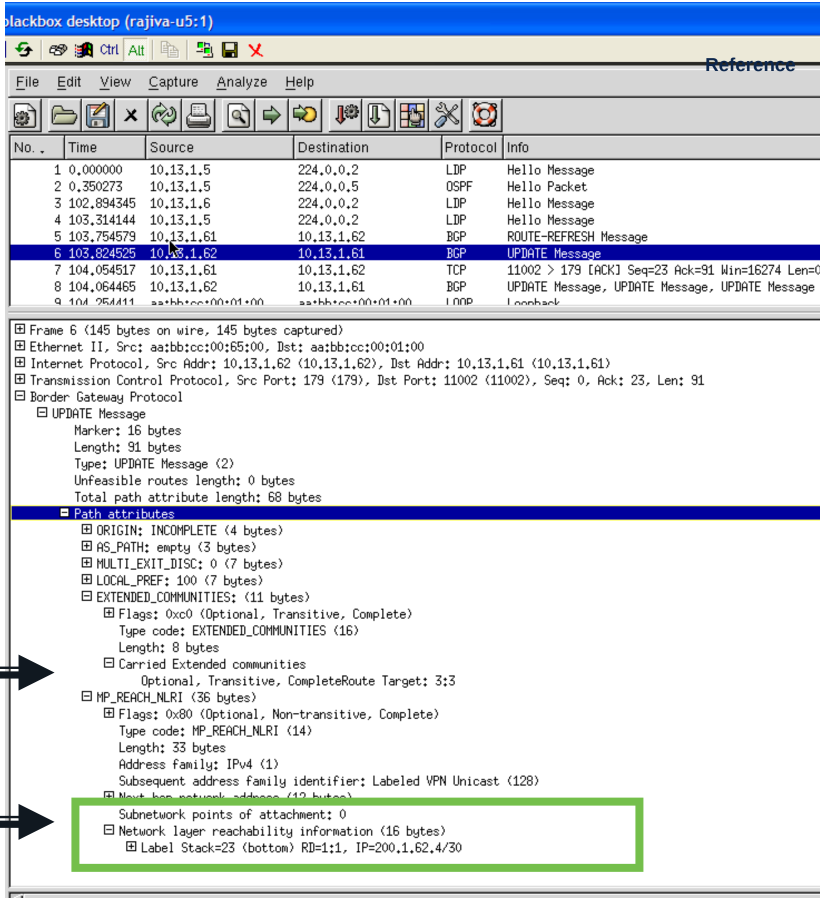

IP/VPN Technology Overview: Control Plane

Route-Distinguisher (rd): 8-byte field

BGP MP_REACH_NLRI

- VPN customer IP prefix is converted into a VPN prefix by appending the RD (1:1, say) to the IP address (200.1.64.0, or 2001:DB8:: say) => 1:1:200.1.64.0 or 1:1:2001:DB8::
  - Makes the customer's IP address unique inside the shared IP/VPN network
- Route Distinguisher (rd) is configured in the VRF at PE
  - RD is not a BGP attribute, just a field in another attribute (MP_REACH_NLRI)

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

* Since 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): 8-byte extended community attribute

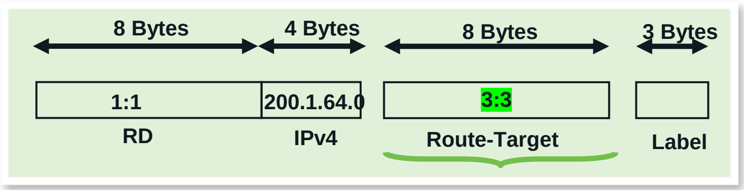

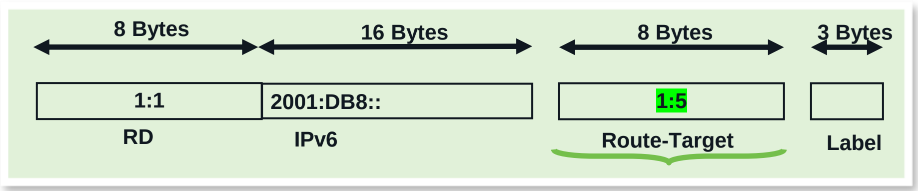

- Route-target (rt) helps PEs color the VPN prefixes
  - Export rt values : attached to VPN routes by PEs in MP-iBGP advertisements
  - Import rt values : used by PEs to identify which VRF(s) keep the received VPN prefixes
- Each VRF should be configured with 1 or more route-targets at PE
  - Export & Import rt must be the same for Any-to-Any topology
  - Export & Import rt must be different for Hub & Spoke topology
- IPv4 and IPv6 address-family RT values are allowed to be different (as shown)

```
IOS_PE#
!
ip vrf green
address-family ipv4
 route-target import 3:3
 route-target export 3:3
!
address-family ipv6
route-target import 1:1
route-target export 1:5
!
```

IP/VPN Technology Overview: Control Plane

Label : 20-bit value

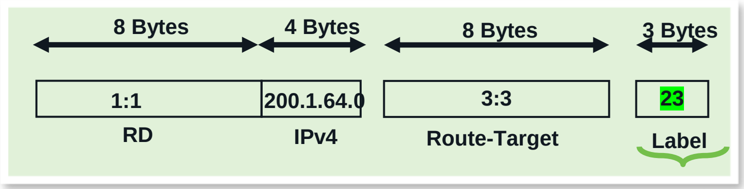

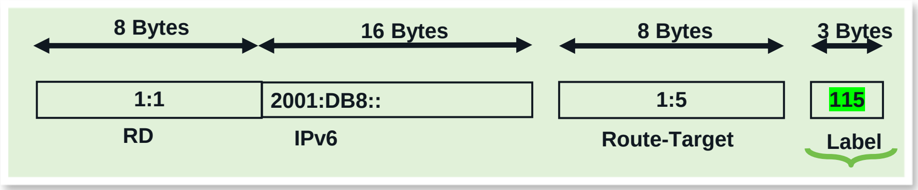

- PE auto-generates & assigns a label for each VPN prefix(es);
  - Next-hop-self towards MP-iBGP neighbors by default i.e. PE sets the NEXT-HOP attribute to its own address (as configured)
  - Label is not an attribute.
- PE addresses used as the BGP next-hops must be uniquely known in IGP
  - CAUTION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

IP/VPN Technology Overview: Control Plane

Putting it all together

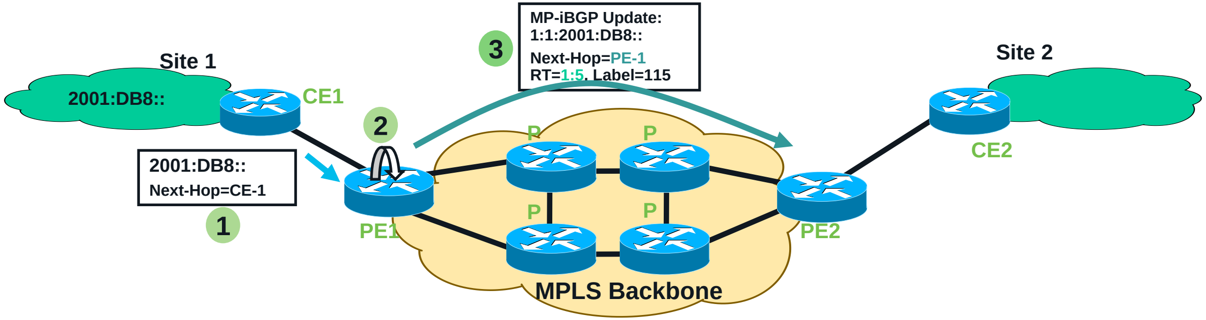

1. PE1 receives an IPv6 (or IPv4) update (eBGP/OSPF/ISIS/RIP/EIGRP)
2. PE1 translates it into VPNv6(v4) address & sends MP-iBGP UPDATE message
   - Associates the RT values (export RT =1:5, say) per VRF green configuration
   - Rewrites next-hop attribute to its IP address (usually loopback0 int)
   - Assigns a label (115, say); Installs it in the MPLS forwarding table.
3. PE1 sends MP-iBGP update to other PE routers

IP/VPN Technology Overview: Control Plane

Putting it all together

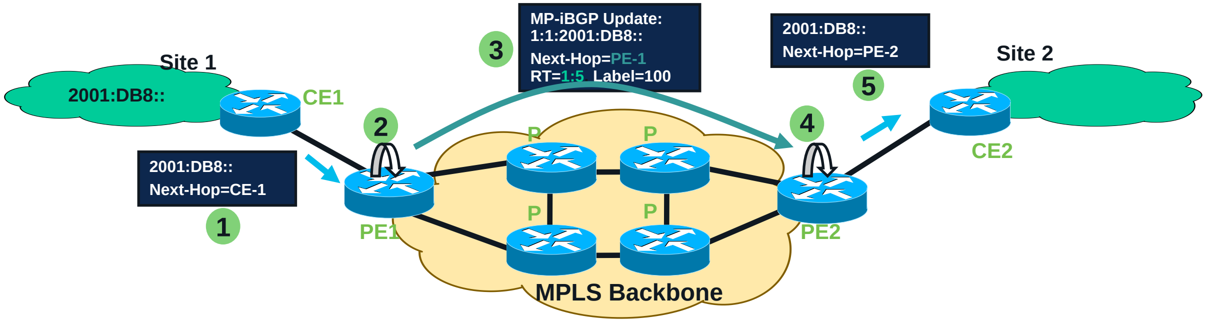

4 PE2 receives and checks whether the RT=1:5 is locally configured as ‘import RT’ within any VRF, if yes, then

- PE2 translates VPN prefix back to IP prefix
- PE2 updates its VRF CEF Table (green) with IP prefix - **2001:DB8::** along with label=115

5 PE2 advertises this IP prefix to CE2 (using whatever routing protocol)

Control Plane is now ready

IP/VPN Technology Overview

Forwarding Plane

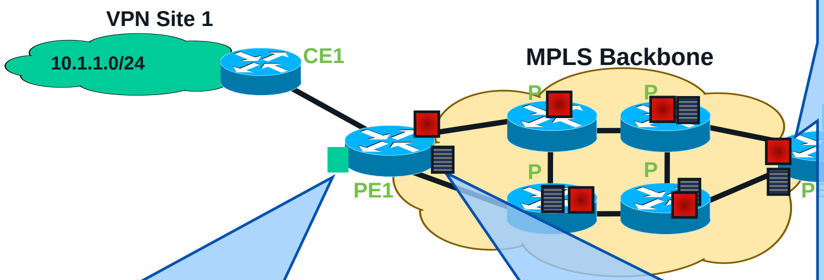

MPLS Forwarding Table

- Stores labels for PE/P routes i.e. next-hops
- Label learned SR/IGP or LDP, RSVP or BGP

```
IOS:show mpls forwarding
NX-OS: show mpls forwarding
IOS-XR: show mpls forwarding
```

Customer/VPN 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> ipv6|ipv4
```

Global CEF Forwarding Table

- Stores PE routes i.e. next-hops with labels
- Next-hop i.e. PE routes learned via IGP
- Label learned SR/IGP or LDP, RSVP or BGP

```
IOS:show ip cef
NX-OS: show forwarding ipv6|ipv4
IOS-XR: show cef ipv6|ip4
```

IP/VPN Technology Overview: Forwarding Plane

Packet Forwarding

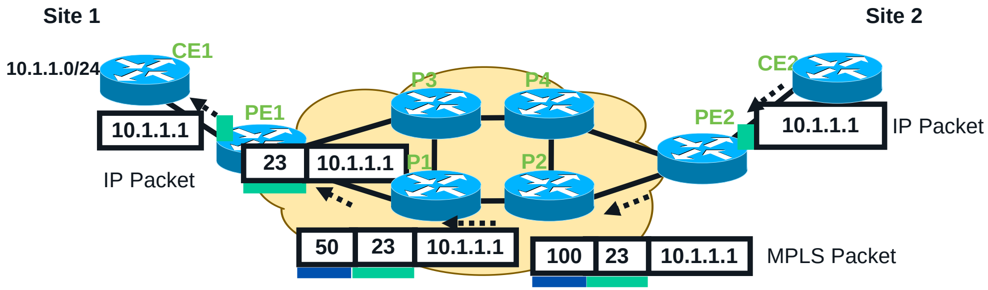

- PE2 imposes two labels (in 2 MPLS headers) for each IP packet going towards remote site
  - Outer label 100 for PE1 address (learned via SR/IGP or LDP or RSVP or static..)
  - Inner label 23 for VPN address (learned via BGP)
- P2 swaps the outer label (100 with 50) per its MPLS forwarding table
- P1 does the Penultimate Hop Popping (PHP) i.e. removes the outer label 50
- PE1 removes label 23, retrieves IP packet and forwards it to CE1.

Reference

IP/VPN Technology Overview: Forwarding Plane

MPLS IP/VPN Packet Capture

- Visualize an MPLS VPN Packet on the wire (PE-P or P-P)
- 2 MPLS headers

Ethernet Header
Outer MPLS header
Inner MPLS Header
IP Header

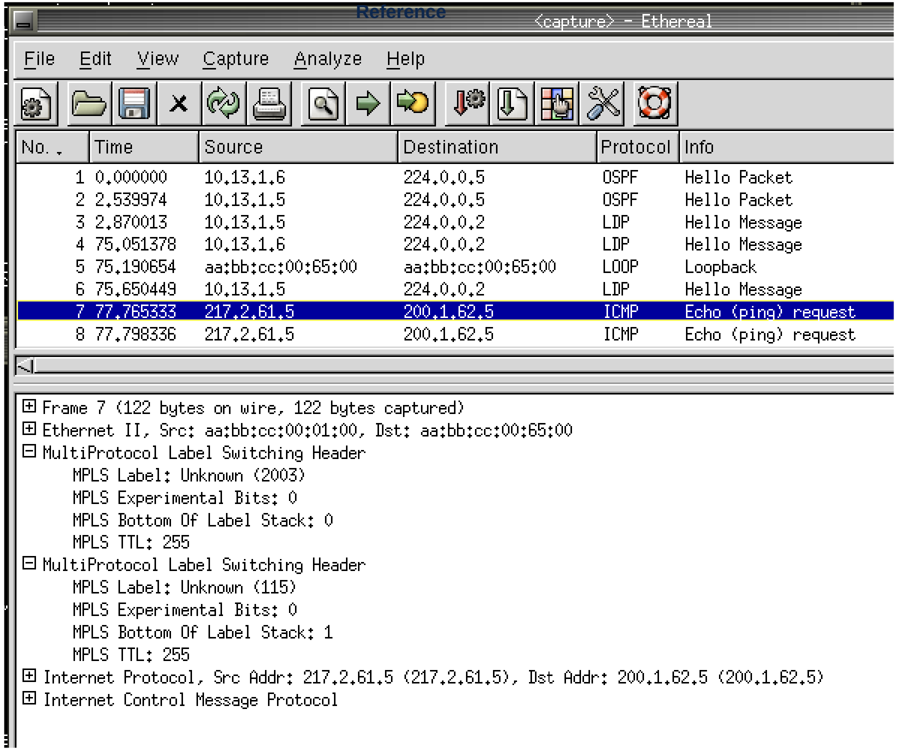

Note: The MPLS values & IP addresses to not refer to the previous examples, sorry. ☹

Agenda

- IP/VPN Overview
  - Technology Overview
  - Configuration Overview (reference only)
- Best Practices
- Use-Cases
- Conclusion

MPLS based IP/VPN Sample Configuration (IOS)

Reference

VRF Definition

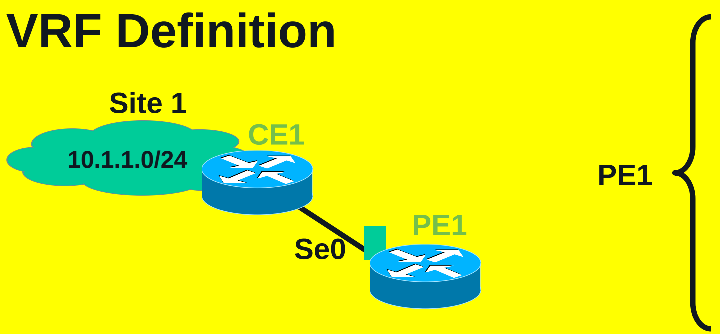


```
vrf definition VPN-A
 rd 1:1
 address-family ipv4
  route-target export  100:1
  route-target import  100:1
 address-family ipv6
  route-target export  100:1
  route-target import  100:1
```

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

IPv4 VPN only config

Both IPv6 VPN
And IPv4 VPN

VRF to Interface Association

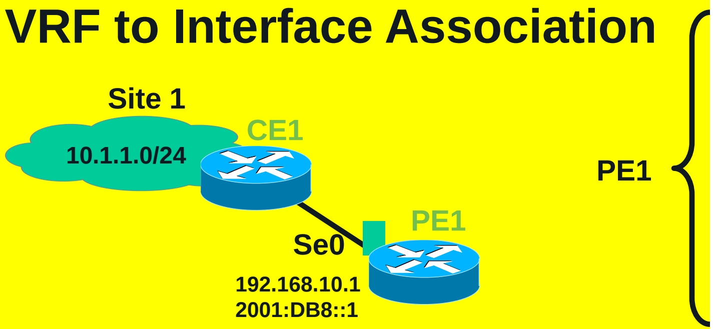


```
interface Serial0
 ip address 192.168.10.1/24
 ipv6 address 2001:DB8::1/124
 ip vrf forwarding VPN-A
```

MPLS based IP/VPN Sample Configuration (IOS)

Reference

VRF Definition

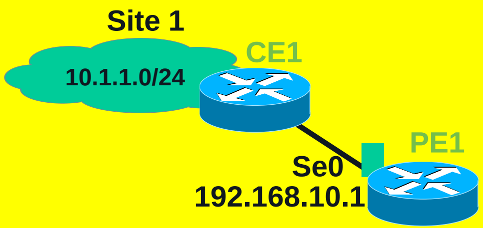

PE1

```
ip vrf VPN-A
 rd 1:1
 address-family ipv4
  route-target export  100:1
  route-target import  100:1
 address-family ipv6
  route-target export  100:1
  route-target import  100:1
```

```
interface Serial0
 ip address 192.168.10.1/24
 ipv6 address 2001:DB8::2/124
 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
 address-family ipv6
  route-target export  100:1
  route-target import  100:1
```

PE-P Configuration

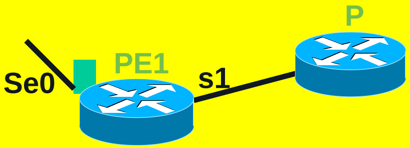

PE1

```
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

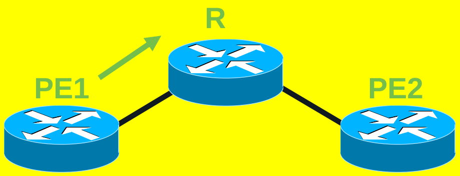

PE 1

```
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
!
 address-family vpnv6
  neighbor 1.2.3.4 activate
  neighbor 1.2.3.4 send-community both
!
```

RR: MP-IBGP Config

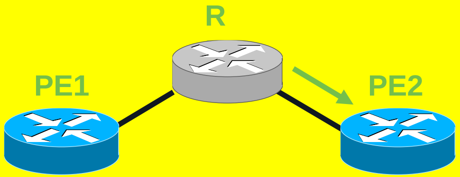

RR

```
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 | vpnv6
 neighbor 1.2.3.6 route-reflector- client
 neighbor 1.2.3.6 activate
!
```

Config Shows Both IPv6 VPN And IPv4 VPN

Reference

MPLS based IP/VPN Sample Configuration (IOS)

PE-CE Routing: BGP

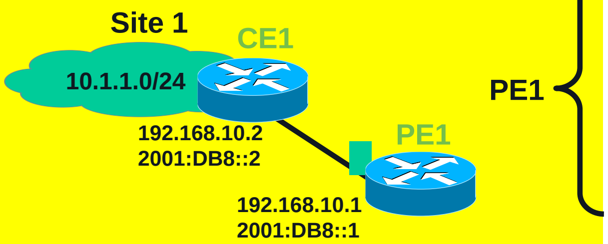


```
router bgp 1
!
address-family ipv4 vrf VPN-A
  neighbor 192.168.10.2 remote-as 2
  neighbor 192.168.10.2 activate
!
```

```
router bgp 1
!
address-family ipv6 vrf VPN-A
  neighbor 192.168.10.2 remote-as 2
  neighbor 192.168.10.2 activate
!
```

PE-CE Routing: OSPF

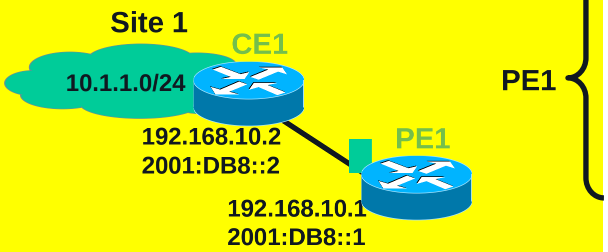


```
router ospf 2 unicast vrf VPN-A
 network 192.168.10.0 0.0.0.255 area 0
 redistribute bgp 1 subnets
!
```

```
router ospfv3
!
address-family unicast vrf VPN-A
 router-id 2001:DB8::2
 redistribute bgp 1 subnets
!
interface Serial0
 ospfv3 2 ipv6 area 0
 redistribute bgp 1 subnets
!
```

MPLS based IP/VPN Sample Configuration (IOS)

Reference

PE-CE Routing: RIP

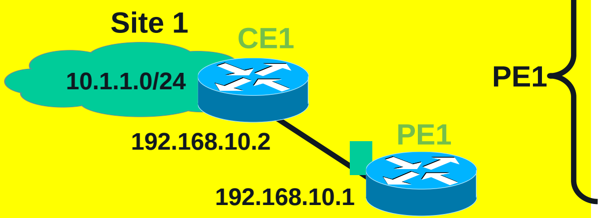


```
router rip
!
 address-family ipv4 vrf VPN-A
  version 2
  no auto-summary
  network 192.168.10.0
  redistribute bgp 1 metric transparent
!
```

```
ipv6 rip vrf-mode enable
!
ipv6 router rip XYZ
  redistribute bgp 1
!
interface Serial0/0
  ipv6 vrf VPN-A XYZ enable
!
```

PE-CE Routing: EIGRP

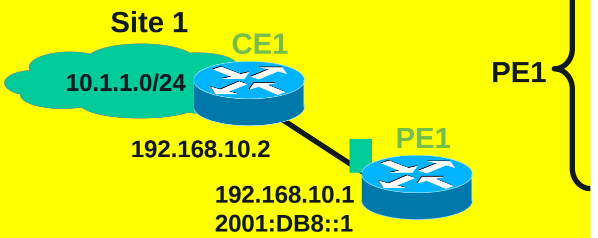


```
router eigrp 1
!
 address-family ipv4 vrf VPN-A
  no auto-summary
  network 192.168.10.0 0.0.0.255
  autonomous-system 10
  redistribute bgp 1 metric 100000 100
255 1 1500
!
```

```
router eigrp XYZ
  address-family ipv6 vrf VPN-A
autonomous-system 1
  af-interface Serial0/0
!
```

MPLS based IP/VPN Sample Configuration (IOS)

Reference

PE-CE Routing: Static

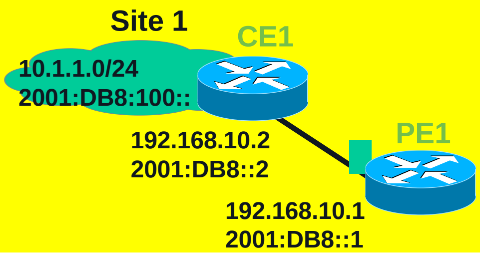

PE1

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

ipv6 route vrf VPN-A 2001:DB8:100::/48 2001:DB8::2
```

If PE-CE Protocol Is Non-BGP (Such as RIP), then Redistribution of VPN Routes from MP-IBGP Is Required (Shown Below for RIP) -

PE-CE: MB-iBGP Routes to VPN

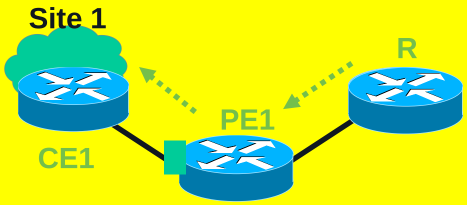

PE1

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


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-CE (Route Distribution)

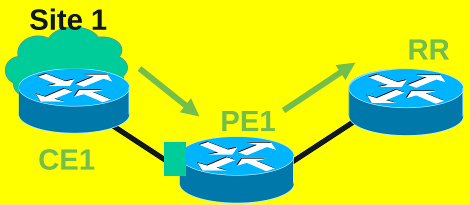


```
router bgp 1
 neighbor 1.2.3.4 remote-as 1
 neighbor 1.2.3.4 update-source loopback 0

 address-family ipv4|ipv6 vrf VPN-A
  redistribute {rip|connected|static|eigrp|ospf}
```

- For hands-on learning, please attend the lab sessions:
  - LTRMPL-2104 Implementing MPLS in SP Networks (Intro Level)
  - LTRMPL-2105 Implementing MPLS in SP Networks (Advanced Level)
- Having familiarized with IOS based config, let's peek through IOS-XR and NX-OS config for VPNs

IP/VPN Deployment Scenarios:

6. IPv6 VPN Service

Supported in IOS, NXOS and IOS-XR

```
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
 !
```

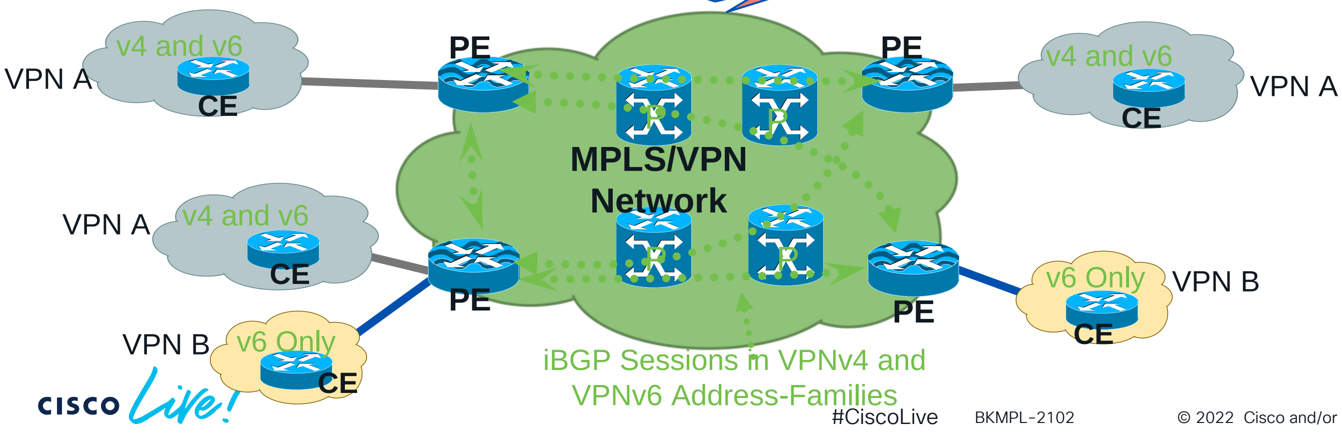

iBGP Sessions in VPNv4 and VPNv6 Address-Families

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

Reference

VRF Definition

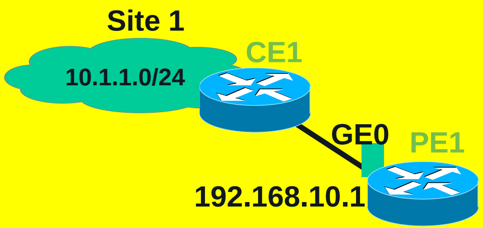

PE 1

```
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
```

```
vrf VPN-A
 address-family ipv6 unicast
   import route-target 100:1
   export route-target 100:1
!
router bgp 1
 vrf VPN-A
  rd 1:1
```

```
Interface GE0
 ipv4 address 192.168.10.1 255.255.255.0
 vrf VPN-A
```

PE-P Configuration

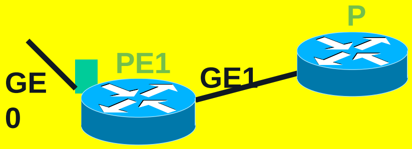

PE 1

```
mpls ip
 int GE1
!
```

```
router ospf 1
 area 0
  interface GE1
```


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

Reference

PE: MP-IBGP Config

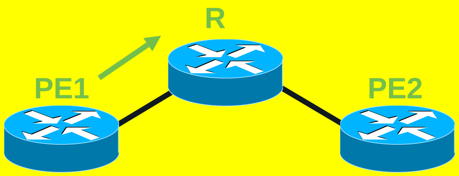

PE 1

```
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
      send-community extended
!
```

RR: MP-IBGP Config

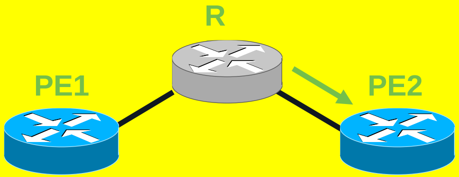

RR

```
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
      send-community extended
      route-reflector-client
!
```

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

Reference

PE-CE Routing: BGP

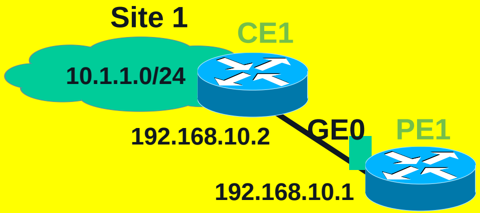

PE1

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

PE-CE Routing: OSPF

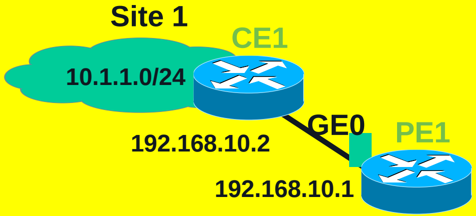

PE1

```
router ospf 2
 vrf VPN-A
  address-family ipv4 unicast
  redistribute bgp 1
  !
  area 0
   interface GE0
 !
```

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

Reference

PE-CE Routing: RIP

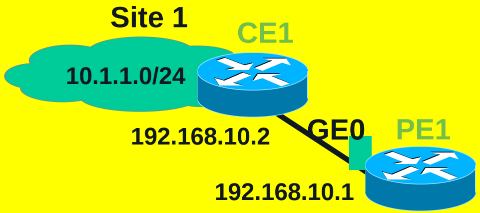

PE1

```
router rip
 vrf VPN-A
  interface GE0
  redistribute bgp 1
 !
```

PE-CE Routing: EIGRP

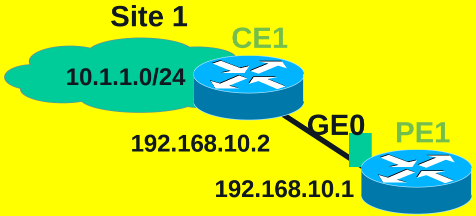

PE1

```
router eigrp 1
 vrf VPN-A
 address-family ipv4
  as 10
  default-metric 100000 100 255 1 1500
  interface GE0
  redistribute bgp 1
```

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

Reference

PE-CE Routing: Static

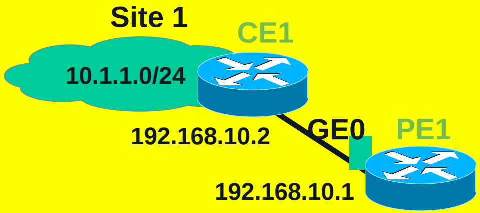

PE1

```
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-PE (Route Distribution)

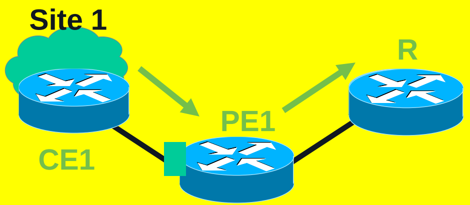

PE1

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


MPLS based IP/VPN Sample Config (NX-OS)

Reference

VRF Definition

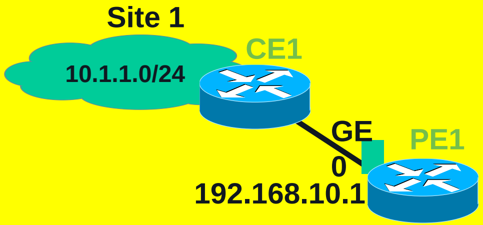

PE 1

```
vrf context VPN-A
 rd 1:1
   address-family ipv4 unicast
      route-target import 1:1
      route-target export 1:1
```

```
Interface GE0
 ip address 192.168.10.1 255.255.255.0
 vrf member VPN-A
```

```
vrf context VPN-A
 rd 1:1
   address-family ipv6 unicast
      route-target import 1:1
      route-target export 1:1
```

PE-P Configuration

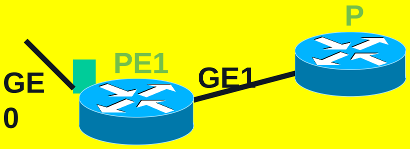

PE 1

```
Interface GE1
 ip address 130.130.1.1 255.255.255.252
 mpls ip
 ip ospf 1 area 0
```

```
router ospf 1
```

MPLS based IP/VPN Sample Config (NX-OS)

Reference

PE: MP-IBGP Config

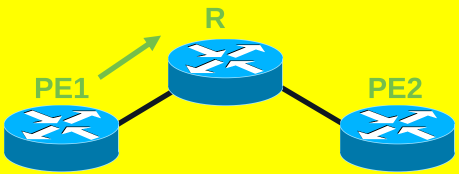

PE 1

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

RR: MP-IBGP Config

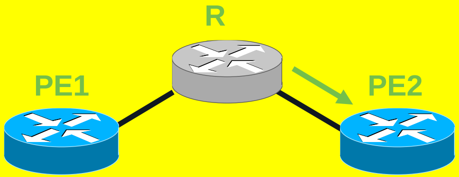

RR

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

MPLS based IP/VPN Sample Config (NX-OS)

Reference

PE-CE Routing: BGP

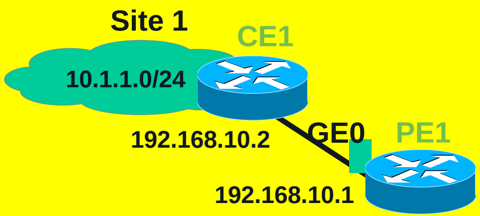

PE1

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

PE-CE Routing: OSPF

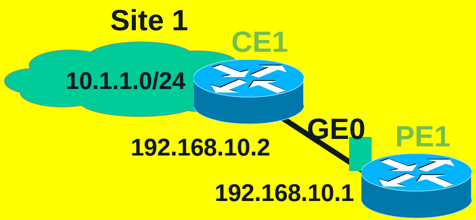

PE1

```
router ospf 2
 vrf VPN-A
  address-family ipv4 unicast
    redistribute bgp 1 route-map name
!
interface GE1
 ip address 192.168.10.1/24
 ip router ospf 2 area 0
```

Reference

MPLS based IP/VPN Sample Config (NX-OS)

PE-CE Routing: RIP

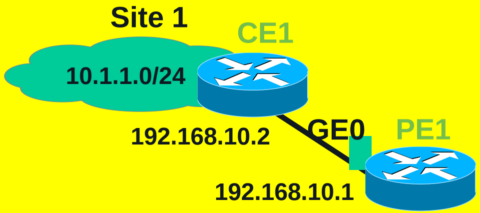

PE1

```
router rip ripxyz1
 vrf VPN-A
  address-family ipv4 unicast
  redistribute bgp 1 route-map name
!
interface GE0
 vrf member vpn1
 ip router rip ripxyz1
```

PE-CE Routing: EIGRP

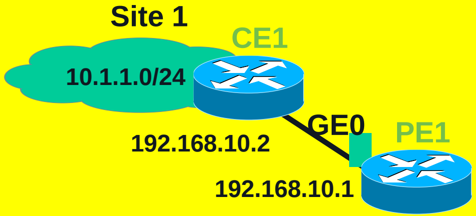

PE1

```
router eigrp 100
 vrf VPN-A
 address-family ipv4
 redistribute bgp 1 route-map name
!
interface GE0
 vrf member vpn1
 ip router eigrp 100
 site-of-origin 1:11
```

MPLS based IP/VPN Sample Config (NX-OS)

PE-CE Routing: Static

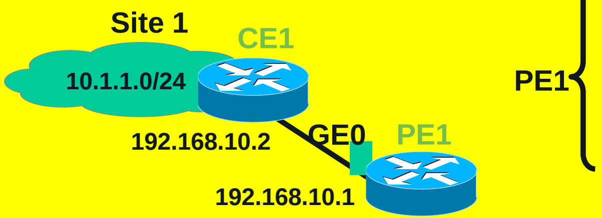


```
vrf context VPN-A
  ip route 10.1.1.0/8 192.168.10.2
```


MPLS based IP/VPN Sample Config (NX-OS)

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)

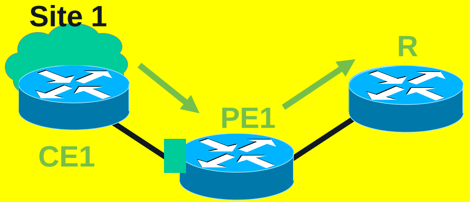

PE1

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

Agenda

- IP/VPN Overview
- Use-Cases Summary
- Best Practices
- Conclusion

Use-Cases

1. SP – Business VPN Service, Mobile Backhaul
2. SP – Internal Usage (e.g. IT), Mobile Backhaul
3. Enterprise – Campus Virtualization/Segmentation
4. Data Center – Multi-Tenancy
5. Data Center – Cloud/Virtualization/Hypervisor

Use-Case #1

SP – Business VPN Services, Mobile Backhaul

- SPs can use IP/VPN to offer L3 site-to-site connectivity to Enterprises/SMB customers'
- SPs can even offer Remote Access integrated with L3VPN

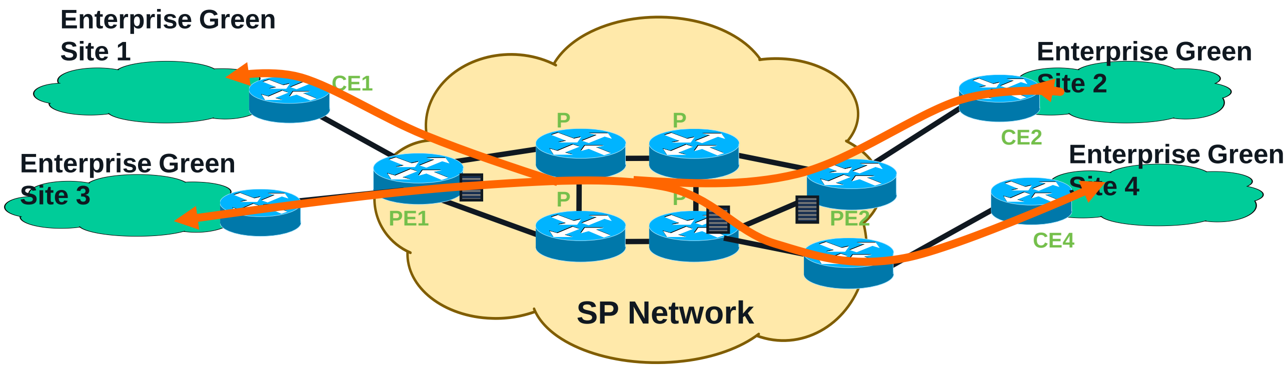

Use-Case #2

SP – Internal Usage (e.g. IT, Mobile Backhaul)

- SP/ISPs can overlay its Enterprise and/or IT WAN connectivity over its MPLS network (that is used to offer L3VPN services to its customers)

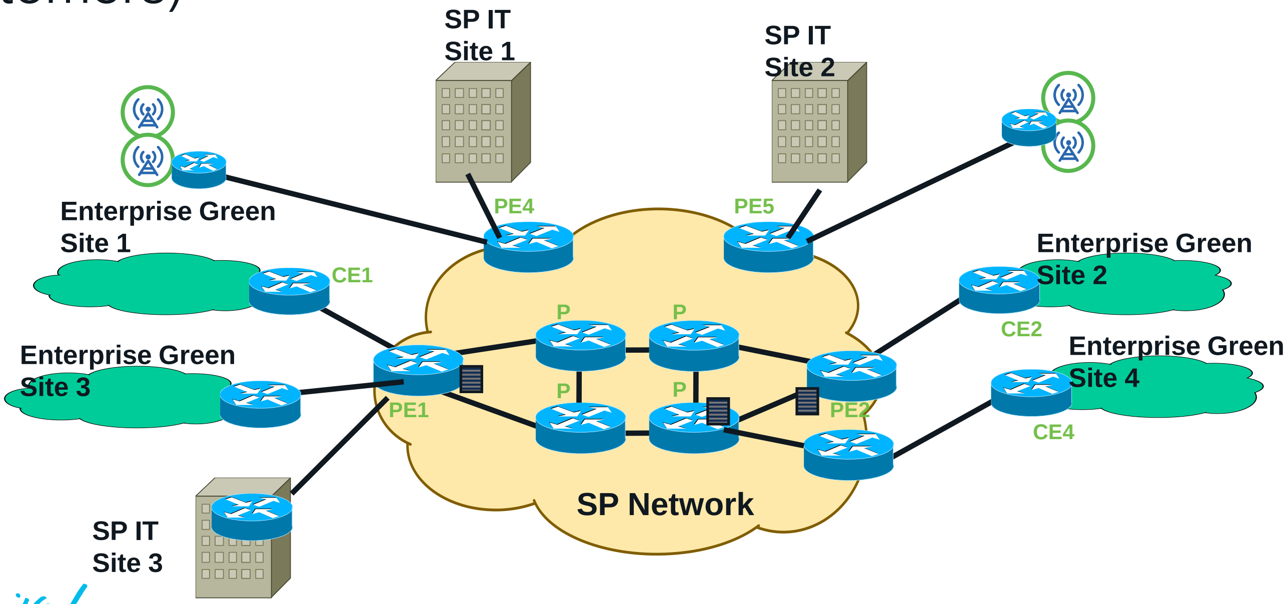

Use-Case#3

Enterprise – Campus Segmentation/Virtualization

- IP/VPN can be used to create multiple logical topologies in the Campus
  - Allows the use of unique security policies per logical domain
  - Provides traffic isolation per application, group, service etc. per logical domain
- IP/VPN segmentation in the Campus can also be extended over the WAN

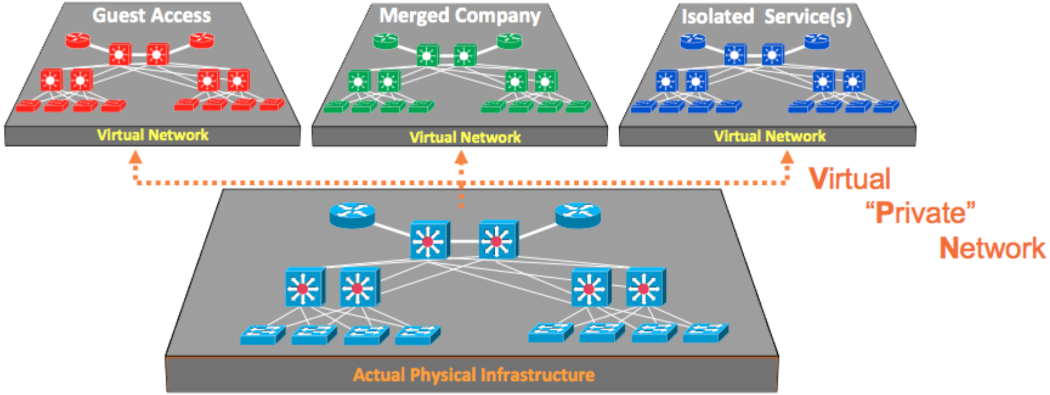

Use-Case#3

Enterprise – Campus Segmentation/Virtualization

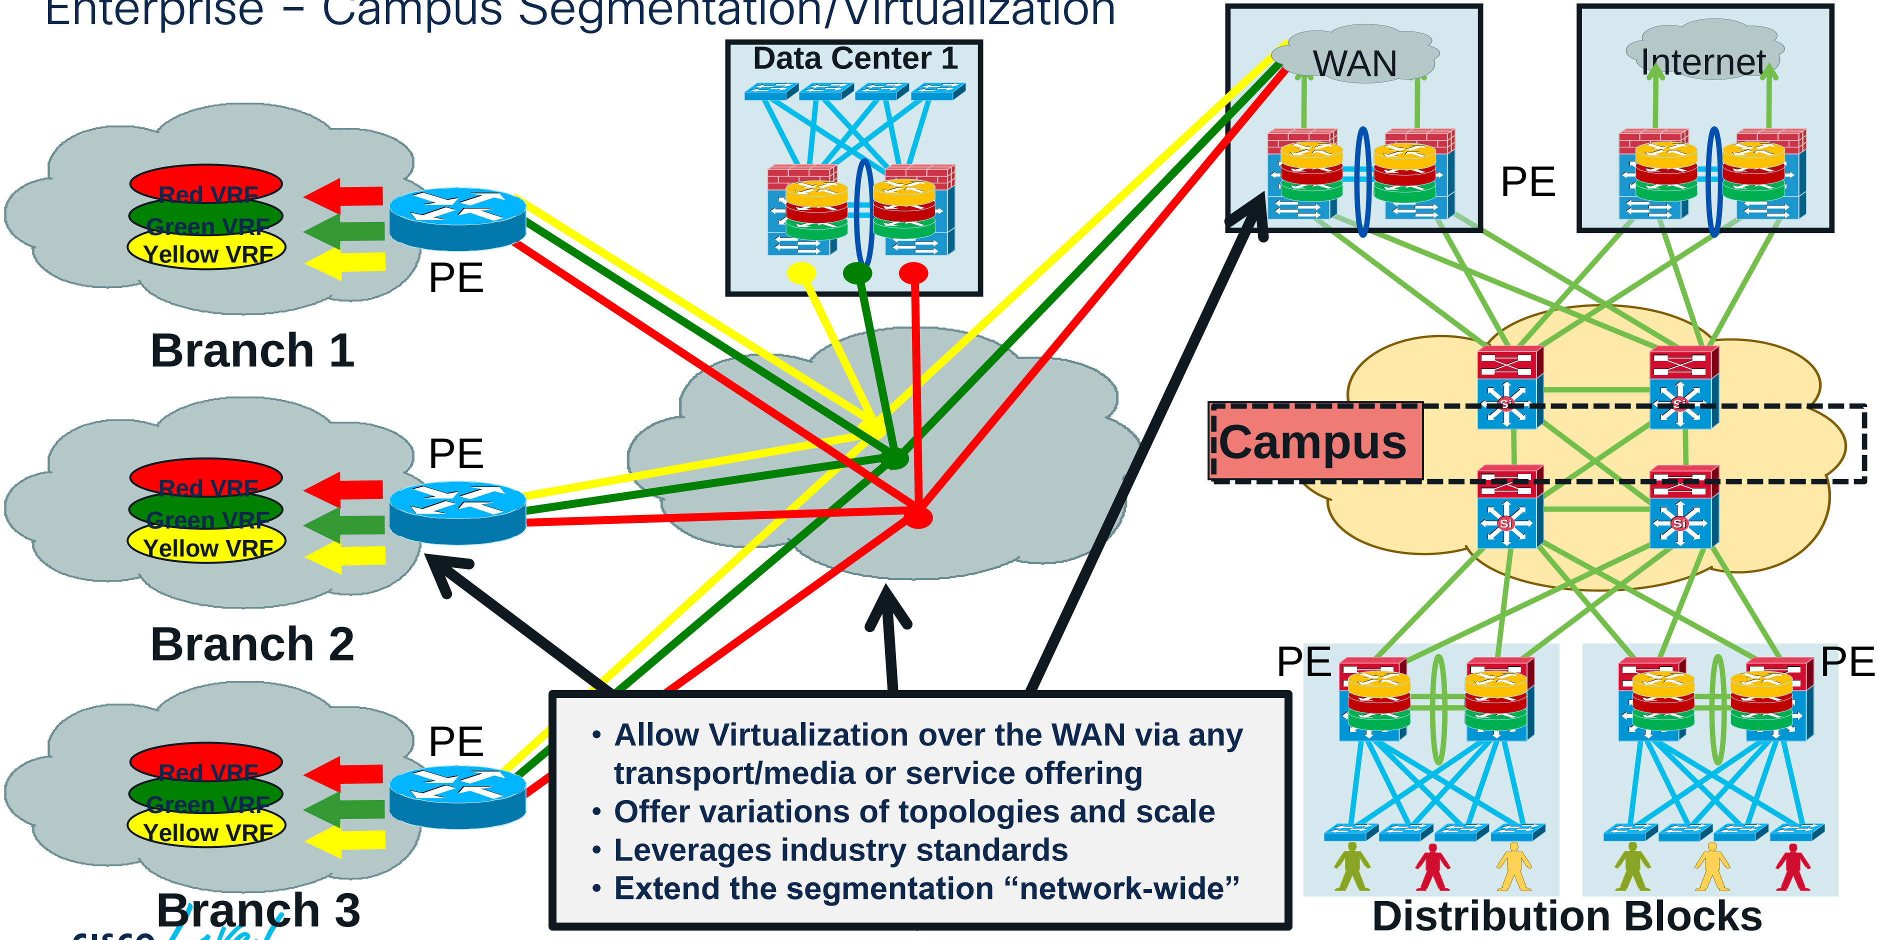

- **Allow Virtualization over the WAN via any transport/media or service offering**
- **Offer variations of topologies and scale**
- **Leverages industry standards**
- **Extend the segmentation "network-wide"**

Eliminates the need for VXLAN

Use-Case#4

Data Center – Multi-Tenancy

- IP/VPN can be used by " Cloud or Hosted DC" providers for multi-tenancy
  - Data Center services to B2B customers
- MPLS upto TOR/Leaf;
  - Segment Routing could be used
- MPLS PE function on TOR / Leaf Device
- CE function on VMs or Bare Metal
- Layer2 between PE and CE

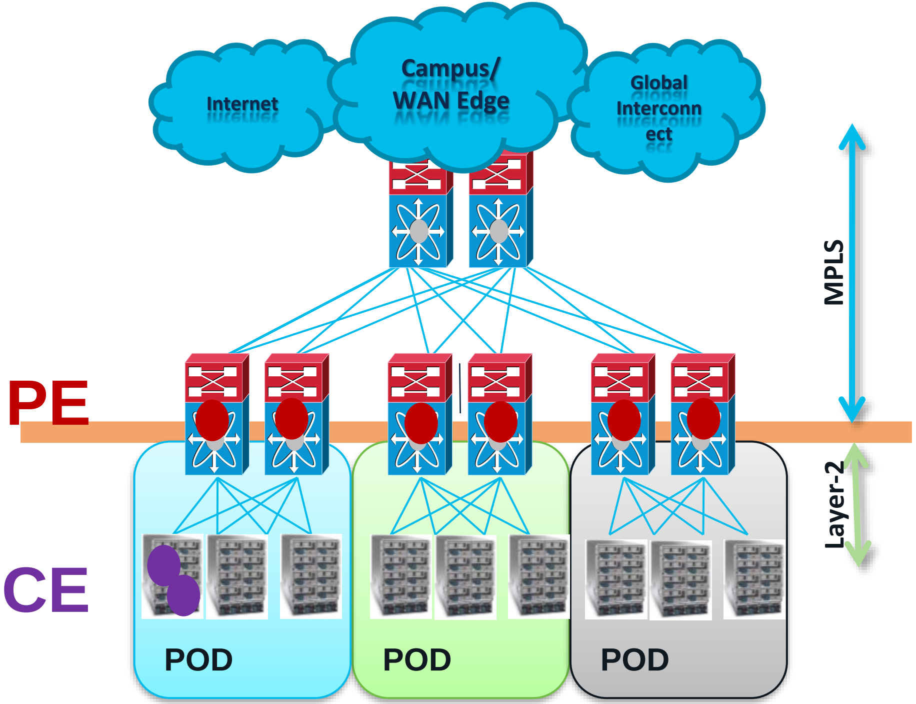

Use-Case#4

Data Center – Multi-Tenancy

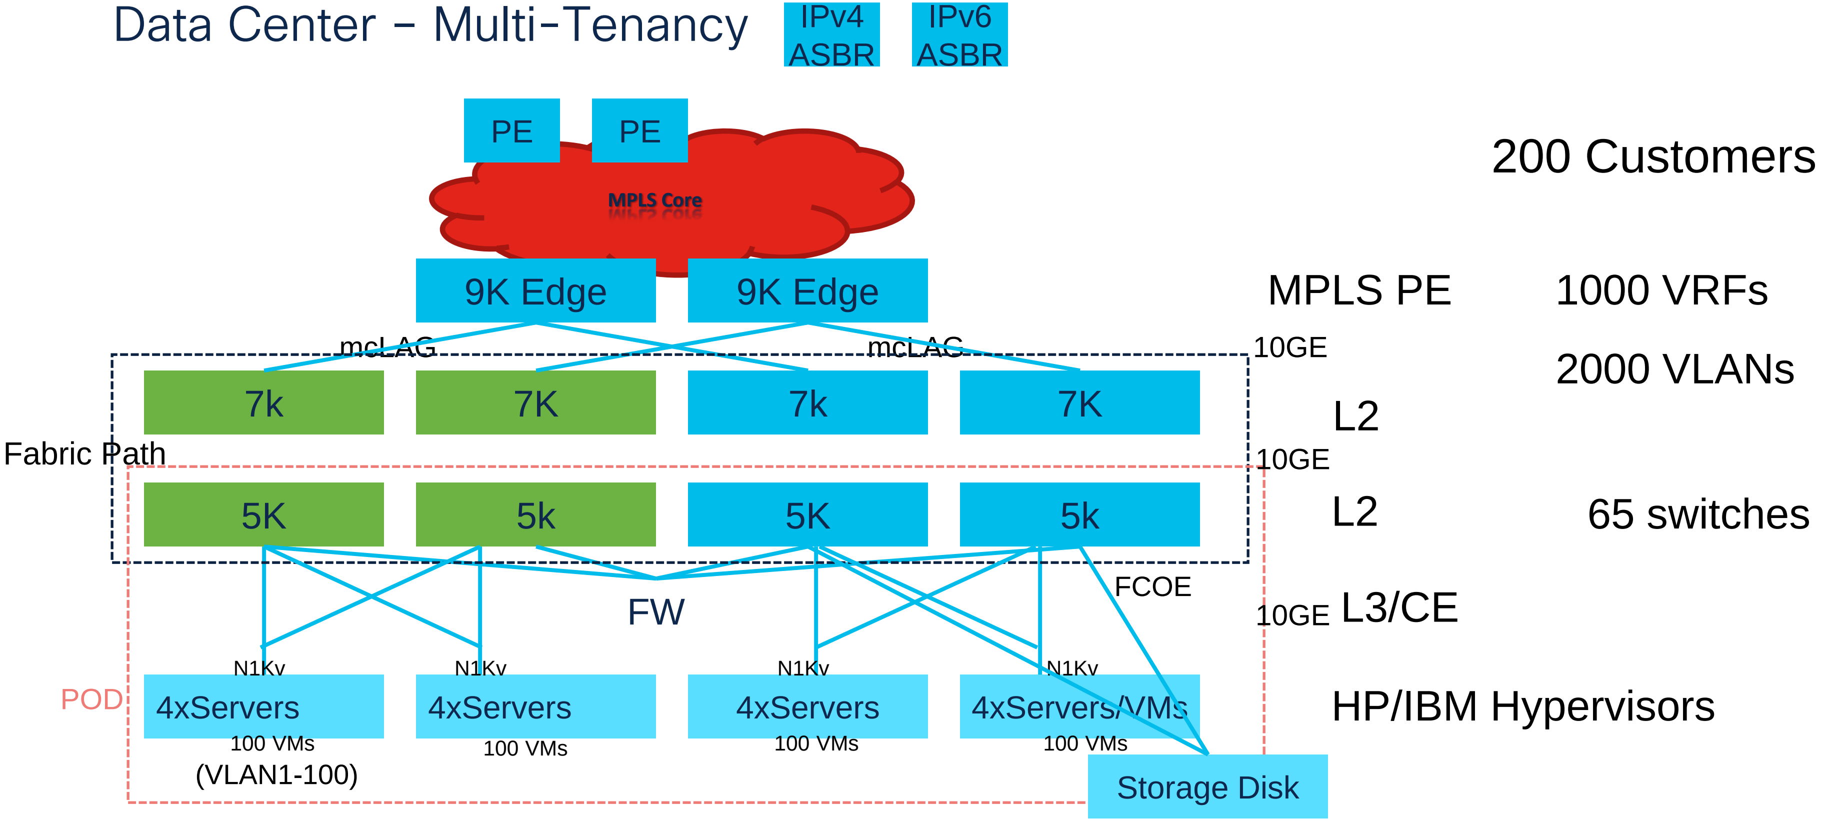

VM mobility is restricted to the POD (shown above)

Eliminates the need for VXLAN

Use-Case#5

Data Center – Cloud / Virtualization

- MPLS in Data Center (Underlay)
- MPLS based IP/VPN as Overlay
- MPLS upto x86 Host;
  - Segment Routing could be used
- MPLS PE function on virtual Router (VM) or Virtual Forwarder (VM or Container)
  - SDN Control Plane and Data Plane Separation in case of latter
- CE function on VMs or Bare Metal
- Layer2 between PE and CE

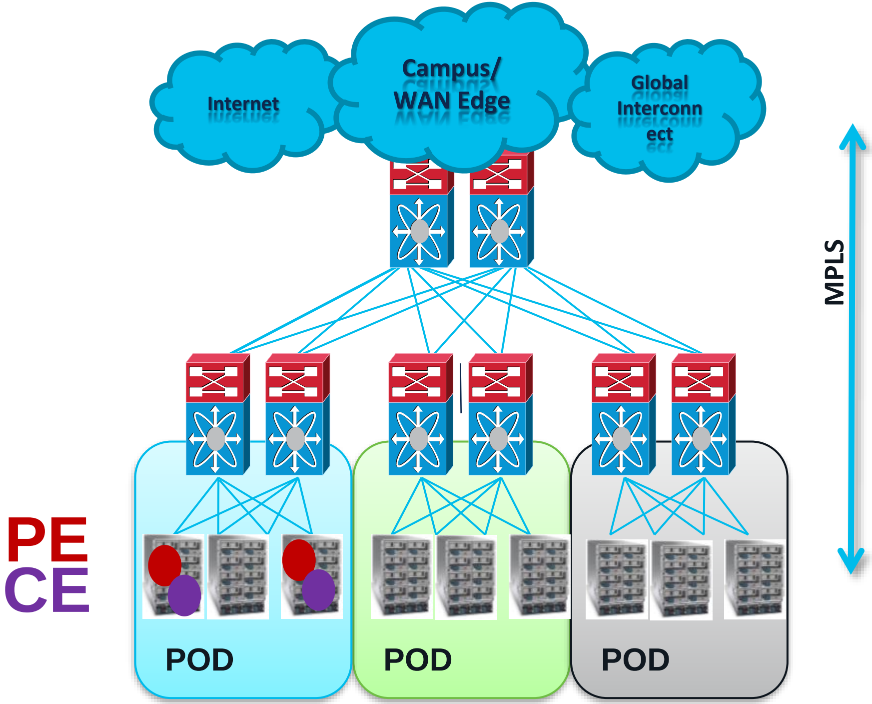

Please see BRKMPL-2115 for MPLS in DC/Cloud Details

Agenda

- IP/VPN Overview
- Use-Cases Summary
- 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 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.
   - Helpful to avoid add-path, shadow RR 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 (IPv6 and IPv4):
   - PIC Core
   - PIC Edge
   - Best-external advertisement
   - Next-hop tracking (ON by default)
8. Consider RT-constraint for PE & RR scalability (millions of routes)
9. Consider 'BGP slow peer' for PE or RR – faster BGP convergence
10. Use a dedicated VPN for CE Management

Agenda

- IP/VPN Overview
- Use-Cases Summary
- Best Practices
- Conclusion

Conclusion

- IP/VPN is the most optimal L3VPN technology
  - Any-to-any, Partial-mesh, Hub-and-Spoke topologies
  - IPv6 or IPv4 or both

- Various IP/VPN deployment scenarios for additional value/revenue

- IP/VPN paves the way for virtualization & Cloud Services
  - Benefits SPs, Enterprises, Data Centers

Technical Session Surveys

- Attendees who fill out a minimum of four session surveys and the overall event survey will get Cisco Live branded socks!
- Attendees will also earn 100 points in the Cisco Live Game for every survey completed.
- These points help you get on the leaderboard and increase your chances of winning daily and grand prizes.

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

CISCO Live!

Thank you

#CiscoLive

Supplemental Material

Agenda

- IP/VPN Overview
- IP/VPN Deployment Scenarios
- Best Practices
- Use-Cases
- Conclusion

1. Multihoming & Load-sharing
2. Hub and Spoke
3. Extranet
4. Internet Access
5. IP/VPN over IP Transport
6. Multi-VRF CE

IP/VPN Deployment Scenarios:

1. Multi-homing & Loadsharing of VPN Traffic

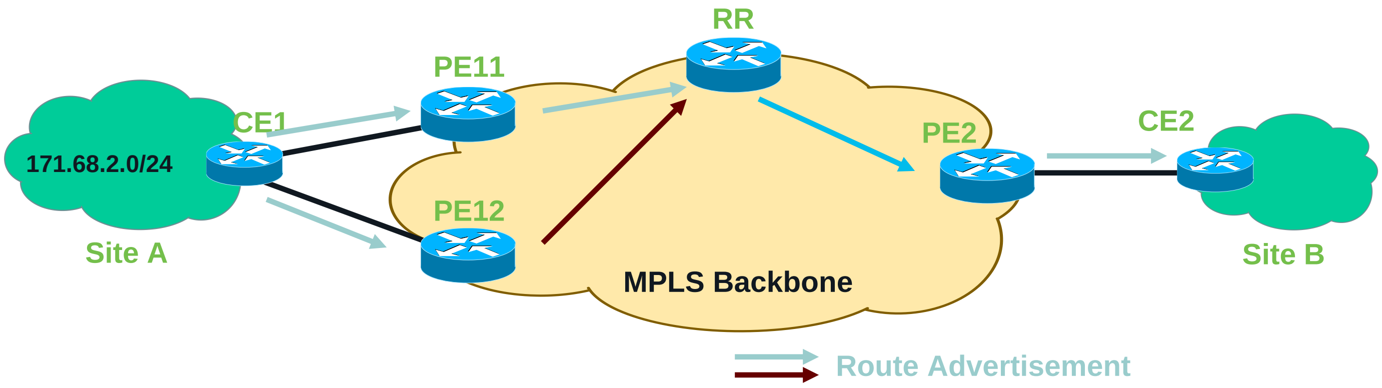

- VPN sites (such as Site A) could be multihomed
- VPN sites need the traffic to (the site A) be loadshared

IP/VPN Deployment Scenarios:

1. Multi-homing & Loadsharing of VPN Traffic

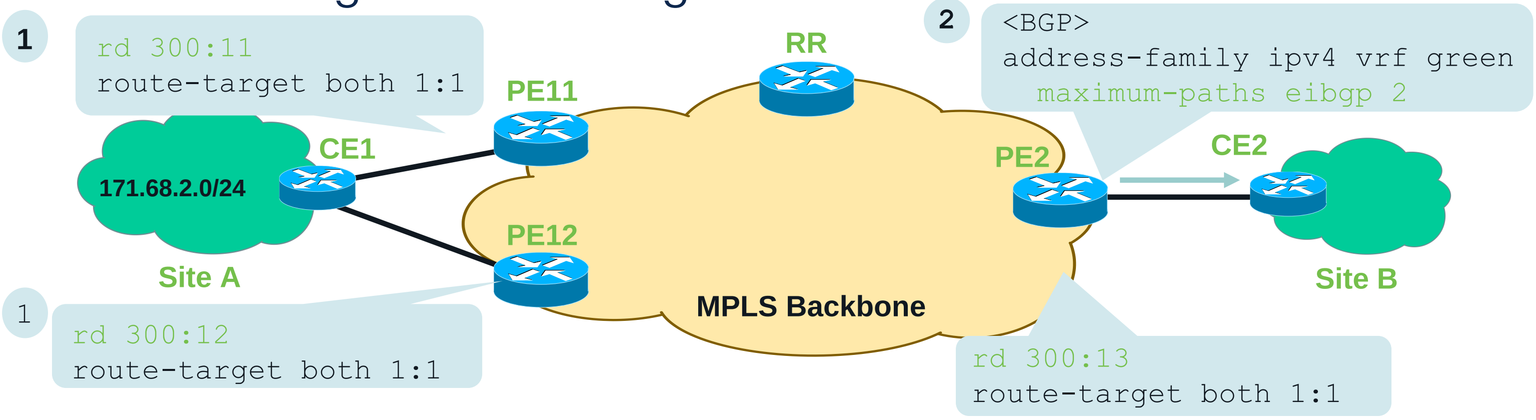

- Configure unique RD per VRF per PE for multi-homed site/interfaces
  - Assuming RR exists
- Enable eiBGP multipath within the relevant BGP VRF address-family at remote PE routers such as PE2 (why PE2?).

Supported in IOS, and IOS-XR

IP/VPN Deployment Scenarios:

1. VPN Fast Convergence—PE-CE Link Failure

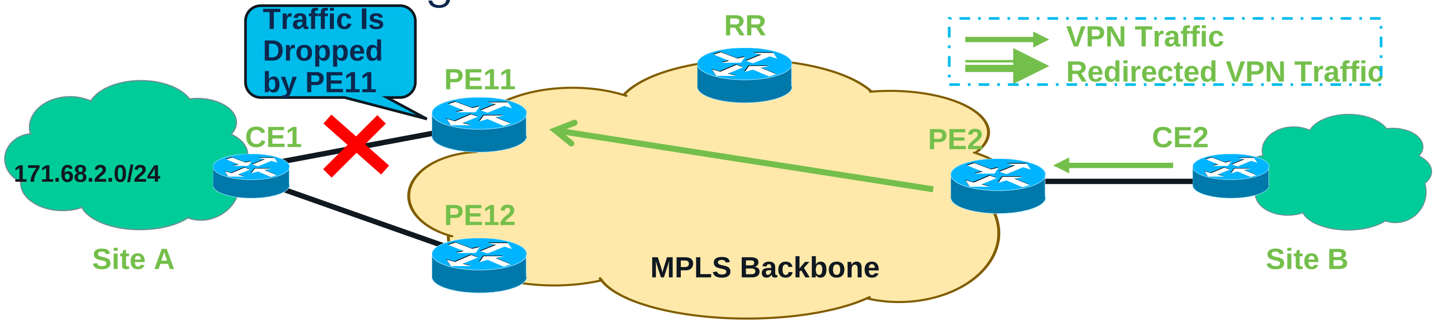

- What if PE11-CE link fails?
  - Wait for BGP convergence (~seconds)

IP/VPN Deployment Scenarios:

1. VPN Fast Convergence—PE-CE Link Failure – PIC Edge Feature

Supported in IOS, and IOS-XR 3.4

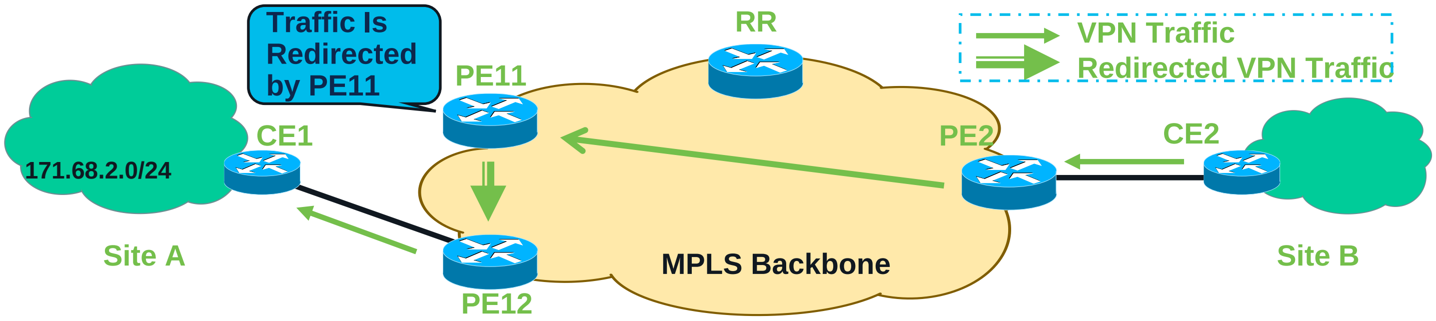

- BGP PIC Edge feature provides fast convergence (~msec) .
  - PE11 temporarily redirects the CE1 bound traffic to PE12 until BGP has converged

- BGP PIC Edge is independent of whether multipath is enabled on PE2 or not

Agenda

- IP/VPN Overview
- IP/VPN Deployment Scenarios
- Best Practices
- Use-Cases
- Conclusion

1. Multihoming & Load-sharing
2. Hub and Spoke
3. Extranet
4. Internet Access
5. IP/VPN over IP Transport
6. Multi-VRF CE

IP/VPN Deployment Scenarios:

2. Hub and Spoke Service

- Many VPN deployments require hub and spoke topology
  - Spoke to spoke communication via Hub site only
  - Example: ATM Machines to HQ, Router Management traffic to NMS/DC
- Despite MPLS based IP/VPN's implicit any-to-any, i.e. full-mesh connectivity, hub and spoke service can easily be offered
  - Uses different import and export of route-target (RT) values
  - Requires unique RD per VRF per PE
- Independent of PE-CE routing protocol per site

IP/VPN Deployment Scenarios:

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 VPN Hub site advertises default or summary routes towards the Spoke sites, otherwise use Option#2

* HDVRF Feature Is Discussed Later

Supported in IOS, NXOS and IOS-XR

IP/VPN Deployment Scenarios:

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

Import and Export RT Values Must Be Different

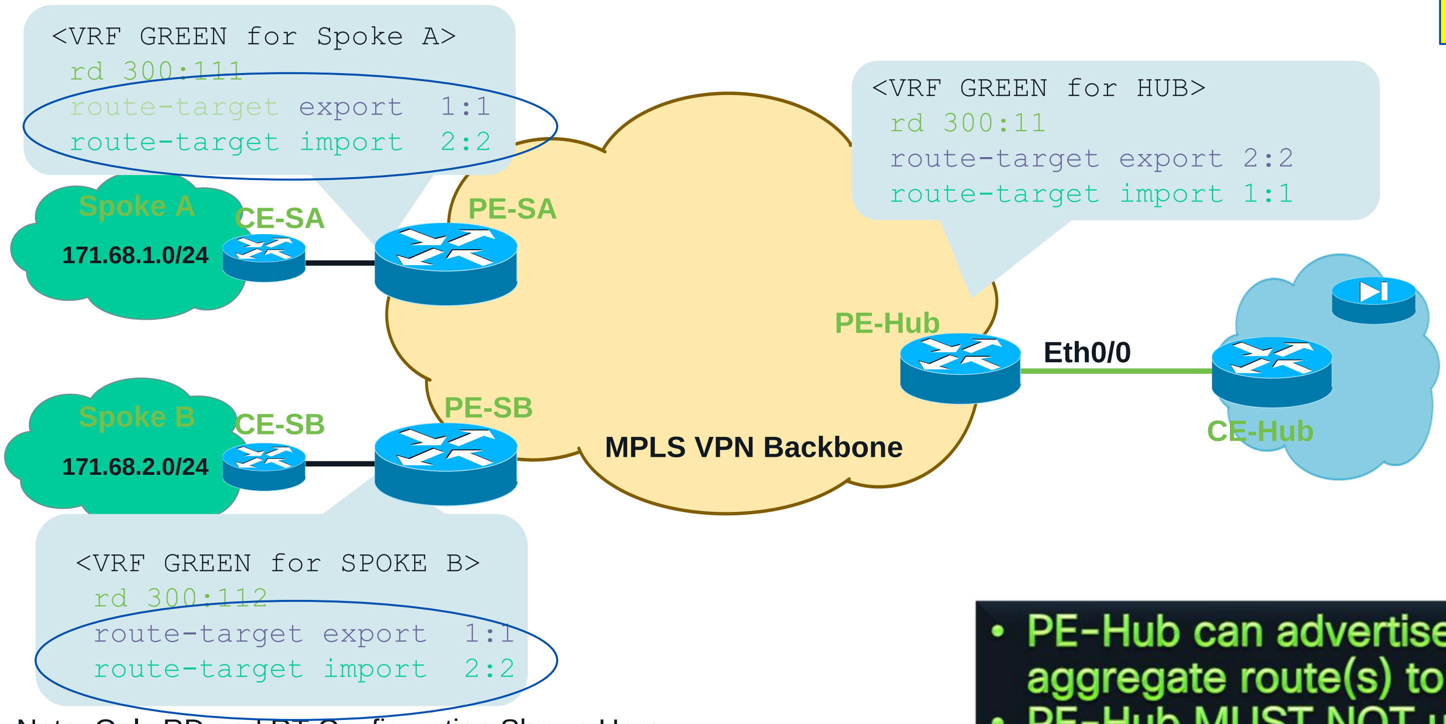

Note: Only RD and RT Configuration Shown Here

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

Supported in IOS, NXOS and IOS-XR

IP/VPN Deployment Scenarios:

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

Import and Export RT Values Must Be Different

```
<VRF GREEN for Spoke A>
rd 300:111
 route-target export  1:1
 route-target import  2:2
```

```
<VRF IN for Hub>
 rd 300:11
 route-target import  1:1
```

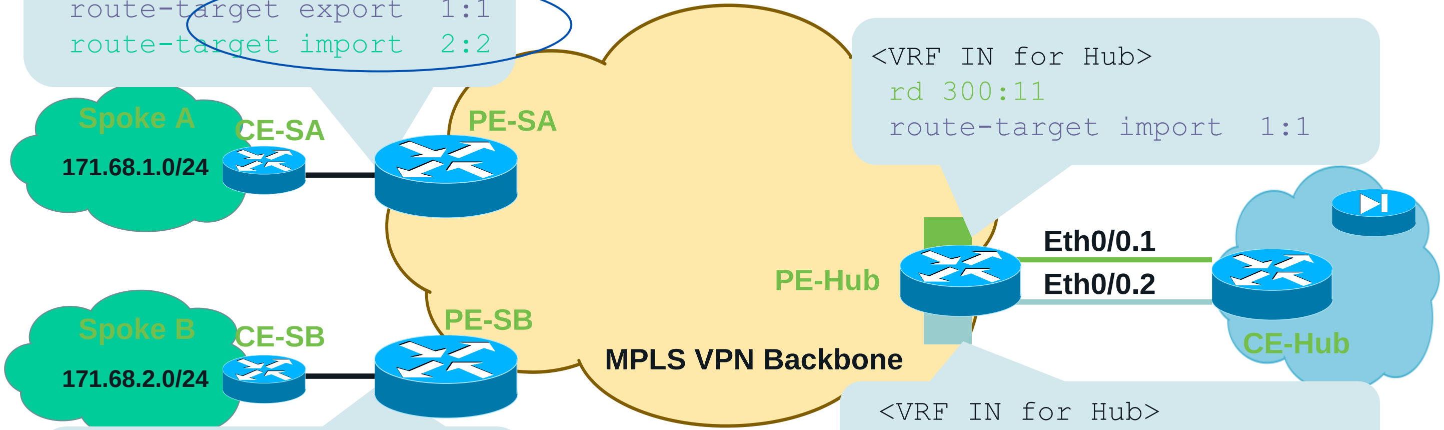


```
<VRF IN for Hub>
 rd 300:12
 route-target export  2:2
```

```
<VRF GREEN for Spoke B>
 rd 300:112
 route-target export  1:1
 route-target import  2:2
```

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

Note: Only RD and RT Configuration Shown Here

Supported in IOS, NXOS and IOS-XR

IP/VPN Deployment Scenarios:

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 Deployment Scenarios:

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

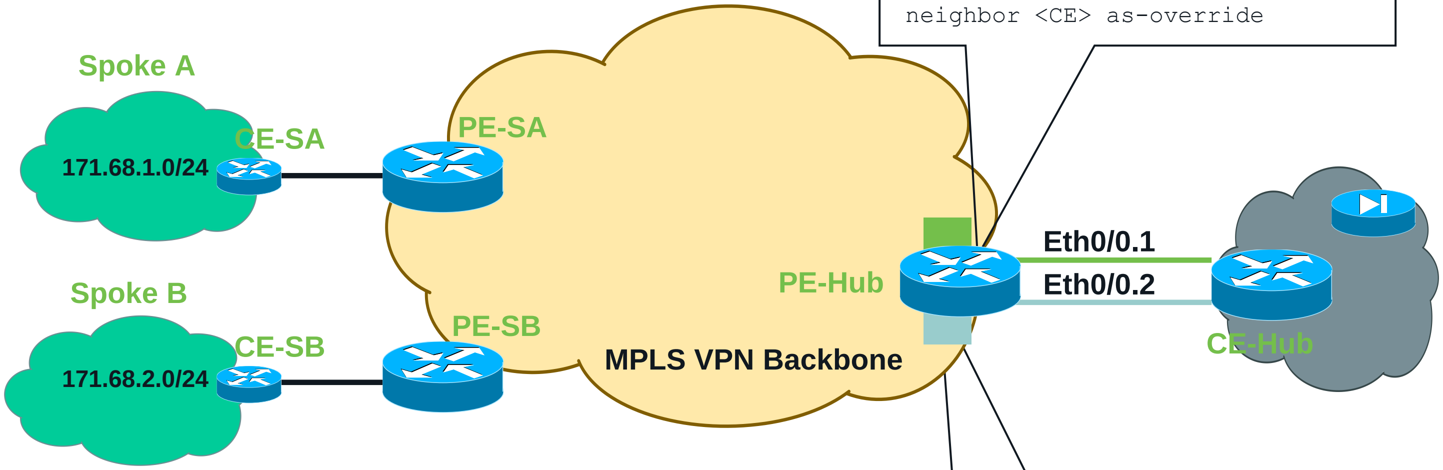


```
<BGP>
address-family ipv4 vrf HUB-IN
 neighbor <CE> as-override
```

```
<BGP>
 address-family ipv4 vrf HUB-OUT
 neighbor <CE> allowas-in 2
```

Supported in IOS, NXOS and IOS-XR

IP/VPN Deployment Scenarios:

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

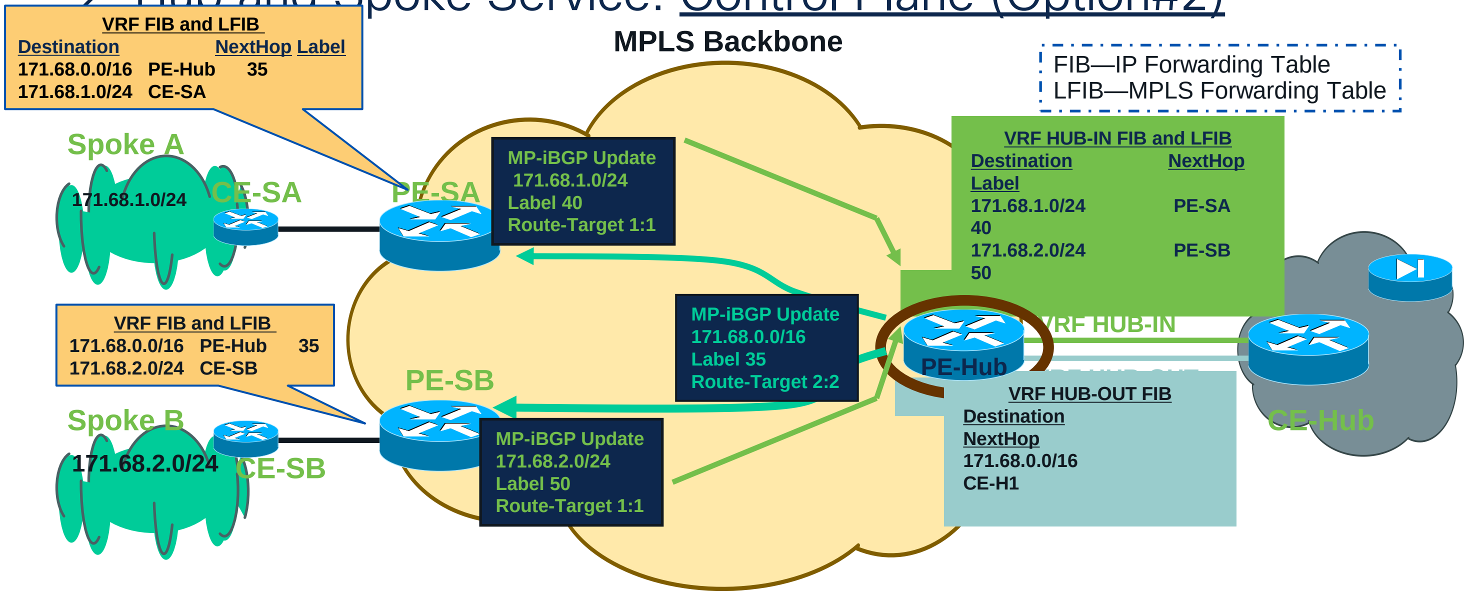

- Two VRFs at the PE-Hub:
  - VRF HUB-IN to learn every spoke routes from remote PEs
  - VRF HUB-OUT to advertise spoke routes or summary 171.68.0.0/16 routes to remote PEs

Supported in IOS, NXOS and IOS-XR

IP/VPN Deployment Scenarios:

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

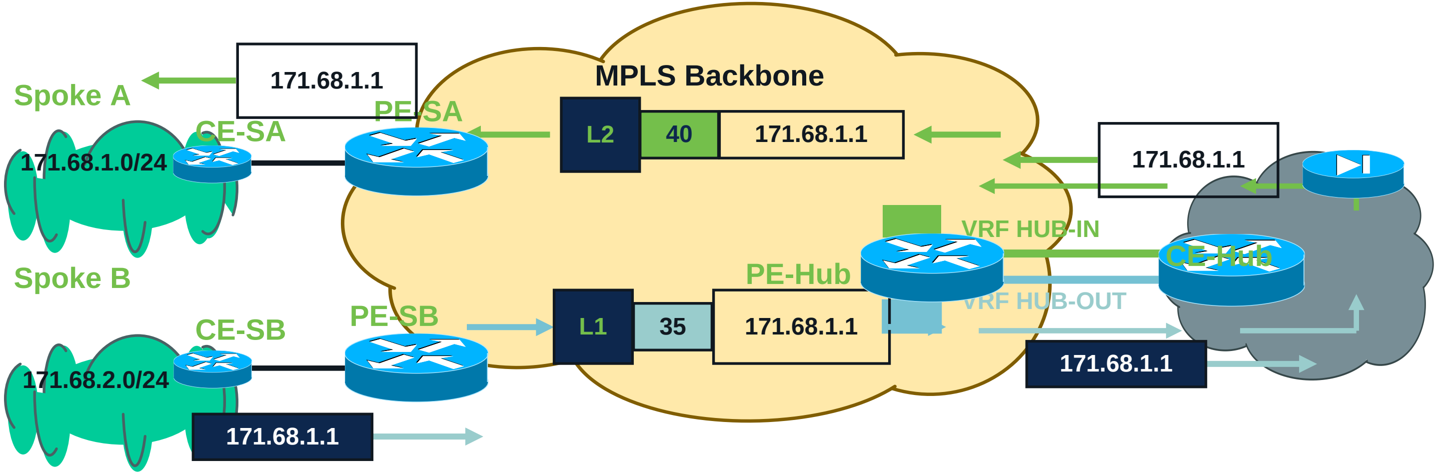

L1 Is the Label to Get to PE-Hub

L2 Is the Label to Get to PE-SA

Supported in IOS and IOS-XR 3.6

IP/VPN Deployment Scenarios:

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 ☹

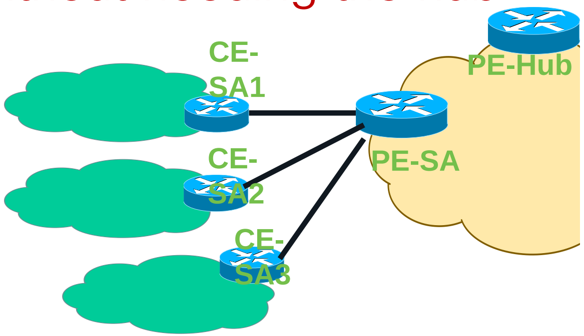

- Half-duplex VRF is the answer
  - Uses two VRFs on the PE (spoke) router :
    - A VRF for spoke->hub communication (e.g. upstream)
    - A VRF for spoke<-hub communication (e.g. downstream)

Note: 12.2(33) SRE. XE 3.0S Support Any Interface Type (Eth, Ser, POS, Virtual-Access, etc.)

IP/VPN Deployment Scenarios:

2. Hub and Spoke Service: Half-Duplex VRF

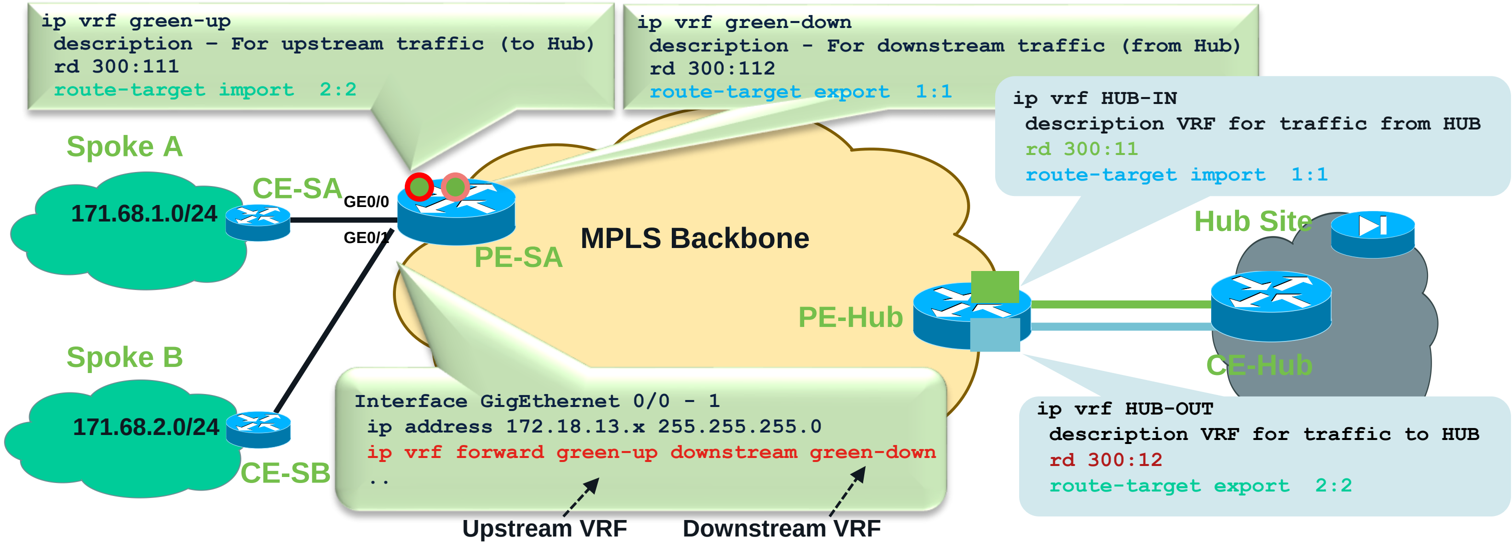

1. PE-SA installs the Spoke routes only in downstream VRF i.e. green-down
2. PE-SA installs the Hub routes only in upstream VRF i.e. green-up
3. PE-SA forwards the incoming IP traffic (from Spokes) using upstream VRF i.e. green-up routing table.
4. PE-SA forwards the incoming MPLS traffic (from Hub)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

IP/VPN Deployment Scenarios:

2. Hub and Spoke Service: Half-Duplex VRF on every PE

Supported in IOS

TBV

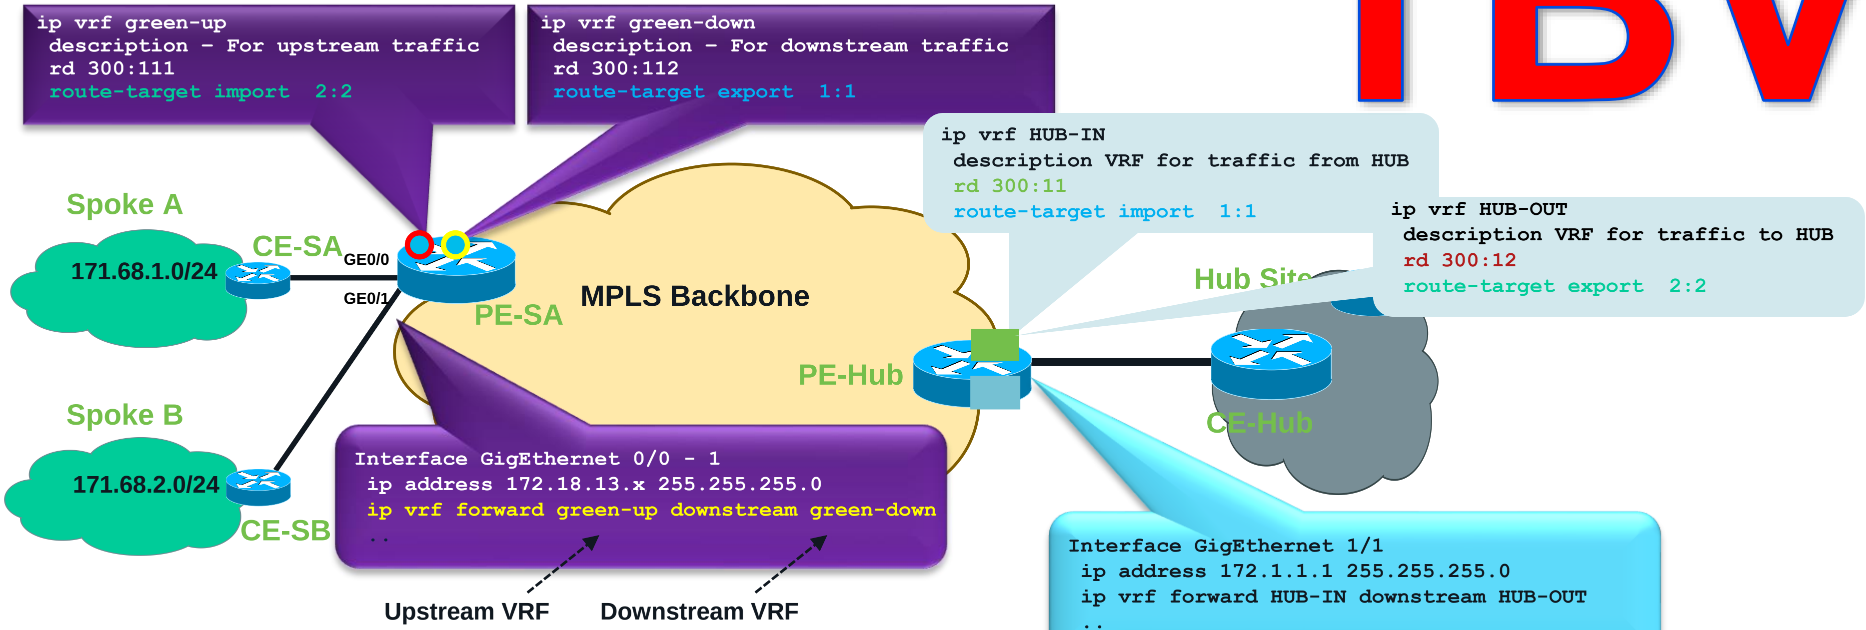

- Single PE-CE interface on Hub

Agenda

- IP/VPN Overview
- IP/VPN Deployment Scenarios
- Best Practices
- Use-Cases
- Conclusion

1. Multihoming & Load-sharing
2. Hub and Spoke
3. Extranet
4. Internet Access
5. IP/VPN over IP Transport
6. Multi-VRF CE

IP/VPN Deployment Scenarios

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.
- Implemented by sharing import and export route-target (RT) values within the VRFs of extranets.
  - Export-map or import-map may be used for advanced extranet.

IP/VPN Deployment Scenarios

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

Supported in IOS, NXOS and IOS-XR

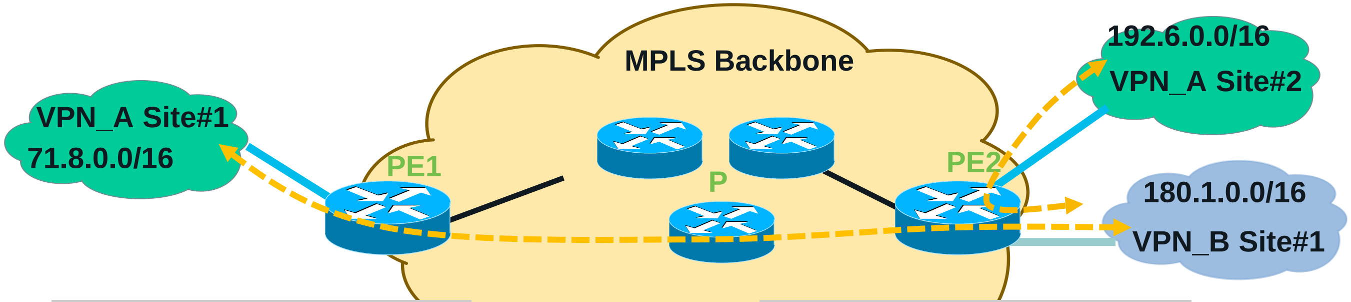


```
<VRF for VPN_A>
route-target import 3000:111
route-target export 3000:111
route-target import 3000:222
```

```
<VRF for VPN_B>
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

IP/VPN Deployment Scenarios

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

Supported in IOS, NXOS and IOS-XR

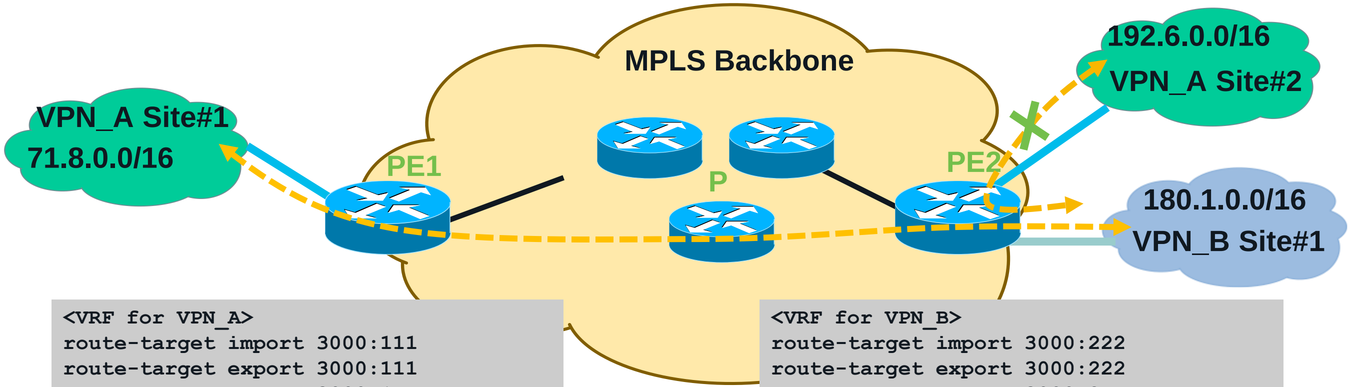


```
<VRF for VPN_A>
route-target import 3000:111
route-target export 3000:111
route-target import 3000:1
import map VPN_A_Import
export map VPN_A_Export
!
route-map VPN_A_Export permit 10
 match ip address 1
 set extcommunity rt  3000:2 additive
!
route-map VPN_A_Import permit 10
 match ip address 2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

```
<VRF for VPN_B>
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 Deployment Scenarios
- Best Practices
- Use-Cases
- Conclusion

1. Multihoming & Load-sharing
2. Hub and Spoke
3. Extranet
4. Internet Access
5. IP/VPN over IP Transport
6. Multi-VRF CE

IP/VPN Deployment Scenarios

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

IP/VPN Deployment Scenarios

4. Internet Access: Design Options

Three 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

IP/VPN Deployment Scenarios
4. Internet Access: Design Options

- VRF specific default route
- Static default route to move traffic from VRF to Internet (global routing table)
- Static routes for VPN customers to move traffic from Internet (global routing table) to VRF
- Works well, but doesn’t scale well (limited to default routing)

- Separate PE-CE Interface
- Besides VRF interface, a global interface also connect to each VPN site
- May use eBGP on the global interface, if dynamic routing pr internet routes are needed
- Works well and scales well, despite the operational overhead

- Extranet with Internet-VRF
- Internet routes inside a dedicated VRF (e.g. Internet-VRF)
- Extranet between Internet-VRF and Customer VRFs that need internet access
-

IP/VPN Deployment Scenarios: Internet Access

Supported in IOS

4.1 Option#1: VRF Specific Default Route

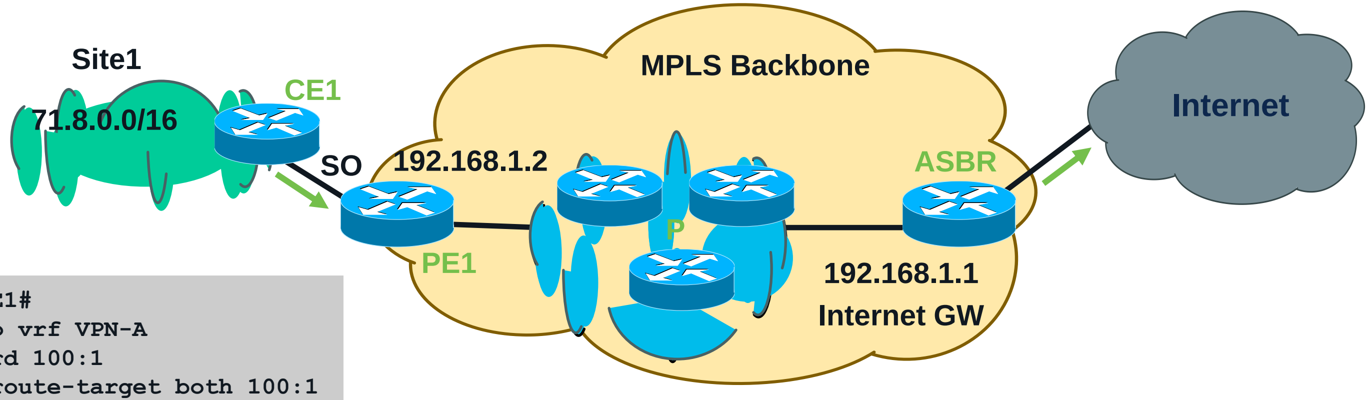


```
PE1#
ip vrf VPN-A
 rd 100:1
 route-target both 100:1
```

```
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 Deployment Scenarios: Internet Access

Supported in IOS,

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

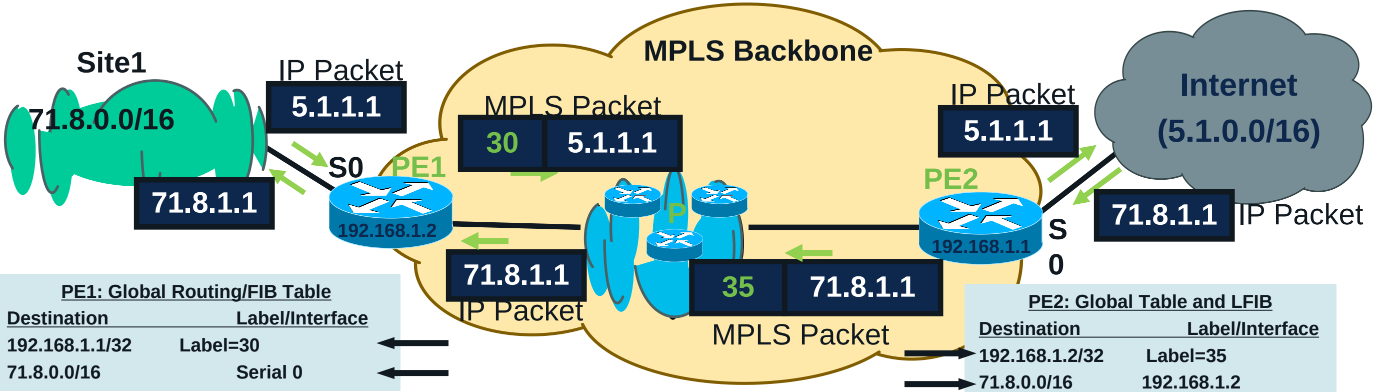

PE1: Global Routing/FIB Table

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)

IP/VPN Deployment Scenarios: Internet Access

Supported in IOS, NXOS and IOS-XR

4.2 Option#2: Separate PE-CE Subinterfaces

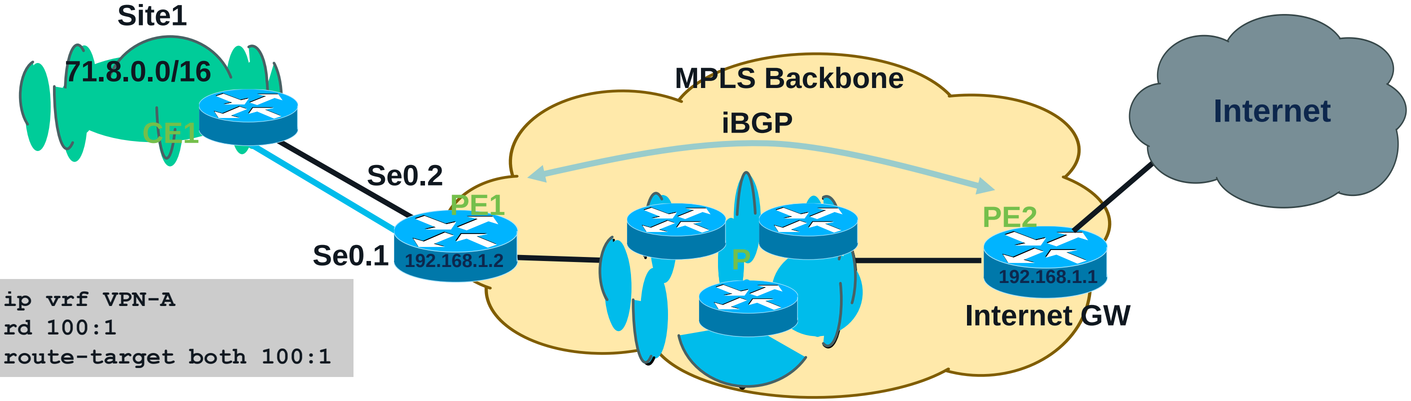


```
ip vrf VPN-A
rd 100:1
route-target both 100:1
```

```
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.

IP/VPN Deployment Scenarios: Internet Access

Supported in IOS, NXOS and IOS-XR

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

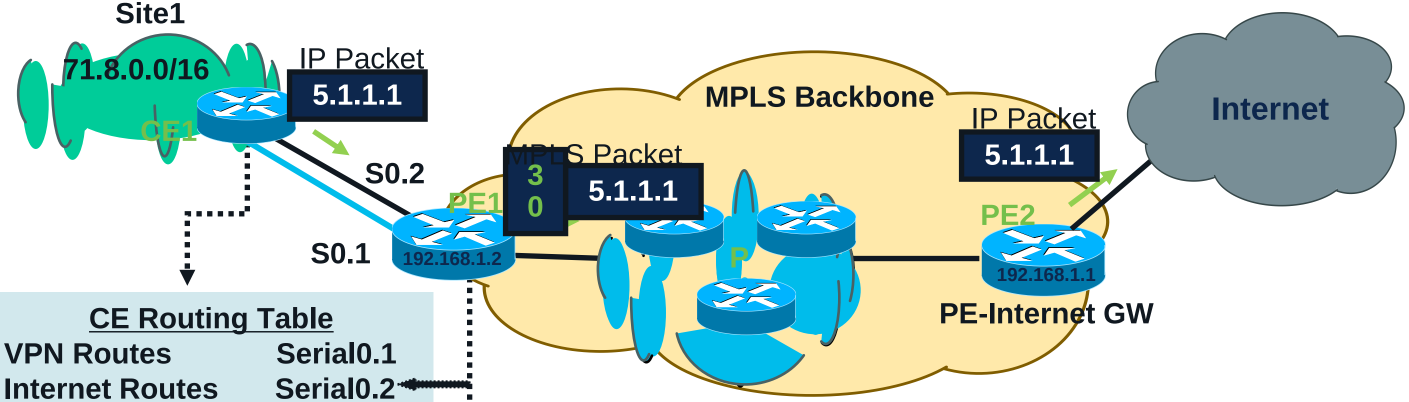

CE Routing Table

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 Deployment Scenarios: Internet Access

4.3 Option#3: Extranet with Internet

- 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 Deployment Scenarios: Internet Access

4.3 Option#3: VPN Extranet with Global (Internet) Table

- Export an IPv6/v4 prefix from VRF to Global routing table
- Import a VPNv6/v4 prefix from Global routing table into VRF
- Advertise imported prefixes to the CE router

```
VRF red
  import ipv4 unicast map foo (10.5/16)      # from global
  export ipv4 unicast map bar (192.34/16)      # to global
```

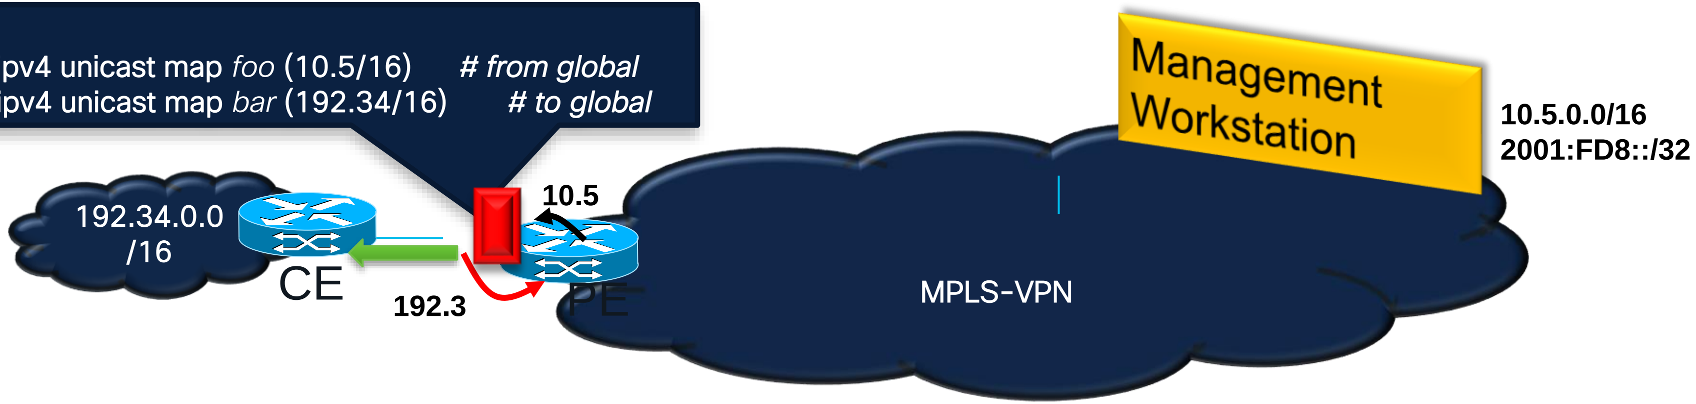

VRF <-> Global Route Leaking

VRF <-> Global Route Leaking eBGP (CE) and iBGP (PE) Advertisement

IOS-XR 4.3.1
IOS-XE 3.10

- Export an IPv6/v4 prefix from VRF to Global routing table
- Import a VPNv6/v4 prefix from Global routing table into VRF
- Advertise imported prefixes to the CE router and optionally PE

```
VRF red
  import ipv4 unicast map foo (10.5/16) export     # from global
  export ipv4 unicast map bar (192.34/16)           # to global
```

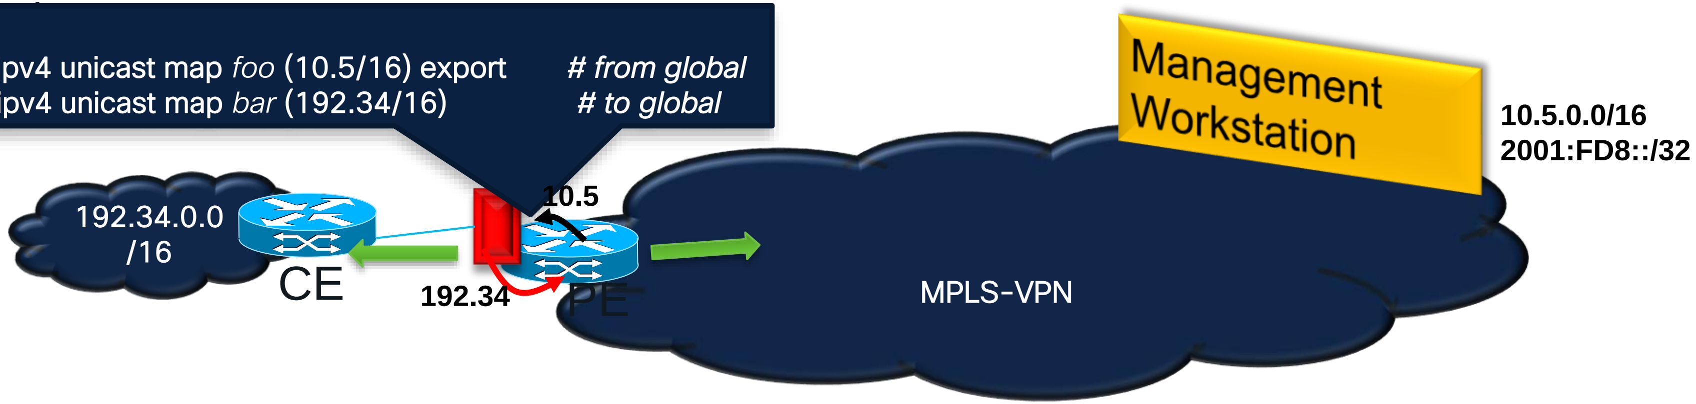

VRF <-> Global Route Leaking

Supported in IOS,

IP/VPN Deployment Scenarios: 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 Deployment Scenarios
- Best Practices
- Use-Cases
- Conclusion

1. Multihoming & Load-sharing
2. Hub and Spoke
3. Extranet
4. Internet Access
5. IP/VPN over IP Transport
6. Multi-VRF CE

Supported in IOS, NXOS and IOS-XR

IP/VPN Deployment Scenarios:

5. Providing MPLS/VPN over IP Transport

- MPLS/VPN (rfc2547) can also be deployed using IP transport
  - No MPLS needed in the core
- PE-to-PE IP tunnel is used, instead of MPLS tunnel, for sending MPLS/VPN packets
  - MPLS labels are still allocated for VPN prefixes by PE routers and used only by the PE routers
  - MPLS/VPN packet is encapsulated inside an IP header
- IP tunnel could be point-to-point or Multipoint GRE encapsulation based.

http://www.cisco.com/en/US/docs/ios/interface/configuration/guide/ir_mplsvpnomgre.html

Supported in IOS, NXOS and IOS-XR

IP/VPN Deployment Scenarios:

5. Providing MPLS/VPN over IP Transport

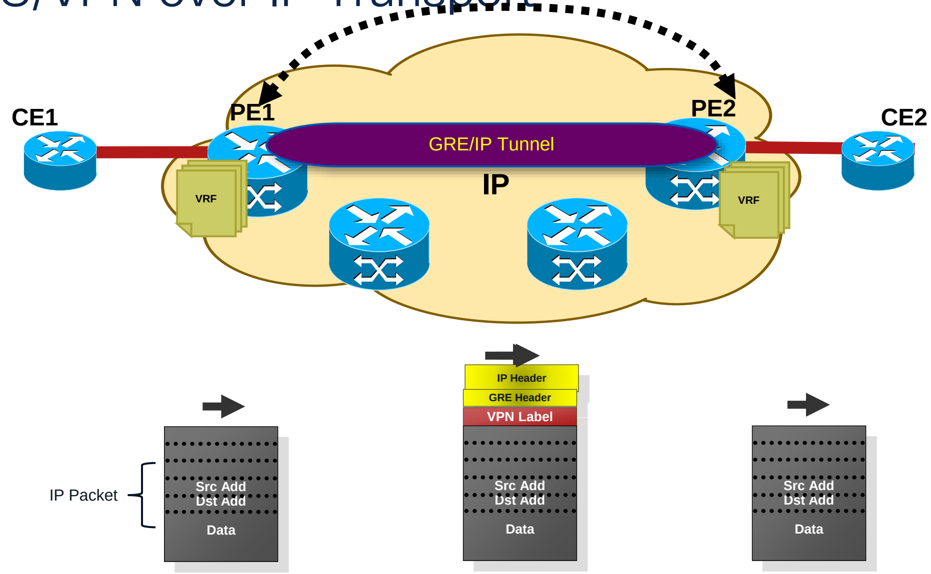

- GRE/IP header and VPN label imposed on VPN traffic by PE1
- VPN traffic is forwarded towards egress PE using IP forwarding
- Egress PE2 decapsulates, and uses VPN label to forward packet to CE2

Source -- http://www.cisco.com/en/US/docs/ios/interface/configuration/guide/ir_mplsvpnomgre.html

Agenda

- IP/VPN Overview
- IP/VPN Deployment Scenarios
- Best Practices
- Use-Cases
- Conclusion

1. Multihoming & Load-sharing
2. Hub and Spoke
3. Extranet
4. Internet Access
5. IP/VPN over IP Transport
6. IPv6 VPN
7. Multi-VRF CE

IP/VPN Deployment Scenarios:

6. IPv6 VPN Service

Supported in IOS, NXOS and IOS-XR

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

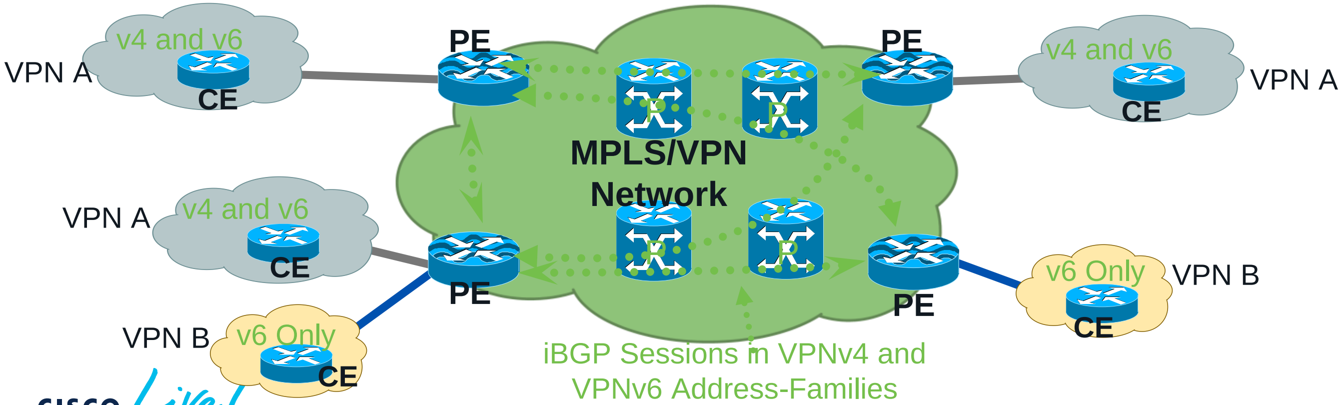

Agenda

- IP/VPN Overview
- IP/VPN Deployment Scenarios
- Best Practices
- Use-Cases
- Conclusion

1. Multihoming & Load-sharing
2. Hub and Spoke
3. Extranet
4. Internet Access
5. IP/VPN over IP Transport
6. Site Segmentation

Supported in IOS, NXOS and IOS-XR

IP/VPN Deployment Scenarios:

7. Providing Multiple VPNs inside VPN Site

- Is it possible for a CE router to keep multiple customer connections separated ?
  - Yes, “multi-VRF CE” a.k.a. vrf-lite can be used
- “Multi-VRF CE” provides multiple virtual routing tables (and forwarding tables) per customer at the CE router
  - Not a feature but an application based on VRF implementation
  - Any routing protocol that is supported by normal VRF can be used in a multi-VRF CE implementation
- No MPLS functionality needed on CE, no label exchange between CE and any router (including PE) ☺

IP/VPN Deployment Scenarios:

7. Multi-VRF CE aka VRF-Lite

Supported in IOS, NXOS and IOS-XR

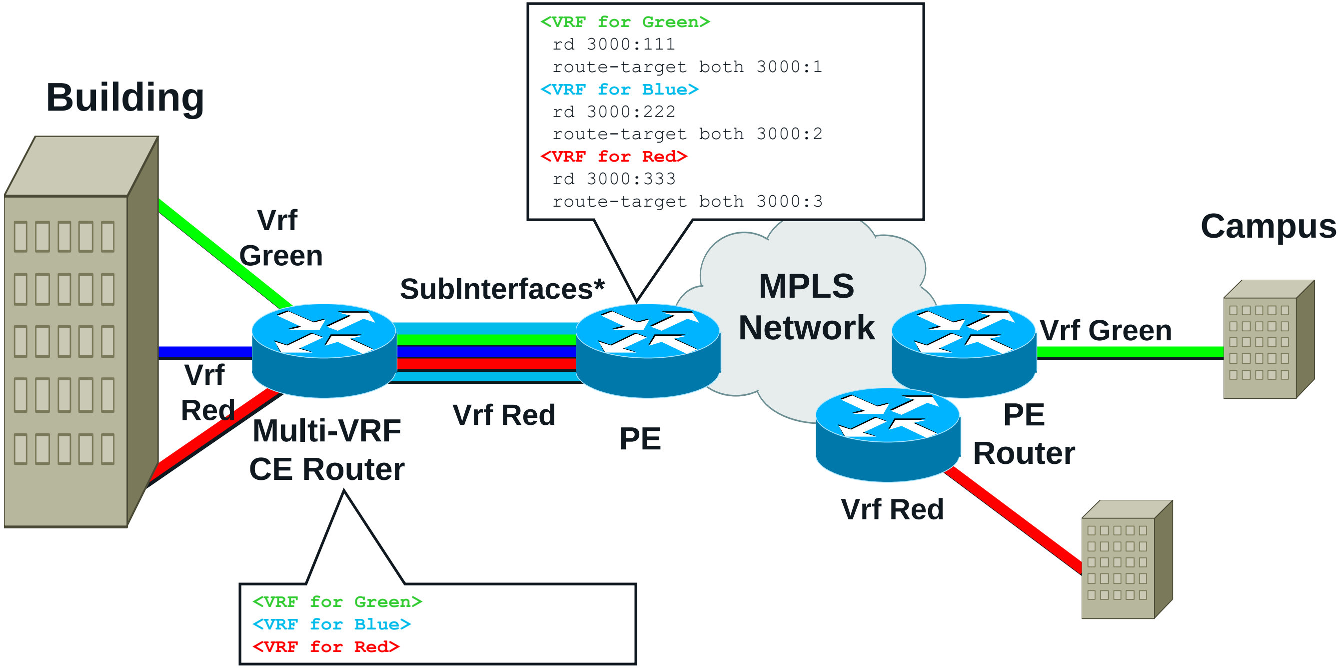

One of Deployment Models for VRF-Lite is Campus Virtualization:=

Extending IP/VPN to CE

*SubInterfaces –Any Interface Type that Supports Sub Interfaces = Ethernet Vlan, Frame Relay, ATM VCs

IP/VPN Deployment Scenarios:

6. IPv6 VPN Service

Supported in IOS, NXOS and IOS-XR

```
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
 !
```

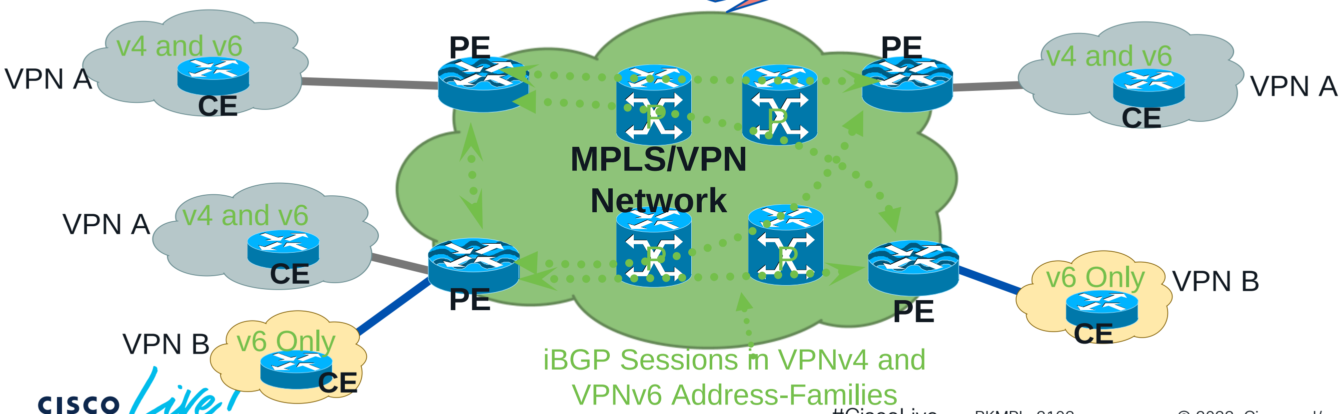

Technical Session Surveys

- Attendees who fill out a minimum of four session surveys and the overall event survey will get Cisco Live branded socks!
- Attendees will also earn 100 points in the Cisco Live Game for every survey completed.
- These points help you get on the leaderboard and increase your chances of winning daily and grand prizes.

Cisco Learning and Certifications

From technology training and team development to Cisco certifications and learning plans, let us help you empower your business and career. www.cisco.com/go/certs

Pay for Learning with Cisco Learning Credits

(CLCs) are prepaid training vouchers redeemed directly with Cisco.

Learn

Cisco U.
IT learning hub that guides teams and learners toward their goals

Cisco Digital Learning
Subscription-based product, technology, and certification training

Cisco Modeling Labs
Network simulation platform for design, testing, and troubleshooting

Cisco Learning Network
Resource community portal for certifications and learning

Train

Cisco Training Bootcamps
Intensive team & individual automation and technology training programs

Cisco Learning Partner Program
Authorized training partners supporting Cisco technology and career certifications

Cisco Instructor-led and Virtual Instructor-led training
Accelerated curriculum of product, technology, and certification courses

Certify

Cisco Certifications and Specialist Certifications
Award-winning certification program empowers students and IT Professionals to advance their technical careers

Cisco Guided Study Groups
180-day certification prep program with learning and support

Cisco Continuing Education Program
Recertification training options for Cisco certified individuals

Here at the event? Visit us at **The Learning and Certifications lounge at the World of Solutions**

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

The bridge to possible

Thank you

cisco Live!

ALL IN

#CiscoLive