
TOMORROW
starts here.
CISCO
Cisco live!

VRF, MPLS and MPBGP Fundamentals

BRKCRT-2601

Andre Laurent, 3X CCIE/CCDE

#clmel

Cisco live!

Agenda

- Introduction to Virtualisation
- VRF-Lite
- MPLS – BGP Free Core
- MP-BGP
- Conclusion

Introduction

MPLS

The Common Perception

- CE Routers owned by customer
- PE Routers owned by SP
- Customer "peers" to "PE" via IP
- Exchanges routing with SP via routing protocol (or static route)
- Customer relies on SP to advertise routes to reach other customer CEs

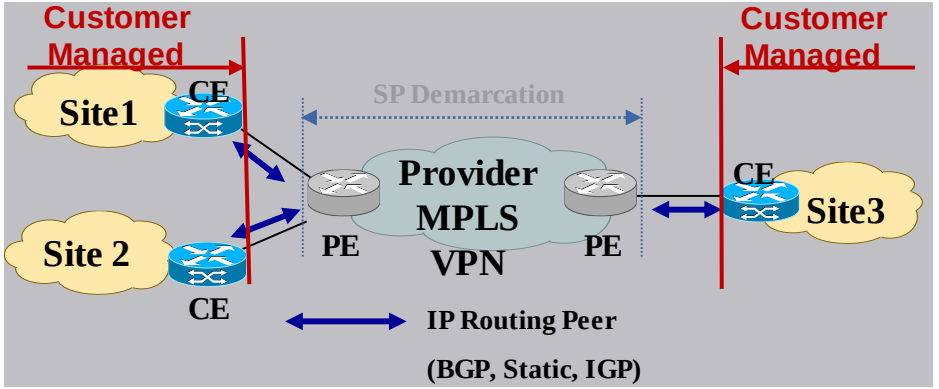

* **No** Labels Are Exchanged with the SP

Enterprise Network Virtualisation

Key Building Blocks

Device Partitioning

“Virtualising” the Routing and Forwarding of the Device

Virtualised Interconnect

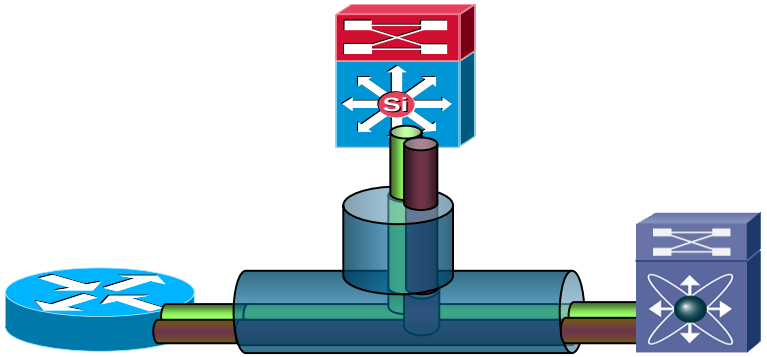

Extending and Maintaining the “Virtualised” Devices/Pools over Any Media

Device Partioning

Layer 2 vs. Layer 3 Virtualisation

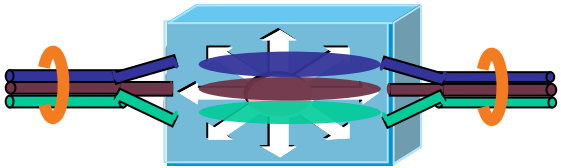

VLAN—Virtual LAN

- Virtualise at Layer 2 forwarding
- Associates to one or more L2 interfaces on switch
- Has its own MAC forwarding table and spanning-tree instance per VLAN
- Interconnect options?

  VLANs are extended via a physical cable or virtual 802.1q trunk

VRF—Virtual Routing and Forwarding

- Virtualise at Layer 3 forwarding
- Associates to one or more Layer 3 interfaces on router/switch
- Each VRF has its own

  Forwarding table (CEF)

  Routing process (RIP, EIGRP, OSPF, BGP)
- Interconnect options (VRF-Lite)?

  802.1q, GRE, sub-interfaces, physical cables, signalling

Path Isolation

Functional Components

Device virtualisation

- Control plane virtualisation
- Data plane virtualisation
- Services virtualisation

Data path virtualisation

- Hop-by-Hop - VRF-Lite End-to-End
- Multi-Hop - VRF-Lite GRE
- MPLS-VPN
- MPLS VPN over IP
- MPLS VPN over DMVPN
- MPLS VPN o GRE/mGRE

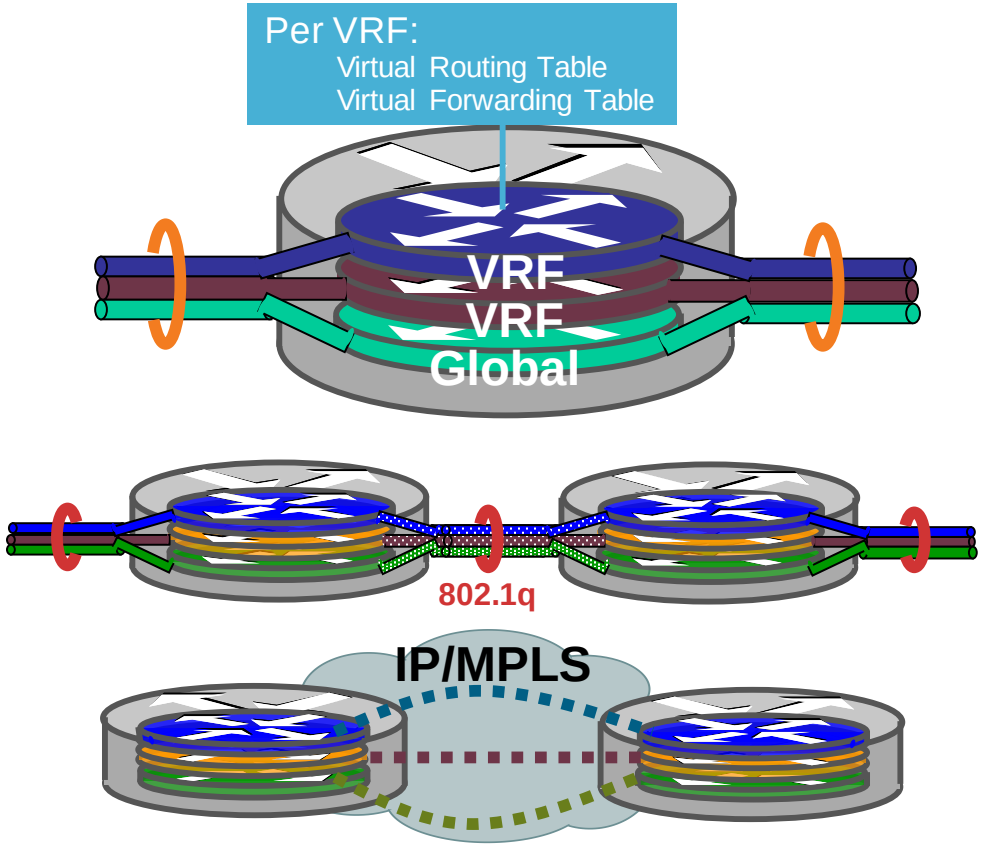


VRF-Lite

What is VRF-Lite?

Functional Components

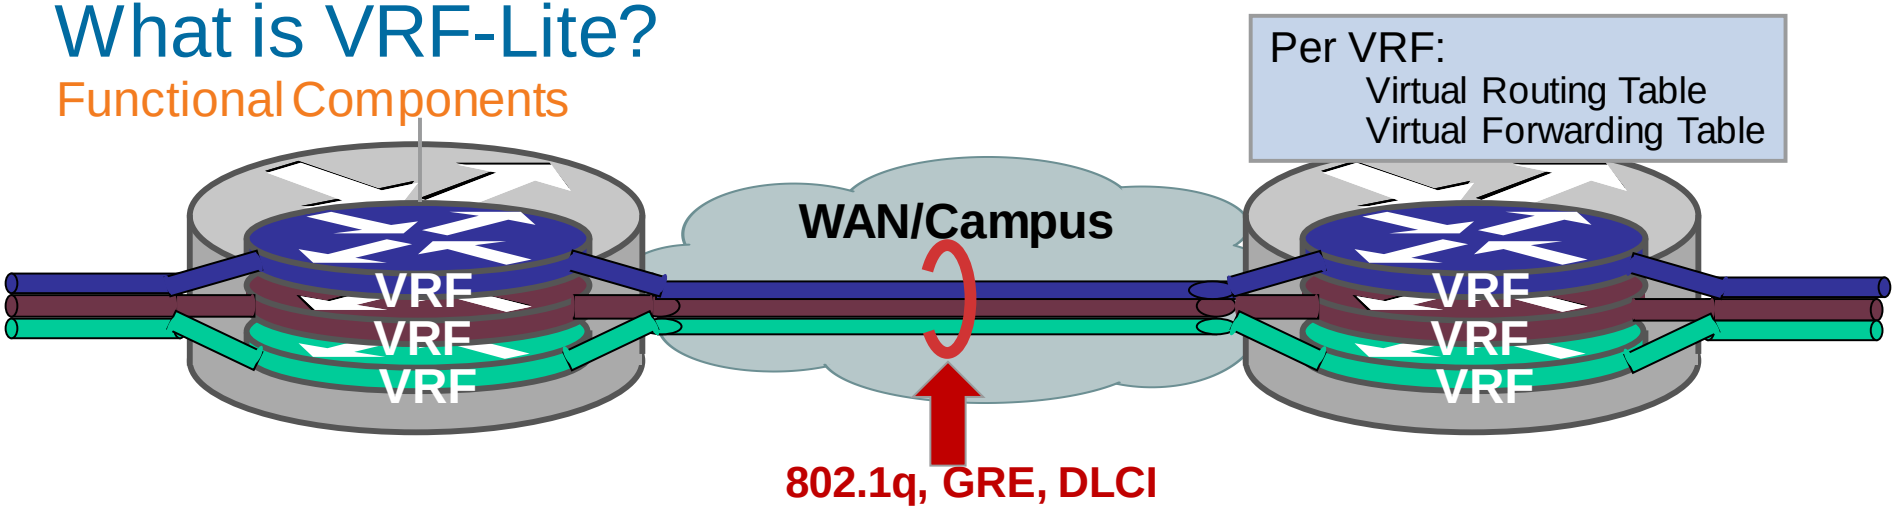

- Defined router supports **routing (RIB), forwarding (FIB), and interface per VRF**
- Leverages "Virtual" encapsulation for separation:
  - Ethernet/802.1Q, GRE, Frame Relay
- The routing protocol is also "VRF aware"
  - RIP/v2, EIGRP, OSPF, BGP, static (per VFR)
- Layer 3 VRF interfaces cannot belong to more than a single VRF

VRF-Lite End-to-End

Target Requirements

- End to End segmentation, per VRF and per interface
- Targets a small number of VRFs as requirement
- Seen frequently in Access → Distribution (vs. end to end)
- No MP-BGP or control plane signalling is required and does not use labels
- No LDP is required (i.e. MPLS)
- Still leverages existing QoS model and supports IP multicast
- Full range of platform support within the Cisco portfolio of switches and routers

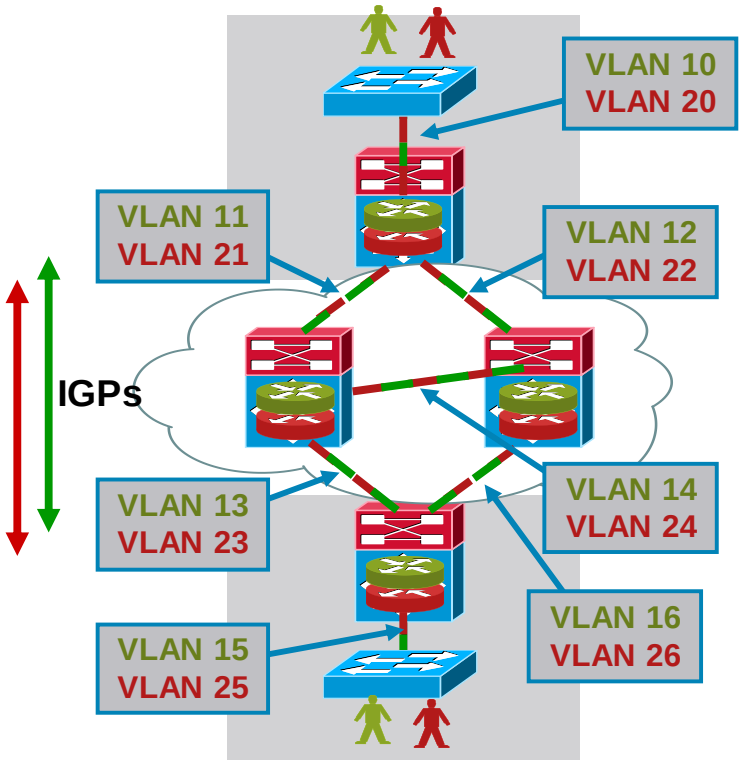

VRF-Lite

Sub-interface Example

Per VRF:
- Virtual Routing Table
- Virtual Forwarding Table
- Locally Significant

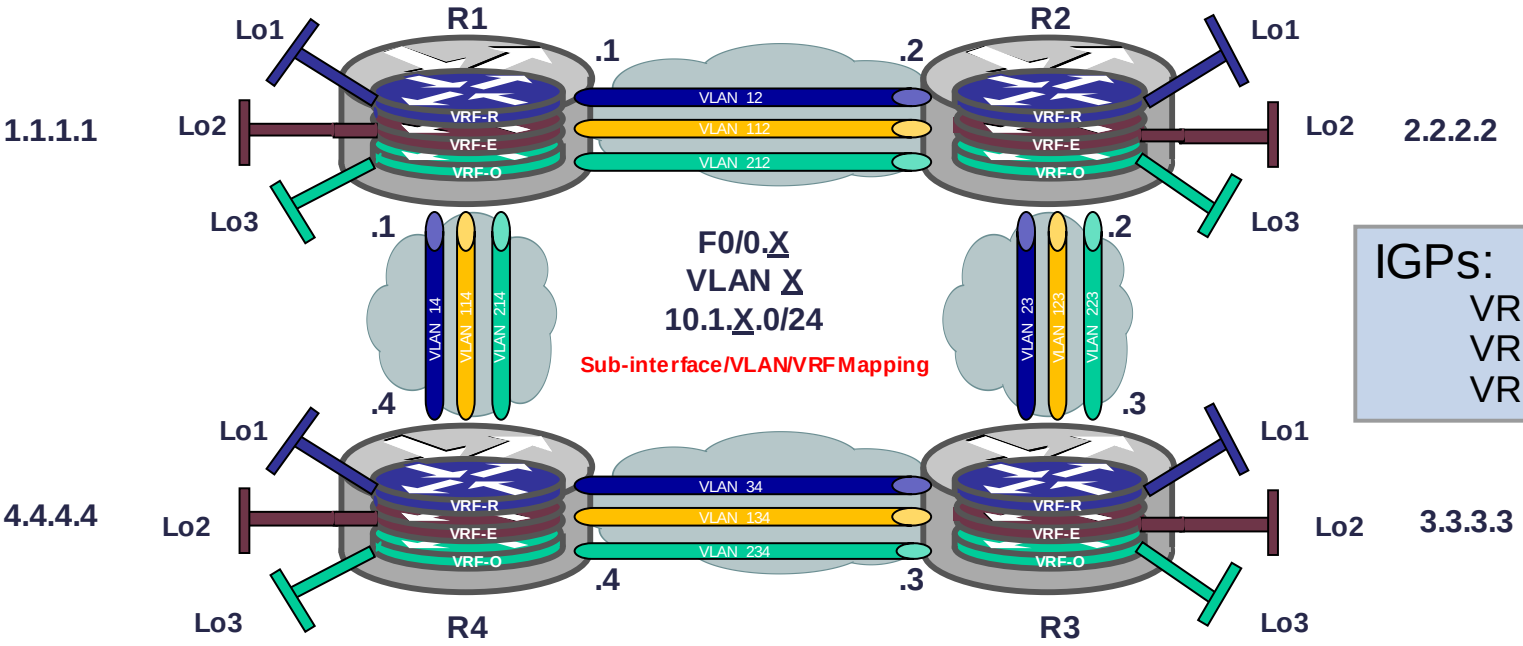

IGPs:
- VRF-R = RIP
- VRF-E = EIGRP
- VRF-O = OSPF

VRF-Lite Sub-interface Configuration

Command Line Interface (CLI) Review

```
ip vrf VRF-R
 rd 1:1

interface FastEthernet0/0.12
 ip vrf forwarding VRF-R

interface Loopback1
 ip vrf forwarding VRF-R

ip vrf VRF-E
 rd 2:2

interface FastEthernet0/0.112
 ip vrf forwarding VRF-E

interface Loopback2
 ip vrf forwarding VRF-R

ip vrf VRF-O
 rd 3:3

interface FastEthernet0/0.212
 ip vrf forwarding VRF-O

interface Loopback3
 ip vrf forwarding VRF-R
```

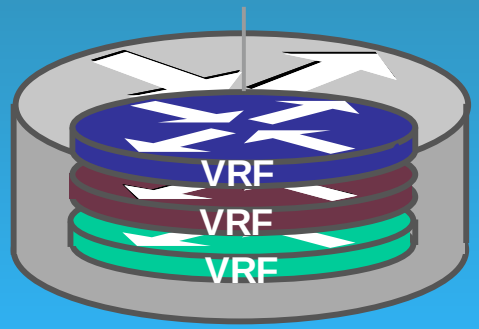

VRF Aware RIP Configuration

Command Line Interface (CLI) Review

```
router rip
 version 2
 network 1.0.0.0
 network 10.0.0.0
 no auto-summary
```

```
router rip
 !
 address-family ipv4 vrf VRF-R
  network 1.0.0.0
  network 10.0.0.0
  no auto-summary
  version 2
 exit-address-family
```

RIP leverages address-family ipv4 vrf _____

Similar to what you already know!

VRF Aware EIGRP Configuration

Command Line Interface (CLI) Review

```
router eigrp 10
 network 1.1.1.1 0.0.0.0
 network 10.1.112.0 0.0.0.255
 no auto-summary
```

```
router eigrp 10
 auto-summary
 !
 address-family ipv4 vrf VRF-E
  network 1.1.1.1 0.0.0.0
  network 10.1.112.0 0.0.0.255
  no auto-summary
  autonomous-system 10
 exit-address-family
```

EIGRP leverages address-family ipv4 vrf _____

Set unique autonomous system number per VRF

Similar to what you already know!

VRF Aware OSPF Configuration

Command Line Interface (CLI) Review

```
router ospf 1
 log-adjacency-changes
 network 1.1.1.1 0.0.0.0 area 1
 network 10.1.212.0 0.0.0.255 area 0
```

```
router ospf 2 vrf VRF-O
 log-adjacency-changes
 network 1.1.1.1 0.0.0.0 area 1
 network 10.1.212.0 0.0.0.255 area 0
```

OSPF leverages vrf _____ after the unique process number

Similar to what you already know!

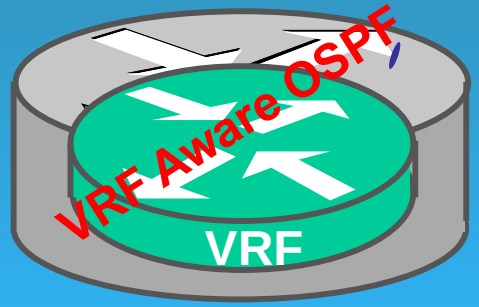


Live Exploration

No Sub-interface Support/No Problem

GRE Example

- VRF Lite can also leverage GRE tunnels as a segmentation technology
- Each VRF uses a unique GRE tunnel
- GRE tunnel interface is “VRF aware”

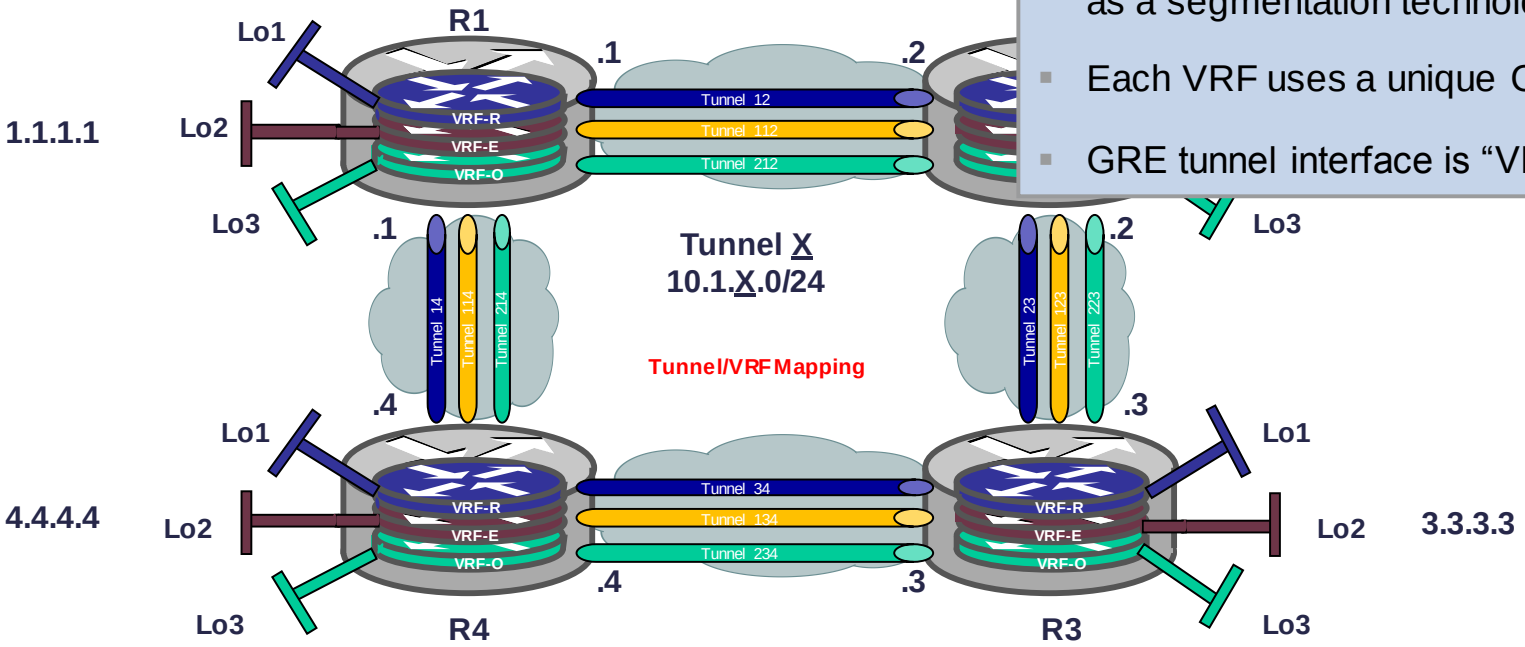

Configuration Note: Each GRE Tunnel Could Require Unique Source/Destination IP (Platform Dependent)

VRF-Lite Tunnel Configuration

Command Line Interface (CLI) Review

```
ip vrf VRF-S
 rd 11:11

interface Loopback11
 ip address 11.11.11.11 255.255.255.255 (Global Routing Table)

interface Tunnel12
 ip vrf forwarding VRF-S
 ip address 10.1.12.1 255.255.255.0
 tunnel source Loopback11
 tunnel destination 22.22.22.22
```

```
ip vrf VRF-E
 rd 22:22

interface Loopback12
 ip address 22.22.22.22 255.255.255.255 (Global Routing Table)

interface Tunnel12
 ip vrf forwarding VRF-S
 ip address 10.1.12.2 255.255.255.0
 tunnel source Loopback22
 tunnel destination 11.11.11.11
```

Similar to what you already know!

```
ip route vrf VRF-S 2.2.2.2 255.255.255.255 10.1.12.2
```


```
ip route vrf VRF-S 1.1.1.1 255.255.255.255 10.1.12.1
```

Layer 2 Serial Link/No Problem

Back-to-Back Frame Relay Example

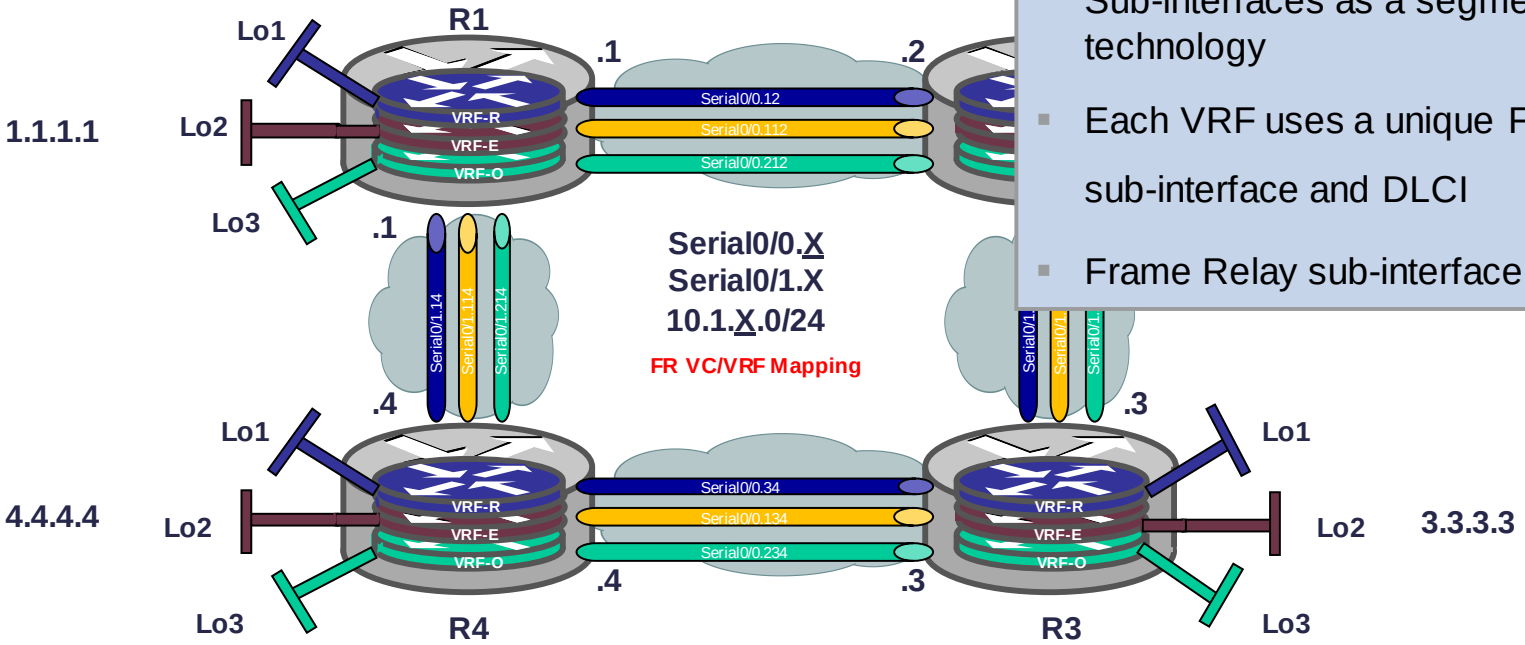

- VRF Lite can also leverage Frame Relay Sub-interfaces as a segmentation technology
- Each VRF uses a unique Frame-Relay sub-interface and DLCI
- Frame Relay sub-interface is “VRF aware”

Configuration Note: Leveraging Back-to-Back Frame-Relay Configuration

VRF-Lite Back-to-Back Frame Relay Configuration

Command Line Interface (CLI) Review

```
ip vrf VRF-B
 rd 111:111

interface Serial0/0
 encapsulation frame-relay

Interface Serial0/0.12
 ip vrf forwarding VRF-B
 ip address 10.1.12.1 255.255.255.0
 frame-relay interface-dlci 201
 no keepalive

ip vrf VRF-B
 rd 222:222

interface Serial0/0
 encapsulation frame-relay

Interface Serial0/0.12
 ip vrf forwarding VRF-B
 ip address 10.1.12.2 255.255.255.0
 frame-relay interface-dlci 201
 no keepalive
```

Similar to what you already know!

```
router bgp 1
 address-family ipv4 vrf VRF-B
  neighbour 10.1.12.2 remote-as 2
  neighbour 10.1.12.2 activate
  no synchronisation
  network 1.1.1.1 mask 255.255.255.255
 exit-address-family
```

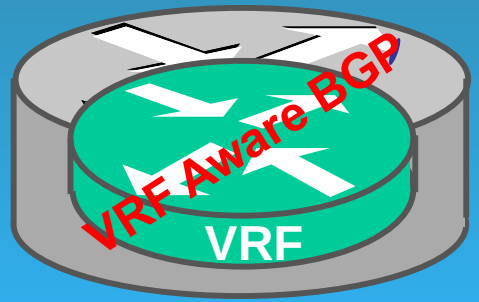


```
router bgp 2
 address-family ipv4 vrf VRF-B
  neighbour 10.1.12.1 remote-as 1
  neighbour 10.1.12.1 activate
  no synchronisation
  network 2.2.2.2 mask 255.255.255.255
 exit-address-family
```


Live Exploration

VRF-Lite

Summary

- Leverages VRF in router (RIB/FIB, interface) and interface for segmentation
- No MPLS, LDP, or BGP required
- Optimal solution when VRF count is small (~ <8)
- Scale usually dependent on routing protocol
- Supports multicast and QoS solutions
- If you understand routing protocols then you already understand PE-CE VRF Routing

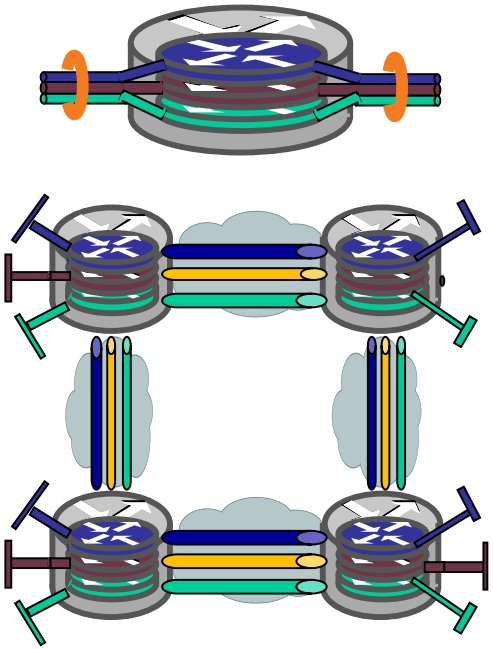

MPLS – BGP Free Core

Cisco live!

BGP

No BGP in Core

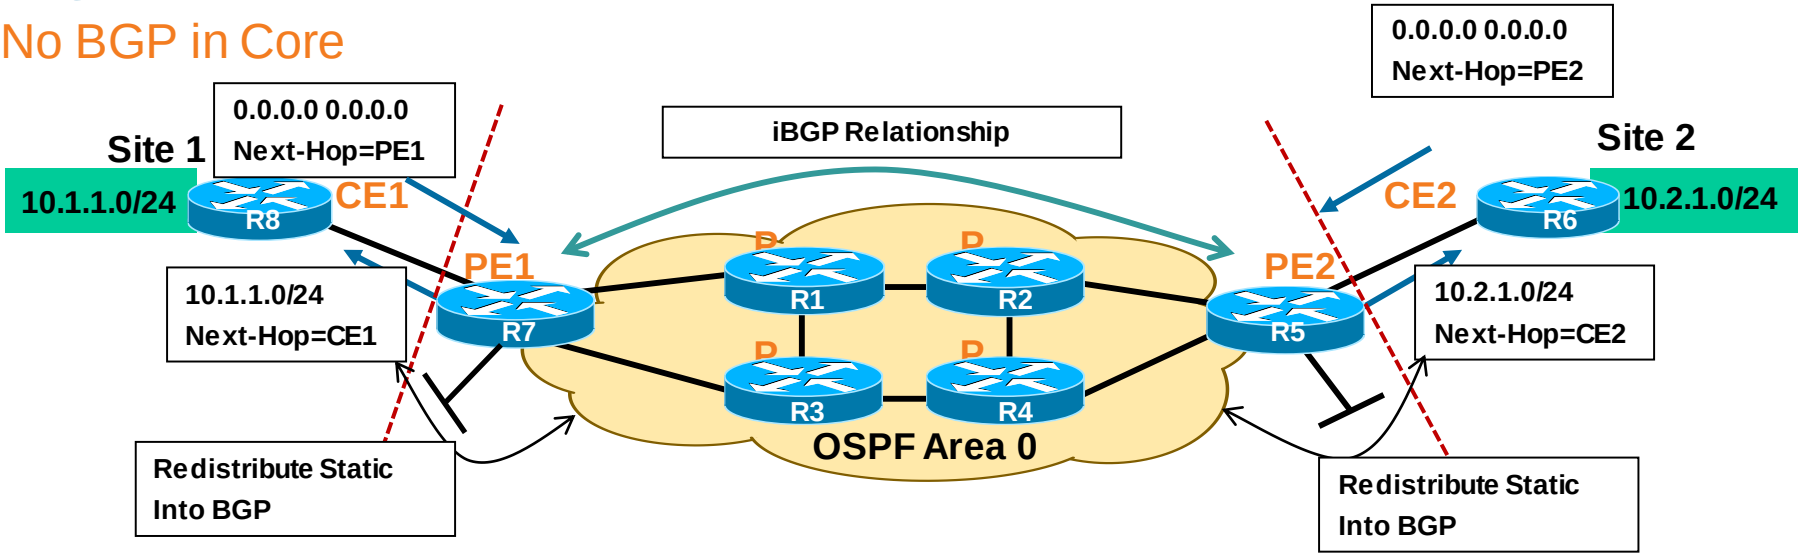

1. Always route towards BGP Next-Hop
2. Routes will be valid on PE Routers
3. Routing will break when reaches the P Routers as they are not participating in BGP

Could run BGP all the way through or redistribute routes into OSPF, but why!

IP Routing

IGP vs. BGP

- Exchange of IP routes for Loopback Reachability
  - OSPF, IS-IS, EIGRP, etc.
- iBGP neighbour peering over IGP transport
- Route towards BGP Next-Hop

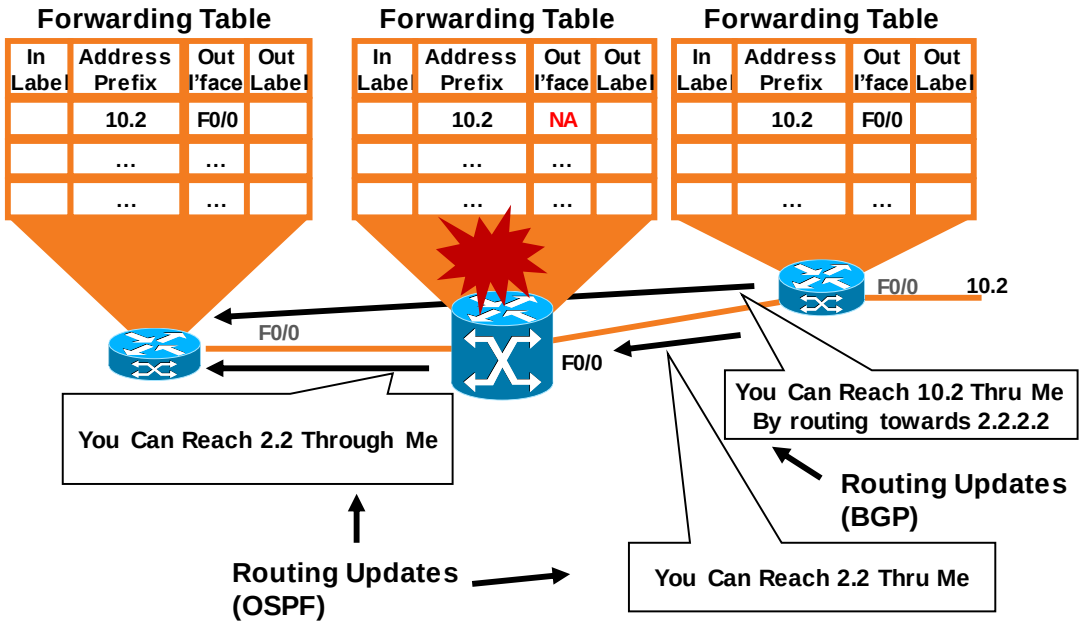

What Is MPLS?

Multi **P**rotocol	Multi-Protocol: The ability to carry any payload Have: IPv4, IPv6, Ethernet, ATM, FR
Label	Uses Labels to tell a node what to do with a packet; separates forwarding (hop by hop behaviour) from routing (control plane)
Switching	Routing == IPv4 or IPv6 lookup. Everything else is Switching.

MPLS Path (LSP) Setup with LDP

Assignment of Remote Labels

- Local label mapping are sent to connected nodes
- Receiving nodes update forwarding table
  - Out label
- LDP label advertisement happens in parallel (downstream unsolicited)

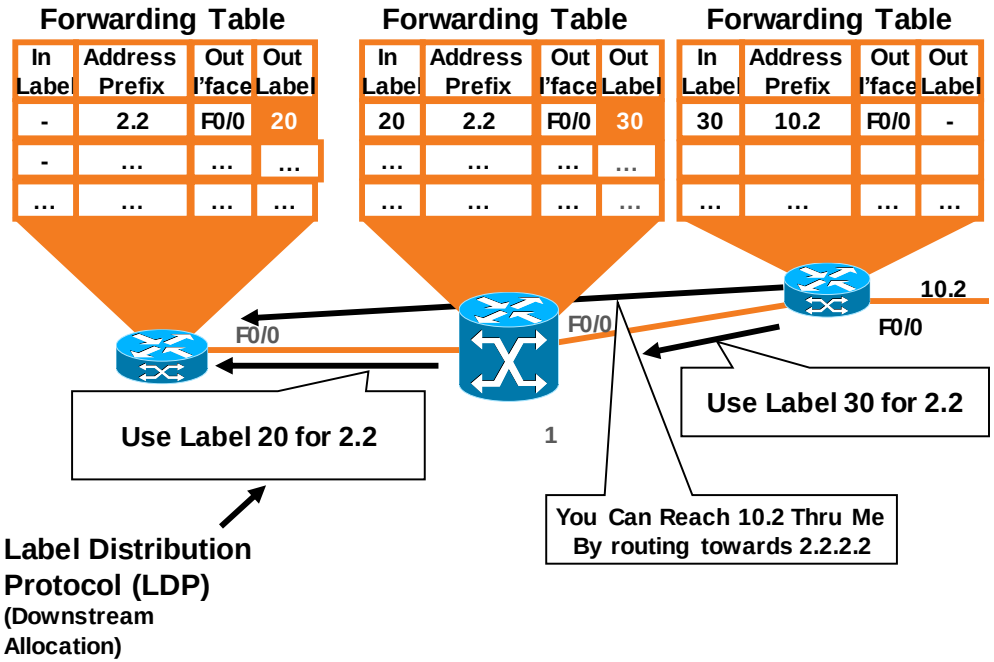

MPLS Traffic Forwarding with LDP

Hop-by-hop Traffic Forwarding Using Labels

- Ingress PE node adds label to packet (push)
  - Via forwarding table
- Downstream node use label for forwarding decision (swap)
  - Outgoing interface
  - Out label
- Egress PE removes label and forwards original packet (pop)

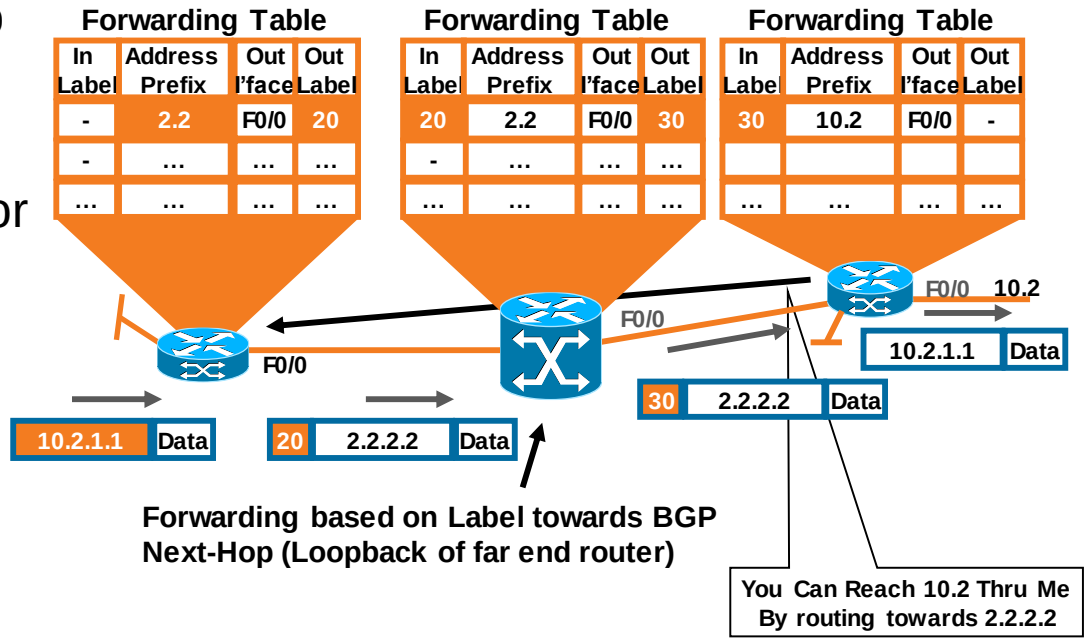

BGP

MPLS in Core

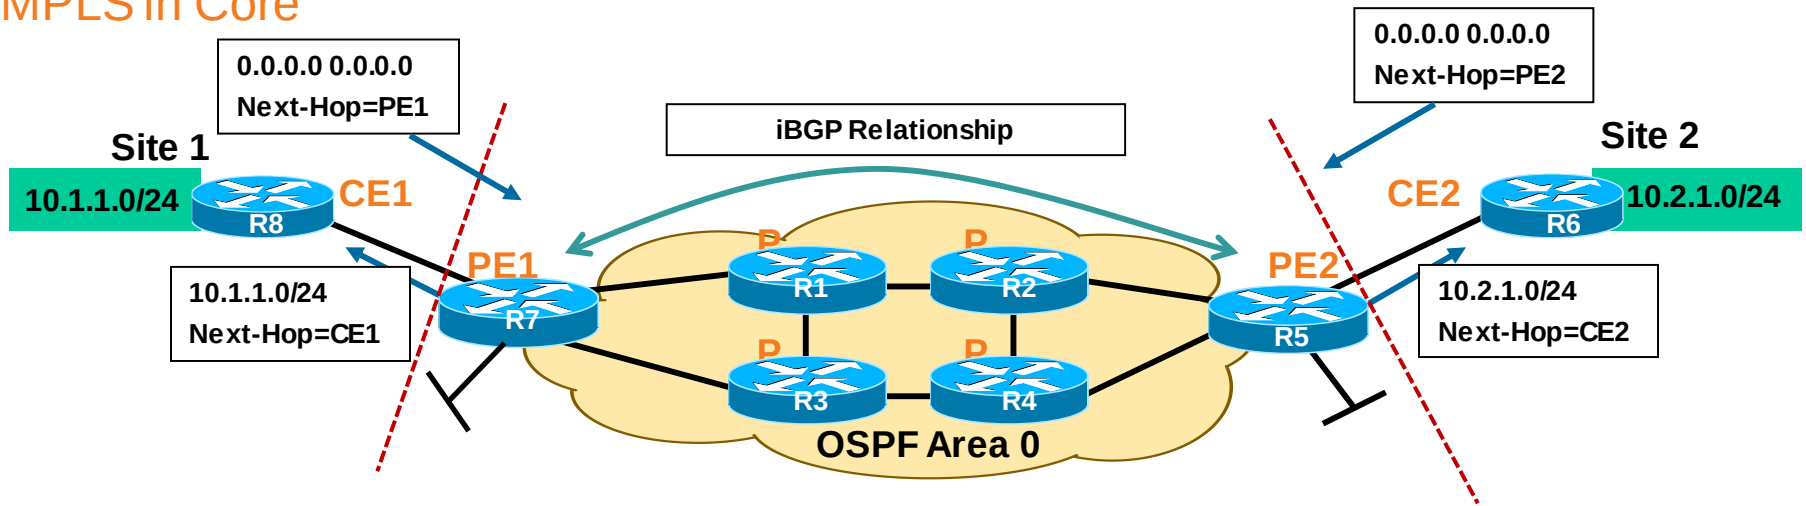

1. Always route towards BGP Next-Hop
2. Routes will be valid on PE Routers
3. Will label switch towards BGP Next-Hop with MPLS enabled

End-to-End BGP and redistribution of routes into OSPF not necessary!


Live Exploration

MP-BGP

Cisco live!

MPLS VPN Technology - Refresher

MPLS VPN Connection Model

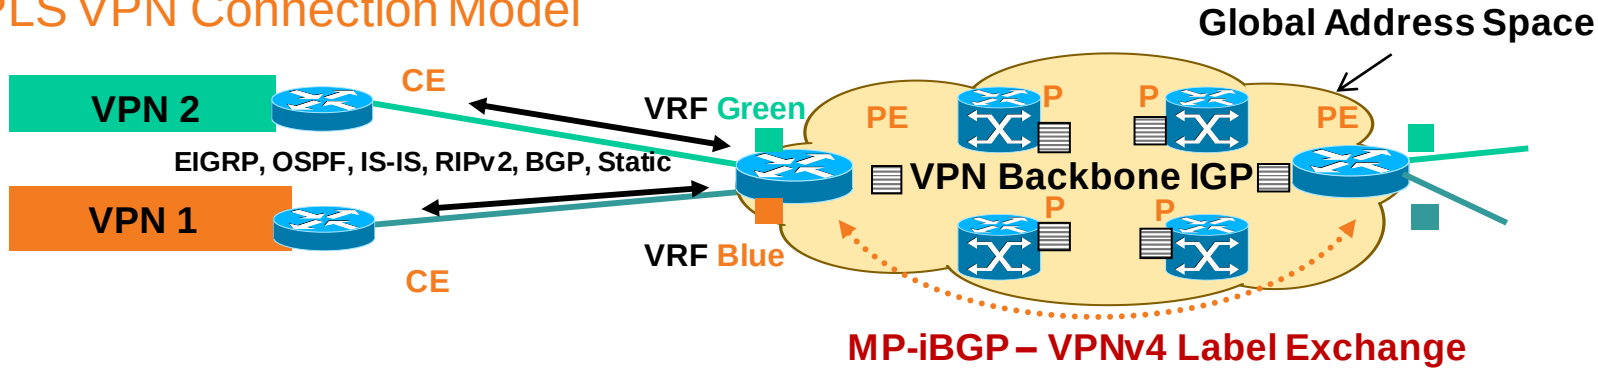

CE Routers

- VRF Associates to one or more interfaces on PE
- Has its own routing table and forwarding table (CEF)
- VRF has its own instance for the routing protocol

  (static, RIP, BGP, EIGRP, OSPF)

PE Routers

- MPLS Edge routers
- MPLS forwarding to P routers
- IGP/BGP – IP to CE routers
- Distributes VPN information through MP-BGP to other PE routers with VPN-IPv4 addresses, extended community, VPN labels

P Routers

- P routers are in the core of the MPLS cloud
- P routers do not need to run BGP
- Do not have knowledge of VPNs
- Switches packets based on labels (push/pop) not IP

Multi Protocol BGP (MPBGP)

Bringing It All Together

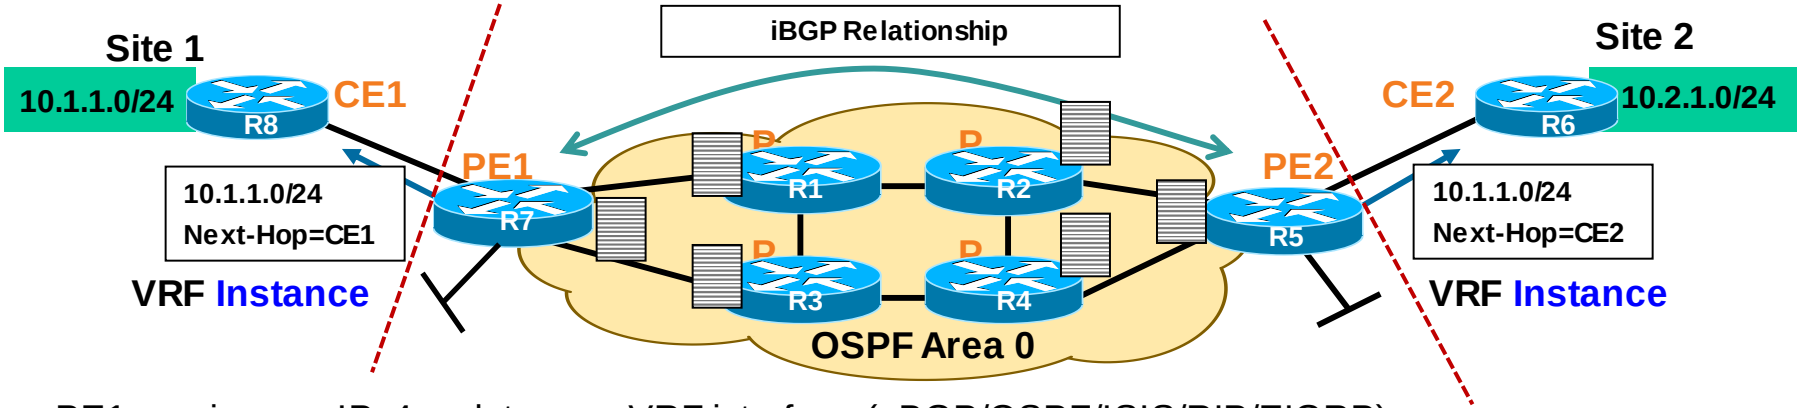

1. PE1 receives an IPv4 update on a VRF interface (eBGP/OSPF/ISIS/RIP/EIGRP)
2. PE1 translates it into VPNv4 address
   - Assigns an RT per VRF configuration
   - Rewrites next-hop attribute to itself
   - Assigns a label based on VRF and/or interface
3. PE1 sends MP-iBGP update to other PE routers

MPLS VPN

Command Line Interface (CLI) Review

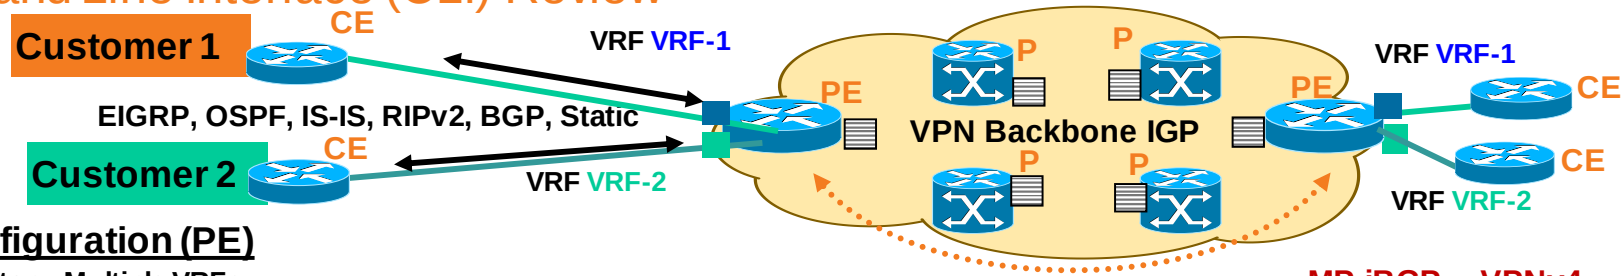

VRF Configuration (PE)

```
! PE Router – Multiple VRFs
ip vrf VRF-1
 rd 65100:10
 route-target import 65102:10
 route-target export 65102:10
ip vrf VRF-2
 rd 65100:20
 route-target import 65102:20
 route-target export 65102:20
!
Interface FastEthernet0/1.10
 ip vrf forwarding VRF-1
Interface FastEthernet0/1.20
 ip vrf forwarding VRF-2
```

MP-iBGP Configuration (PE)

MP-iBGP – VPNv4 Label Exchange

```
! PE router
router bgp 65102
 no bgp default ipv4-unicast
 neighbour 2.2.2.2 remote-as 65102
 !
 address-family vpnv4
 neighbour 2.2.2.2 activate
 neighbour 2.2.2.2 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf VRF-1
 redistribute rip
 exit-address-family
```


Live Exploration

Closing Tips

Personal Insights

- Separate the building blocks of “MPLS” to foster an improved understanding
- Don’t over complicate things
- PE-CE VRF based routing is not much different than regular routing
- MPLS LDP Configuration is pretty simple
- MPBGP and traditional IPv4 BGP configuration is nearly the same
- If routes are not present on CE routers check route-target import/export and redistribution between IPv4 VRF address-families under IGP and BGP
- If routes are present but you are having problems with reachability, check MPLS configuration
- Remember that on PE devices you are living in a VRF world


Q & A
Cisco live!

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2015 T-Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site http://showcase.genie-connect.com/clmelbourne2015
- Visit any Cisco Live Internet Station located throughout the venue

T-Shirts can be collected in the World of Solutions on Friday 20 March 12:00pm - 2:00pm

Learn online with Cisco Live!
Visit us online after the conference for full access to session videos and presentations. www.CiscoLiveAPAC.com

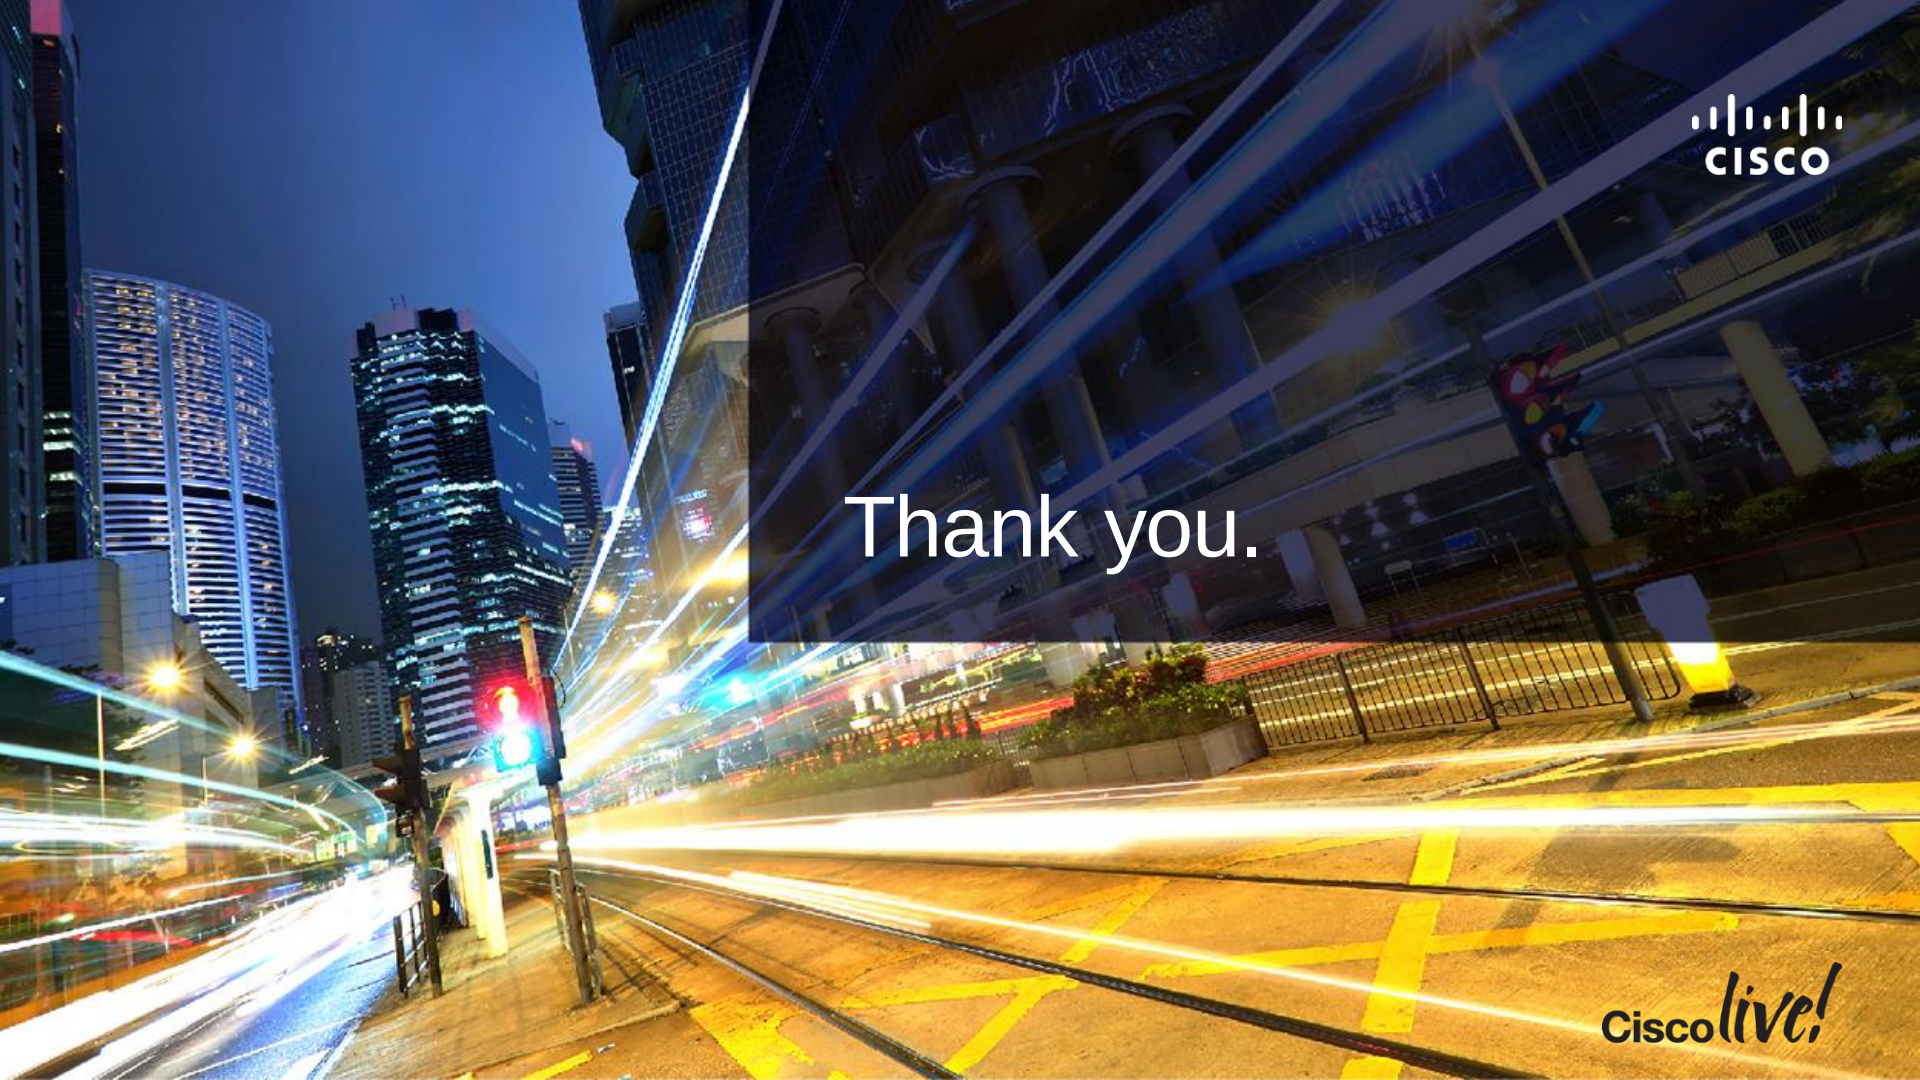
Thank you.
CISCO
Cisco live!


CISCO