
cisco
TOMORROW
starts here.
Cisco live!

L3 VPN over IP Transport, Design and Solutions in the WAN

BRKRST-2045

Craig Hill – Distinguished Systems Engineer

#clmel

Cisco live!

Session Assumptions and Disclaimers

- Participants should have a:
  - Intermediate knowledge of IP routing, IP/GRE tunnels, VRF's, and WAN design fundamentals and technologies
  - Intermediate knowledge of IPSec, DMVPN, GETVPN, MTU considerations
  - Intermediate knowledge of MPLS VPNs operation, MP-BGP, GRE tunnelling, IP QoS
- This discussion will not cover VMware, Virtual Machines, or other server Segmentation technologies
- Data Centre Interconnection (DCI) is an important element in a complete WAN Segmentation infrastructure, but is not a focus in this session nor is Layer 2 Segmentation technologies
- RFC 2547 (BGP/MPLS IP VPNs) is now replaced with RFC 4364.

Agenda

- Introduction - Network Segementation Drivers and Concepts
- WAN Transport Impact on L3 VPN over IP
- Technology Deep-Dive on Advancements in L3 VPN over IP
- QoS, MTU, and Encryption Recommendations
- Recent “Innovations” Evolving in L3 Segmentation
- Summary

Cisco live!

Agenda

- **Introduction - Network Segmentation Drivers and Concepts**
- WAN Transport Impact on L3 VPN over IP
- Technology Deep-Dive on Advancements in L3 VPN over IP
- QoS, MTU, and Encryption Recommendations
- Recent "Innovations" Evolving in L3 Segmentation
- Summary

Cisco live!

Evolution of “Network” Segmentation

…Means Many Things to Many People ☺

- It has evolved a long way from technologies like TDM (1960’s)
- From TDM, ATM/FR Virtual Circuits in the WAN, to…
- VLANs in the Campus, to… Logical/Virtual Routers on routing devices, to…
- Virtual Machines on server clusters in the Data Centre

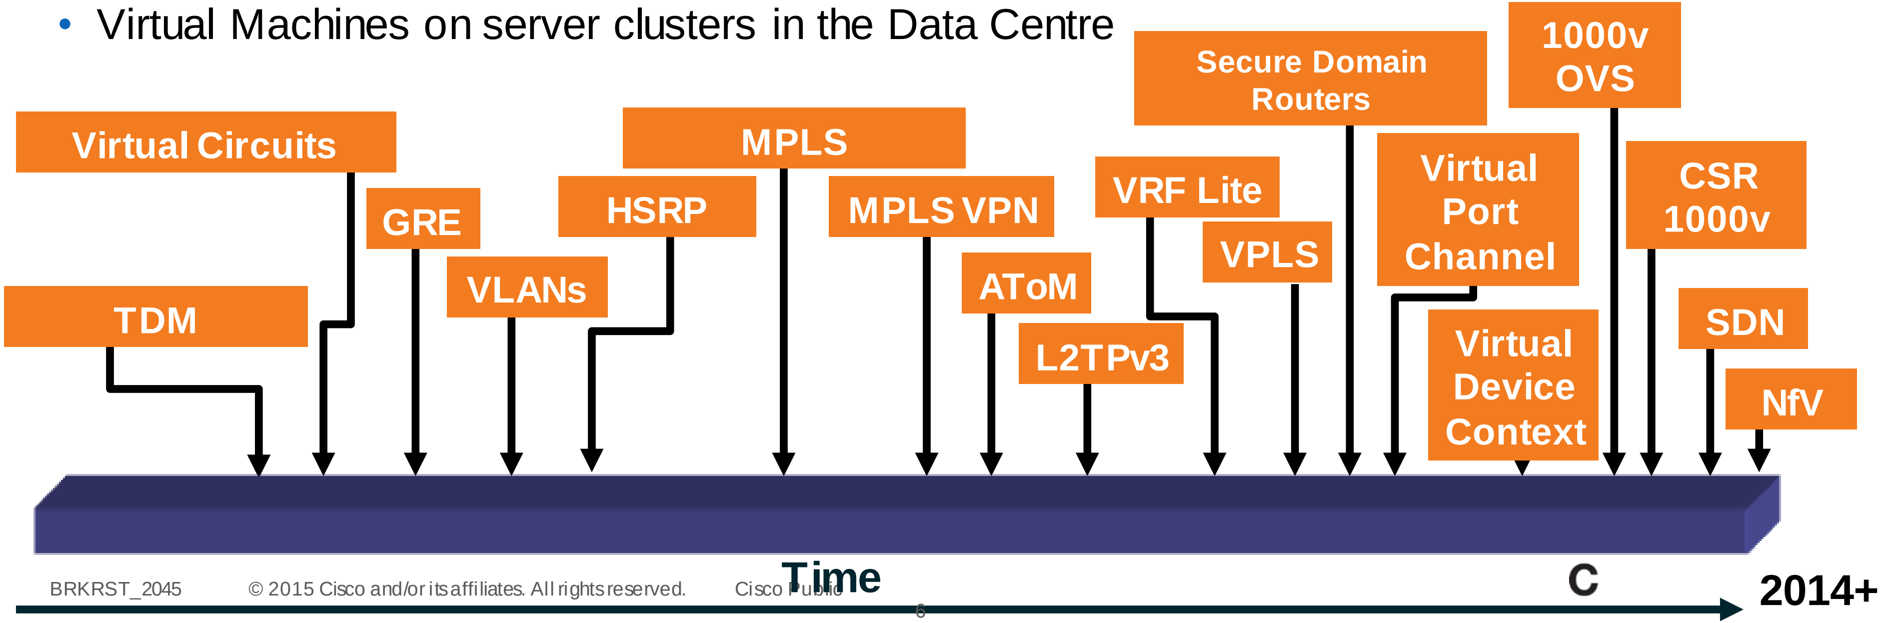

What is Enterprise L3 “Network” Segmentation?

- Giving One physical network the ability to support multiple L3 virtual networks
- End-user perspective is that of being connected to a dedicated network (security, independent set of policies, routing decisions…)
- Maintains Hierarchy, Virtualises devices, data paths, and services

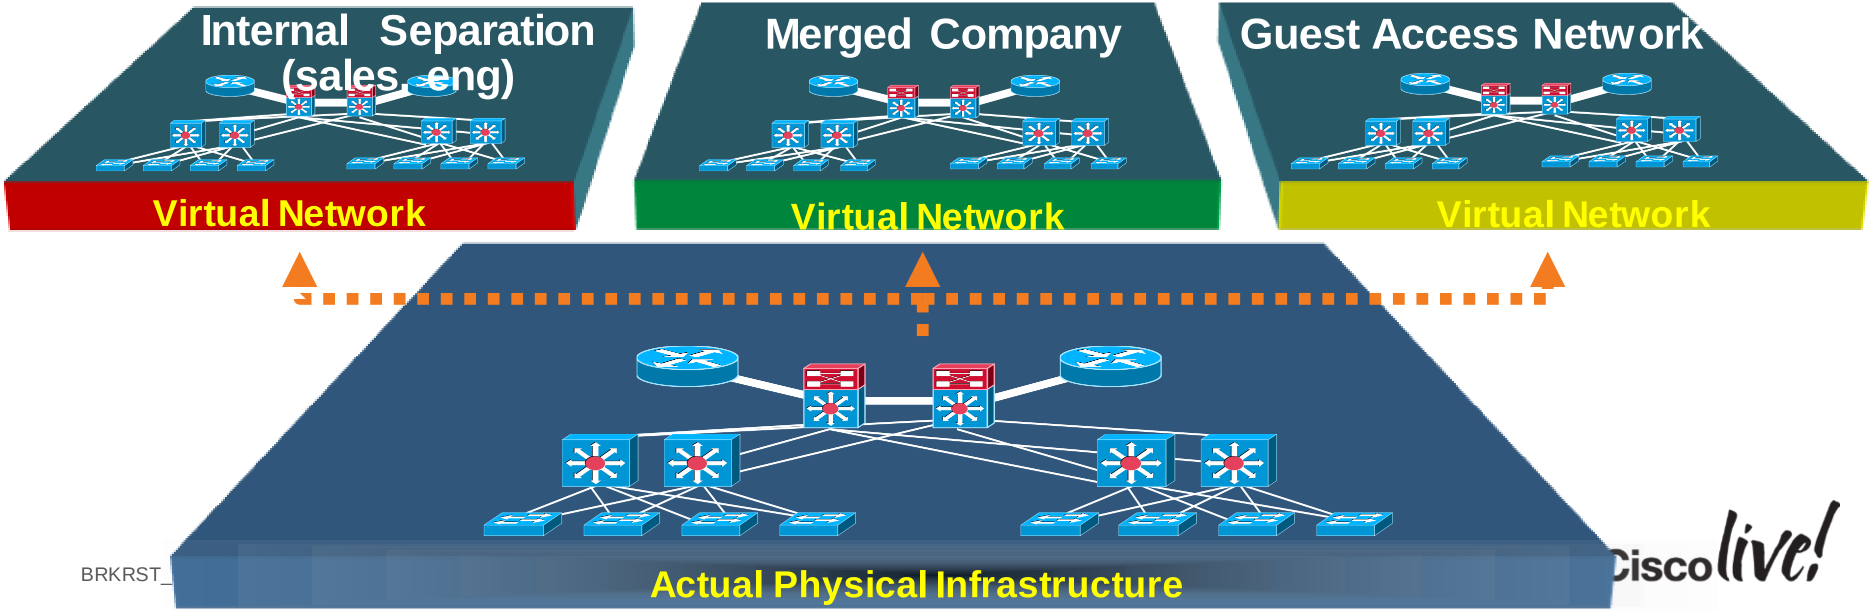

Why L3 Network Segmentation?

Key Drivers and Benefits

- **Cost Reduction**—allowing a single physical network the ability to offer multiple users and virtual networks
- **Simpler OAM**—reducing the amount of network devices needing to be managed and monitored
- **Security**—maintaining segmentation of the network for different departments over a single device/Campus/WAN
- **High Availability**—leverage Segmentation through clustering devices that appear as one (vastly increased uptime)
- **Data Centre Applications**—require maintained separation, end-to-end (i.e. continuity of Segmentation from server-to-campus-to-WAN) , including Multi-tenant DC's for Cloud Computing

L3 Network Segmentation Use Cases

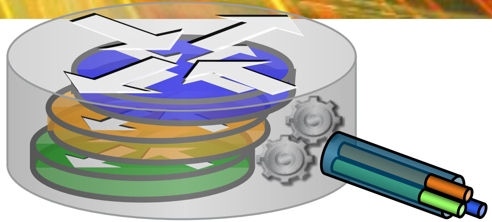

Requirement exists for L3 VPN segmentation within their organisation

- **Multi-Tenant Dwelling requiring Separation**
  - Airports – airlines (United, Delta, etc…) sharing network transport space (physical)
  - Government Facilities – Federal agencies sharing single building/campus
  - Intra Organisation segmentation – Separation of sales, engineering, HR, LoB
  - Company mergers – allowing slow migration for transition, overlapping addressing
  - Data Centre Applications – VM→VLAN→VRF orchestration for segmentation
  - Separation of Facility equipment (IP cameras, badge readers) from the user data
- **Security**
  - Mandates to logically separate varying levels of security enclaves
- **Regulation requirements**
  - Health Care – HIPPA | Financial and Transactional – Sarbanes-Oxley, PCI Compliance
- **Cloud Computing and WAN Orchestration**
  - L3 segmentation (VRF's) are configured dynamically, or part of the automation process, in multi-tenant cloud environments

Enterprise Network Segmentation over the WAN

The Building Blocks – Example Technologies

Device Partitioning

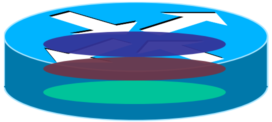

VLANs

VRFs

EVN

(Easy Virtual Network)

VDC (NX-OS)

(Virtual Device Context)

SDR (IOS-XR)

(Secure Domain Routers)

FW Contexts

Device Pooling

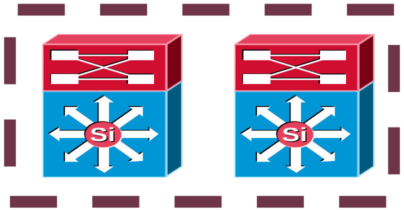

Virtual Sw System (VSS)

Virtual Port Channel (vPC)

HSRP/GLBP

Stackwise

ASR 9000v/nV Clustering

Inter-Chassis Control Protocol (ICCP)

Enterprise Network Segmentation over the WAN

The Building Blocks – Example Technologies

Device Partitioning

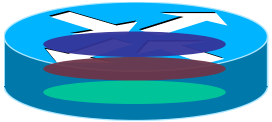

VLANs

VRFs

EVN

(Easy Virtual Network)

VDC (NX-OS)

(Virtual Device Context)

SDR (IOS-XR)

(Secure Domain Routers)

FW Contexts

WAN Interconnect

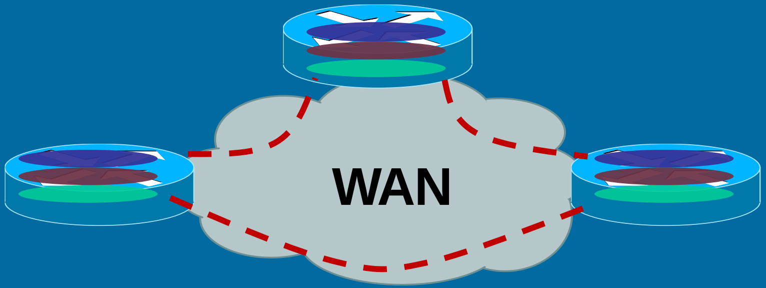

L2 VPNs –PWE3, VPLS, L2 PW over GRE, L2TPv3, OTV (Overlay Transport Segmentation)

Evolving Standards – PBB/E-VPN, VxLAN, NVGRE

L3 VPNs – VRF-Lite, VRF-Lite over GRE, MPLS BGP VPNs, **MPLS BGP VPNs over GRE/mGRE, LISP Multi-tenant**

Device Pooling

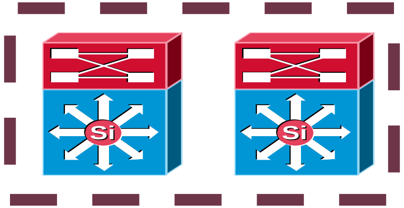

Virtual Sw System (VSS)

Virtual Port Channel (vPC)

HSRP/GLBP

Stackwise

ASR 9000v/nV Clustering

Inter-Chassis Control Protocol (ICCP)

Defining the Virtual Route Forwarding (VRF) Instance

Components, Functions, Uses…

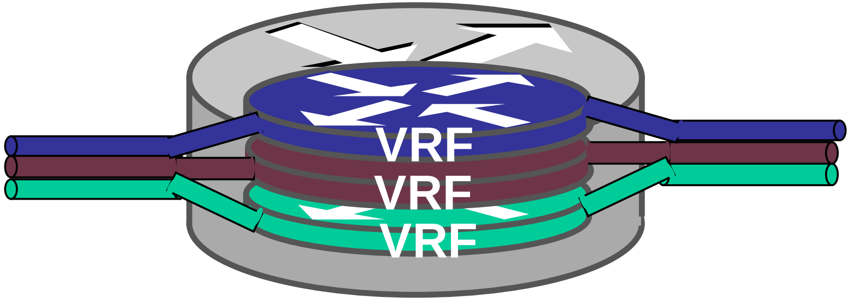

- Associates to one or more interfaces on router (typically a PE)
  Privatise an interface, i.e. colour the interface
- VRF has its own routing table (RIB) and forwarding table (FIB)
- VRF has its own instance for the routing protocols
  (static, RIP, BGP, EIGRP, OSPF)
- Level of segmentation allows overlapping address space

Agenda

- Introduction - Network Segmentation Drivers and Concepts
- **WAN Transport Impact on L3 VPN over IP**
- Technology Deep-Dive on Advancements in L3 VPN over IP
- QoS, MTU, and Encryption Recommendations
- Recent “Innovations” Evolving in L3 Segmentation
- Summary

Transport Options for Private IP VPN Services

WAN Segmentation Models

1. Self Deployed MPLS Backbone Supporting BGP VPNs
2. Self deployed MPLS BGP VPNs “over the top” of an SP Offered IP VPN transport

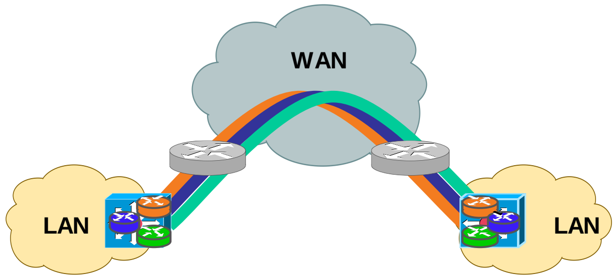

Self Deployed MPLS vs. SP Managed Services

Self Deployed BGP MPLS IP VPN Backbone (RFC 4364)

- Self Deployed offers Service richness and control
- Customer manages and owns:
  - IP routing, provisioning
  - Transport links for PE-P, P-P, PE-CE
  - Full L2, L3 service portfolio
  - SLA's, to "end" customer, QoS
- Customer controls how rapidly services are turned up
- **Allows customer full control E2E**
- **Requires more expertise on the operations team**

BGP MPLS IP VPN Backbone

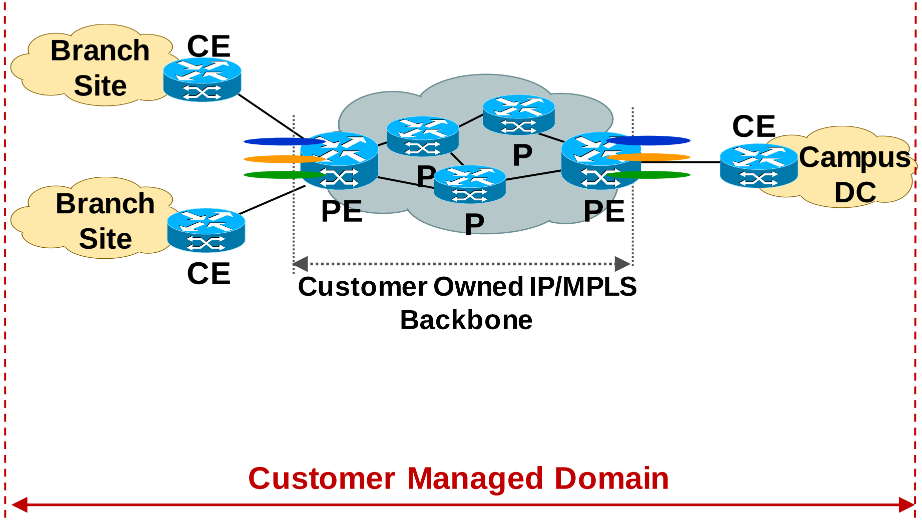

Self Deployed MPLS vs. SP Managed Transport

SP Managed "IP VPN" Transport Services

SP Managed "IP VPN" Service

- **CE Routers owned by customer**
- **PE Routers owned by SP**
- Customer "peers" to "PE" via IP
  - No labels are exchanged with SP PE
  - No end-to-end visibility of other CE's
- Route exchange with SP done via eBGP/static
- Customer relies on SP to advertise their internal routes to all CE's in the VPN for reachability
- SP can offer multiple services: QoS, multicast, IPv6

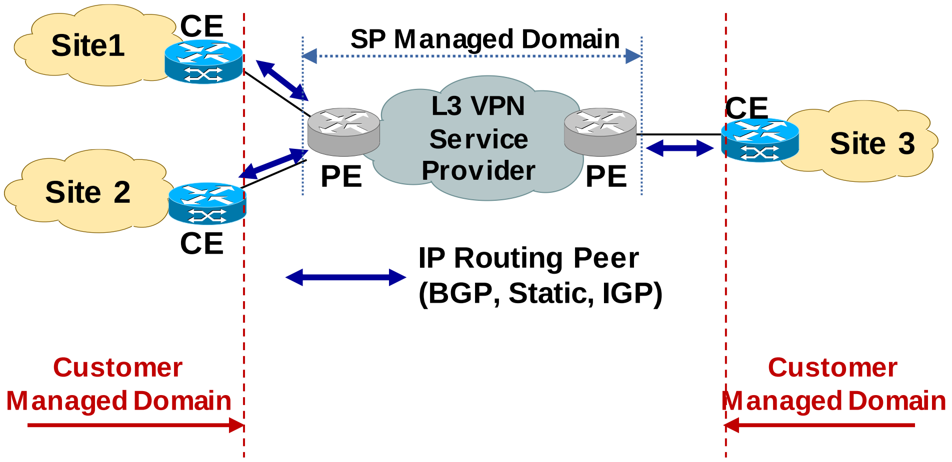

*** No Labels Are Exchanged with the SP**

Self Deployed MPLS vs. SP Managed Transport

MPLS VPN “over the top” of an SP Managed “IP VPN” Transport Service

- **CE customer owned, PE provider owned**
- **Customer enables “PE ” functionality (RFC 4364) on the CE (transparent to SP)**
  - Customer Routing done “Over the Top” of the SP transport
  - Customer IP forwarding encapsulated in GRE, so SP only sees GRE packets
  - Because GRE is used, all traffic can leverage IPSec encryption
  - **Solution must:** **scale, be simple to operate, leverage standards, support QoS, IPSec, be transport independent**

SP Managed “IP VPN” Service

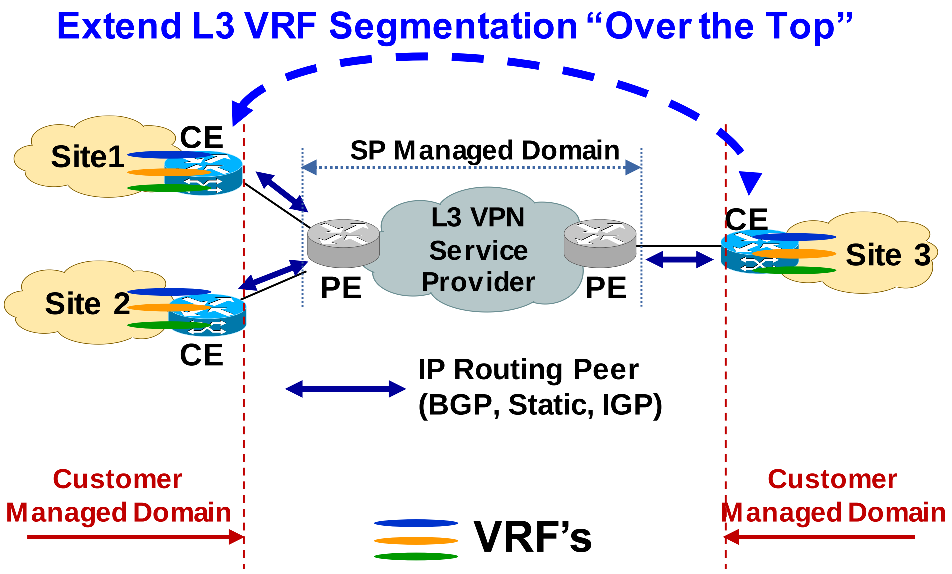

*** No Labels Are Exchanged with the SP**

Key Benefits – Private IP MPLS VPN "Over the Top"

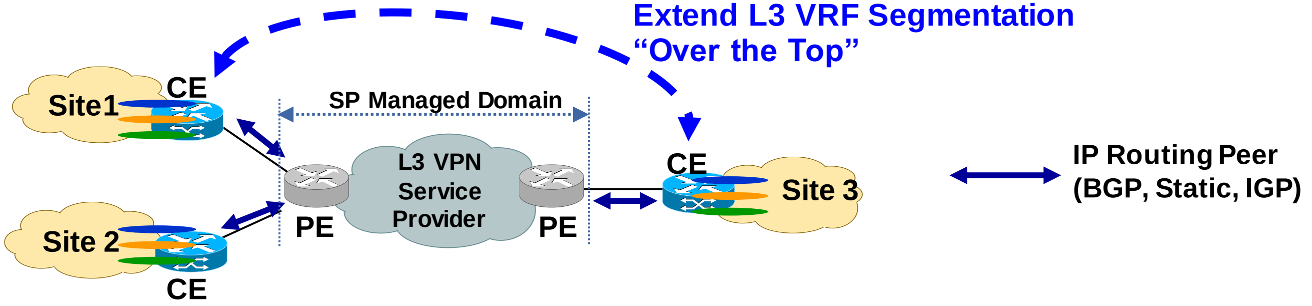

- Allows enterprise to deploy smaller scale MPLS VPN solutions over IP
- VRF changes for end customer goes from days to hours
  - Customer Ex: 30-60 days VRF change in SP | 1 hour VRF change in Private IP VPN Solution
- Can still leverage cost effective L3 transport services from SP
- Can still leverage encryption, QoS, and private BGP AS over the top
- Target Use cases: IPv4 VPN, v6 VPN over v4, align QoS with provider, scale

Agenda

- Introduction - Network Segmentation Drivers and Concepts
- WAN Transport Impact on L3 VPN over IP
- **Technology Deep-Dive on Advancements in L3 VPN over IP**
- QoS, MTU, and Encryption Recommendations
- Recent “Innovations” Evolving in L3 Segmentation
- Summary

Cisco live!

Private IP VPN “Over the Top” Solution Options

Layer 3 VPN Peering – Private IP VPN “Over the Top” Solutions (RFC 4797)

Private VRF Extension Options

Layer 3 IP VPN Transport

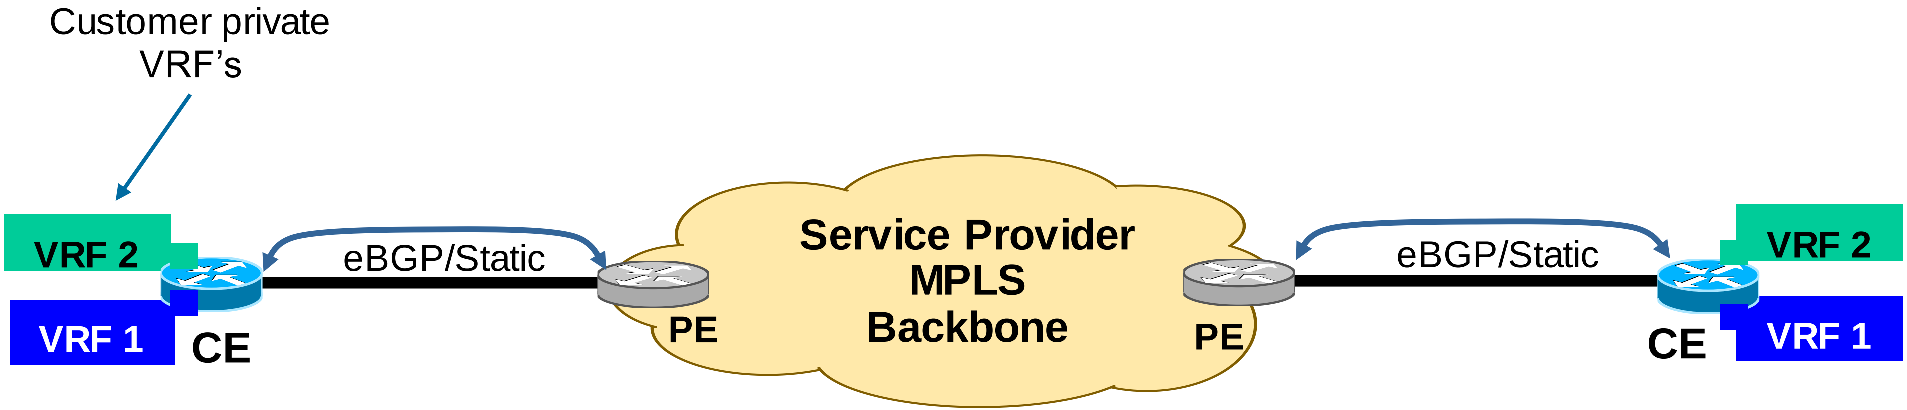

- Back to Back VRF's – VRF-Lite to provider PE
- VRF-Lite over GRE tunnels - CE-to-CE per VRF
- MPLS VPN over IP

MPLS VPN over IP…

Simplifying MPLS VPN over IP Using RFC 4797 Concepts

- Customer may not control the WAN transport Between MPLS networks
  - EXAMPLE: Customers are leveraging IP VPN Service from SP
- Complex to require MPLS label forwarding everywhere in the network
- Customer requires encrypting their PE to PE MPLS traffic
  - No native MPLS encryption exists today
  - Encapsulating MPLS into IP allows use of standard-based IPsec
- Leveraging any IP transport between MPLS PE's is cost effective and simpler

In Summary, the Implementation Strategy Described Enables the Deployment of BGP/MPLS IP VPN Technology in Networks Whose Edge Devices are MPLS and VPN Aware, But Whose Interior Devices Are Not (Source: RFC 4797)

Encapsulation for MPLS in GRE (RFC 4023)

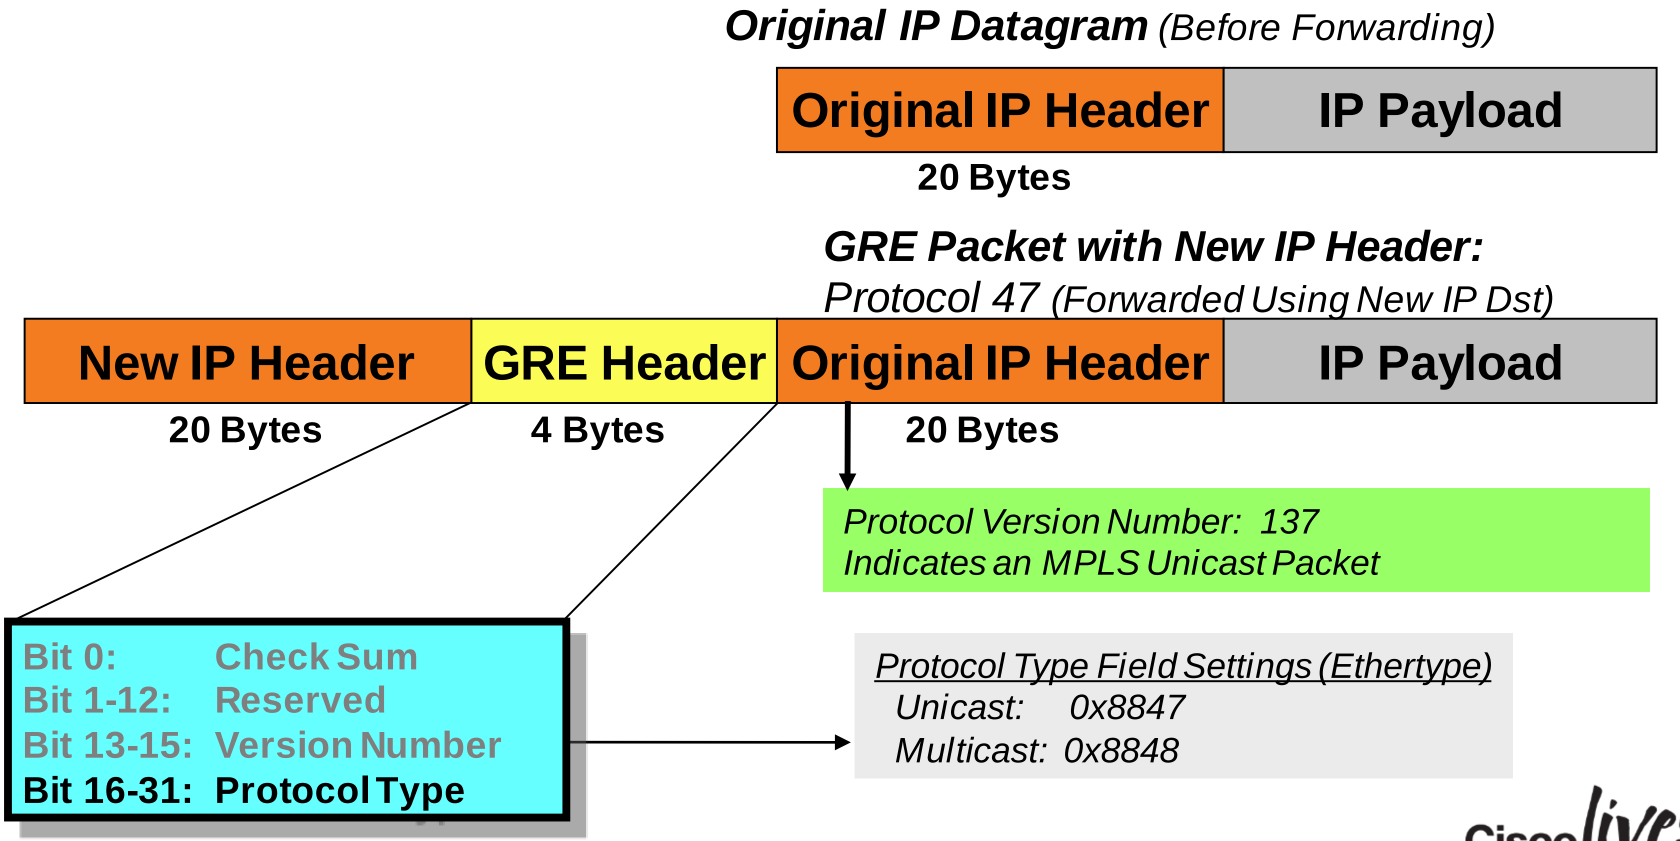

GRE Tunnel Format with MPLS

(Reference: RFC 4023)

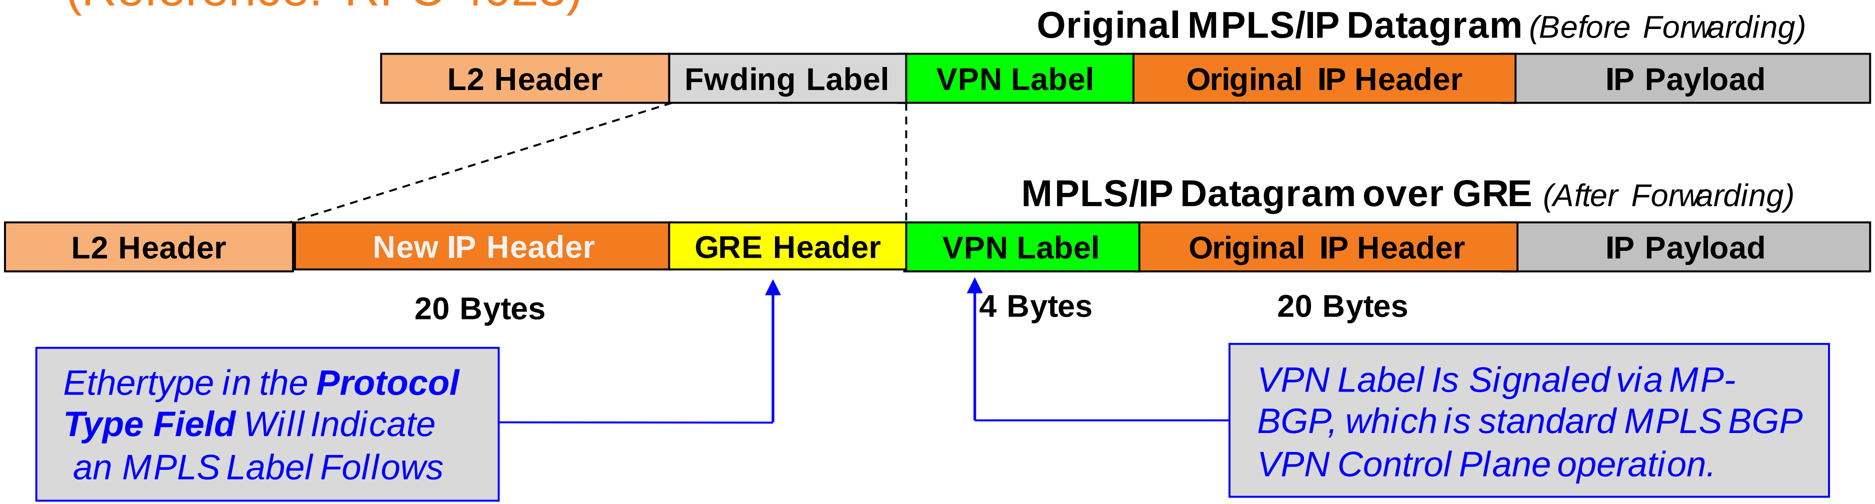

- MPLS Tunnel label (top) is replaced with destination PE’s IP address
- Encapsulation defined in RFC 4023
- Most widely deployed form of MPLS over IP encapsulation

GRE Tunnel Modes

"Stateful" vs. "Stateless" GRE Tunnelling

Point-to-Point GRE

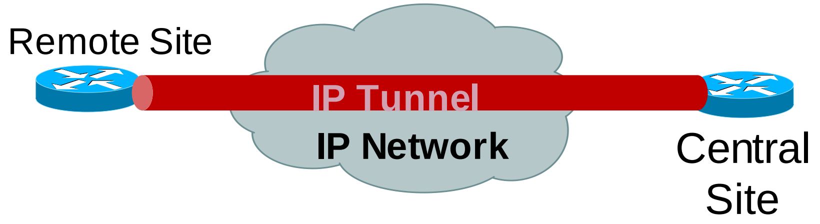

- Source and destination requires manual configuration
- Tunnel end-points are stateful neighbours
- Tunnel destination is explicitly configured
- Creates a logical point-to-point "Tunnel"
- IGP, BGP, and LDP/MPLS run through static tunnel

GRE Tunnel Modes

“Stateful” vs. “Stateless” GRE Tunnelling

<u>Point-to-Point GRE</u>

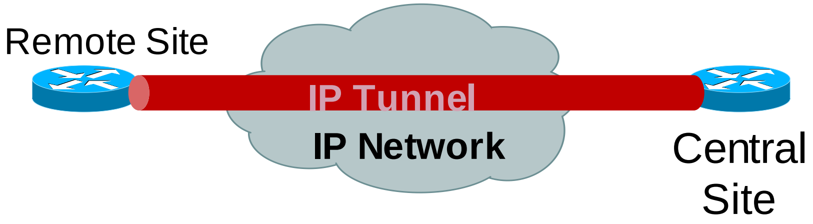

- Source <u>and</u> destination requires manual configuration
- Tunnel end-points are stateful neighbours
- Tunnel destination is explicitly configured
- Creates a logical point-to-point “Tunnel”
- IGP, BGP, and LDP/MPLS run through static tunnel

<u>Multipoint GRE</u>

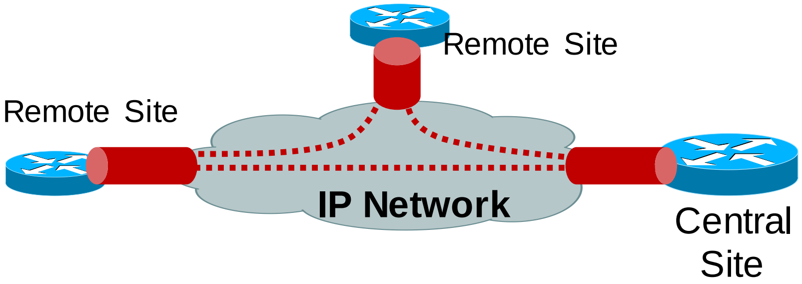

- **<u>Single</u>** multipoint tunnel interface is created per node
- Only the tunnel <u>source</u> is defined
- Tunnel destination is derived dynamically
  - DMVPN – uses NHRP
  - MPLS VPN over mGRE – uses BGP
- Creates an “encapsulation” using IP headers (GRE)

MPLS VPN Technology

Private L3 VPNs "Over the Top"

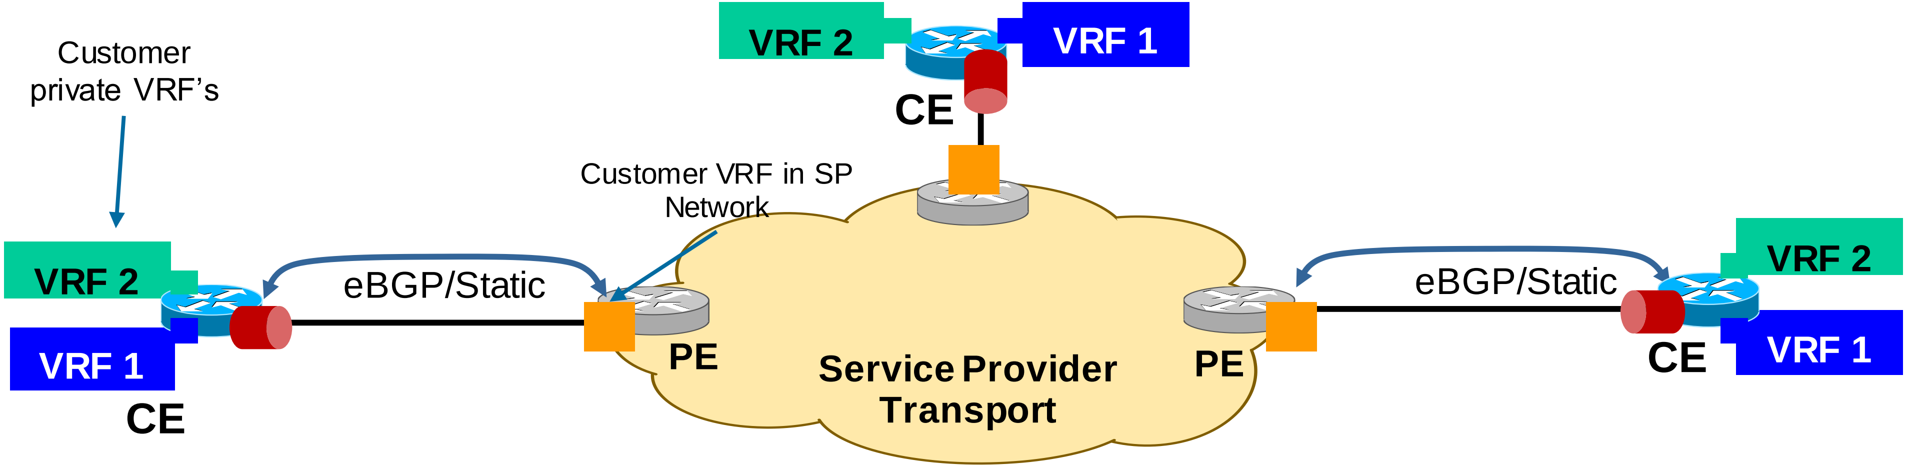

- Basic eBGP/static to peer with SP router

MPLS VPN Technology

Private L3 VPNs "Over the Top"

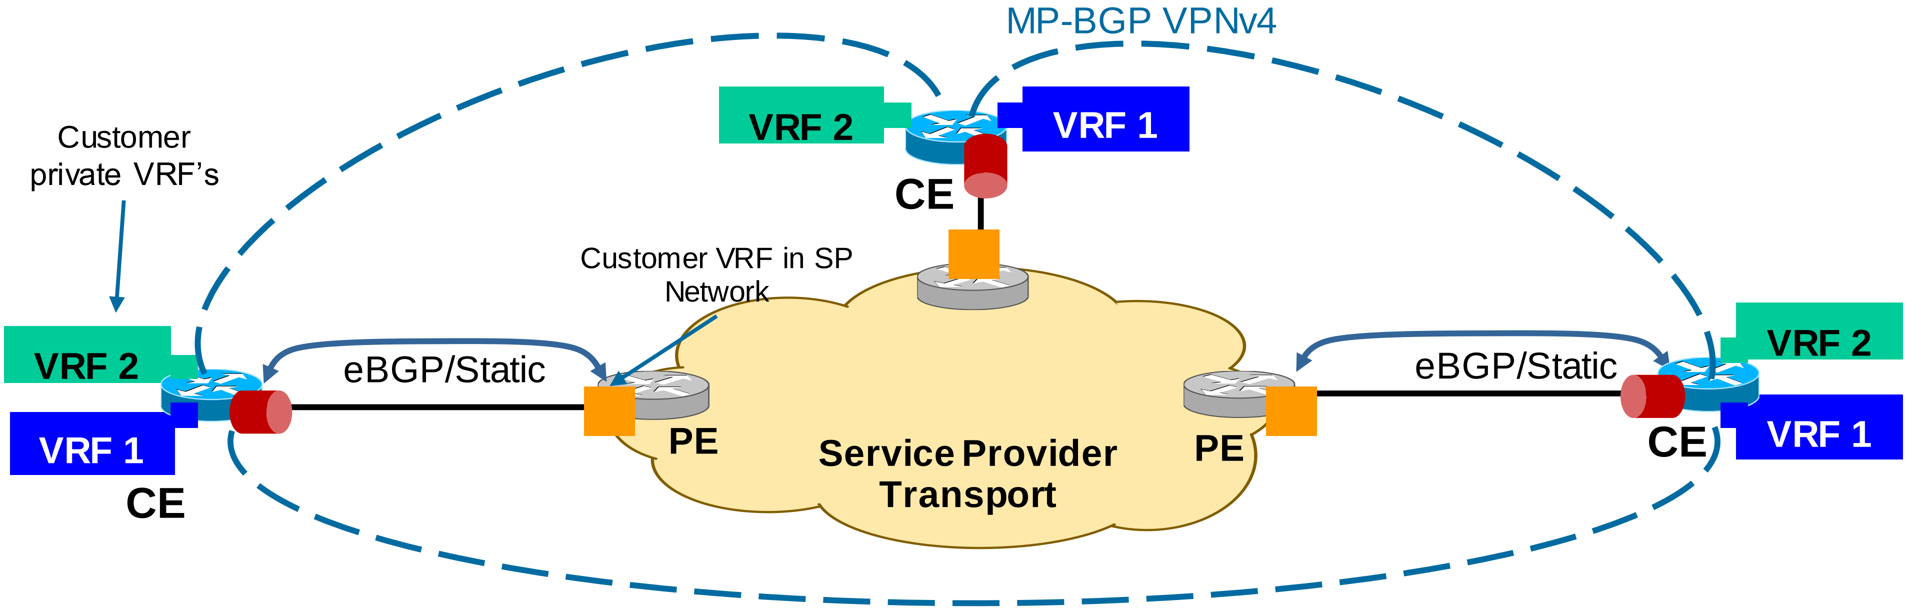

- Basic eBGP/static to peer with SP router
- Run iBGP over the top of the SP between CE routers

MPLS VPN Technology

Private L3 VPNs "Over the Top"

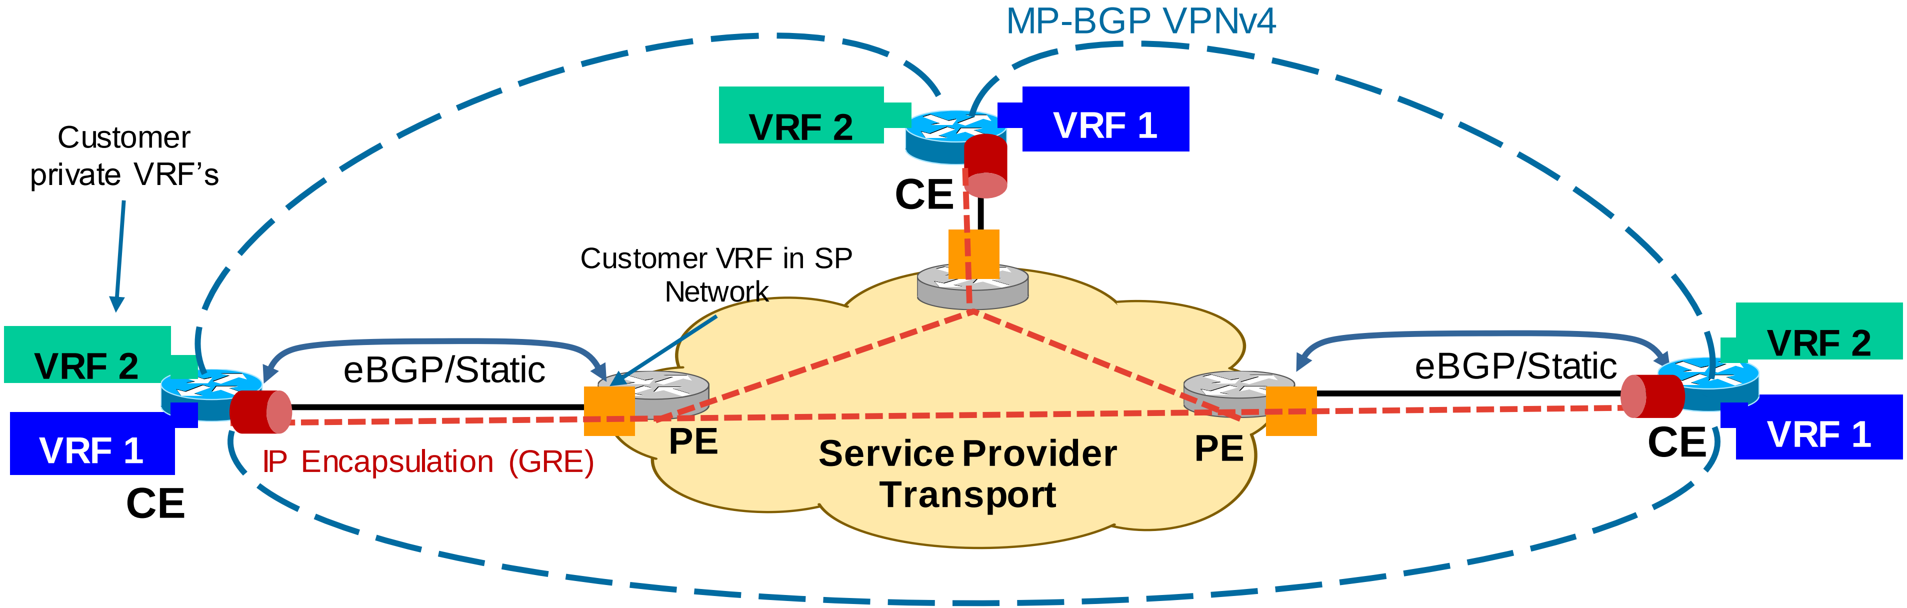

- Basic eBGP/static to peer with SP router
- Run iBGP over the top of the SP between CE routers
- Leverage MPLS VPN over GRE encapsulation for transport
- SP only forwards IP packets (GRE and iBGP) from its data plane view

Advancements in Private IP VPN’s “Over the Top”

Enhancing the L3 VPN Segmentation Portfolio…

- **VRF Lite Options**
  - Leverage Carrier Ethernet E-LINE/E-LAN services
  - Over GRE (GRE tunnel per VRF)
  - Over DMVPN (mGRE interface per VRF)
  - Easy Virtual Networking (EVN) over an E-LINE Carrier Ethernet service
- **L3 MPLS BGP VPN (RFC 4364)**
  - Over L2 transport (PE-PE, P-P, PE-P)… optical, Ethernet, SONET/SDH, etc…
  - Over p2p GRE tunnels
  - Over DMVPN
- **MPLS VPN over Multipoint GRE (mGRE)**

MPLS VPN over Multipoint GRE (mGRE)

MPLS VPNs over Multipoint GRE Using BGP for Dynamic Next-Hop Learning

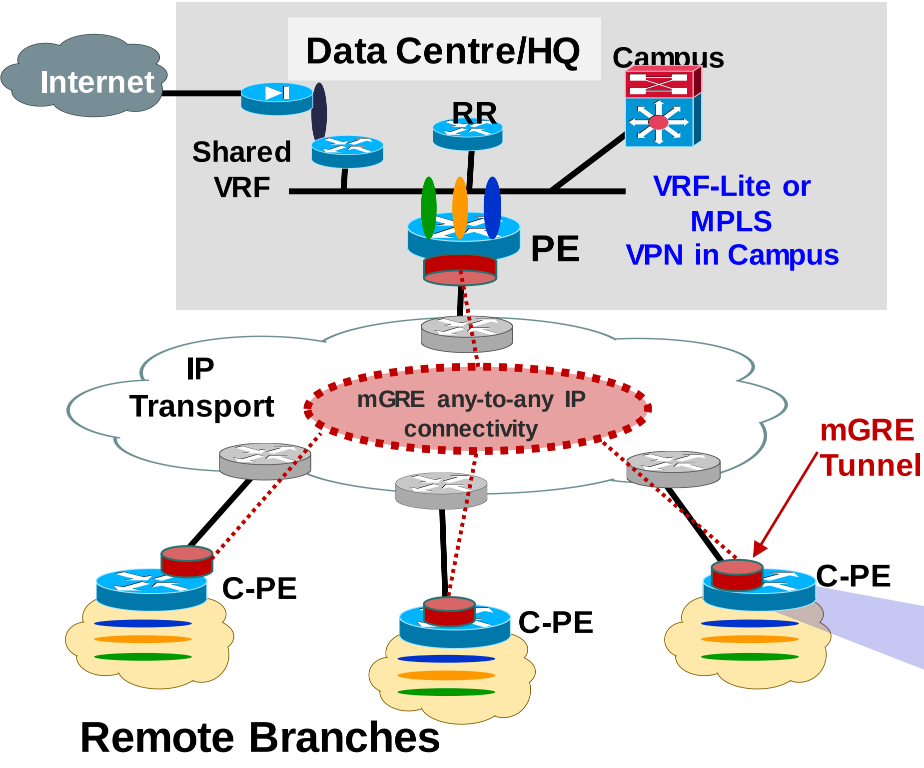

mGRE Interface

- Offers MPLS-VPN over IP
- Inherit spoke-to-spoke communications
- Uses standard RFC 4364 MP-BGP control plane
- Uses standard MPLS over GRE data plane
- Offers dynamic Tunnel Endpoint next-hop via BGP
- Requires only a single IP address for transport over SP network
- Reduces configuration: Requires No LDP, No GRE configuration setup

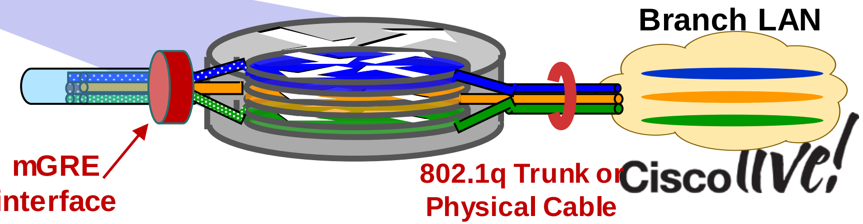

MPLS VPN over Multipoint GRE (mGRE)

Control/Data Plane Example over Service Provider Model

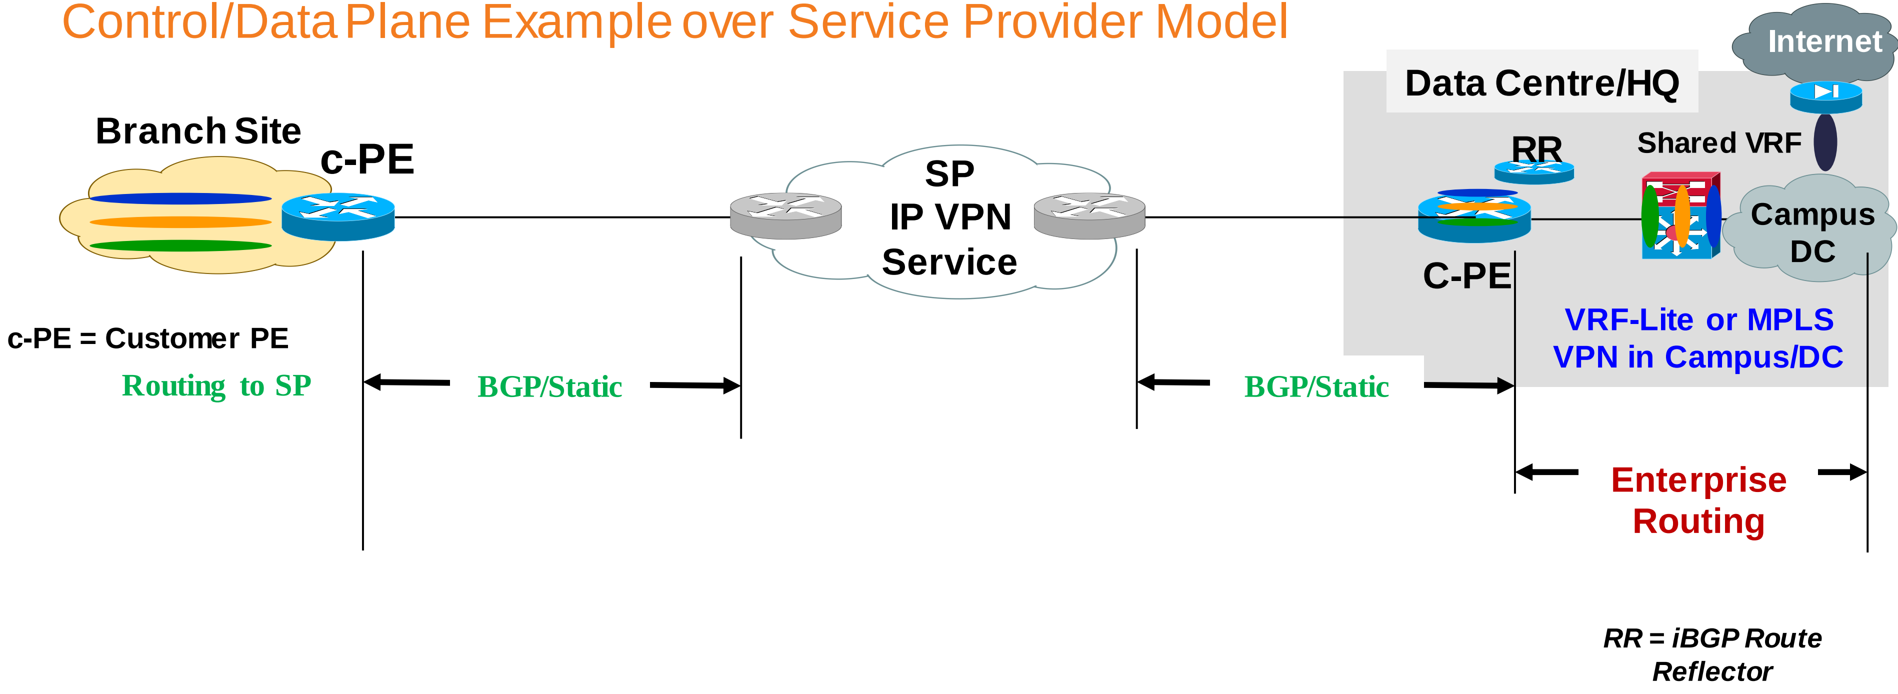

MPLS VPN over Multipoint GRE (mGRE)

Control/Data Plane Example over Service Provider Model

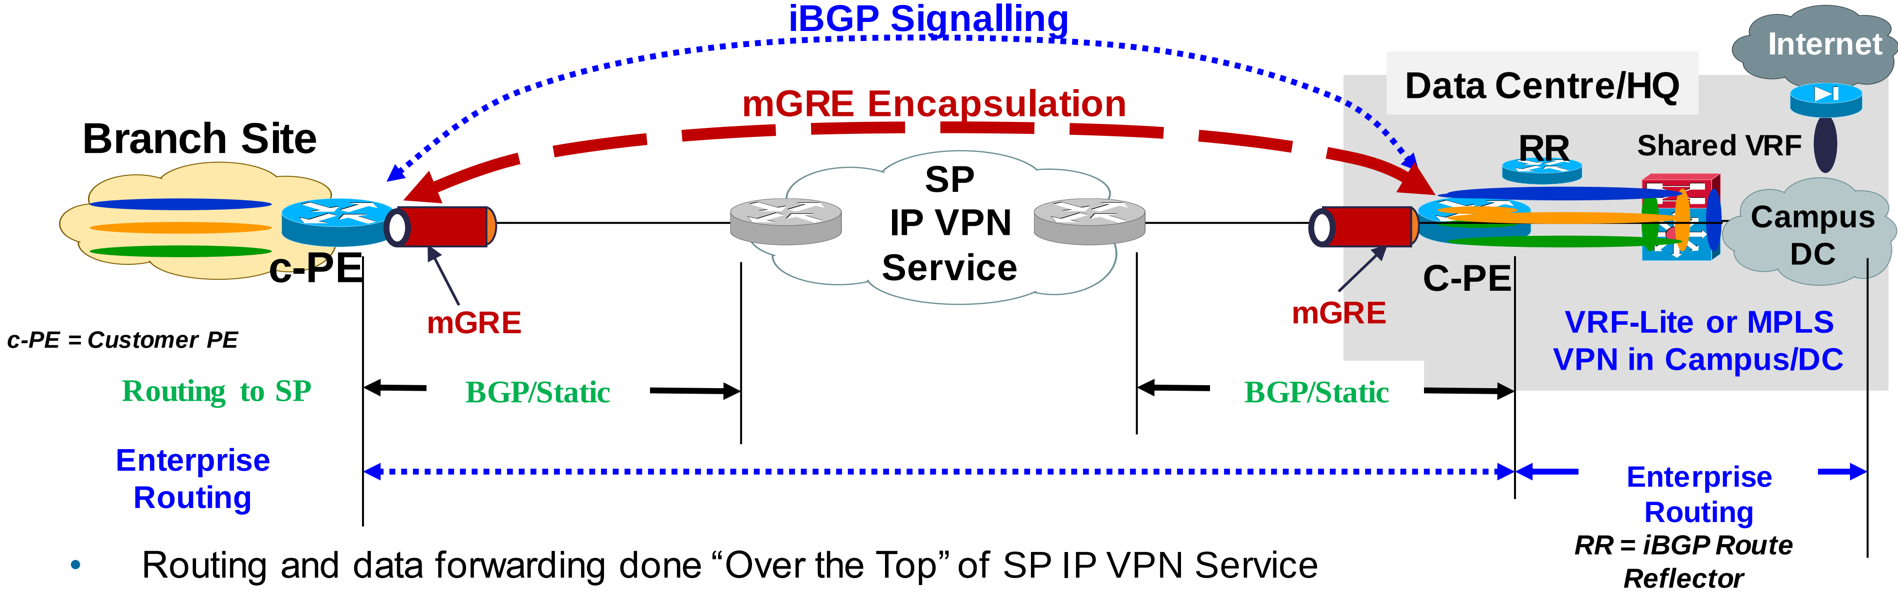

- Routing and data forwarding done "Over the Top" of SP IP VPN Service
- iBGP: (1) Advertise VPNv4 routes, (2) exchange VPN labels
- eBGP: (1) exchange tunnel end point routes with SP (or directly connected)
- Requires advertising a SINGLE IP prefix to SP (e.g. IP tunnel "end points")

MPLS VPN over Multipoint GRE (mGRE)

Control Plane

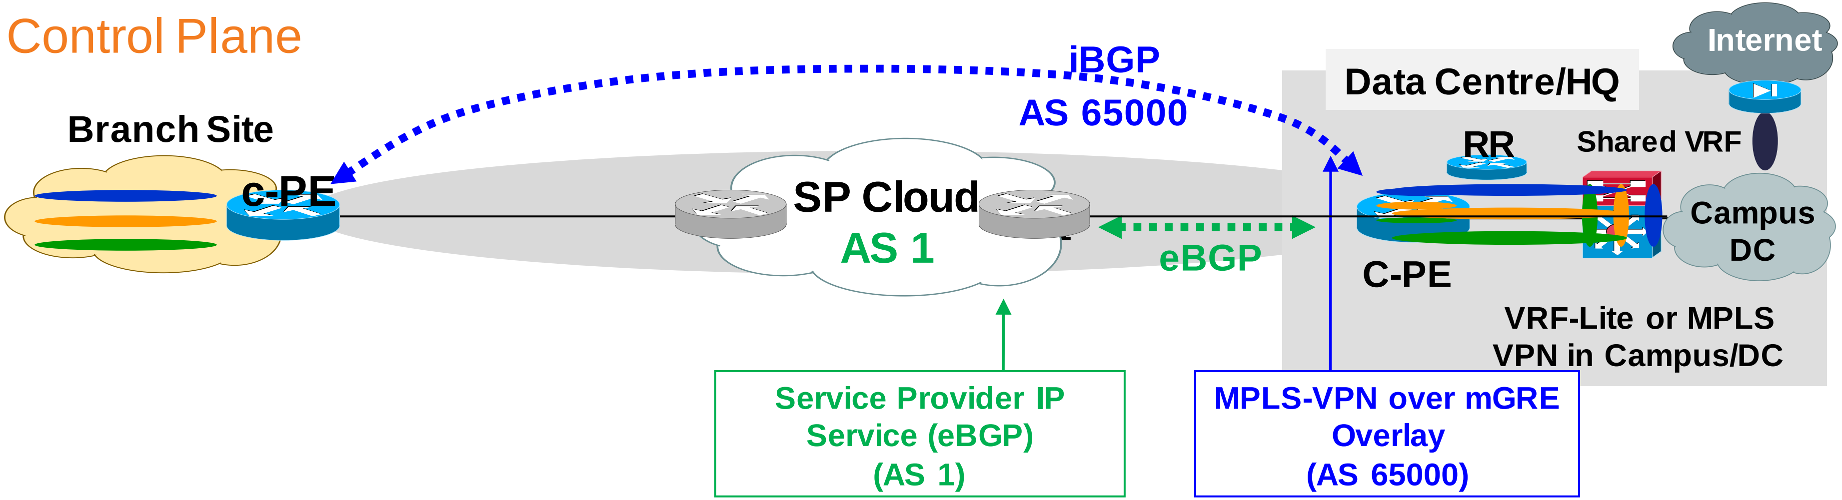

- eBGP (AS 1): used to peer to the SP PE router
- i-BGP (AS 65000): used for MP-BGP and VPNv4 prefix and label exchange
- C-PE for e-BGP appears as CE to the SP
- C-PE for i-BGP functions as a PE in supporting MPLS-VPN over mGRE
- **eBGP used for advertising iBGP next-hop (and mGRE tunnel endpoint) only**

MPLS VPN over mGRE Model

mGRE Interface is Dynamic and De-coupled from Physical Interfaces

- System dynamically configures mGRE tunnel (via tunnel profile)
- mGRE tunnel is decoupled from physical interface
- User traffic is in VRF/VPNv4 of mGRE payload (hidden from provider)
- Only a single IP address (source GRE/BGP-source) advertised to provider

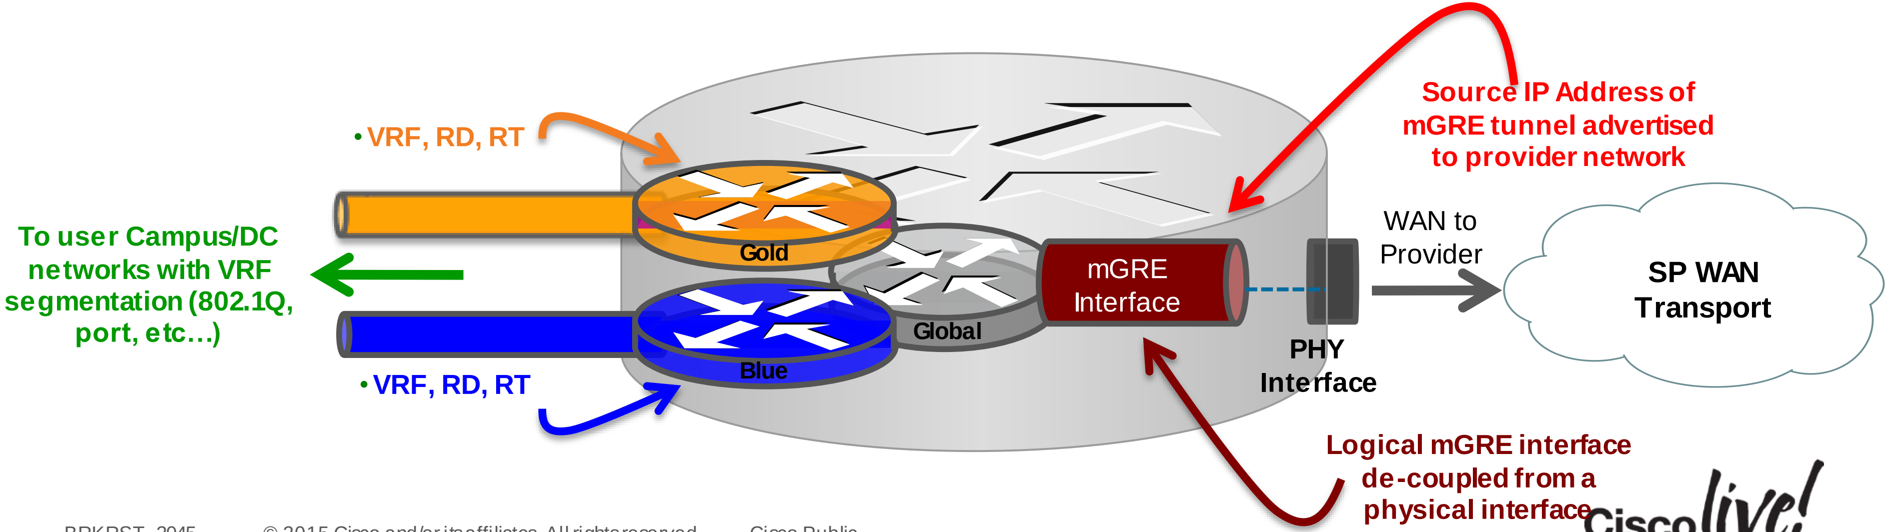

MPLS VPN over Multipoint GRE (mGRE)

Feature Components

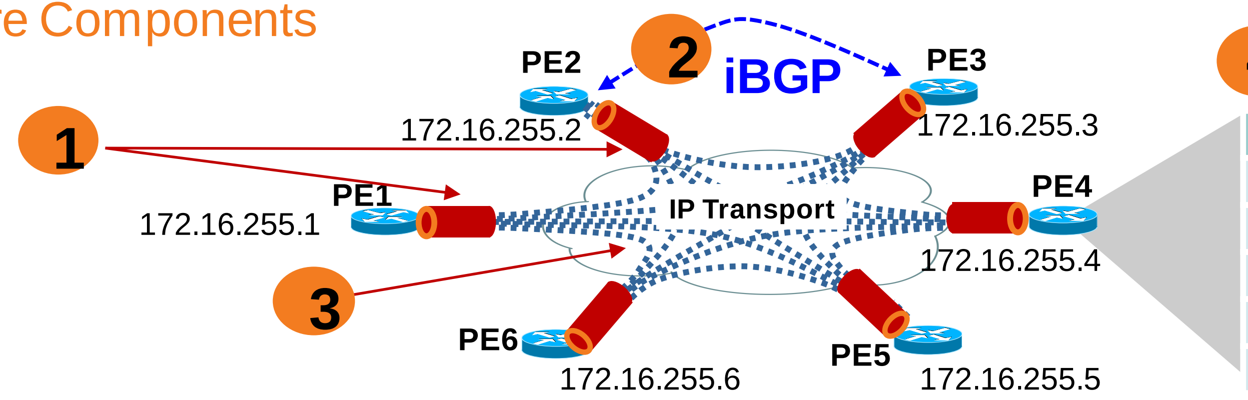

4 **View for PE 4**

Tunnel Endpoint DB
172.16.255.1
172.16.255.2
172.16.255.3
172.16.255.5
172.16.255.6

1. **mGRE** is a multipoint bi-directional GRE tunnel
2. Control Plane leverages RFC 4364 using MP-BGP

   Signalling VPNv4 routes, VPN labels, and building IP next hop (locally)
3. VPNv4 label (VRF) and VPN payload is carried in mGRE tunnel encapsulation
4. New **encapsulation profile** (see next slide) in CLI offers dynamic endpoint discovery:

   (1) Sets IP encapsulation for next-hop

   (2) Installs signaled BGP peer and end-point into "tunnel endpoint database"

Multipoint GRE Interface

MPLS VPN over Multipoint GRE (mGRE)

VPNv4 Configuration Example

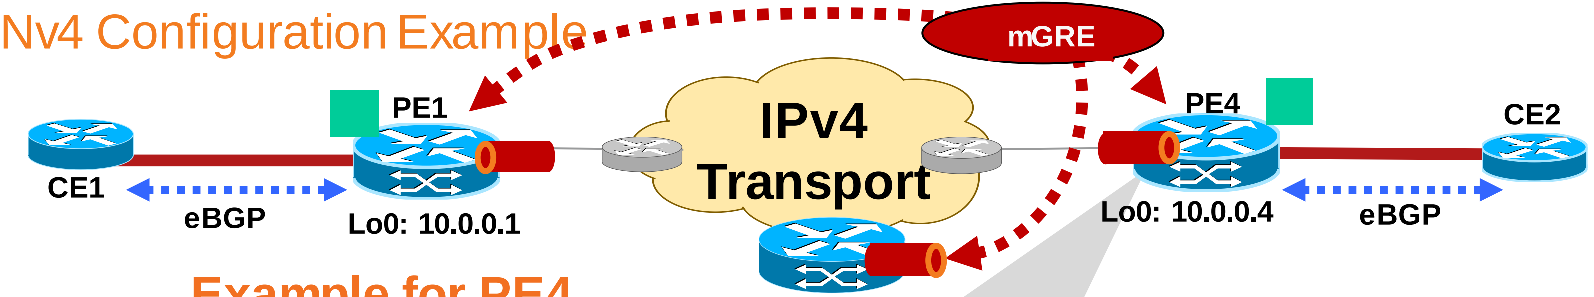

Example for PE4

```
interface Loopback0
 ip address 10.0.0.4 255.255.255.255
!
l3vpn encapsulation ip Cisco
  transport ipv4 source Loopback0

!
router bgp 100
. . .
 address-family vpnv4
  neighbor 10.0.0.1 activate
  neighbor 10.0.0.1 send-community extended
  neighbor 10.0.0.1 route-map next-hop-TED in
 exit-address-family
. . .
!
route-map next-hop-TED permit 10
 set ip next hop encapsulate l3vpn Cisco
```

Sets mGRE Encapsulation "Profile" for BGP Next-Hop

Apply Route-Map to Received Advertisement from Remote iBGP Neighbour

Use IP Encap (GRE) for Next-Hop and Install Prefix in VPN Table as Connected IP Tunnel Interface

MPLS VPN over Multipoint GRE (mGRE)

IPv6 Configuration Example

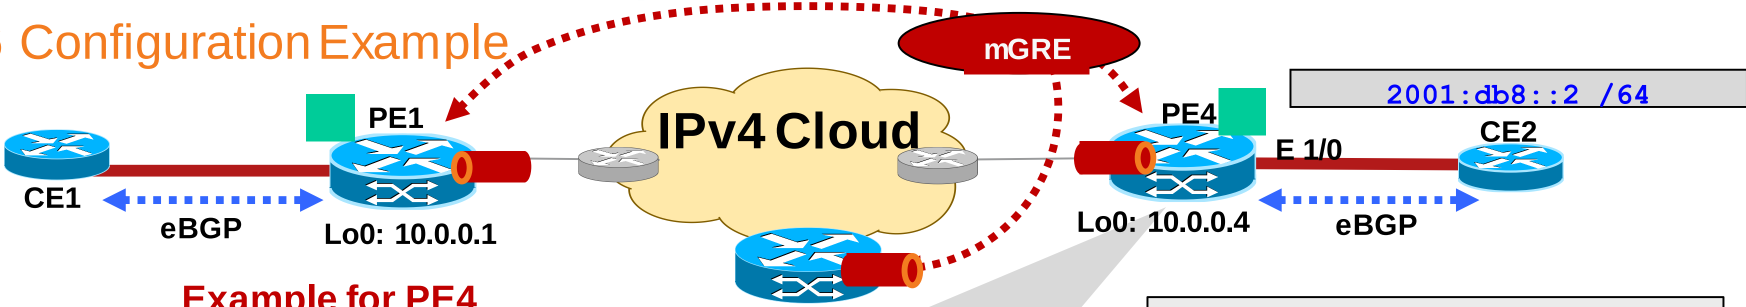

Example for PE4

```
interface Ethernet 1/0
 vrf forwarding green
 ip address 209.165.200.253 255.255.255.224
 ipv6 address 2001:db8:: /64 eui-64
!
router bgp 100
. . .
 address-family vpnv6
  neighbor 10.0.0.1 activate
  neighbor 10.0.0.1 send-community both
  neighbor 10.0.0.1 route-map next-hop-TED in
 exit-address-family
. . .
!
route-map next-hop-TED permit 10
 set ip next-hop encapsulate l3vpn Cisco
 set ipv6 next-hop encapsulate l3vpn Cisco
```

NOTE: Relevant MPLS VPN over mGRE Commands That Are Same for IPv4, Are Not Shown in This IPv6 Example

IPv6 Address Applied to CE2 Facing Interface

Apply Route-Map to Received Advertisement from Remote iBGP Neighbour (Same as vpnv4)

Use IP Encap (GRE) for Next-Hop and Install IPv6 Prefix in VPNv6 Table as Connected Tunnel Interface

MPLS VPN Deployment Considerations for WAN Designs (over IP)

EXAMPLE: MPLS VPN over mGRE (BGP)

Example: 50 – 1000 Sites

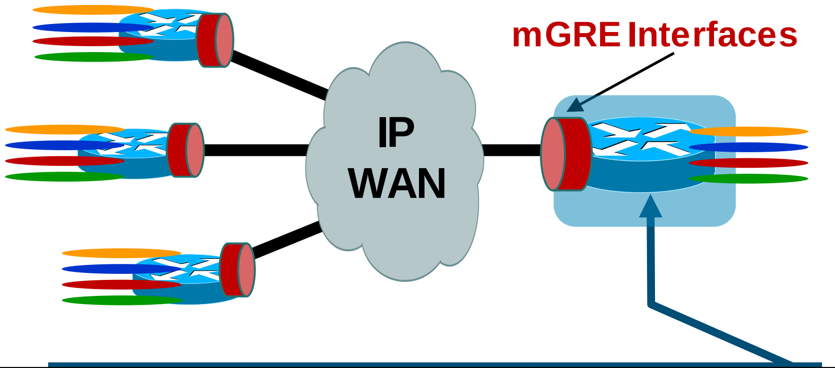

VRFs	Neighbours	GRE Tunnel Interface
50	50	1
100	100	1
250	200	1
500+	1000	1

Key questions to ask yourself:

- How many VRFs will be required at initial deployment? 1 year? 3+ years?
- Are frequent adds/deletes and changes of VRFs required?
- How many locations will the network grow?
- Do I require any-to-any traffic patterns?
- What is the transport? (i.e. is GRE required?)
- Do I have the expertise to manage an MPLS VPN network?

MPLS VPN over Multipoint GRE (mGRE)

Summary and Configuration Notes

- Only requires advertising a single IP prefix to SP for mGRE operation
- Leverages standard MP-BGP control plane (RFC 4364)
- Tunnel endpoint discovery is done via iBGP/route-map
- E-BGP can/is still be used for route exchange (mGRE end-point) with the SP
- Solution requires NO manual configuration of GRE tunnels or LDP
- Supports MVPN and IPv6 per MPLS VPN model (MDT and 6vPE respectfully)
  - MVPN Platform Support today: ISR/G2, SUP-2T (ASR 1000 – FUTURE)
- Supports IPSec for PE-PE encryption (GET VPN or manual SA – Discussed later)
- Platform Support

Today: 7600/12.2(33) SRE, ASR 1000 (3.1.2S), ISR product line (15.1(2)T), 6500/SUP-2T (15.0(1) SY), MWR-2941

Future: IOS-XR Platforms (Future planning)

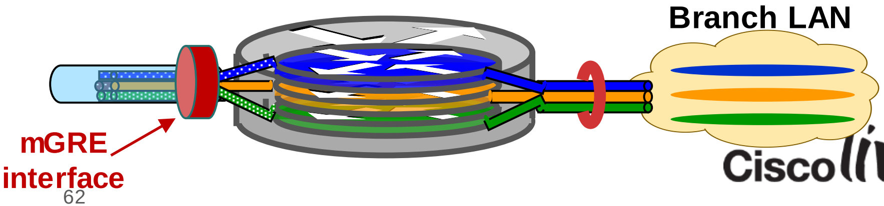

MPLS VPN over mGRE – “Config” and “Show” Examples

MPLS VPN over Multipoint GRE (mGRE)

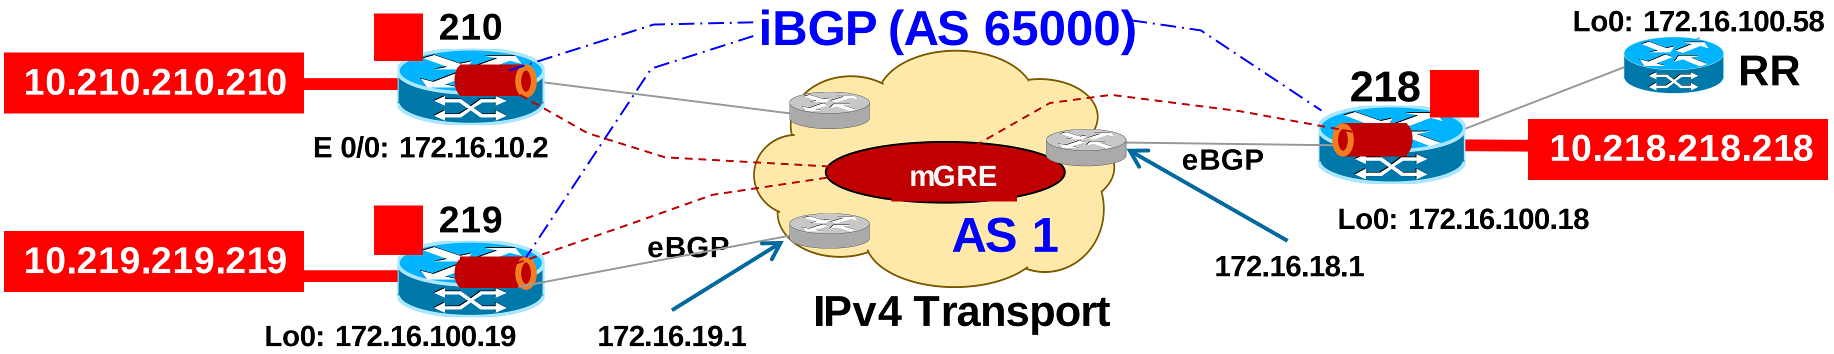


```
!
vrf definition red
 rd 1:1
 route-target export 1:1
 route-target import 1:1
 !
 address-family ipv4
!
interface Loopback0
 ip address 172.16.100.18 255.255.255.255
!
interface Ethernet0/0
 ip address 172.16.18.2 255.255.255.0
 service-policy output parent
!
```

```
!
l3vpn encapsulation ip Cisco
 transport ipv4 source Loopback0
 mpls mtu max

 !
!
route-map mgre-v4 permit 10
 set ip next-hop encapsulate l3vpn Cisco
```

MPLS VPN over Multipoint GRE (mGRE)

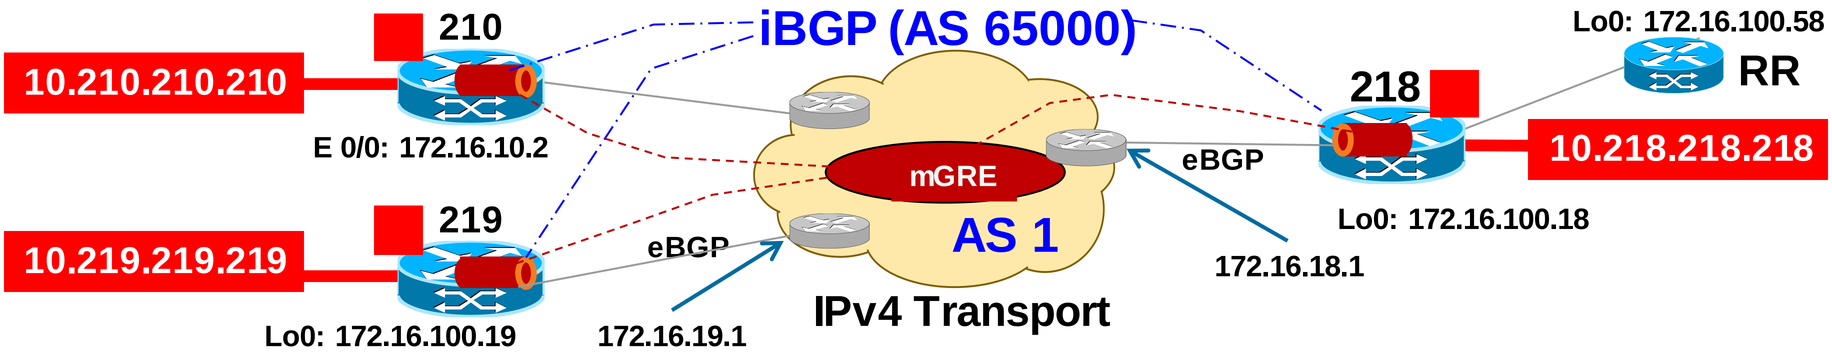


```
!
router bgp 65000
 neighbor 172.16.18.1 remote-as 1
 neighbor 172.16.18.1 update-source Eth 0/0
 neighbor 172.16.100.58 remote-as 65000
 neighbor 172.16.100.58 update-source Loop 0
 !
 address-family ipv4
  network 172.16.100.18 mask 255.255.255.255
  neighbor 172.16.18.1 activate
  neighbor 172.16.18.1 allowas-in 5
  neighbor 172.16.100.58 activate
 exit-address-family
 !
```

```
!
 address-family vpnv4
  neighbor 172.16.100.58 activate
  neighbor 172.16.100.58 send-community ext
  neighbor 172.16.100.58 route-map mgre-v4 in
!

 address-family ipv4 vrf red
 network 10.218.218.218 mask 255.255.255.255
!
```

MPLS VPN over Multipoint GRE (mGRE)

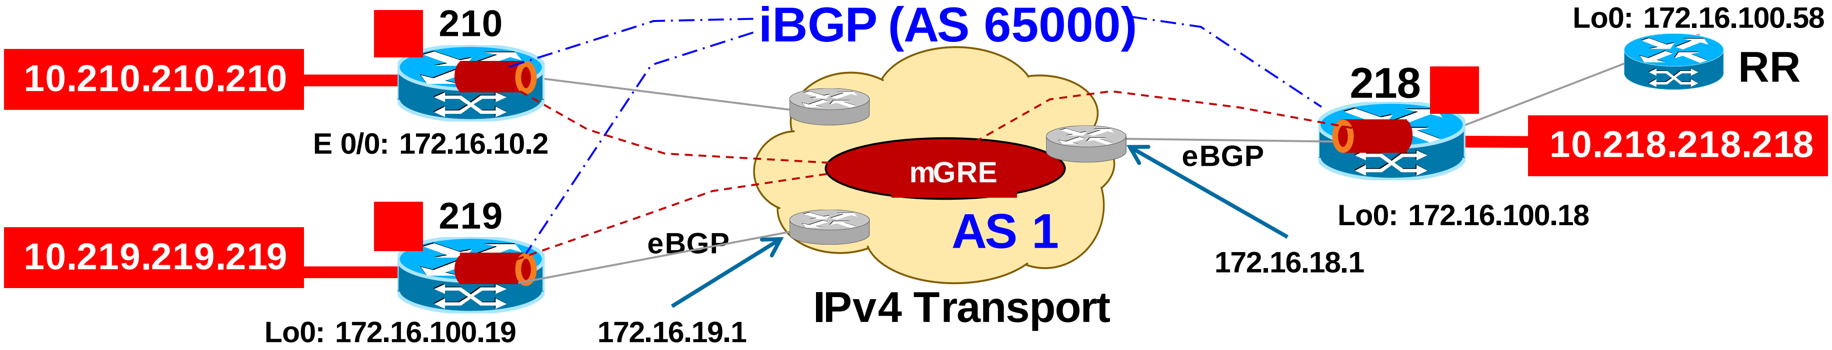


```
218#conf t
Enter configuration commands, one per line.  End with CNTL/Z.

218(config)#l3vpn encapsulation ip Cisco
218(config-l3vpn-encap-ip)#
*%LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0, changed state to up
```

MPLS VPN over Multipoint GRE (mGRE)

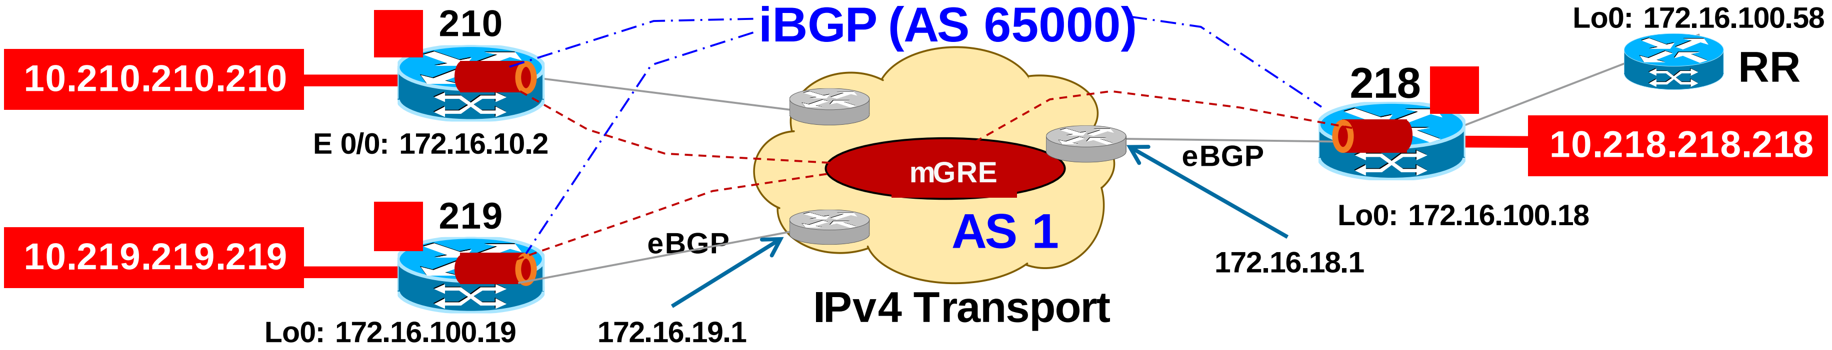


```
218#sh adjacency tunnel 0
Protocol Interface                 Address
IP       Tunnel0                   172.16.10.2(3)
TAG      Tunnel0                   172.16.10.2(3)
IP       Tunnel0                   172.16.100.19(3)
TAG      Tunnel0                   172.16.100.19(3)
```

```
218#sh l3vpn encapsulation ip

 Profile: Cisco
  transport ipv4 source Loopback0
  protocol gre
  payload mpls
   mtu max
  Tunnel Tunnel0 Created [OK]
  Tunnel Linestate [OK]
  Tunnel Transport Source Loopback0 [OK]
```


Cisco live!

MPLS VPN over Multipoint GRE (mGRE)

```
218#sh ip bgp vpnv4 vrf red
BGP table version is 8, local router ID is 172.16.100.18
.....

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf red)
 *>i 10.210.210.210/32
                       172.16.10.2              0    100      0 ?
 *>  10.218.218.218/32
                       0.0.0.0                  0         32768 i
 *>i 10.219.219.219/32
                       172.16.100.19            0    100      0 iD
```

```
218#sh ip route vrf red

Routing Table: red

Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 3 subnets
B        10.210.210.210 [200/0] via 172.16.10.2, 5d15h, Tunnel0
C        10.218.218.218 is directly connected, Loopback218
B        10.219.219.219 [200/0] via 172.16.100.19, 02:20:23, Tunnel0
```

Cisco live!

VRF-Lite over Dynamic Multipoint VPN (DMVPN)

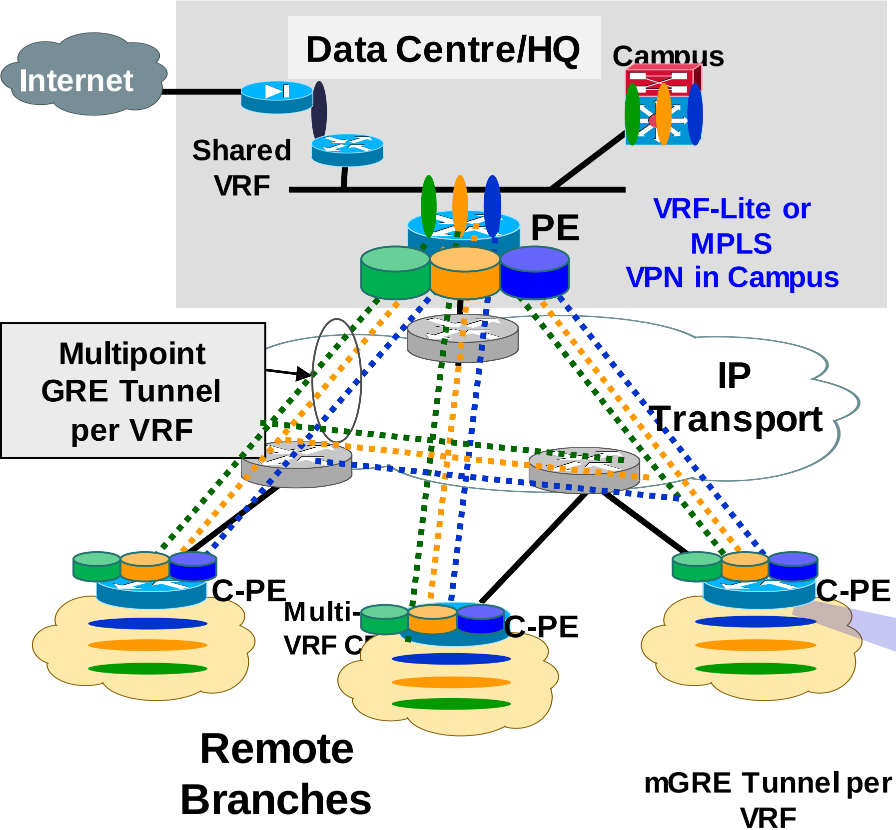

- Allows VRF segmentation over DMVPN framework
- A Multipoint GRE (mGRE) interface is enabled per VRF (1:1)
- Solution allows spoke-to-spoke data forwarding per VRF
- **Deployment Target:** Customers already running DMVPN, but needs to add VRF capabilities to sites

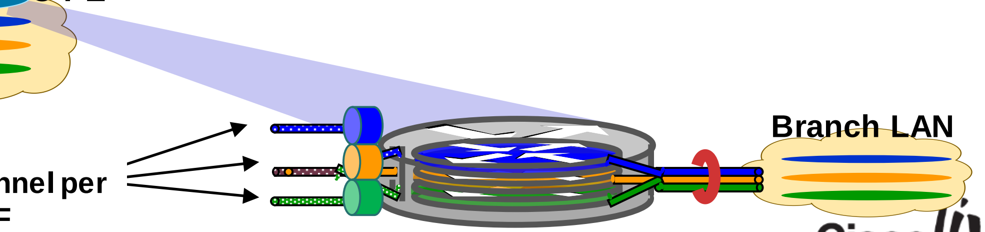

MPLS VPN over Dynamic Multipoint VPN (DMVPN)

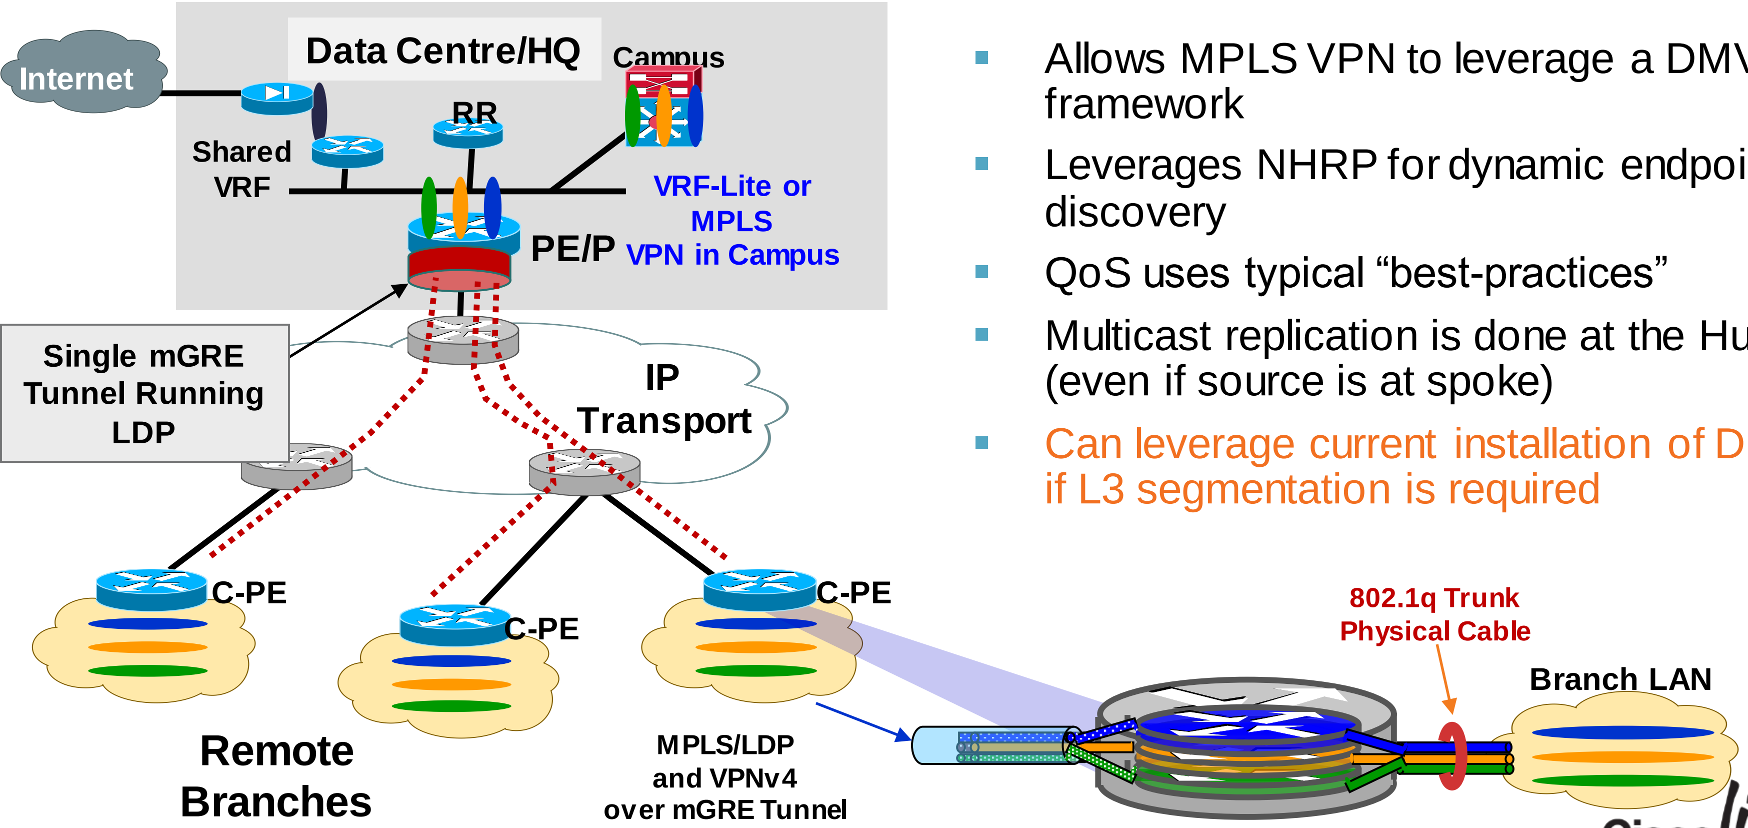

- Allows MPLS VPN to leverage a DMVPN framework
- Leverages NHRP for dynamic endpoint discovery
- QoS uses typical "best-practices"
- Multicast replication is done at the Hub (even if source is at spoke)
- Can leverage current installation of DMVPN if L3 segmentation is required

Inter AS for Campus-to-WAN BGP AS Interconnect

Inter AS Options for MPLS and MPLS-VPN over IP

Other Deployment Model Options

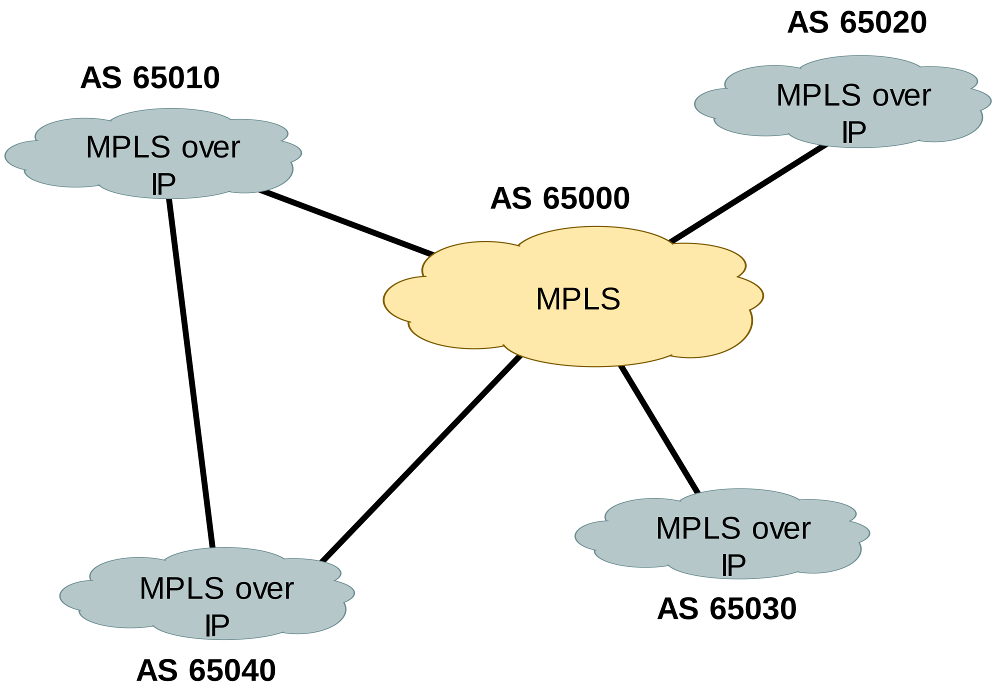

- Requirement is needed to interconnect L3 VPN AS's that exist in the network
- Campus to WAN, WAN to WAN, or WAN to DC
- Each AS is autonomously controlled by unique Ops team, but route exchange is required
- Several options exist for this "Inter AS" capability

Campus-to-WAN Interconnection

Inter AS Option A (Back to Back VRFs)

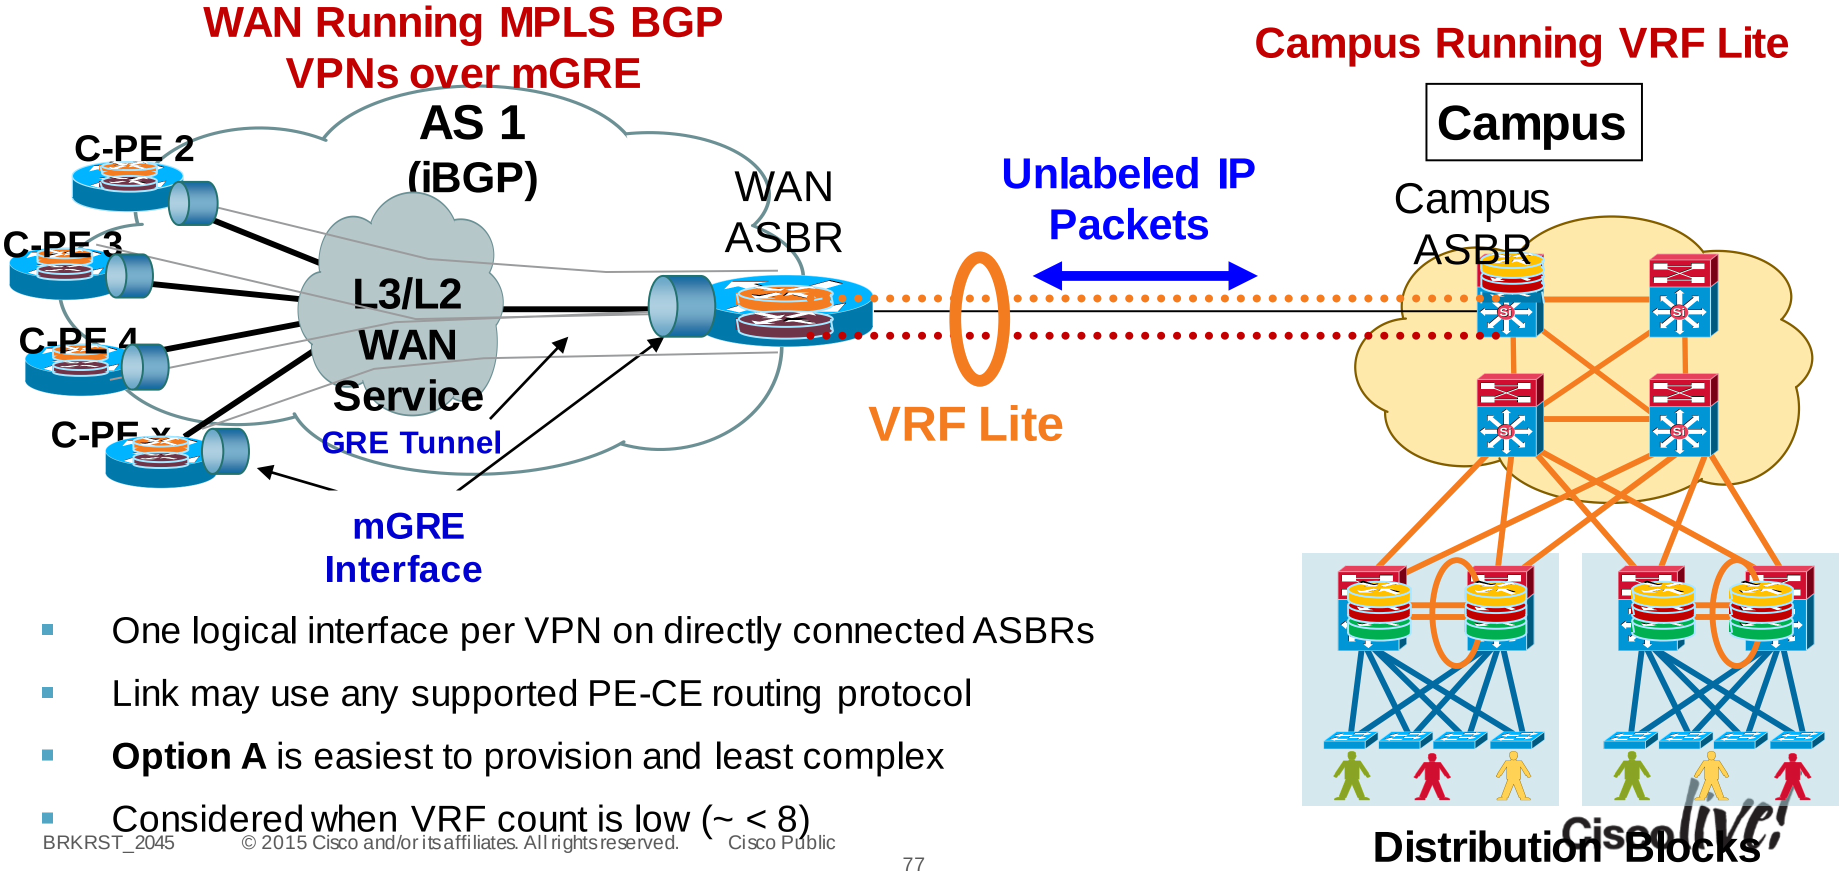

- One logical interface per VPN on directly connected ASBRs
- Link may use any supported PE-CE routing protocol
- **Option A** is easiest to provision and least complex
- Considered when VRF count is low (~ < 8)

Campus-to-WAN Interconnection

Inter AS Option B (Medium/Large VRF Deployments)

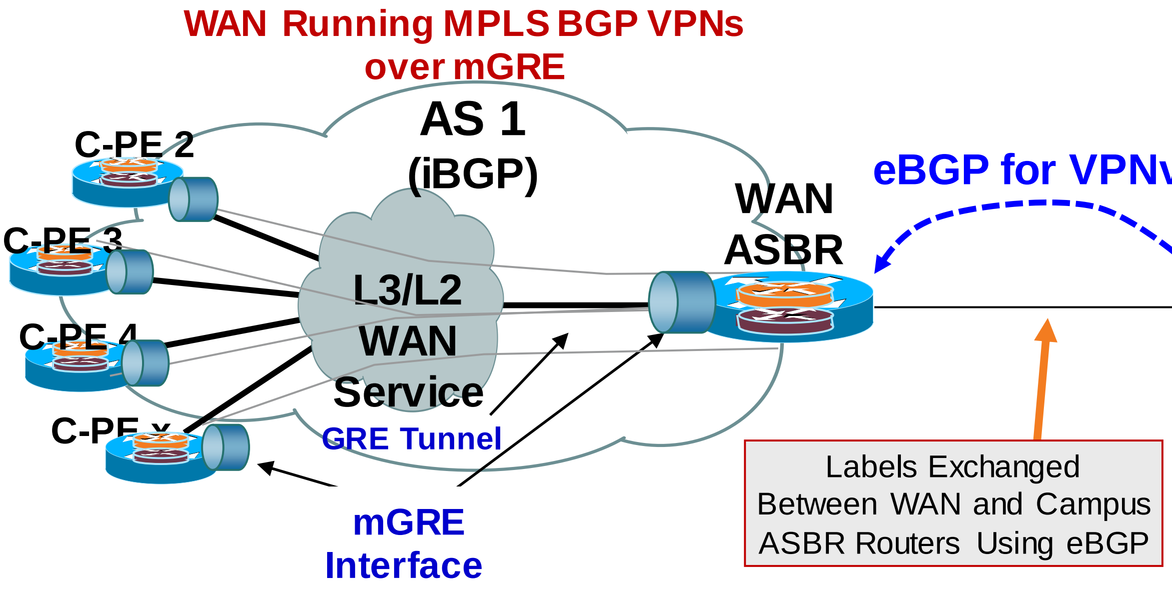

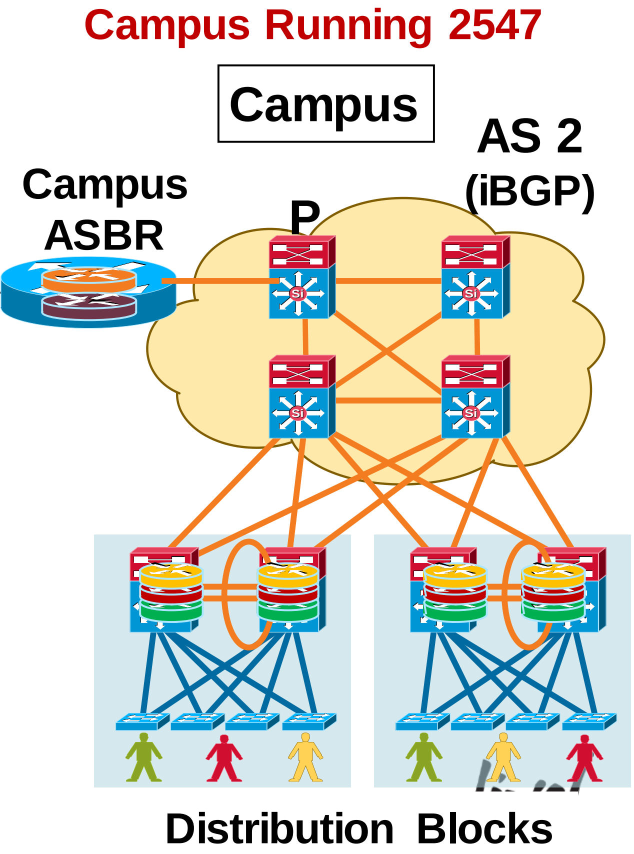

- ASBRs exchange VPN routes using eBGP
- ASBRs hold all VPNv4 routes needing exchange
- Recommended when VRF count is higher (~ >8)
- More complex that Option A, but more flexible

MPLS VPN over mGRE

Inter AS Example

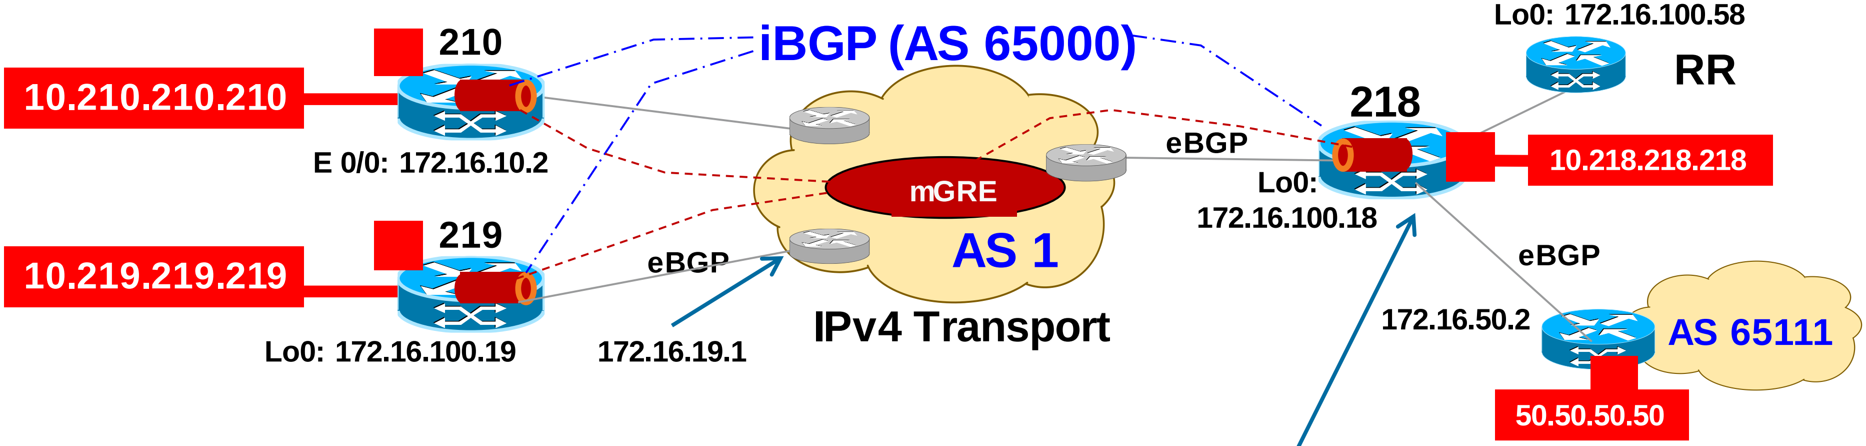

Enable next-hop-self under VPNv4 AF

```
router bgp 65000
......
!
 address-family vpnv4
  neighbor 172.16.100.58 activate
  neighbor 172.16.100.58 send-community ext
  neighbor 172.16.100.58 route-map mgre-v4 in
  neighbor 172.16.100.58 next-hop-self
!

 address-family ipv4 vrf red
 network 10.218.218.218 mask 255.255.255.255
!
```


MPLS VPN over mGRE + Inter AS

Inter AS Example

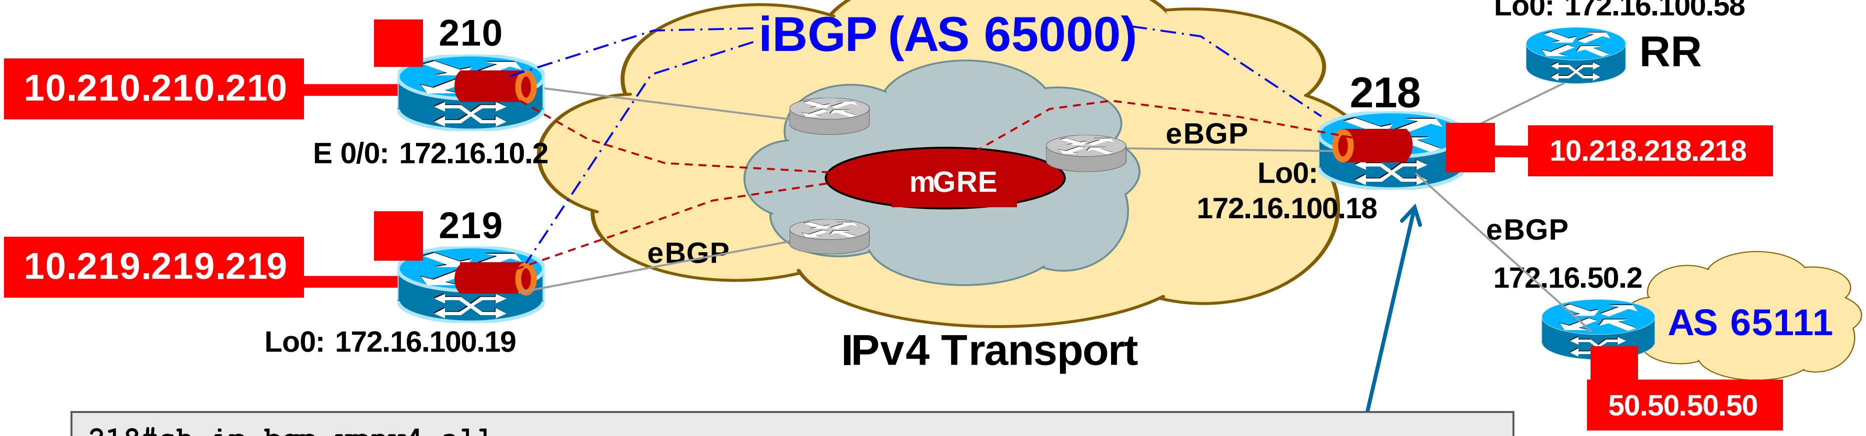


```
218#sh ip bgp vpnv4 all

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf red)

 *>  50.50.50.50/32   172.16.50.2              0             0 65111 ?
```

```
Route-Reflector#sh ip bgp vpnv4 all

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1

 *>i 50.50.50.50/32   172.16.100.18            0    100      0 65111 ?
```

Using Locator ID Separation Protocol (LISP) for L3 Segmentation over the WAN

Enhancing the L3 VPN Segmentation Portfolio…

- **VRF Lite Options**
  - Leverage Carrier Ethernet E-LINE/E-LAN services
  - Over GRE (GRE tunnel per VRF)
  - Over DMVPN (mGRE interface per VRF)
  - Easy Virtual Networking (EVN) over an E-LINE Carrier Ethernet service
- **L3 MPLS BGP VPN (RFC 4364)**
  - Over L2 transport (PE-PE, P-P, PE-P)… optical, Ethernet, SONET/SDH, etc…
  - Over p2p GRE tunnels
  - Over DMVPN
- **MPLS VPN over Multipoint GRE (mGRE)**
- **LISP Multi-Tenancy for L3 Segmentation**

What is LISP? (Locator-ID Separation Protocol)

A Next Generation Routing Architecture – RFC 6830

LISP creates a "Level of indirection" with two namespaces: **EID** and **RLOC**

- **EID (Endpoint Identifier)** is the IP address of a host – just as it is today
- **RLOC (Routing Locator)** is the IP address of the LISP router for the host
- **EID-to-RLOC mapping** is the distributed architecture that maps **EIDs** to **RLOCs**

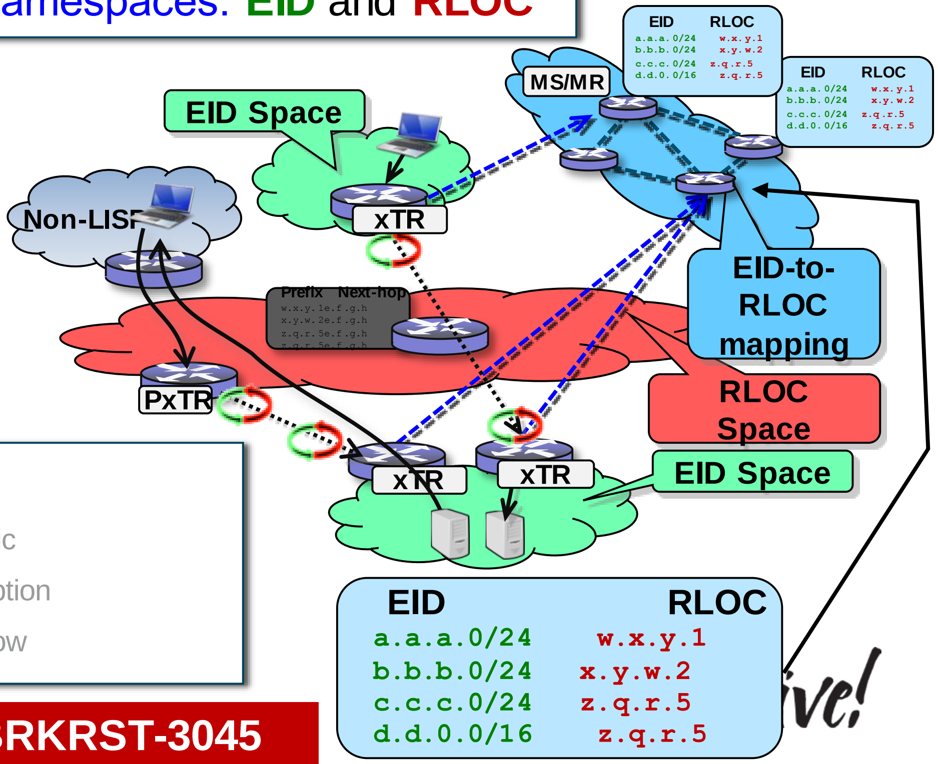

- Network-based solution
- No host changes
- Minimal configuration
- Incrementally deployable
- Support for mobility
- Address Family agnostic
- IPv4 to v6 Transition option
- In Cisco IOS/NX-OS now

More Details on LISP Covered in Session BRKRST-3045

LISP Overview

What do we mean by "location" and "identity"?

Today's Internet Behavior
Loc/ID "overloaded" semantic

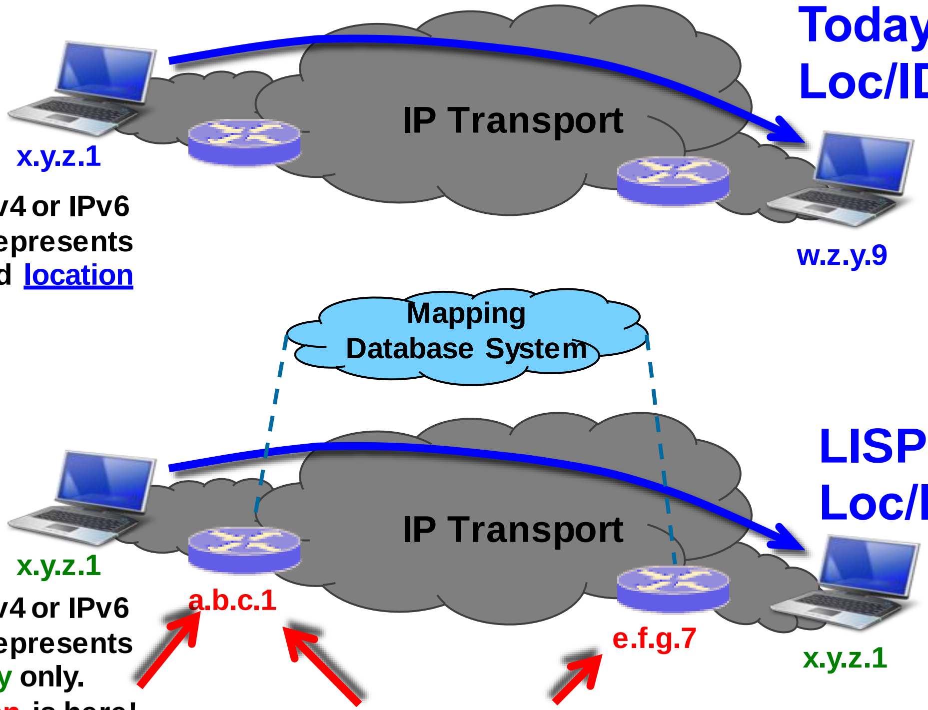

Device IPv4 or IPv6 address represents identity and location

When the device moves, it gets a new IPv4 or IPv6 address for its new identity and location

LISP Behavior
Loc/ID "split"

Device IPv4 or IPv6 address represents identity only.
Its location is here!

Only the location changes

When the device moves, keeps its IPv4 or IPv6 address.
It has the same identity

LISP Operations

LISP Mapping Resolution – DNS Analogy…

LISP "Level of Indirection" is analogous to a DNS lookup

- DNS resolves **IP addresses** for **URLs**

DNS
Name-to-IP
URL Resolution

LISP resolves **locators** for queried **identities**

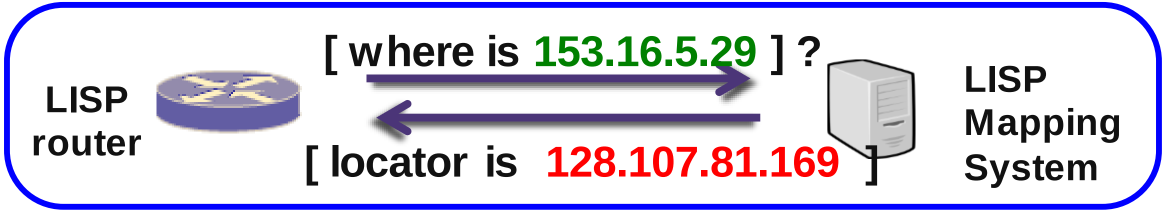

LISP
Identity-to-locator
Mapping Resolution

LISP - Basic Routing Concept

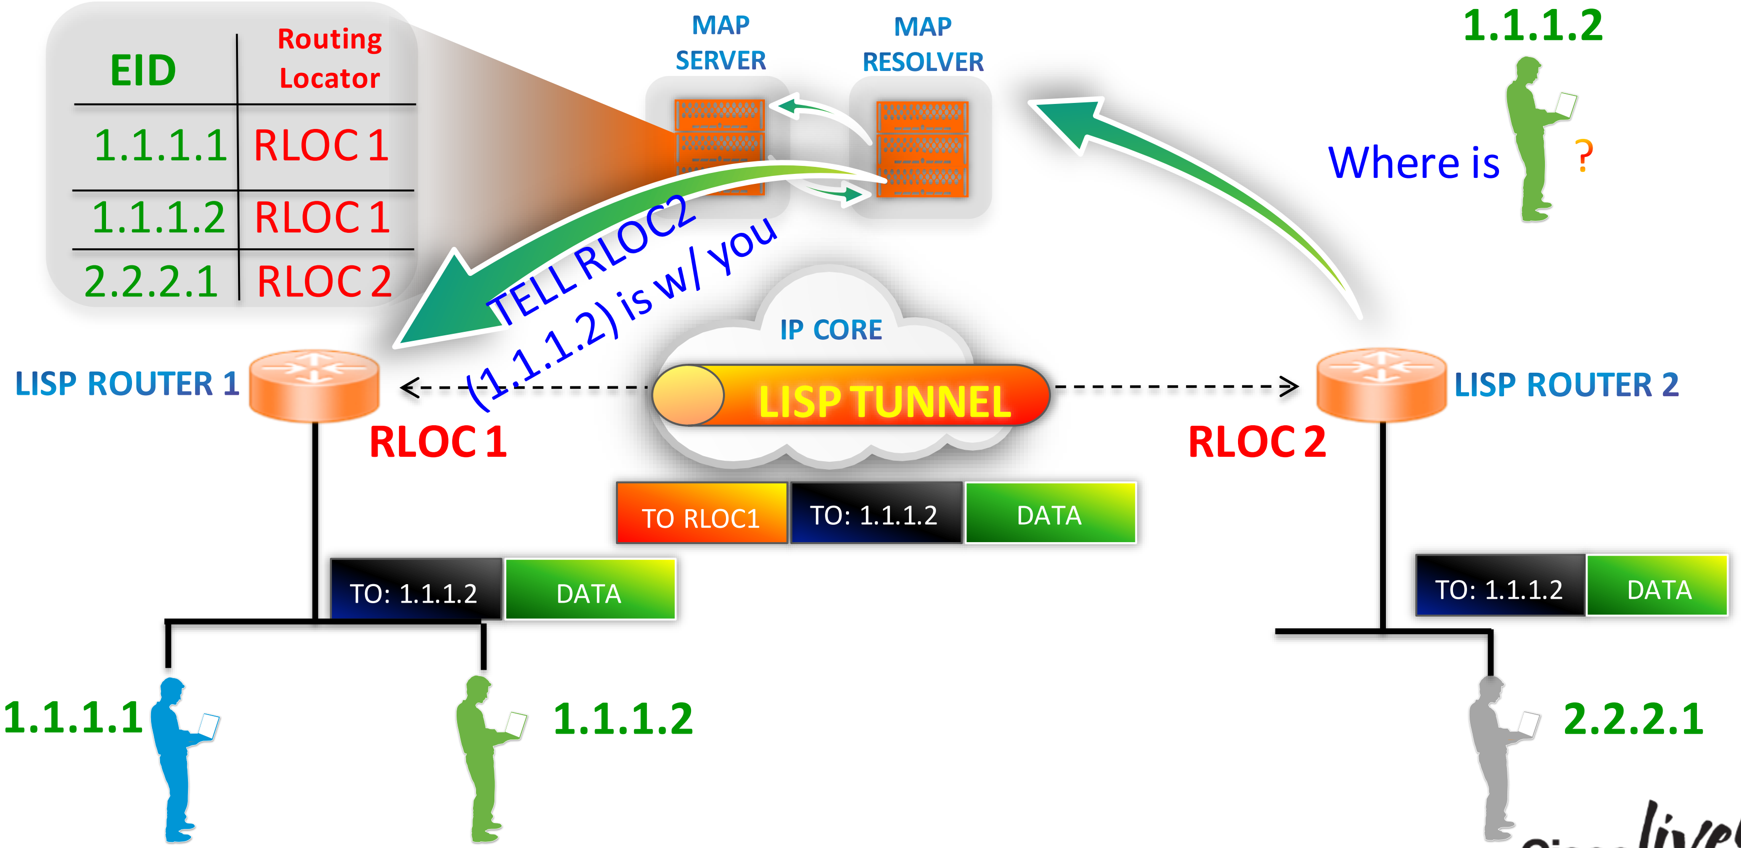

LISP Use Cases

The Five Core LISP Use-Cases

1. Efficient Multi-Homing
2. IPv6 Transition Support
3. **Network Segmentation/Multi-Tenancy**
4. Host/VM Mobility
5. LISP Mobile-Node

LISP Operations

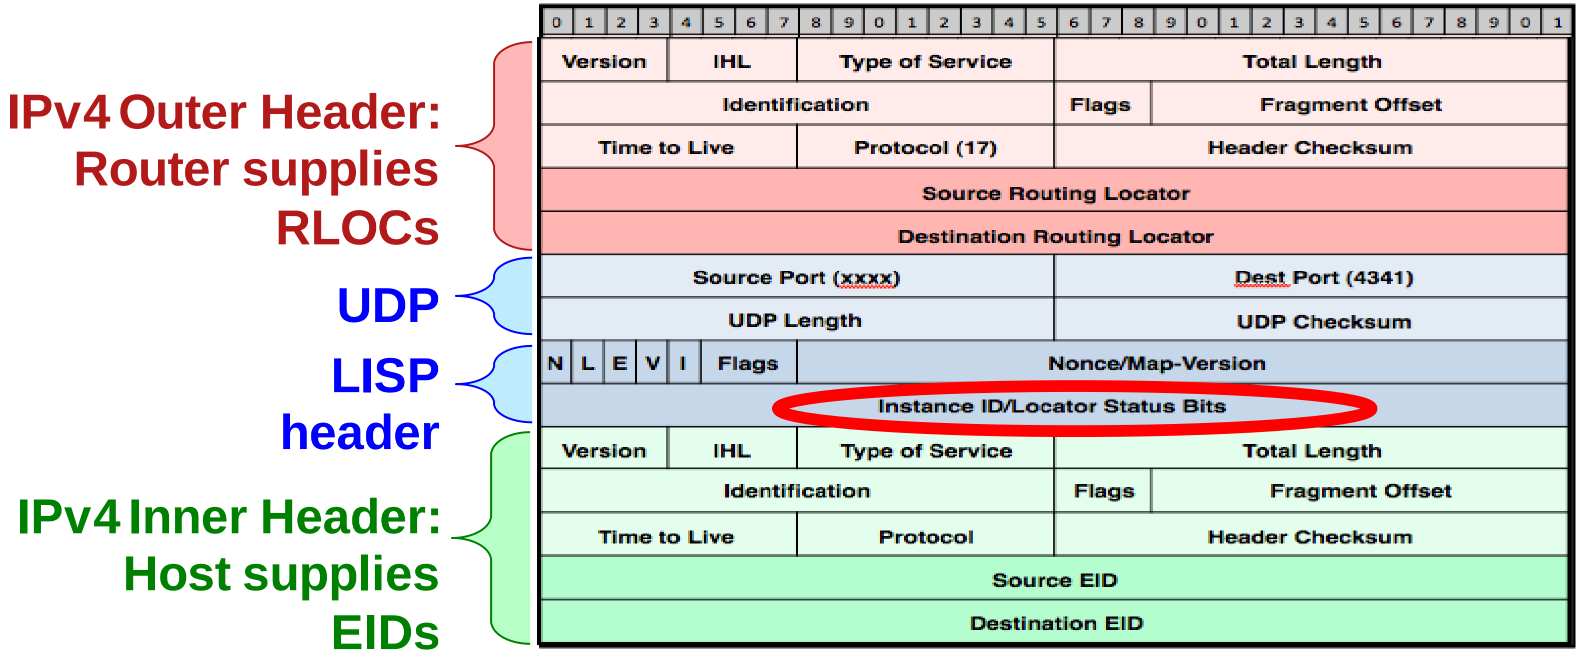

LISP Segmentation/VPN

Efficient Segmentation/Multi-Tenancy Support – Concepts…

- Because LISP considers Segmentation of both EID and RLOC namespaces, two models of operation are defined: Shared and Parallel
- Shared Model
  - Virtualises the EID namespaces
  - Binds an EID namespace privately defined using a VRF to an Instance-ID
  - Uses a common (shared) RLOC (locator) address space
  - The Mapping System is also part of the locator namespaces and is shared
- Parallel Model
  - Virtualises the RLOC (locator) namespaces
  - One or more EID instances may share a virtualised RLOC namespace
  - A Mapping System must also be part of each locator namespaces

LISP Segmentation/VPN

Efficient Segmentation/Multi-Tenancy Support – Shared Model…

- **Shared Model** – at the device level (think MPLS/MPLS-VPN…)
  - Multiple EID-prefixes are allocated privately using VRFs
  - EID lookups are in the VRF associated with an Instance-ID
  - All RLOC lookups are in a single table – (default/global or RLOC VRF)
  - The Mapping System is part of the locator address space and is shared

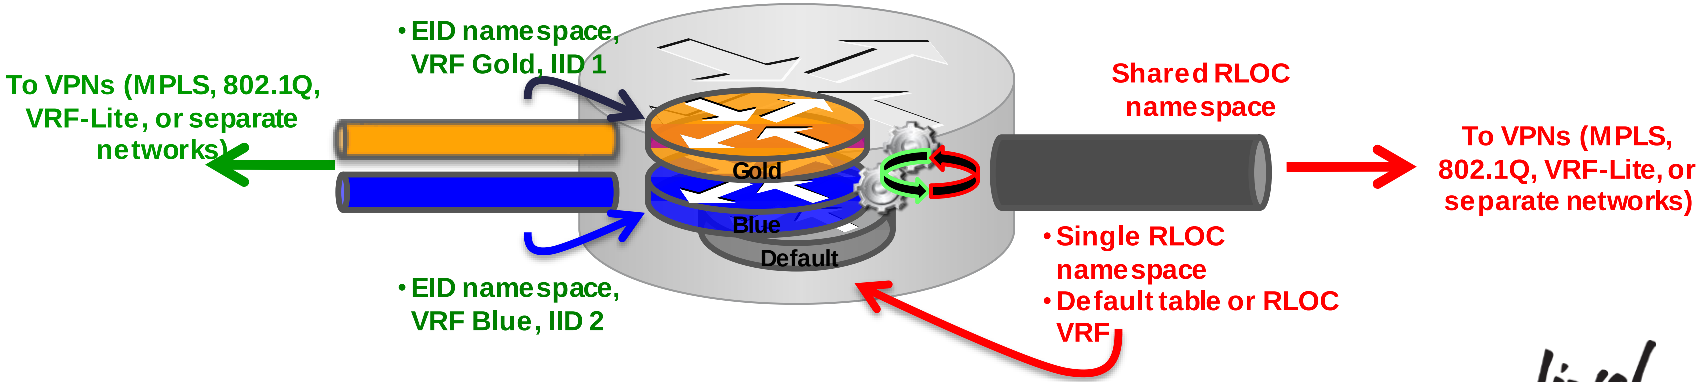

LISP Segmentation/VPN

Efficient Segmentation/Multi-Tenancy Support – Parallel Model…

- **Parallel Model** – at the device level (think VRF-Lite…)
  - Multiple EID-prefixes are allocated privately using VRFs
  - EID lookups are in the VRF associated with an Instance-ID
  - RLOC lookups are in the VRF associated with the locator table
  - A Mapping System must be part of each locator address space

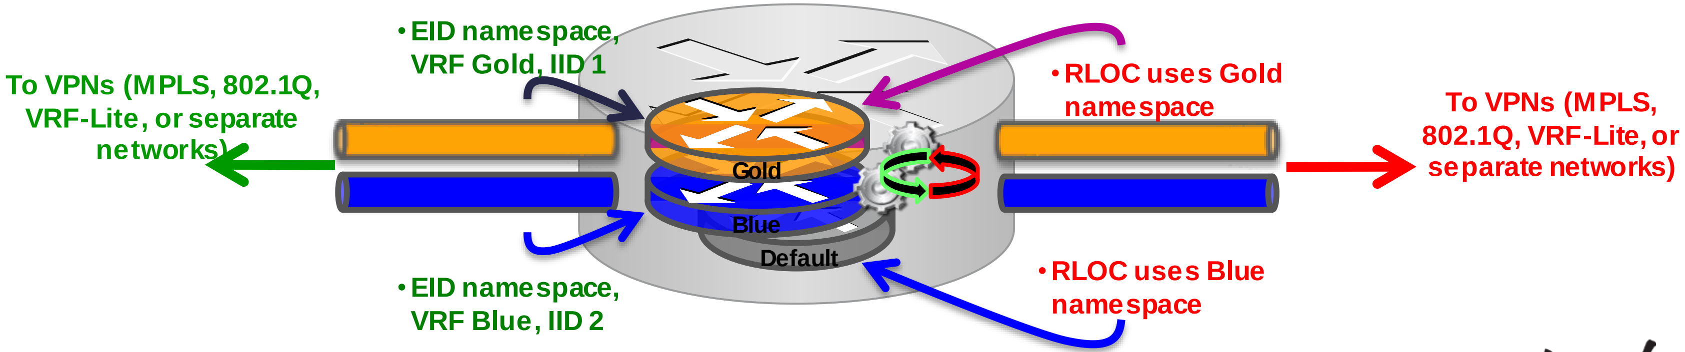

LISP in Enterprise WAN/Branch

Leverage LISP Framework for WAN Branch Backhaul

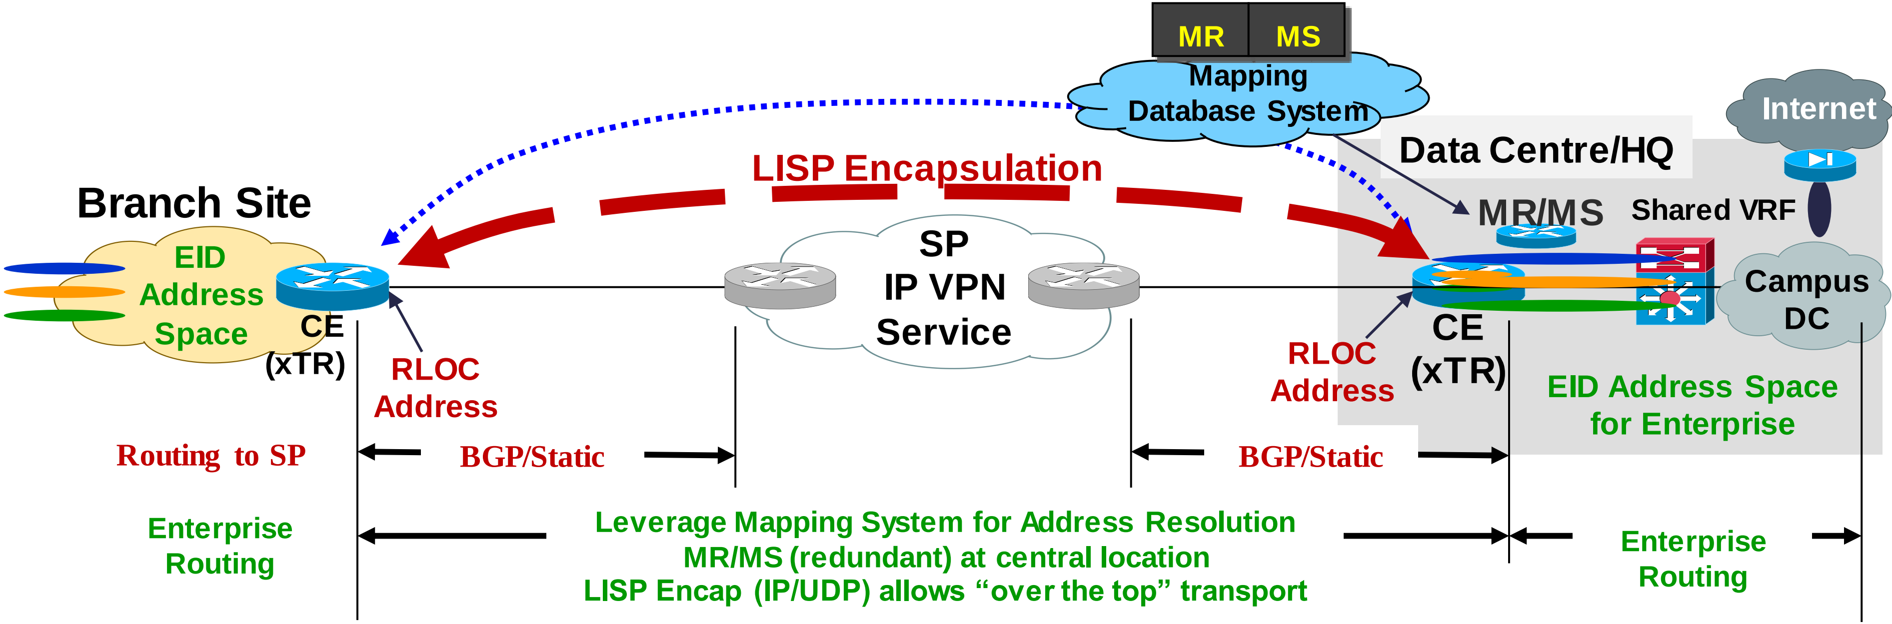

- Allows network segmentation on xTR (viewed as CE in L3 VPN model)
- PE routers require minimal routes (RLOC address only, which only SP knows)
- VRF Segmentation is applied to CE/xTR
- Offers another "over the top" Segmentation solution (VRF capabilities
- Can leverage GET VPN for additional data security (IPSec)

MR = Map Resolver
MS = Map Server

	MPLS VPN over mGRE	LISP Segmentation
IPv6 Transition	Y	Y
Segmentation	VRF	VRF
VRF Identifier	VPN Label	Instance ID
Scale	1000+	1000+
Multi-Homing	Y (BGP/IGP recursion)	Y (simple)
Spoke to Spoke (w/ Virt)	Y (Y)	Y (Y)
Tunneless IP (encap)	Y (RFC 4023)	Y (native IP/UDP)
Manual Tunnel config	N	N
Single IP address sent to provider?	Y (mGRE source IP)	Y (RLOC)
Control Plane	RFC 4364 i/eBGP (RR)	Map DB
Encryption Support	Y (GET)	Y (GET)
Route Learning	BGP (Push)	MR/MS (Pull)
Convergence	Sub-second (BGP PIC)	seconds
Load Balance over multiple links	N (limited)	Y
MVPN Support	Y	Y
Route Distribution Model	PUSH (BGP advertisement)	PULL (on-demand only)

Agenda

- Introduction - Network Segmentation Drivers and Concepts
- WAN Transport Impact on L3 VPN over IP
- Technology Deep-Dive on Advancements in L3 VPN over IP
- **QoS, MTU, and Encryption Recommendations**
- Recent "Innovations" Evolving in L3 Segmentation
- Summary

Securing L3 VPN Solutions over the WAN with GET VPN

Group Encrypted Transport (GET) VPN

Public/Private WAN

Before: IPSec P2P Tunnels

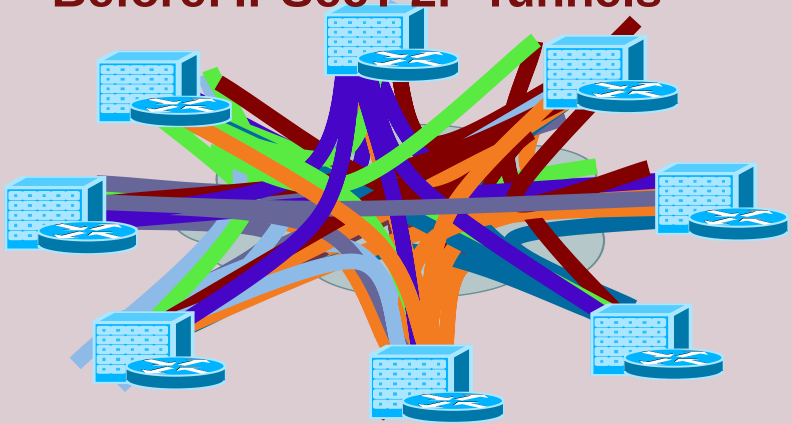

- Scalability—an issue (N^2 problem)
- Overlay routing
- Any-to-any instant connectivity can't be done to scale
- Limited QoS
- Inefficient Multicast replication

Private WAN

After: Tunnel-Less VPN

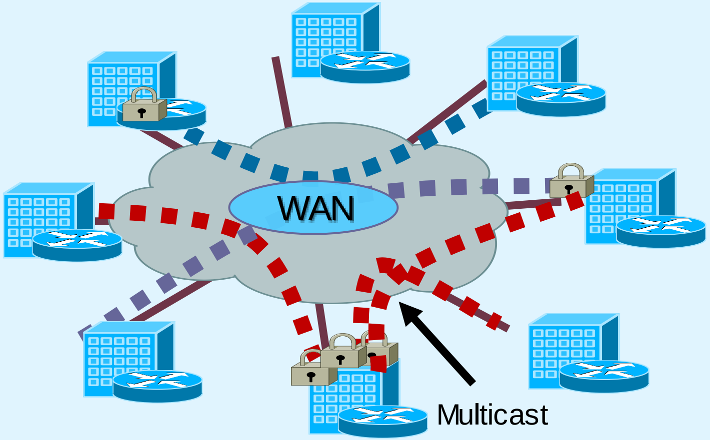

- Scalable architecture for any-to-any connectivity and encryption
- No overlays—native routing
- Any-to-any instant connectivity
- Enhanced QoS
- Efficient Multicast replication

Combining Technologies into Secure L3 Segmentation

Leverage MPLS VPN over mGRE + GET VPN Encryption

C-PE = Customer "owned" PE
GM = Group Member

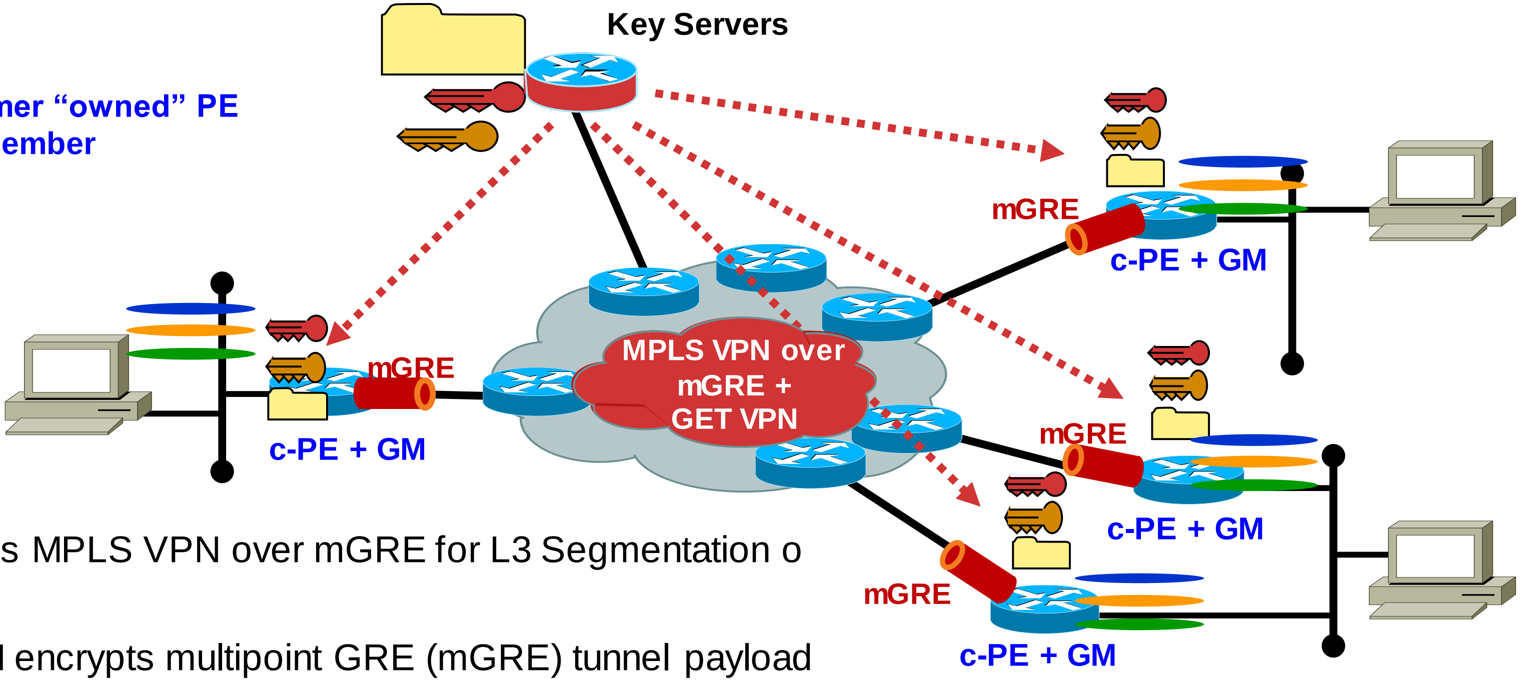

- C-PE runs MPLS VPN over mGRE for L3 Segmentation o IP
- GETVPN encrypts multipoint GRE (mGRE) tunnel payload
- Payload of VPNv4 (VRF) traffic is encrypted
- Leverage simplicity of MPLS VPN over mGRE + GETVPN

MPLS VPN over mGRE + GET VPN - White Paper

http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns431/ns658/white_paper_c11-726689.html

Cisco live!

White Paper

Secure Extension of Community of Interests Across Wide Area Networks

Authors

Mark "Mitch" Mitchiner
Solutions Architect
U.S. Federal Area
mmitchin@cisco.com

Craig Hill
Distinguished Systems Engineer
U.S. Federal Area
crhill@cisco.com

Abstract

This paper examines how recent network-based virtualization technology can be used to simplify community of interest (COI) deployment and operations within Department of Defense (DoD), Intelligence Community (IC), and secure enterprise networks.

The primary innovations addressed in this paper are Multiprotocol Label Switching (MPLS) over multipoint GRE (mGRE), combined with Group Encrypted Transport (GET) Virtual Private Network (VPN) technology while utilizing Next Generation Encryption ([NGE], also known as Suite B). These technologies, when combined as an architectural framework, address some of the major scaling, deployment, and operational challenges common in secure Wide Area Networks (WANs) today when Layer 3 network virtualization is required.

http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns431/ns658/white_paper_c11-726689.pdf

QoS Considerations for L3 Segmentation over the WAN

QoS with GRE, MPLS over GRE

GRE Header

GRE Header with ToS Reflection

MPLS over GRE Header with ToS Reflection

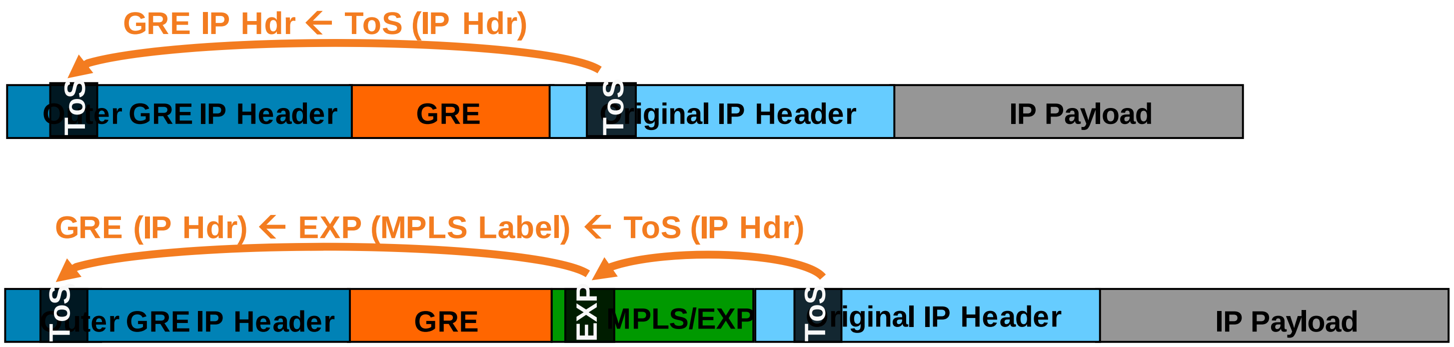

- Router will copy original ToS marking to outer GRE header
- For MPLS over GRE, the EXP marking is copied to the outer header of the GRE tunnel
- This allows the IPv4 “transport” to perform QoS on the multi-encapsulated packet

Caveats:

- **Traffic originating on the router (SNMP, pak_priority for routing, etc…), could have different behavior**

QoS Deployment Models in a Virtualised Environment

- **Aggregate Model**

  A common QoS strategy is used for all VRFs

  – **i.e. same marking for voice, video, critical data, best effort… regardless of the VRF the traffic is sourced from or destined too.**

  Allows identical QoS strategy to be used with/without Segmentation

- **Prioritised VRF Model**

  Traffic in a VRF(s) are prioritised over other VRFs

  Example: Prioritise "production" traffic over "Guest" access

  More complex. Could leverage PBR with MPLS-TE to accomplish this

Aggregate vs. Prioritised Model

Following the **"Aggregate Model"** Allows the Identical QoS Strategy to Be Used With/Without Network Segmentation

QoS Deployment with Network Segmentation

Point-to-Cloud Example - Hierarchical QoS + MPLS VPN over mGRE

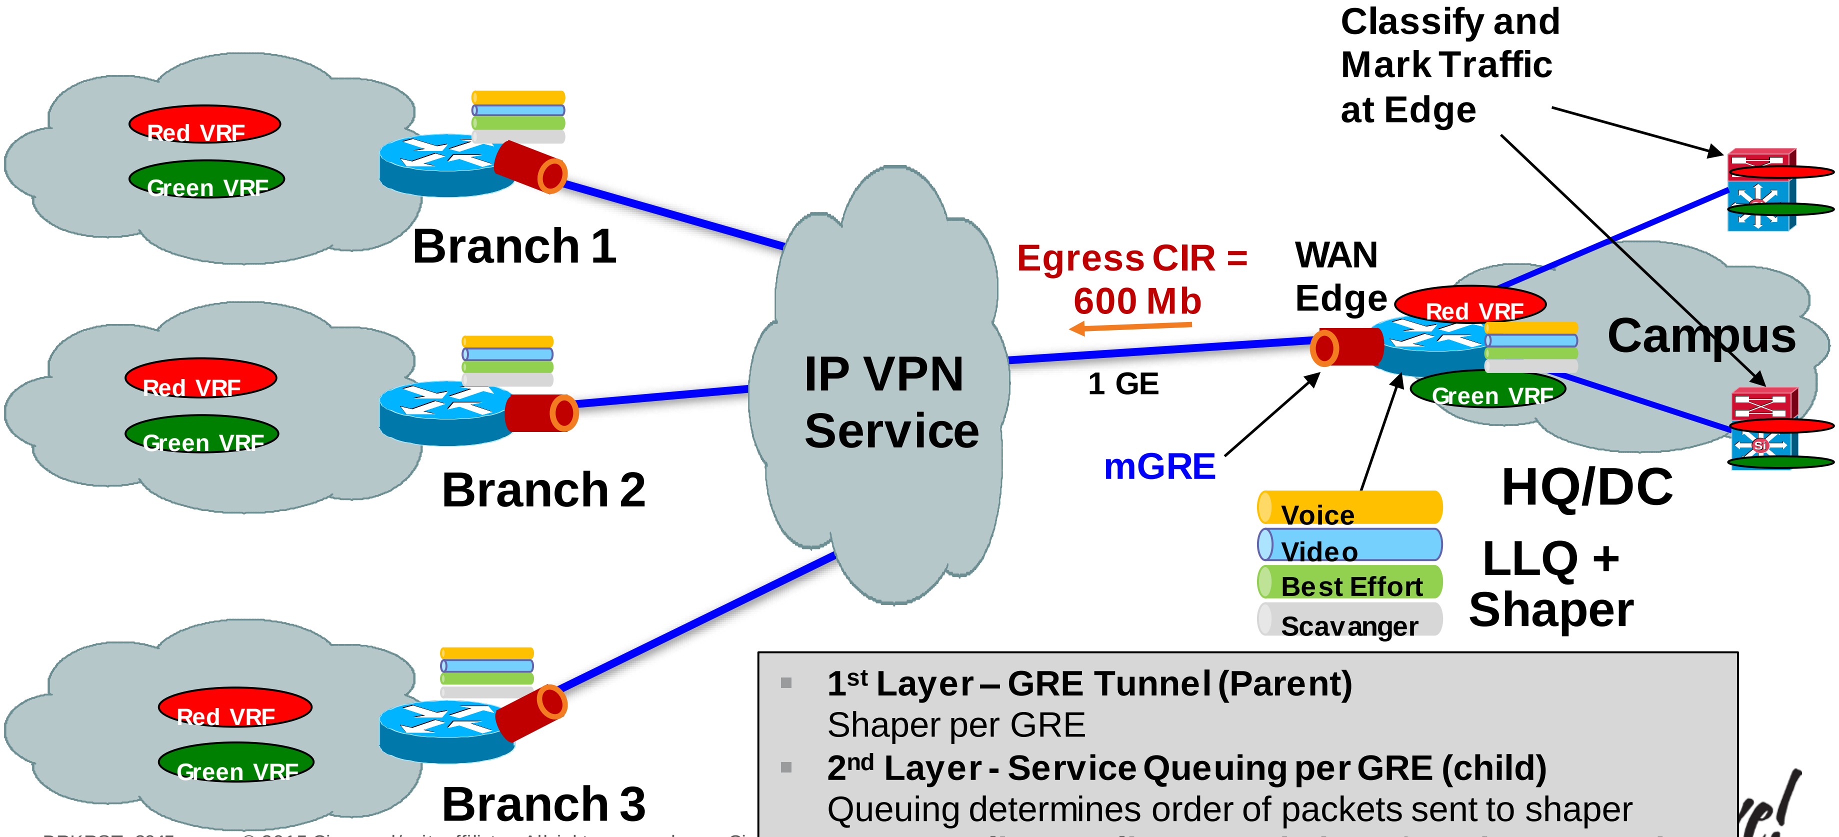

- 1st Layer – GRE Tunnel (Parent)
  Shaper per GRE
- 2nd Layer - Service Queuing per GRE (child)
  Queuing determines order of packets sent to shaper
- H-QoS policy applies to main interface (not mGRE)

Hierarchical QoS Example

H-QoS Policy on Interface to SP, Shaper = CIR

Two MQC Levels

```
Policy-map PARENT
   class class-default
    shape average 600000000   [600 Mbps shaper]
    service-policy output CHILD

Policy-map CHILD
    class Voice
      police cir percent 10
    class Video
      police cir percent 20
    class Scav
        bandwidth remaining ratio 1
    class class-default
     bandwidth remaining ratio 9

Interface gigabitethernet 0/1.100
  service-policy output PARENT
```

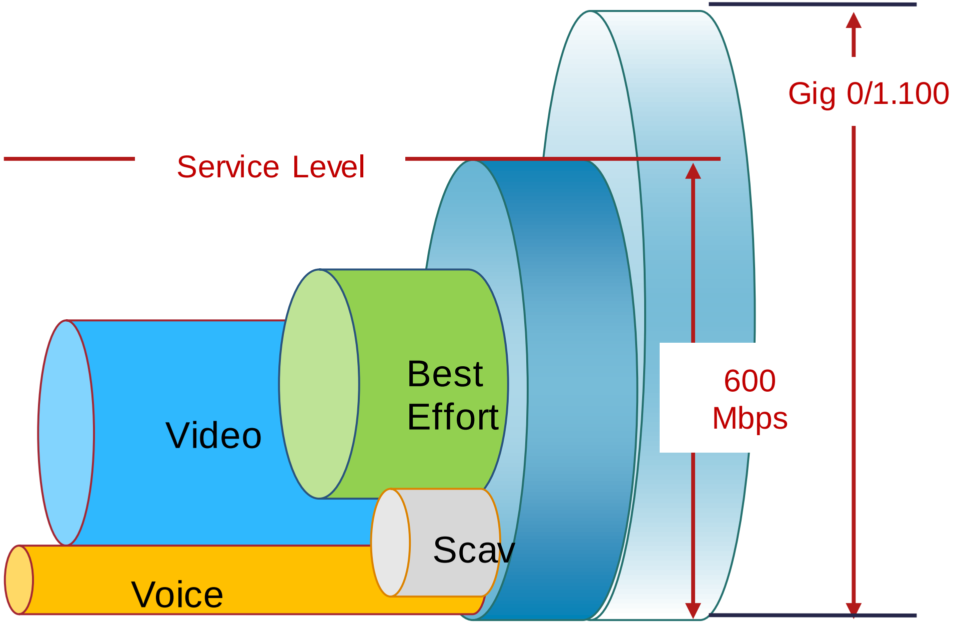

MTU Considerations with GRE Tunnels

- Fragmentation is unavoidable in some cases
- The use of GRE tunnels increase the chances of MTU issues (i.e. fragmentation) due to the increase in IP packet size GRE adds
- Main Issue: The performance impact to the router when the GRE tunnel destination router must re-assemble fragmented GRE packets
- Common Cases where fragmentation occurs?:
  - Customer does not control end to end IP path (some segment is < MTU)
  - Router generates an ICMP message, but the ICMP message gets blocked by a router or firewall (between the router and the sender). Most Common!! ☹

MTU Recommendations

- ✓ Avoid fragmentation ☺ (if at all possible)
- ✓ Consider "tunnel path-mtu-discovery" command to allow the GRE interface to copy DF=1 to GRE header, and run PMTUD on GRE
- ✓ Set "ip mtu" on the GRE to allow for MPLS label overhead (4-bytes)
  - ✓ If using IPSec, "ip mtu 1400" is recommended
- ✓ Configure ip tcp adjust-mss for assist with TCP host segment overhead
- ✓ MTU Setting options:
  - ✓ Setting the MTU on the physical interface larger than the IP MTU

```
interface Ethernet 1/0
. . .
 mtu 1500
```

  - ✓ Set IP MTU to GRE default (1476) + MPLS service label (4)

```
interface Tunnel0
. . .
 ip mtu 1472
```

- ✓ Best to fragment prior to encapsulation, than after encapsulation, as this forces the "host" to do packet reassembly (vs. the remote router)

MTU Recommendations

- ✓ Multipoint GRE (mGRE) interfaces are "stateless"
- ✓ "tunnel path-mtu-discovery" command is not supported on mGRE interfaces (defaults to DF=0 for MPLS VPN o mGRE)
- ✓ For the MPLS VPN over mGRE Feature, "ip mtu" is automatically configured to allow for GRE overhead (24-bytes) (and GRE tunnel key if applied)

```
interface Tunnel 0
. . .
Tunnel protocol/transport multi-GRE/IP
    Key disabled, sequencing disabled
    Checksumming of packets disabled
  Tunnel TTL 255, Fast tunneling enabled
  Tunnel transport MTU 1476 bytes
```

IP MTU Defaults to 1476 When MPLS VPN over mGRE Is Used

- ✓ Configure ip tcp adjust-mss for assist with TCP hosts (inside interface)
- ✓ MTU Setting options:
  - ✓ Setting the MTU on the physical interface larger than the IP MTU
- ✓ Best to fragment prior to encapsulation, than after encap, as remote router (GRE dest) must reassemble GRE tunnel packets

IP MTU Technical White Paper:
http://www.cisco.com/en/US/tech/tk827/tk369/technologies_white_paper09186a00800d6979.shtml

Agenda

- Introduction - Network Segmentation Drivers and Concepts
- WAN Transport Impact on L3 VPN over IP
- Technology Deep-Dive on Advancements in L3 VPN over IP
- QoS, MTU, and Encryption Recommendations
- **Recent "Innovations" Evolving in L3 Segmentation**
- Summary

Innovations Worth Investigating Further

- IWAN 3.0 Solutions
  - Leverage Intelligent overlay networks for latency based routing
- VRF Aware Services Interface (VASI)
  - (in backup slides)
- EIGRP Over The Top
- Leveraging SDN for WAN Automation Provisioning
  - Using WAN Automation Engine (WAE) in self deployed MPLS networks
- Flex VPN in Virtualised Networking Environments

Agenda

- Introduction - Network Segmentation Drivers and Concepts
- WAN Transport Impact on L3 VPN over IP
- Technology Deep-Dive on Advancements in L3 VPN over IP
- QoS, MTU, and Encryption Recommendations
- Recent “Innovations” Evolving in L3 Segmentation
- **Summary**

Summary – Key Takeaways…

WAN Segmentation - Key Takeaways

- The ability for an enterprise to extend Layer 3 (L3) Segmentation technologies over the WAN is critical for today's applications
- The ability to transport VRF-Lite and MPLS-VPN over IP allows flexible transport options, including ability to encrypt segmented traffic
- Understanding key network criteria (topology, traffic patterns, VRFs, scale, expansion) is vital to choosing the "optimal" solution for extending Segmentation over the WAN
- MPLS VPN over mGRE offers simpler, and more scalable, deployment, eliminating LDP, manual GRE, for the WAN
- Understand the options for QoS, GET VPN in mGRE environments, and the impact of MTU and available tools in IOS for MTU discovery
- Begin to understand Cisco innovations (MPLS VPN over mGRE, EVN, LISP Segmentation) and how they can help simplify network Segmentation in the WAN for future designs
- **Leverage the technology, but "Keep it Simple" when possible ☺**

Continue Your Education

- Demos in the Cisco Campus
- Walk-in Self-Paced Labs
- Meet the Expert 1:1 meetings

Recommended Reading

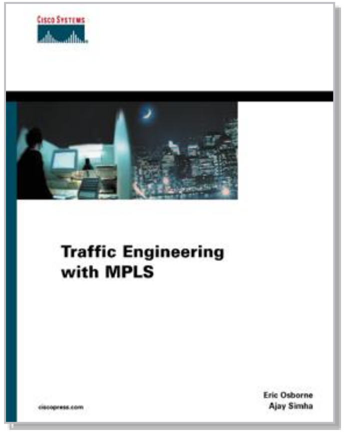

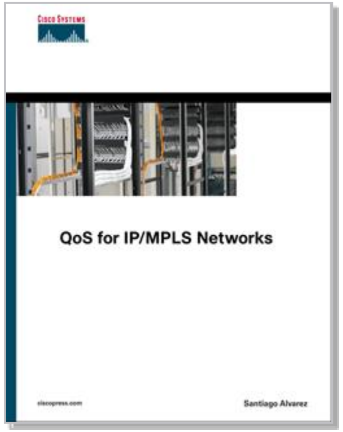

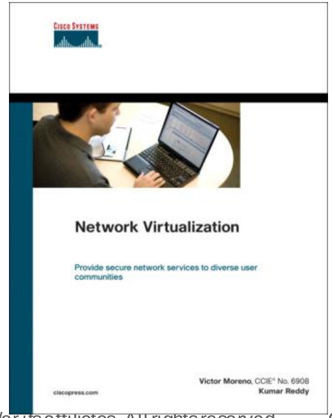

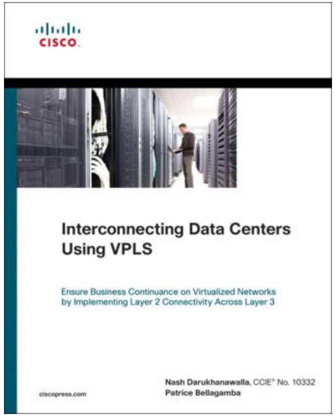

Q & A

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

Cisco live!

Thank you.