# TOMORROW starts here.

# Network Virtualisation Design Concepts over the WAN

BRKRST-2045

Craig Hill

Distinguished Systems Engineer

Cisco live!

# Session Assumptions & 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
  - Basic knowledge of MPLS VPNs operation, MP-BGP, GRE tunnelling, IP QoS
- This discussion will not cover VMware, Virtual Machines, or other server Virtualisation technologies
- Data Centre Interconnection (DCI) is an important element in a complete WAN Virtualisation infrastructure, but is not a focus in this session nor is Layer 2 Virtualisation technologies
- RFC 2547 (BGP/MPLS IP VPNs) is referenced frequently for MPLS VPN.  This is for familiarity only.  RFC 2547 is now replaced with RFC 4364.

# Agenda

- Introduction - Network Virtualisation Drivers and Concepts
- SP WAN Transport Service Impact on L3 Virtualiation Solution Choices
- Technology and Deployment Deep-Dive for L3 Virtualised WAN
- Integrating Encryption into L3 Virtualisation Solutions
- Recent “Innovations” Evolving in L3 Virtualisation
- QoS/H-QoS and MTU Considerations for L3 Virtualisation Deployments
- Summary and Wrap Up

# Agenda

- **Introduction - Network Virtualisation Drivers and Concepts**
- SP WAN Transport Service Impact on L3 Virtualisation Solution Choices
- Technology and Deployment Deep-Dive for L3 Virtualised WAN
- Integrating Encryption into L3 Virtualisation Solutions
- Recent “Innovations” Evolving in L3 Virtualisation
- QoS/H-QoS and MTU Considerations for L3 Virtualisation Deployments
- Summary and Wrap Up

# Evolution of “Network” Virtualisation

…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

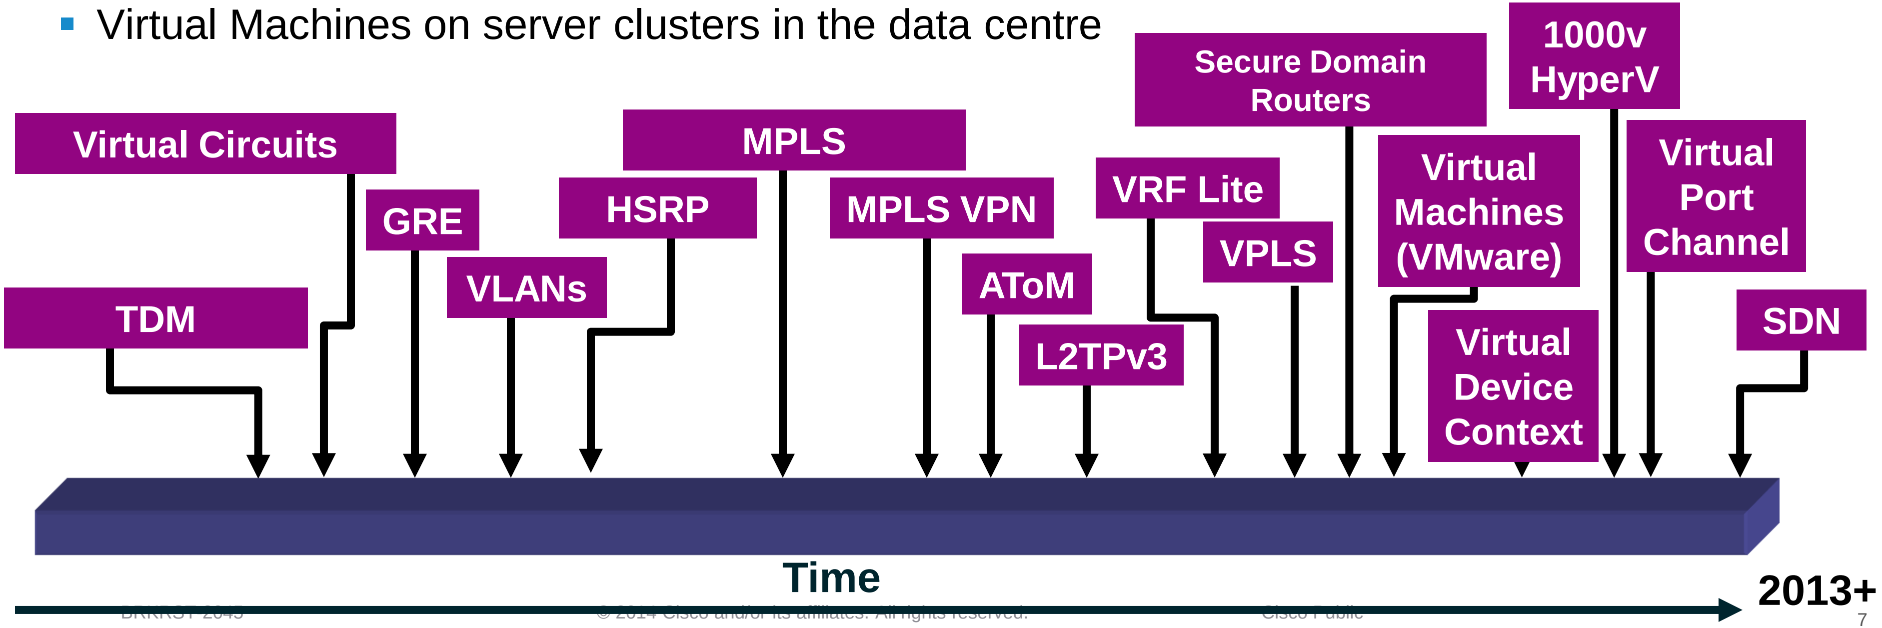

# What Is Enterprise L3 “Network” Virtualisation?

- 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

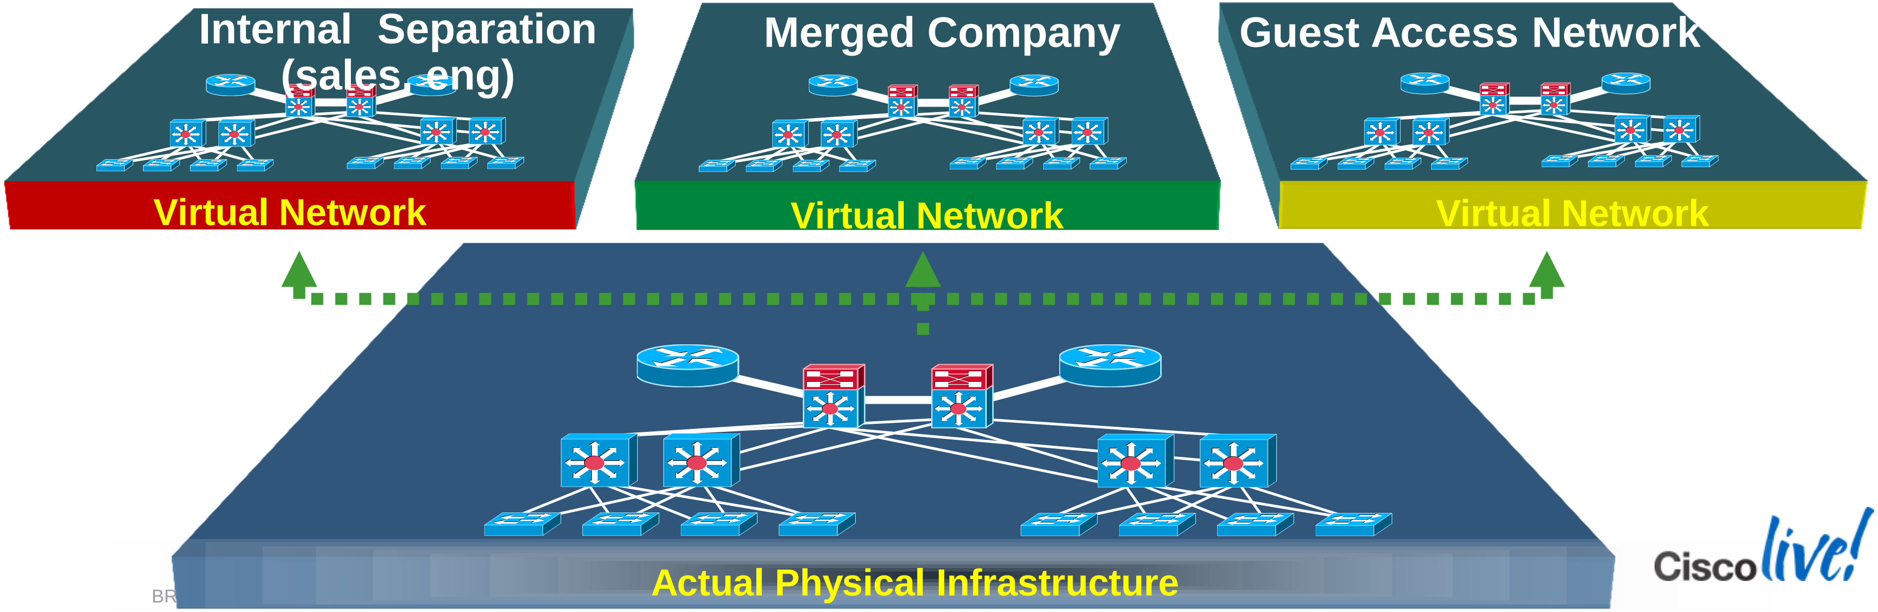

# Why Network Virtualisation?

## Key Benefits

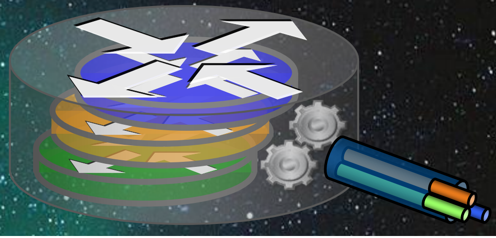

- **Cost Reduction**—allowing a single physical network the ability to offer multiple users and virtual networks
- **Simpler OAM**—reducing the amount of network devices needing to be managed and monitored
- **Security**—maintaining segmentation of the network for different departments over a single device/Campus/WAN
- **High Availability**—leverage virtualisation through clustering devices that appear as one (vastly increased uptime)
- **Data Centre Applications**—require maintained separation, end-to-end (i.e. continuity of Virtualisation from server-to-campus-to-WAN) , including <u>Multi-tenant DC's for Cloud Computing</u>
- **Common Use Cases**
  - Guest Access, Airports, Cloud Computing IaaS, Physical Security Separation, Company Mergers
  - Regulation/Compliance – Health Care (HIPPA), Credit Card (PCI)

# Network Virtualisation Use Cases

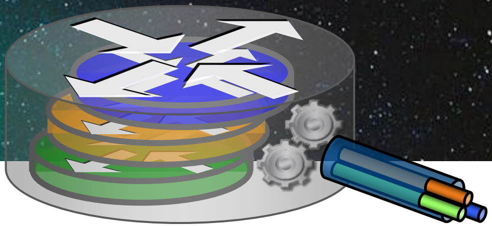

- **Multi-Tenant Dwelling requiring Separation**
  - Airports – airlines (United, Delta, etc…) sharing
  - Government Facilities – Federal agencies sharing single building/campus
  - Intra organisation segmentation – Separation of sales, engineering, HR
  - Company mergers – allowing slow migration for transition, overlapping addressing
  - Data Centre Applications – VM→VLAN→VRF orchestration for segmentation

- **Security**
  - Mandates to logically separate varying levels of security enclaves

- **Regulation requirements**
  - Health Care – HIPPA
  - Financial and Transactional – Sarbanes-Oxley, PCI Compliance

  

# Multi Tenant Cloud and DC

- Add multitenant
- Even if the VRF’s are configured dynamically, or part of the automation process, they are required in multi tenant cloud environments

# Enterprise Network Virtualisation

## Key Building Blocks

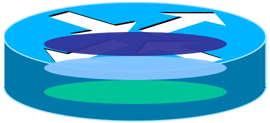

"Virtualising" the Routing and Forwarding of the Device

## Virtualised Interconnect

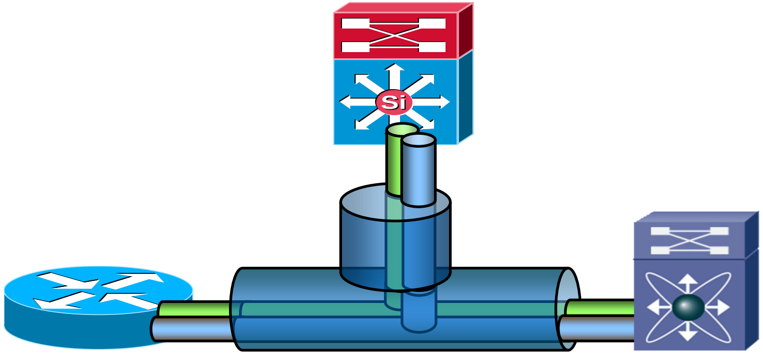


Extending and Maintaining the "Virtualised" Devices/Pools over Any Media

## Device Pooling

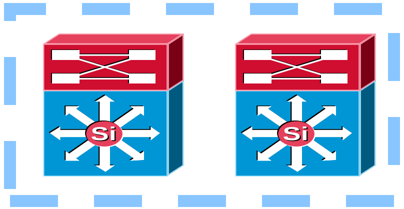


"Virtualising" Multiple Devices to Function as a Single Device

# Enterprise Network Virtualisation

## The Building Blocks – Example Technologies

### Device Partitioning

- **VLANs**
- **VRFs**
- EVN (Easy Virtual Network)
- VDC (Virtual Device Context)
- SDR (Secure Domain Routers)
- FW Contexts
- VASI (VRF Aware Service Int)

### Virtualized Interconnect




- L3 VPNs – MPLS VPNs, GRE, VRF-Lite, MPLS services (L2/L3) over GRE
- L2 VPNs - AToM, Unified I/O, VLAN trunks
- Evolving – TRILL, 802.1ah, 802.1af

### Device Pooling




- Virtual Sw System (VSS)
- Virtual Port Channel (vPC)
- HSRP/GLBP
- Stackwise
- ASR 9000v/nV Clustering
- Inter-Chassis Control Protocol (ICCP)

# Enterprise Network Virtualisation

## The Building Blocks – Example Technologies

### Device Partitioning



VLANs

VRFs
VDCs

SDR (XR)

FW Contexts

### Virtualized Interconnect



L3 VPNs – MPLS VPNs, GRE, VRF-Lite, MPLS services (L2/L3) over GRE

L2 VPNs - AToM, Unified I/O, VLAN trunks

### Device Pooling

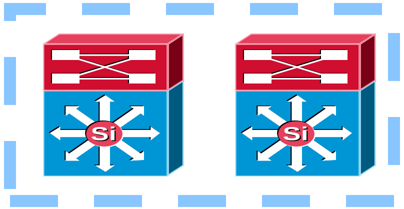

Virtual Sw System (VSS)

Virtual Port Channel (vPC)

HSRP/GLBP

Stackwise

ASR 9000v/nV Clustering

Inter-Chassis Control Protocol (ICCP)

# Enterprise Network Virtualisation over the WAN

## The Building Blocks – Example Technologies

### Device Partitioning

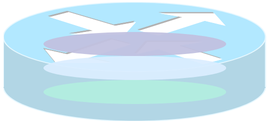

VLANs

VRFs

EVN

(Easy Virtual Network)

VDC (NX-OS)

(Virtual Device Context)

SDR (IOS-XR)

(Secure Domain Routers)

FW Contexts

### Virtualised Interconnect

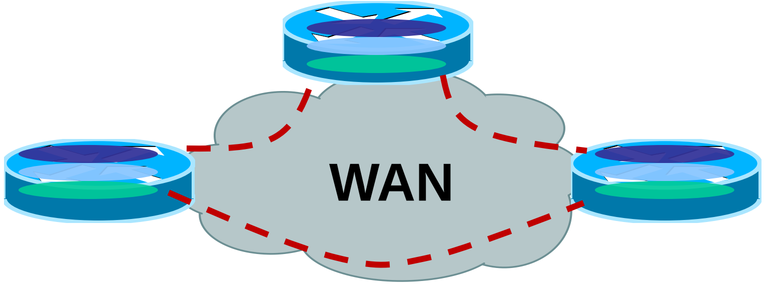


**<u>L3 VPNs</u> – MPLS VPNs, MPLS VPN over GRE/mGRE, VRF-Lite, VRF-Lite over IP, LISP Multi-tenant**

<u>L2 VPNs</u> –PWE3, VPLS, L2 VPN over IP, L2TPv3, OTV (Overlay Transport Virtualisation), FabricPath/L2MP

<u>Evolving Standards</u> – TRILL, Fat-PW, MPLS-TP, PBB/E-VPN, VxLAN, NVGRE

### Device Pooling

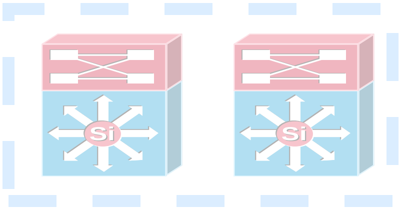


Virtual Sw System (VSS)

Virtual Port Channel (vPC)

HSRP/GLBP

Stackwise

ASR 9000v/nV Clustering

Inter-Chassis Control Protocol (ICCP)

# Agenda

- Introduction - Network Virtualisation Drivers and Concepts
- **SP WAN Transport Service Impact on L3 Virtualisation Solution Choices**
- Technology and Deployment Deep-Dive for L3 Virtualised WAN
- Integrating Encryption into L3 Virtualisation Solutions
- Recent “Innovations” Evolving in L3 Virtualisation
- QoS/H-QoS and MTU Considerations for L3 Virtualisation Deployments
- Summary and Wrap Up

# Today's WAN Transport Options

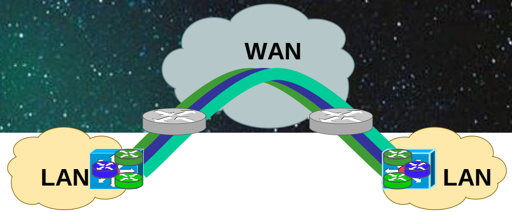


## Topologies

- Point-point, multi-point
- Full/partial mesh
- Hub/Spoke or Multi-Tier

## Media

- Serial, ATM/FR, OC-x
- Dark fibre, Lambda
- Ethernet

## SP VPN Offerings

- L2 – Ethernet (p2p, p2mp)
- L3 – Private IP VPN

## SP Transport

- L3 – Public (Internet)
- L3 – Broadband/WiFi/3G/4G

# Primary Transport Options Utilised in Enterprises
(With/Without L3 WAN Virtualisation)

## 1 SP Managed “Ethernet” Service

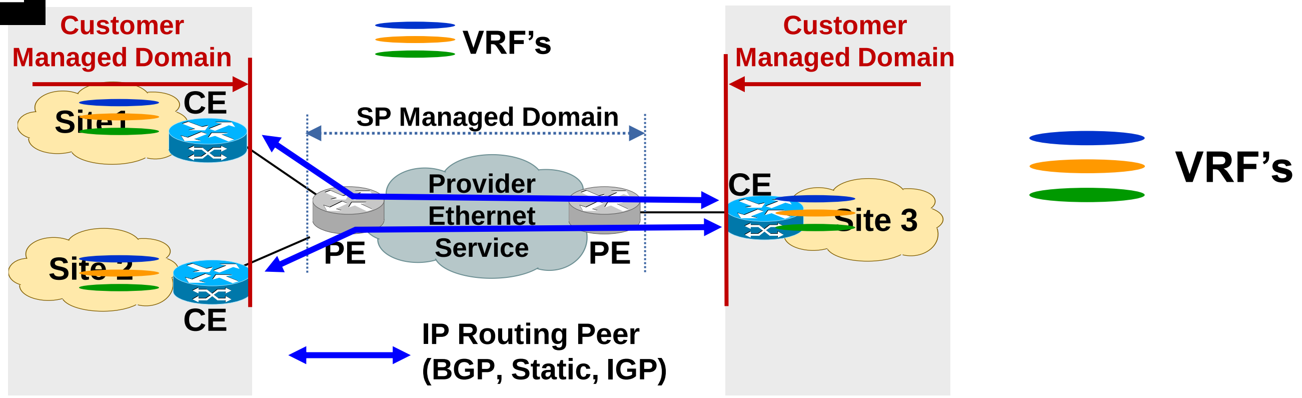


## 2 SP Managed “IP VPN” Service

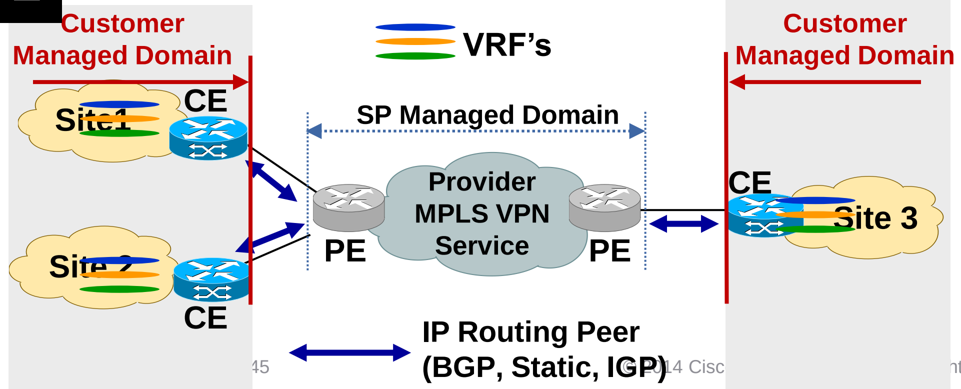


## 3 Self Deployed IP/MPLS Backbone

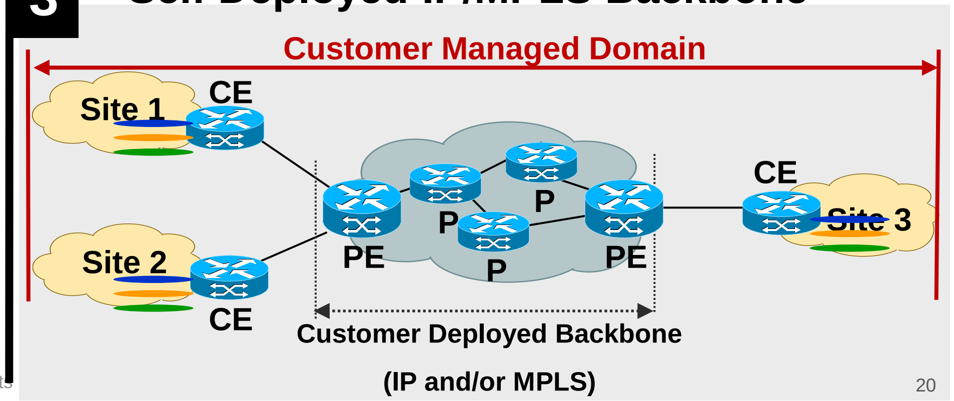

# Ethernet Virtual Connection (EVC)

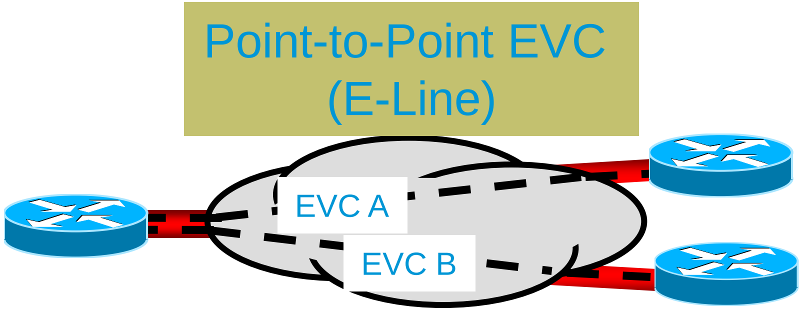


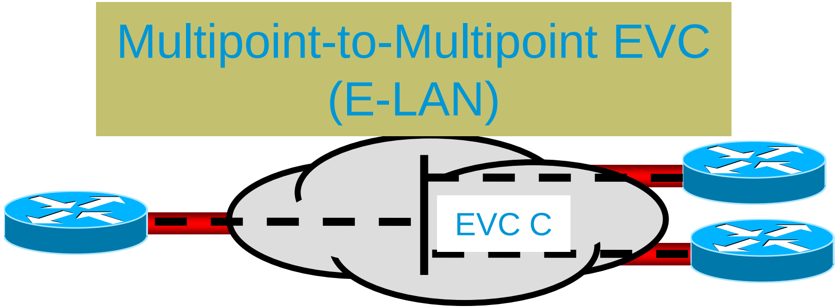


- Metro Ethernet Forum (EF) Service Types:
  - E-LINE: associated to an Point-to-Point EVC
  - E-LAN: associated to an Multipoint EVC

# Deployment Options - Self Deployed vs. SP Managed

## SP Offered Ethernet Service (Layer 2 Service) - Customer owns CE

- **CE Routers owned by customer**
- **PE Routers owned by SP**
- Customer leverages E-LINE/E-LAN Ethernet Service
  - VLAN or port-mode
  - Point-to-point (PW), multi-point (VPLS)
- Routing controlled and managed by end customer
- SP service offers P2P or P2MP service transport
  - Other P2P options include:  T1/E1, T3/E3, ATM/FR PVC, OC-x, CH OC-x

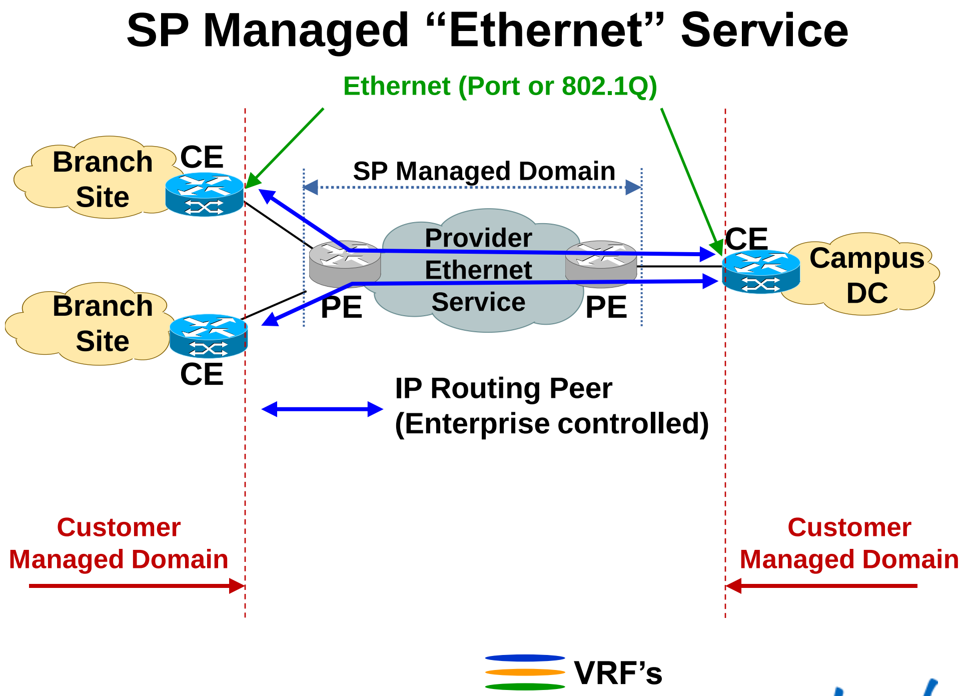

# Deployment Options - Self Deployed vs. SP Managed

SP Offered IP VPN Service (Layer 3 Service) - Customer owns CE

## 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

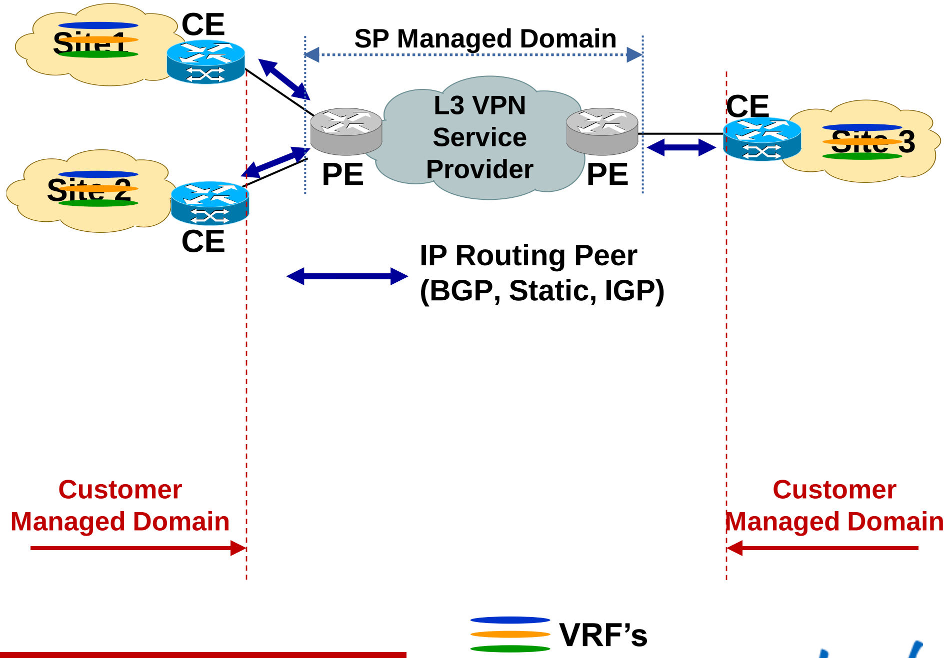


* **No** Labels Are Exchanged with the SP

# Deployment Options - Self Deployed vs. SP Managed

## Customer Deploys Their Own Internal MPLS VPN Network – Controls E2E

- 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**

### Self Deployed IP Backbone

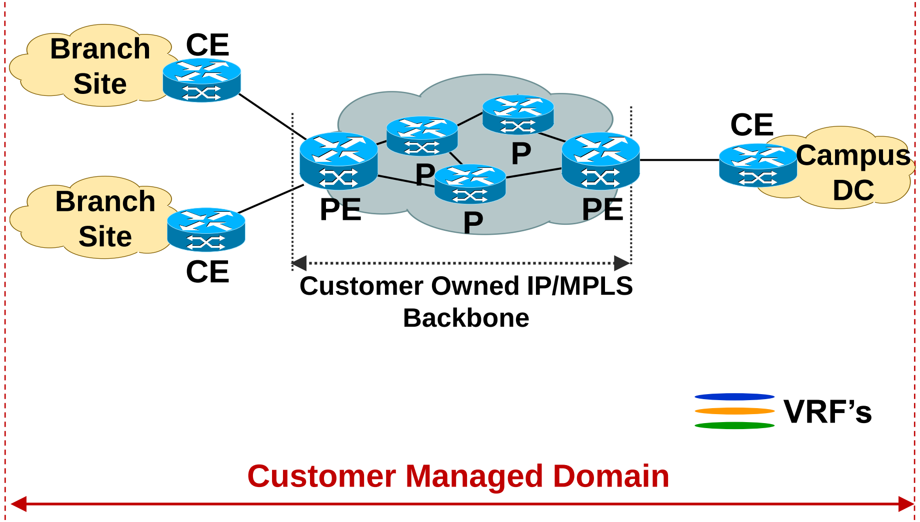

# Primary Transport Options Utilised in Enterprises

(With/Without L3 WAN Virtualisation)

## 1 Self Deployed IP/MPLS Backbone

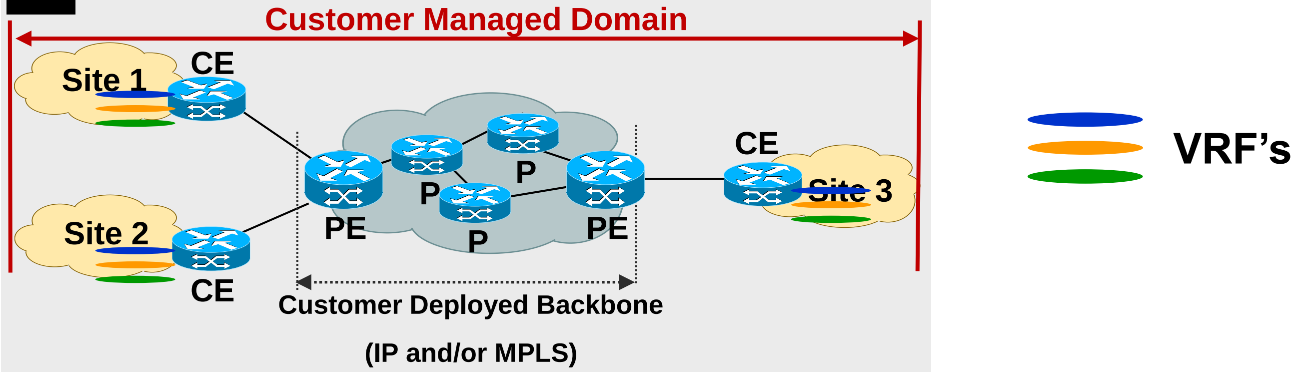


## 2 SP Managed "Ethernet" Service

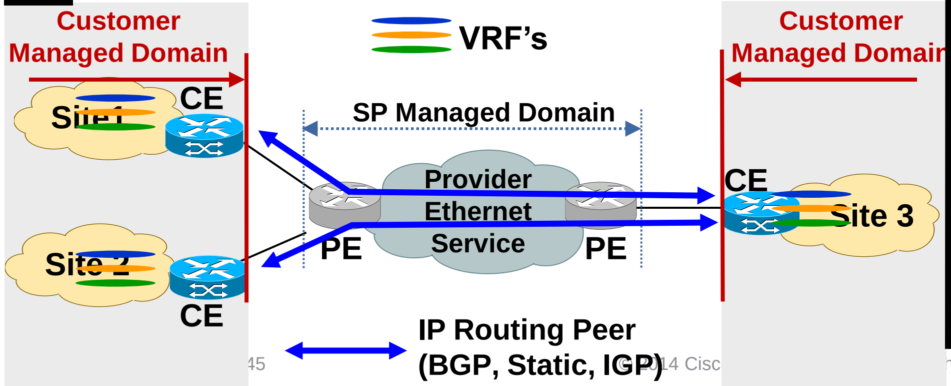


## 3 SP Managed "IP VPN" Service

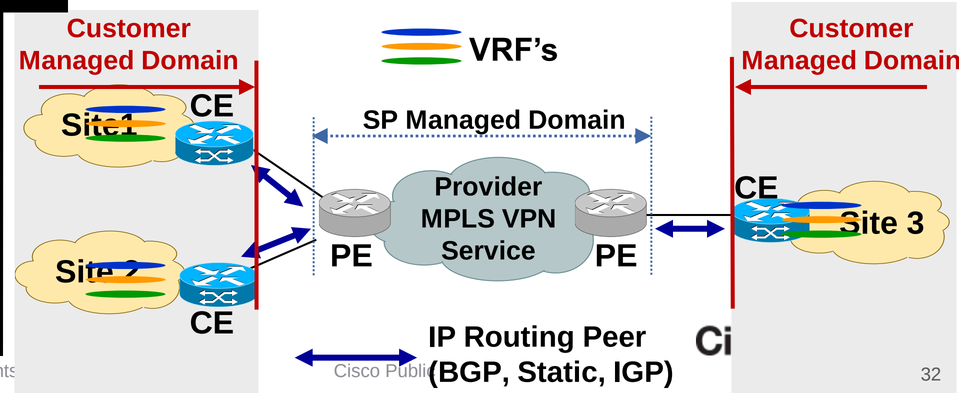

# Agenda

- Introduction - Network Virtualisation Drivers and Concepts
- SP WAN Transport Service Impact on L3 Virtualisation Solution Choices
- Technology and Deployment Deep-Dive for L3 Virtualised WAN
  - **MPLS VPN “101” for Self Deployed Solution**
  - L3 Virtualisation Solution Options over L2 Services
  - L3 Virtualisation Solution Options over L3 Services
- Integrating Encryption into L3 Virtualisation Solutions
- Recent “Innovations” Evolving in L3 Virtualisation
- QoS/H-QoS and MTU Considerations for L3 Virtualisation Deployments
- Summary and Wrap Up




# L3 Virtualisation using VRF-Lite and MPLS BGP VPN (RFC 4364)

# MPLS: Large Scale Virtualisation Enabler in the WAN

## Allows Vast Network “Virtualisation” Capabilities over WAN

- **Layer 3 VPN/Segmentation**
  - VPN (RFC 4364)
  - Provides Any-to-Any connectivity
- Maximise Link Utilisation with Selective Routing/Path Manipulation
  - Traffic Engineering
  - Optimisation of bandwidth and protection using Fast-ReRoute (FRR)
- **Layer 2 VPN/Transport**
  - AToM (Any Transport over MPLS) i.e. “pseudo-wire”
  - Layer-2 transport: Ethernet, ATM/FR, HDLC/PPP, interworking
  - Layer-2 VPN: VPLS for bridged L2 domains over MPLS
- QoS Capabilities
  - Diffserv, Diffserv aware Traffic Engineering (DS-TE)
- Bandwidth Protection Services
  - Combination of TE, Diffserv, DS-TE, and FRR
- IP Multicast (per VPN/VRF)
- Transport of IPv6 over an IPv4 (Global Routing Table)
- Unified Control Plane (Generalised MPLS)

**Key Virtualisation Mechanisms over an IP Infrastructure**

- **L3 VPN**
- **L2 VPN (P2P)**
- **L2 VPN (P2MP)**
- **QoS**
- **IPv6, MVPN**

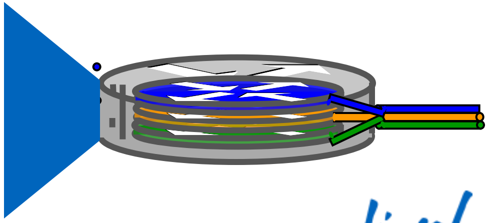

# What is a VRF?

## Virtual Routing and Forwarding Instance

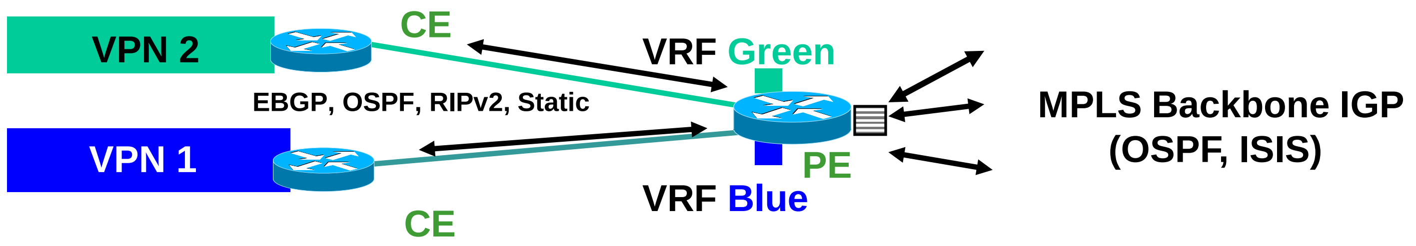


- Associates to one or more interfaces on PE
  - Privatise an interface i.e., colour the interface
- Has its own routing table and forwarding table (CEF)
- VRF has its own instance for the routing protocol
  - (static, RIP, BGP, EIGRP, OSPF)
- CE router runs standard routing software
- Allows overlapping address space

# MPLS Label Encapsulations

## Applicable When Using MPLS over Layer 2 Transport

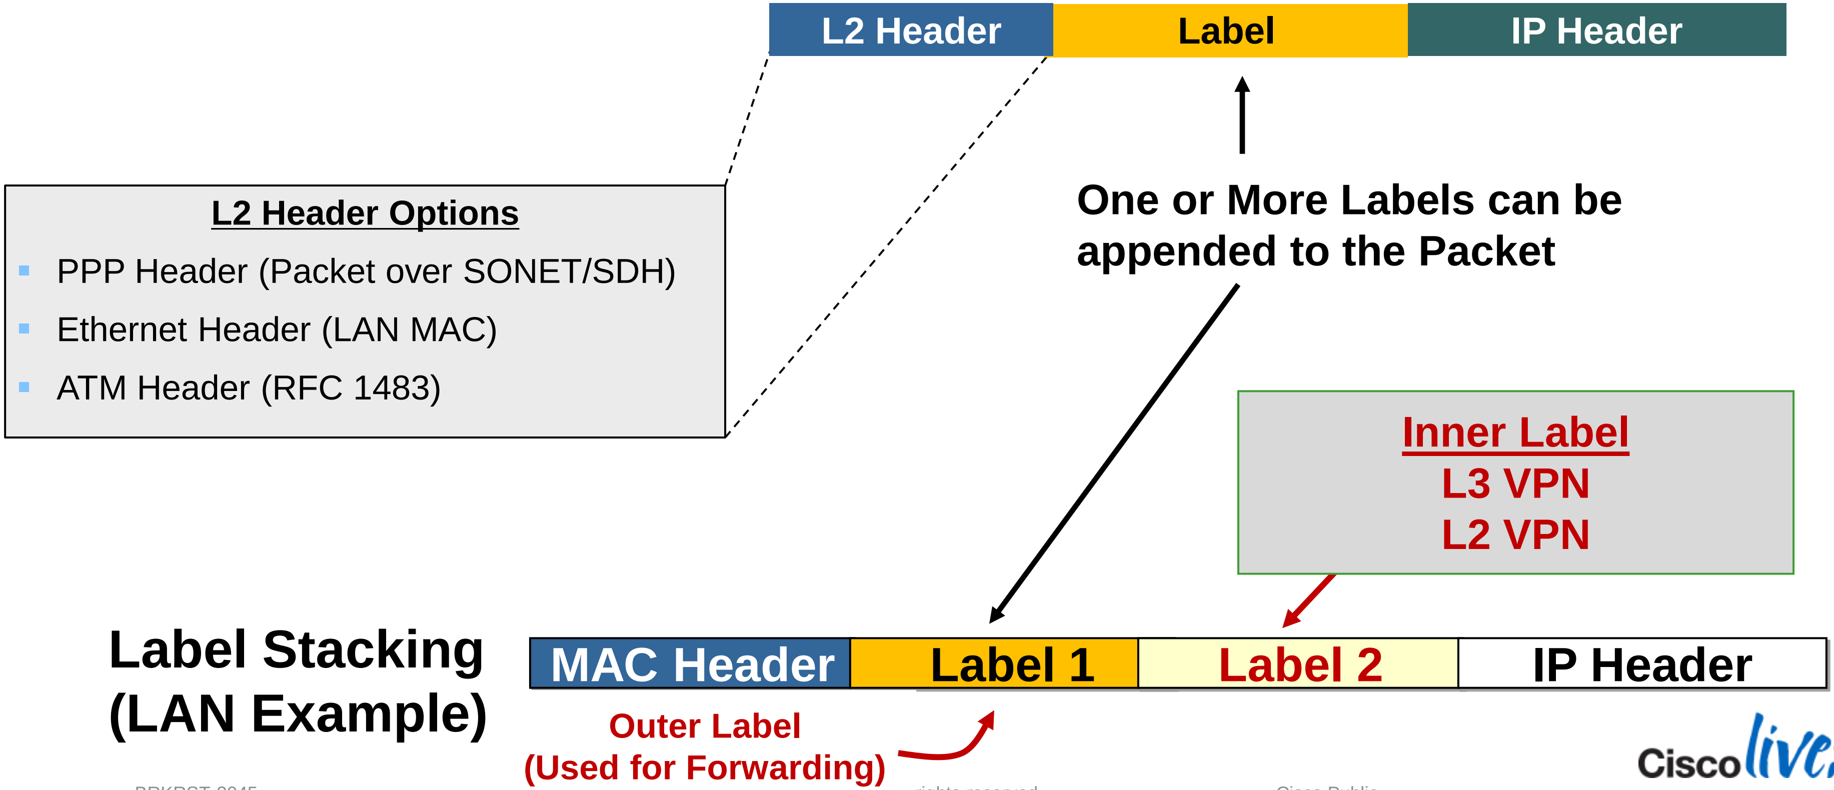

# MPLS VPN Technology—Refresher

## Control Plane

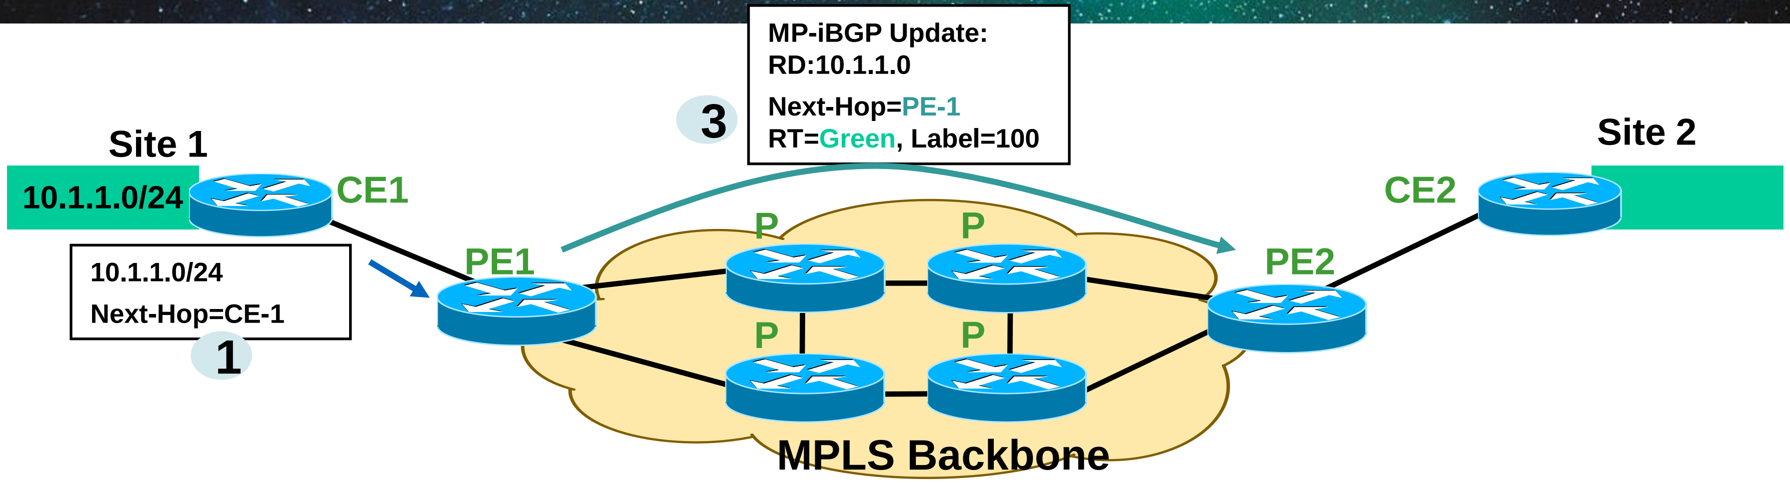


1. PE1 receives an IPv4 update (eBGP/OSPF/ISIS/RIP/EIGRP)
2. PE1 translates it into VPNv4 address
   - Assigns an RT per VRF configuration
   - Rewrites next-hop attribute to itself
   - Assigns a label based on VRF and/or interface
3. PE1 sends MP-iBGP update to other PE routers

# MPLS VPN Technology—Refresher

## Control Plane

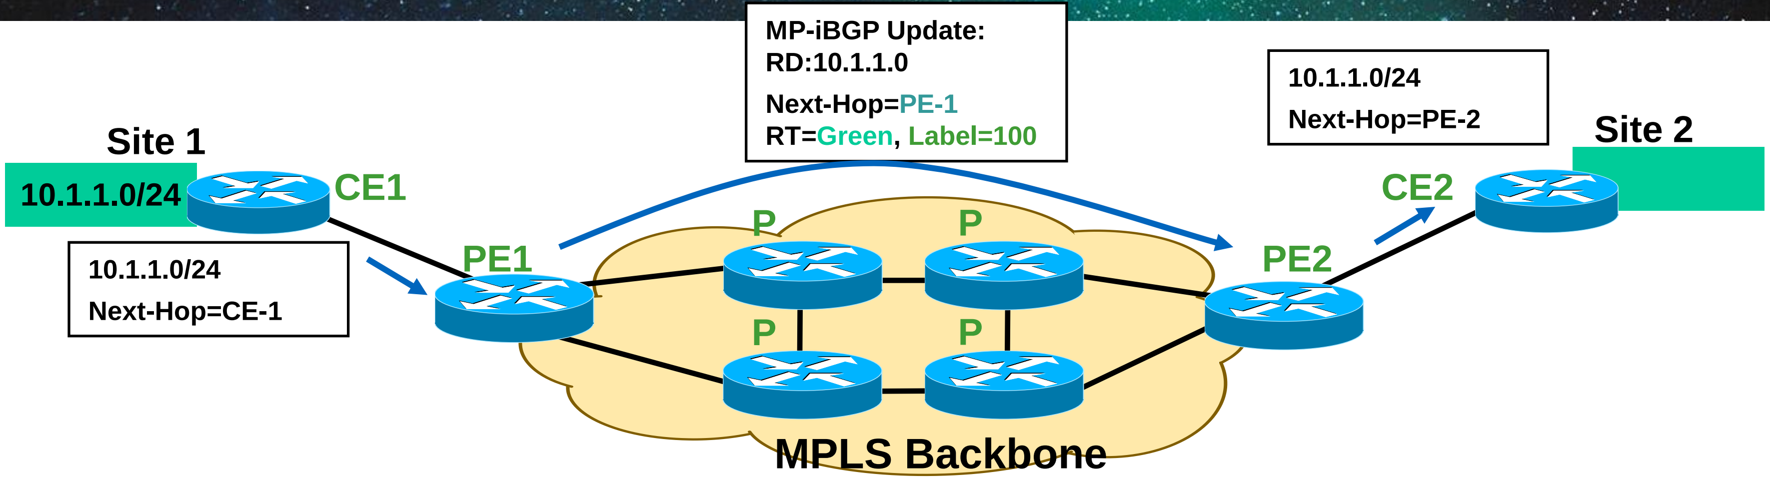


4. PE2 receives and checks whether the RT=green (40:103, say) is locally configured within any VRF, if yes, then
5. PE2 translates VPNv4 prefix back into IPv4 prefix,
   - Installs the prefix into the VRF routing table
   - Updates the VRF CEF table with label=100 for 10.1.1.0/24
   - Advertise this IPv4 prefix to CE2 (using EBGP/RIP/OSPF/ISIS/EIGRP)

# MPLS VPN Technology—Refresher

## Forwarding Plane

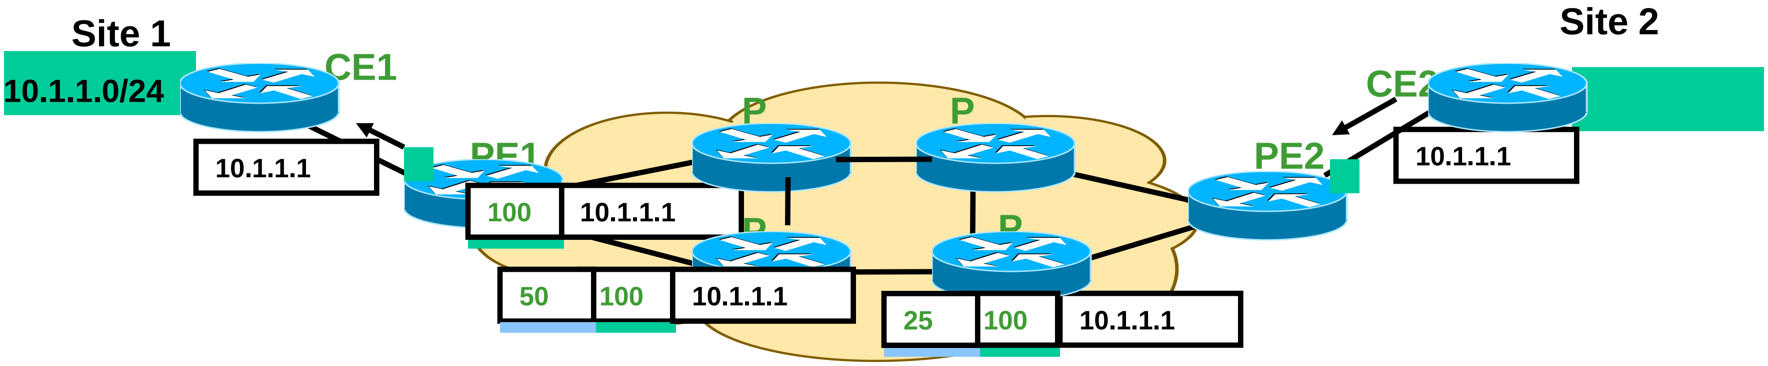


- PE2 imposes TWO labels for each packet going to the VPN destination 10.1.1.1
- The top label is LDP learned and derived from an IGP route
  - Represents LSP to PE address (exit point of a VPN route)
- The second label is learned via MP-BGP
  - Corresponds to the VPN address

# Self Deployed MPLS VPN

## Configuration Example (IOS)

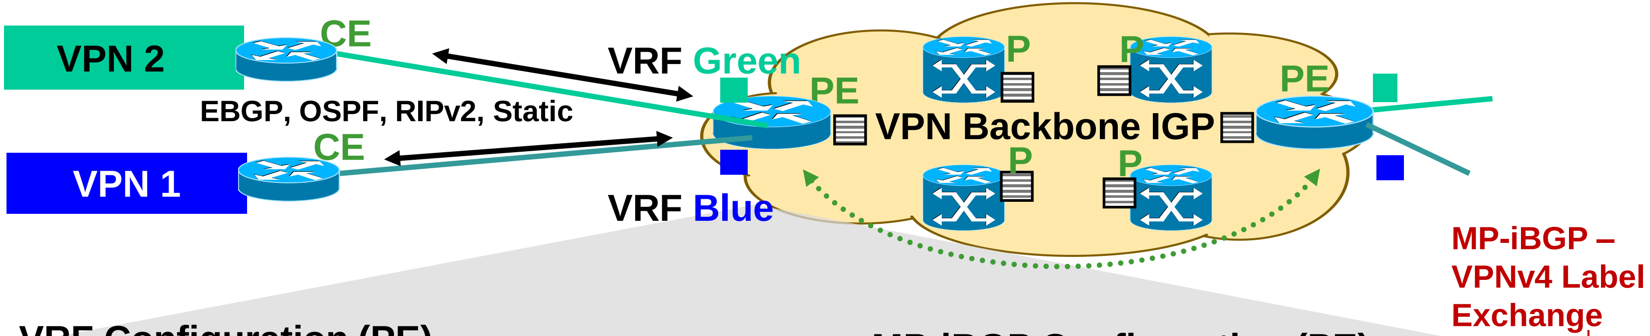


**VRF Configuration (PE)**

```
! PE Router – Multiple VRFs
ip vrf blue
 rd 65100:10
 route-target import 65100:10
 route-target export 65100:10
ip vrf green
 rd 65100:20
 route-target import 65100:20
 route-target export 65100:20
!
interface GigabitEthernet0/1
 ip vrf forwarding blue
interface GigabitEthernet0/2
 ip vrf forwarding green
```

**MP-iBGP Configuration (PE)**

```
! PE router
router bgp 65100
 neighbor 192.168.100.4 remote-as 65100
 !
 address-family vpnv4
 neighbor 192.168.100.4 activate
 neighbor 192.168.100.4 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf blue
 neighbor 172.20.10.1 remote-as 65111
 neighbor 172.20.10.1 activate
 exit-address-family
```

# Self deployed MPLS VPN – End to End Control

## Summary and Deployment Targets

- Targets large-scale VRF's and customers wanting control!
- Leverages standard based L2 transports (no overlay) in the WAN (ATM, SONET/SDH, Ethernet, dark fibre/lambda's)
- Target customers usually function as an "internal Service Provider" for their company/agency
- Allows full deployment of MPLS services
  - L2 VPN (PW, VPLS), QoS, Multicast/mVPN, IPv6, MPLS TE, TE-FRR
- Offers tight control for QoS Service Level requirements
- Offers rapid deployment for Virtualisation "turn up"
- Massively scalable but does require a higher level of Operational expertise

# What is VRF Lite?

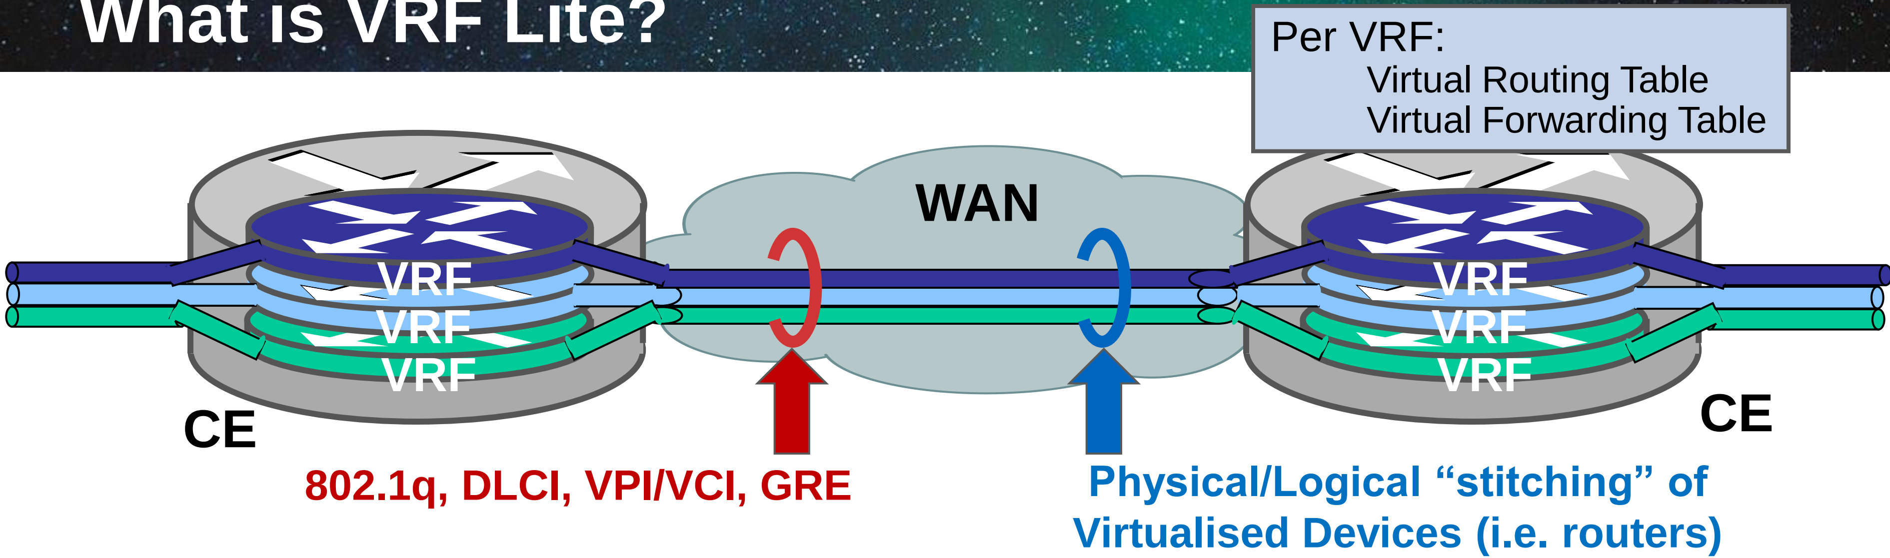


- Defines router supports **routing (RIB), forwarding (FIB), and interface per VRF !!**
- Leverages "Virtual" encapsulation for separation:
  - ATM VCs, Frame Relay, Ethernet/802.1Q
- The routing protocol is also "VRF aware"
  - EIGRP, OSPF, BGP, RIP/v2, static (per VFR)
- Layer 3 VRF interfaces cannot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

# MPLS VPN + VRF-Lite Technology

## MPLS VPN Using Separate "CE" router per VRF

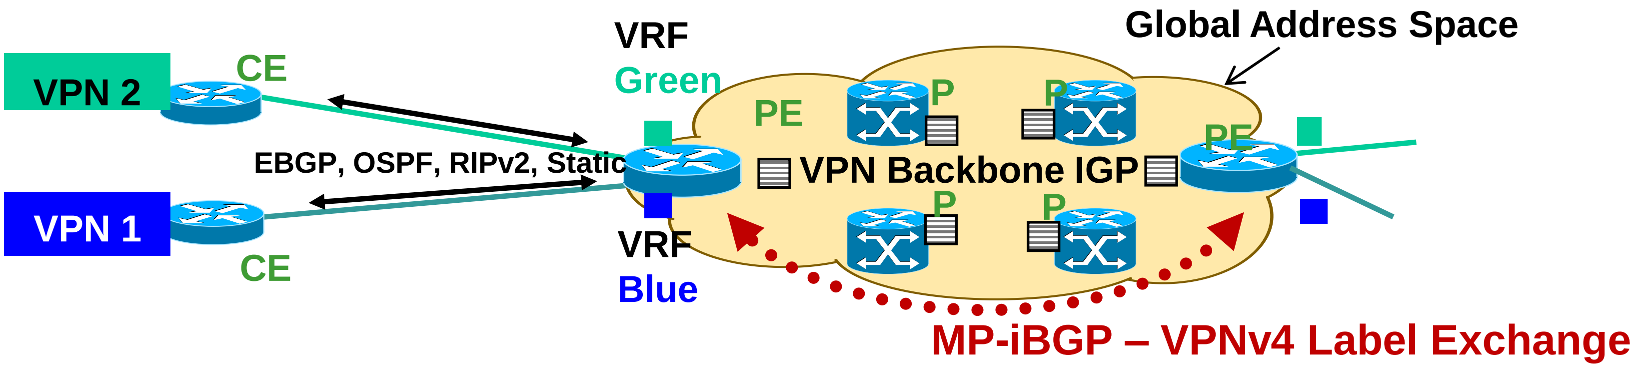

# MPLS VPN + VRF-Lite Technology

## Combining MPLS L3 VPN + VRF-Lite (PE-CE)

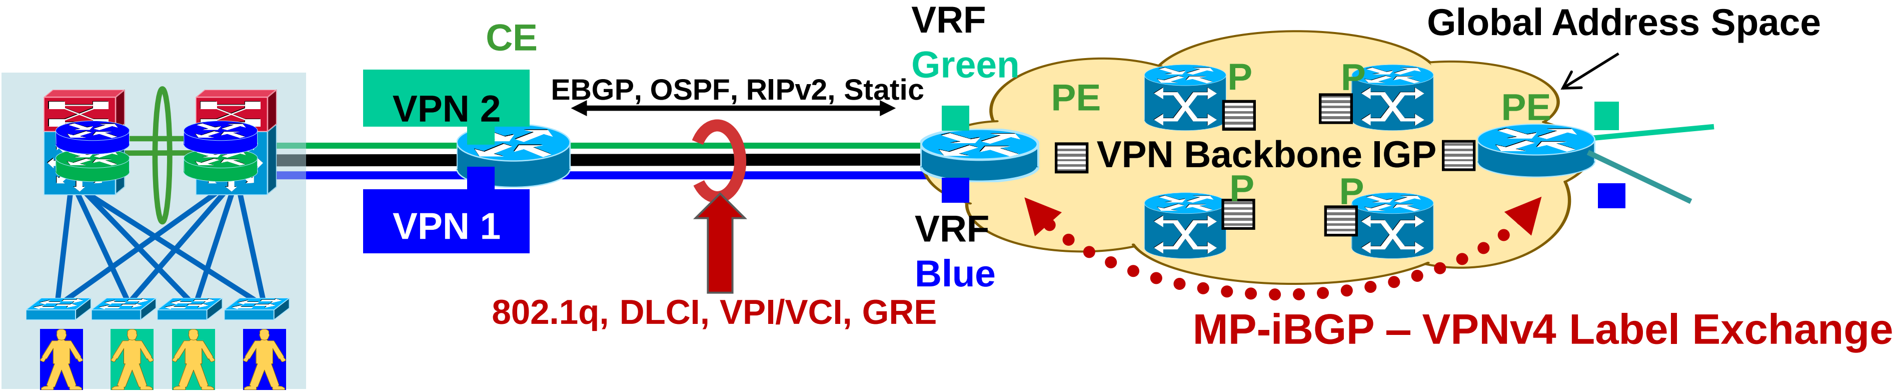


- MPLS VPN backbone remains the same
- Leverage VRF-Lite CE to PE
- CE to PE can be "local" (fibre, copper) or remote (WAN, Metro service)
- If WAN, transport will dictate technology chosen for CE – PE
  - Ethernet service (802.1Q), WAN (DLCI), IP WAN (GRE)
- VRF has its own instance for the routing protocol
  - (static, RIP, BGP, EIGRP, OSPF)

# Agenda

- Introduction - Network Virtualisation Drivers and Concepts
- SP WAN Transport Service Impact on L3 Virtualisation Solution Choices
- Technology and Deployment Deep-Dive for L3 Virtualised WAN
  - MPLS VPN “101” for Self Deployed Solution
  - **L3 Virtualisation Solution Options over L2 Services**
  - L3 Virtualisation Solution Options over L3 Services
- Integrating Encryption into L3 Virtualisation Solutions
- Recent “Innovations” Evolving in L3 Virtualisation
- QoS/H-QoS and MTU Considerations for L3 Virtualisation Deployments
- Summary and Wrap Up




# L3 Virtualisation over Ethernet Transport Services

# Deployment Options - Self Deployed vs. SP Managed

## SP Offered Ethernet Service (Layer 2 Service) - Customer owns CE

- **CE Routers owned by customer**
- **PE Routers owned by SP**
- Customer leverages E-LINE/E-LAN Ethernet Service
  - VLAN or port-mode
  - Point-to-point (PW), multi-point (VPLS)
- Routing controlled and managed by end customer
- SP service offers P2P or P2MP service transport
  - Other P2P options include: T1/E1, T3/E3, ATM/FR PVC, OC-x, CH OC-x

**SP Managed "Ethernet" Service**

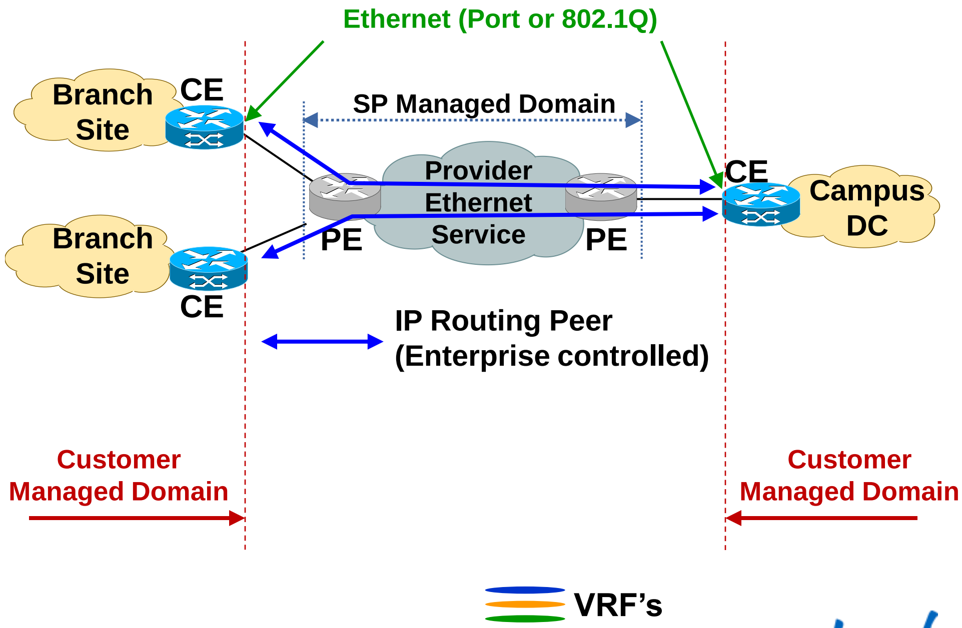

# Ethernet Virtual Connection (EVC)

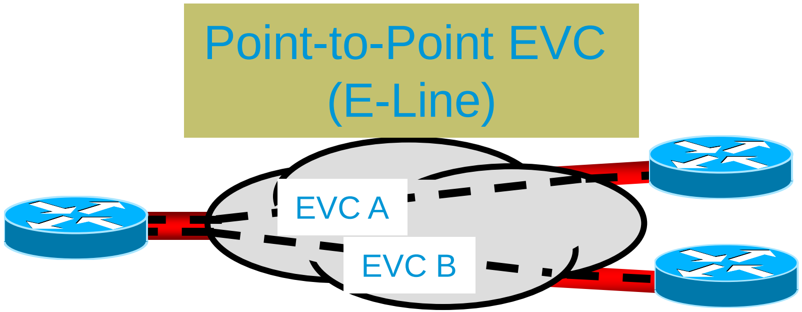


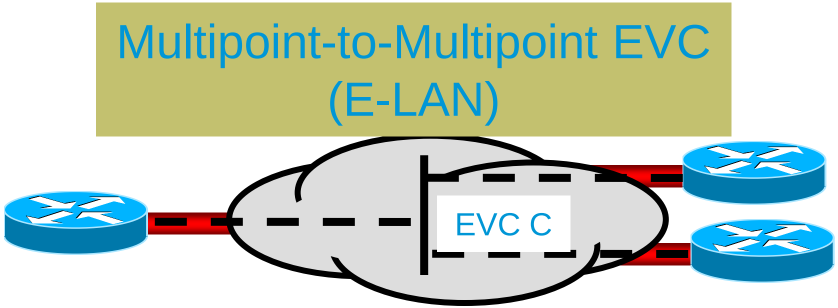


- Metro Ethernet Forum (EF) Service Types:
  - E-LINE: associated to an Point-to-Point EVC
  - E-LAN: associated to an Multipoint EVC

# Deployment Options - Self Deployed vs. SP Managed

SP Offered Ethernet Service (Layer 2 Service) - Customer owns CE

## SP Managed “Ethernet” Service

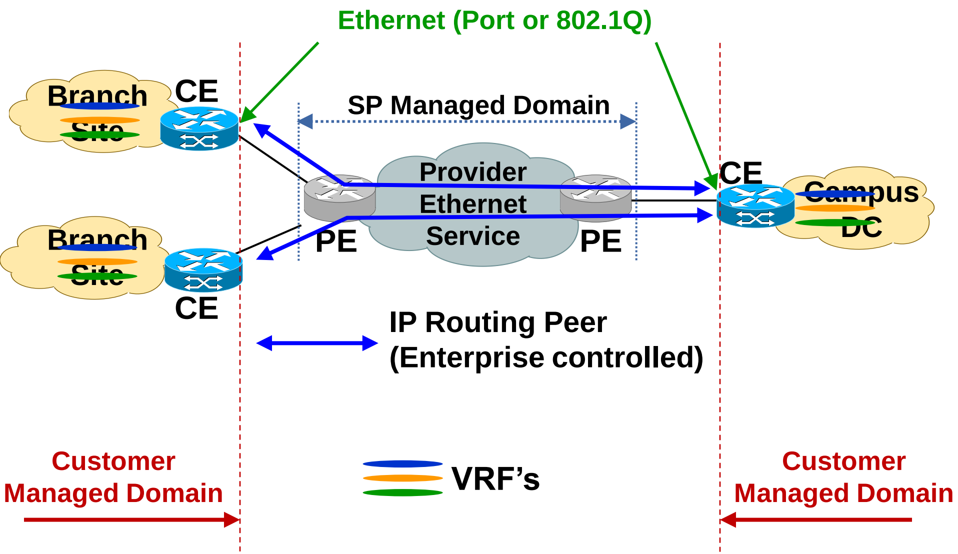


- Add text

# Deployment Options - Self Deployed vs. SP Managed

## SP Offered Ethernet Service (Layer 2 Service) - Customer owns CE

### SP Managed “Ethernet” Service

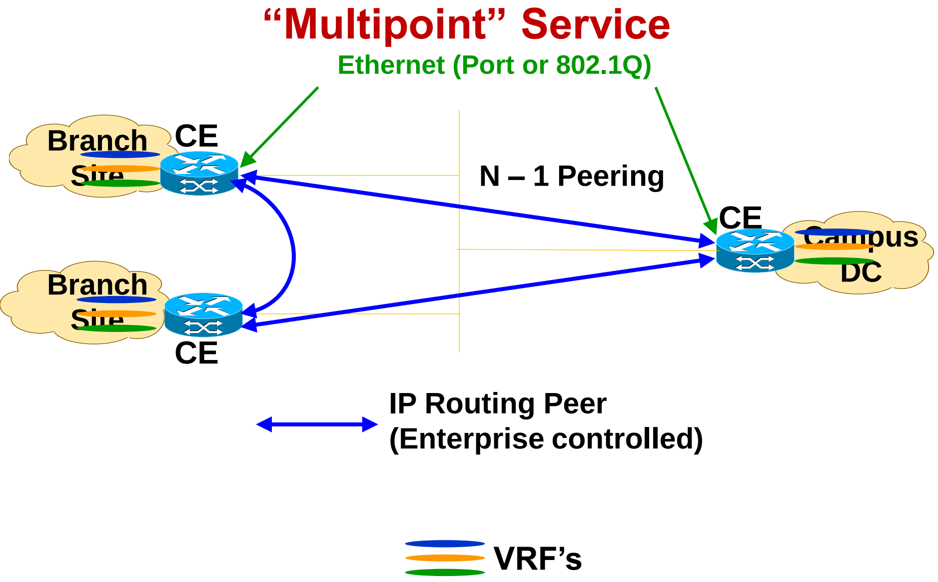


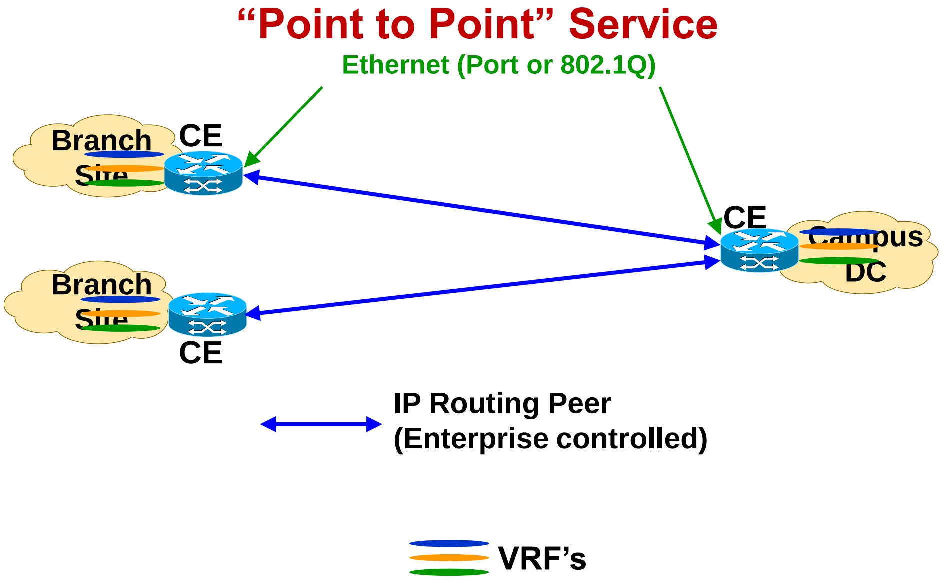


- Add text

# VRF-Lite over Layer 2 Transport (Point to Point Ethernet Example)

Extend Virtualisation over WAN L2 Service

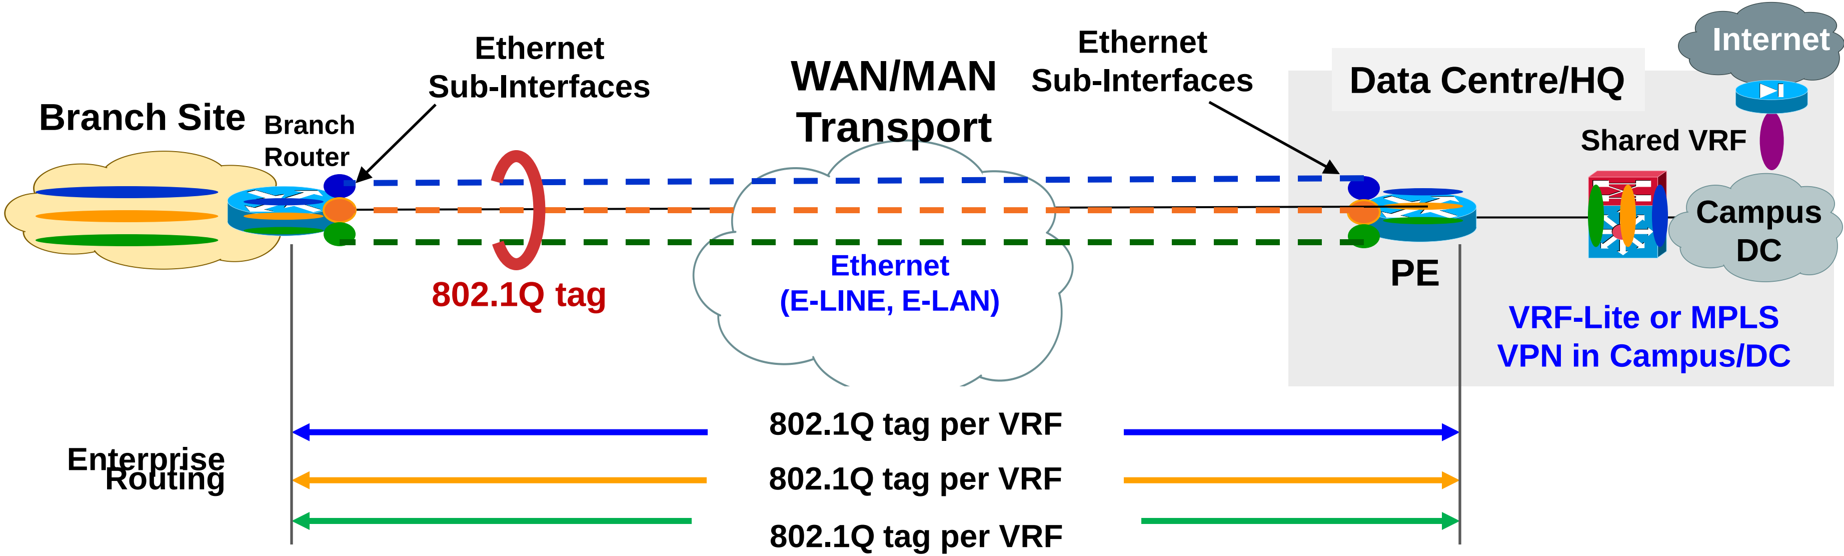


- Each Ethernet interface (or serial) leverages a sub-interface
- Unique DLCI (frame relay) or 802.1Q tag (Ethernet) per VRF
- IGP process created per VRF in both Branch/Campus
- Offers virtualised segmentation within a single “physical” interface

# VRF-Lite Options over an Ethernet Transport

## 802.1Q or Q-in-Q + MPLS Option

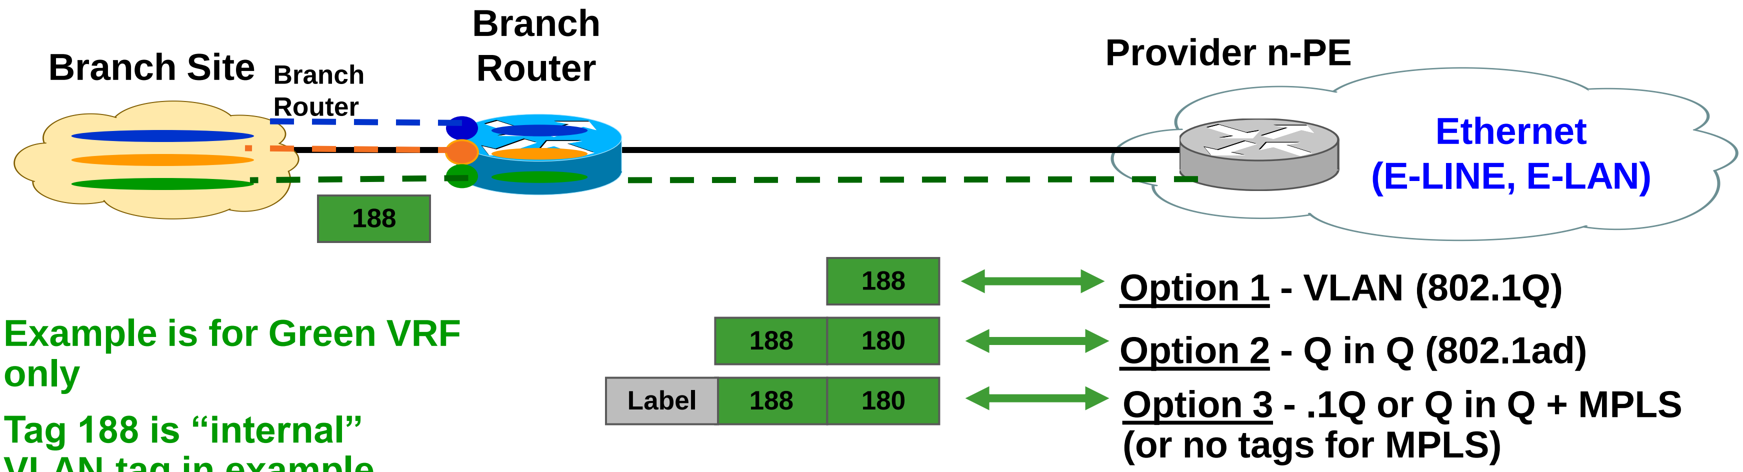


- Example is for Green VRF only
- Tag 188 is “internal” VLAN tag in example

- **Option 1** – configure single sub-interface/802.1Q tag to provider
- **Option 2** – leverage Q-in-Q (802.1ad) to send required tag to SP, but hide customers tag
- **Option 3** – Leverage Option 1 or 2 + MPLS (MPLS could be run without the need for .1Q tags, or any Q-in-Q)

# VRF-Lite over Layer 2 Transport

Example:  Ethernet Service (Point to Point)

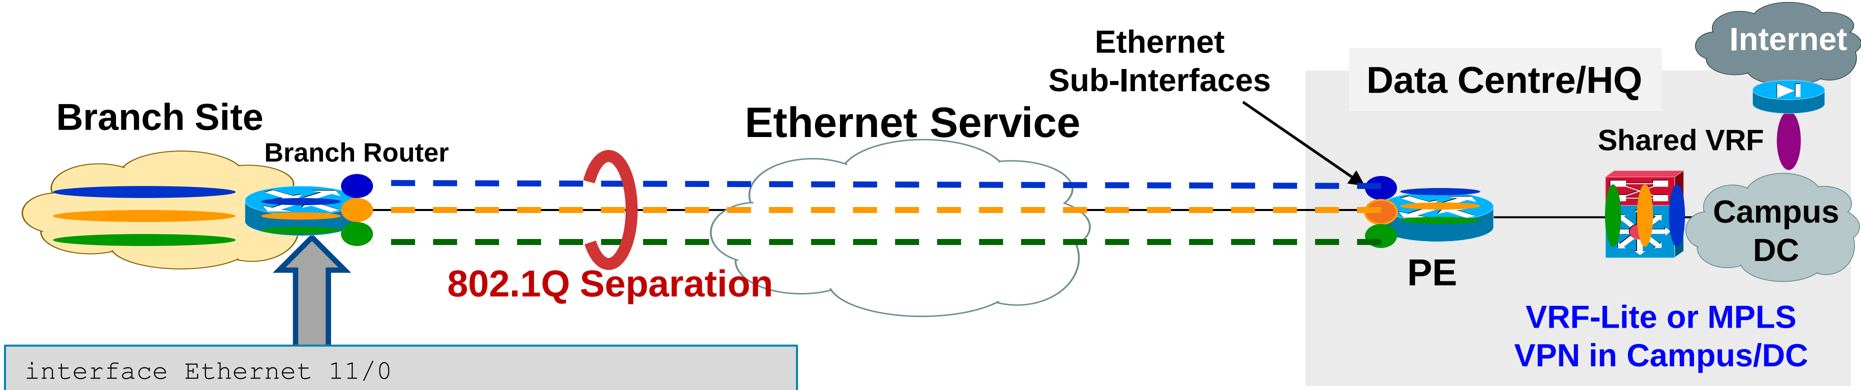


**Option exists for .1Q or "Q-in-Q" on sub-interface**

```
interface Ethernet 11/0
!
interface Serial 11/0.1
 encapsulation dot1Q 100 second-dot1q 1000
 ip vrf forwarding blue
 ip address 192.168.51.2 255.255.255.252
!
interface Ethernet 11/0.2
 encapsulation dot1Q 200
 ip vrf forwarding green
 ip address 192.168.61.2 255.255.255.252
!
interface Ethernet 11/0.3
 encapsulation dot1Q 300
 ip vrf forwarding yellow
 ip address 192.168.71.2 255.255.255.252
```

```
router ospf 10 vrf blue
 log-adjacency-changes
 network 192.168.51.0 0.0.0.255 area 0
!
router ospf 20 vrf green
 log-adjacency-changes
 network 192.168.61.0 0.0.0.255 area 0
!
router ospf 100 vrf yellow
 log-adjacency-changes
 redistribute bgp 65000 subnets
 network 192.168.71.0 0.0.0.255 area 0
```

# VRF-Lite over Layer 2 Transport

## Example: T1/T3, OC-x, FR/ATM (Point to Point)

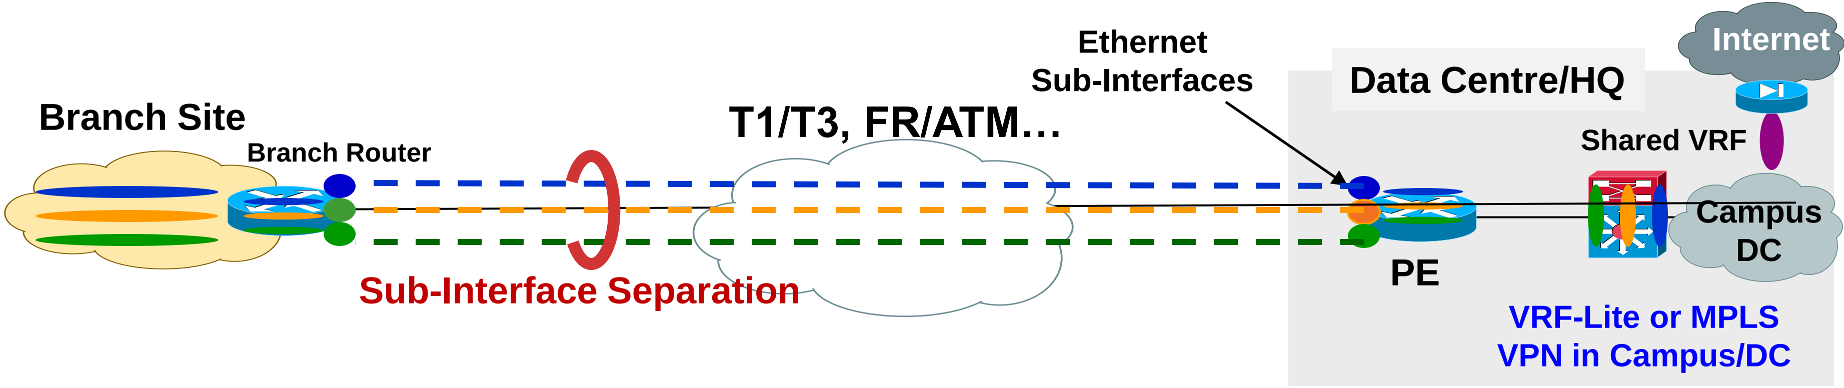


**Configuration Note:**

- Frame Relay encapsulation can be used to virtualise a leased line (E1/T1, E3/T3,…)
- Enabling Frame Relay encap allows the use of sub-interfaces
- Then VRF forwarding can be enabled per sub-interface
- **Allows VRF-Lite over leased-line**

# VRF-Lite Considerations in WAN Deployments

## Is VRF-Lite the Best Fit for My Network?

**Example: 4 Sites with 4 VRFs**

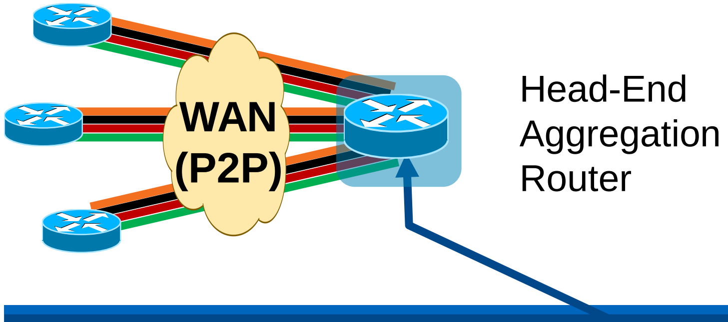


| Virtual Networks | Neighbours | VRF Sub-interfaces |
|---|---|---|
| 4 | 3 | 12 |
| 10 | 3 | 30 |
| 20 | 3 | 60 |
| 30 | 3 | 90 |

**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 and/or locations required?
- How much (locations) will the network grow?
- Does my team have the expertise to manage a more complex MPLS VPN network, if that is the best solution?

# Design Considerations in WAN Deployments

## VRF-Lite vs. MPLS BGP VPN (RFC 4364)

**Example: 4 Sites with 4 VRFs**

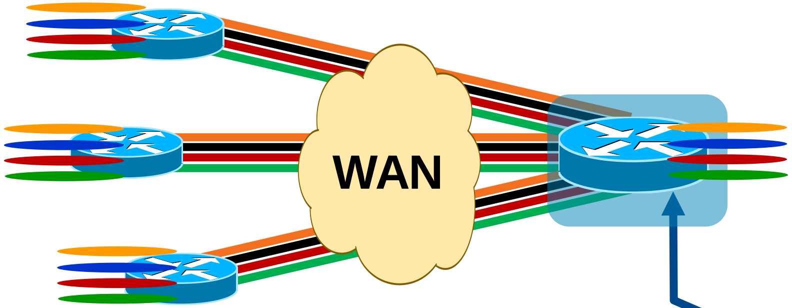


| VRFs | Neighbours | GRE Tunnels (1 per VRF) |
|---|---|---|
| 4 | 3 | 12 |
| 10 | 3 | 30 |
| 20 | 3 | 60 |
| 30 | 3 | 90 |

**Example: 4 Sites with 4 VRFs**

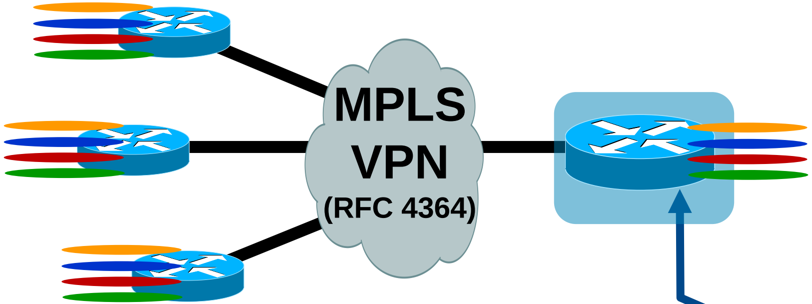


| VRFs | Neighbours | Interfaces to the WAN |
|---|---|---|
| 4 | 3 | 1 |
| 10 | 3 | 1 |
| 20 | 3 | 1 |
| 30 | 3 | 1 |

# Agenda

- Introduction - Network Virtualisation Drivers and Concepts
- SP WAN Transport Service Impact on L3 Virtualisation Solution Choices
- Technology and Deployment Deep-Dive for L3 Virtualised WAN
  - MPLS VPN “101” for Self Deployed Solution
  - L3 Virtualisation Solution Options over L2 Services
  - **L3 Virtualisation Solution Options over L3 Services**
- Integrating Encryption into L3 Virtualisation Solutions
- Recent “Innovations” Evolving in L3 Virtualisation
- QoS/H-QoS and MTU Considerations for L3 Virtualisation Deployments
- Summary and Wrap Up




# L3 Virtualisation Solutions over IP Transport Services

# L3 WAN Virtualisation Deployment – IP VPN Transport

“IP VPN” Service Offering (PE→CE Model)

## SP Managed “IP VPN” Service

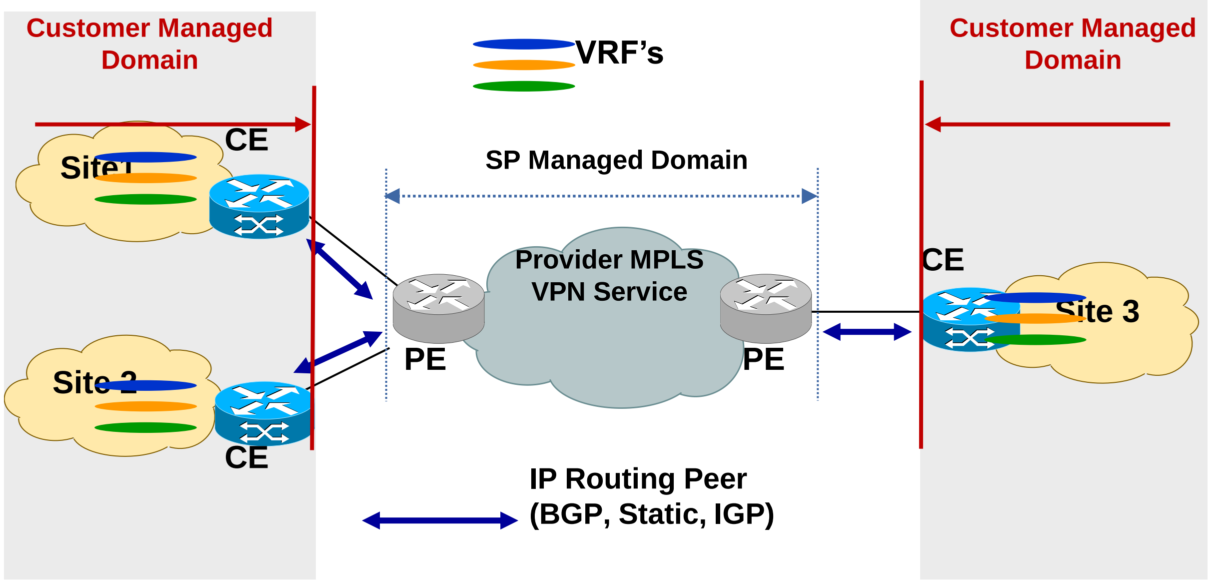


- **<u>L3 Virtualisation Options for “IP VPN Service”:</u>**
  - Back to Back VRFs (to SP PE)
  - Carrie Supporting Carrier (CsC) with RFC 3107
  - IP “Over the Top” (MPLS VPN or VRF-Lite over IP)

# Layer 3 VPN Peering – Private IP VPN “Over the Top” Solutions

# MPLS VPN Technology

## MPLS VPN over IP Encapsulation ("Over the Top")

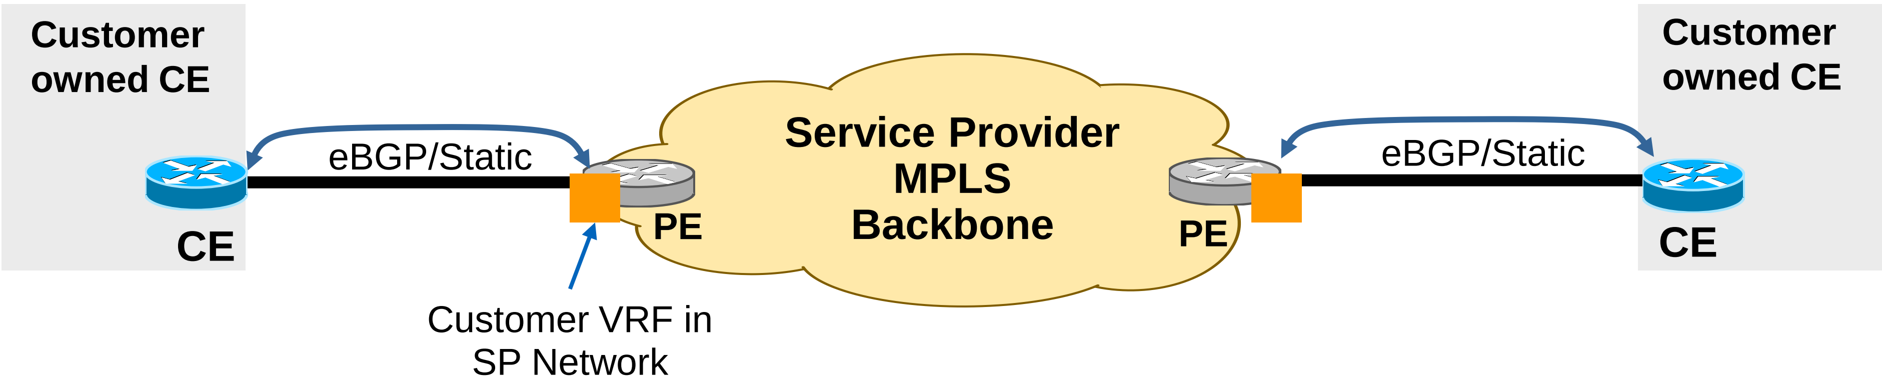

# Why Do We Need MPLS VPN over IP?

- Not all “transport/transit” networks are MPLS
  - i.e. MPLS is not available for transport on every network
- IP is the only Transit Offered Between MPLS Islands (i.e. networks)
- Customers are leveraging IP VPN Service from SP
- Customer uses “external” IP encryption units (i.e. device does not support MPLS)
- MPLS packets require encryption (no native MPLS encryption exists)

**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)**

# MPLS VPN Technology

## MPLS VPN or VRF-Lite over IP Encapsulation ("Over the Top")

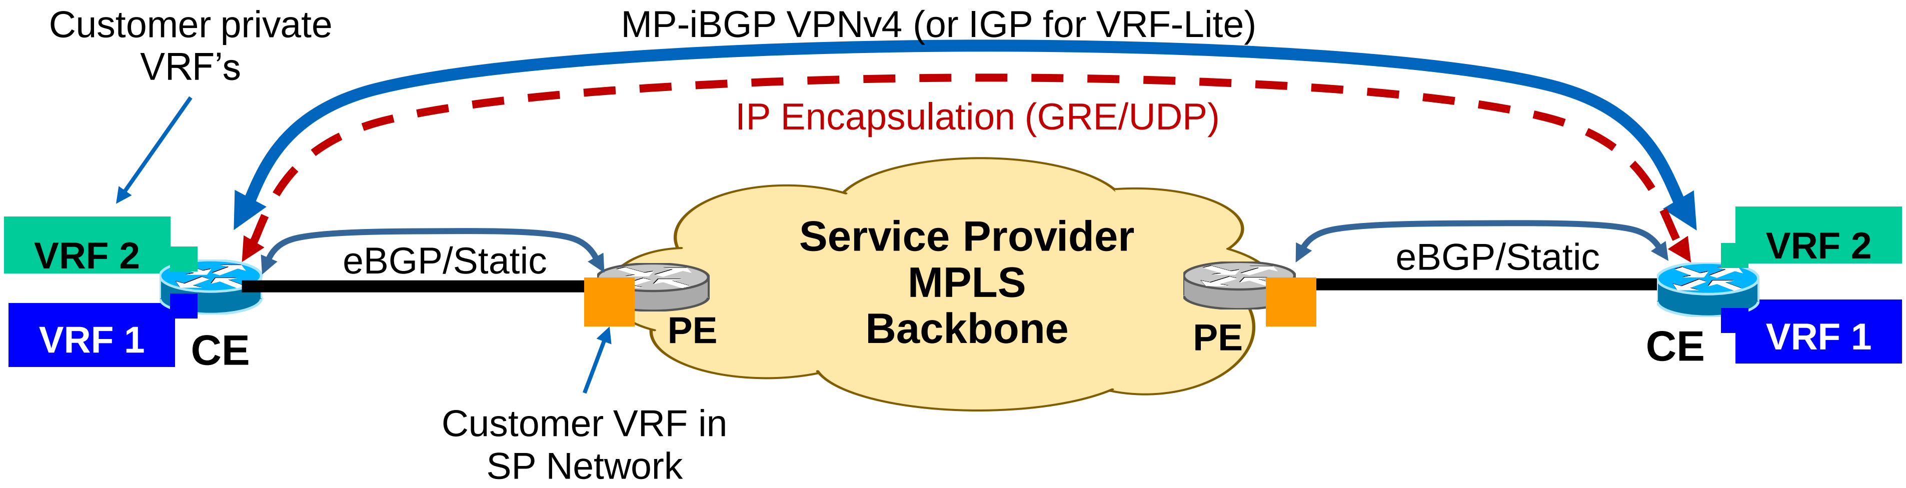

# MPLS VPN Technology

## MPLS VPN or VRF-Lite over IP Encapsulation (“Over the Top”)

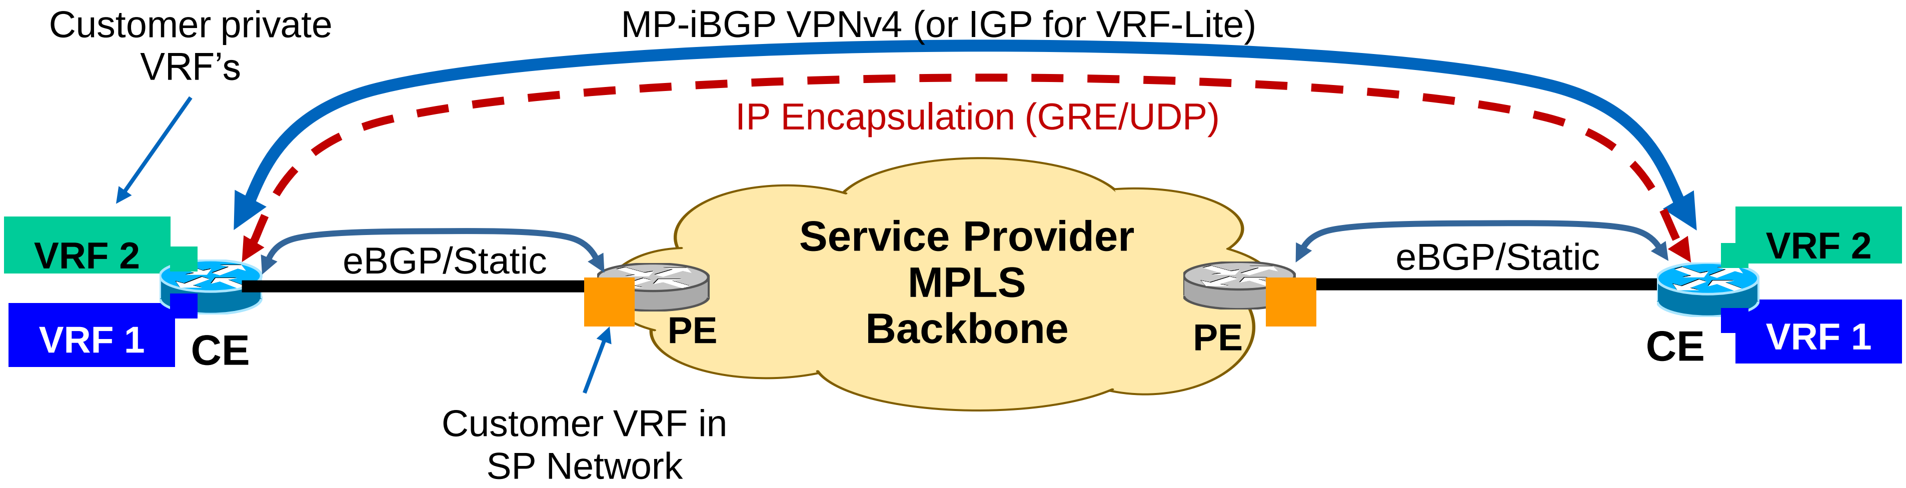


- MPLS VPN or VRF-Lite over IP Encapsulation
- Routing and data forwarding done “Over the Top” of the SP transport
- Enterprise routing - exchanged either inside IP tunnel, and/or over the top (BGP)
- Routing to SP – BGP/static and minimal (typically IP tunnel end-points)
- Multicast can be supported either (1) leveraging the SP service, or (2) inside the IP tunnel

# MPLS VPN Technology

## MPLS VPN or VRF-Lite over IP Encapsulation ("Over the Top")

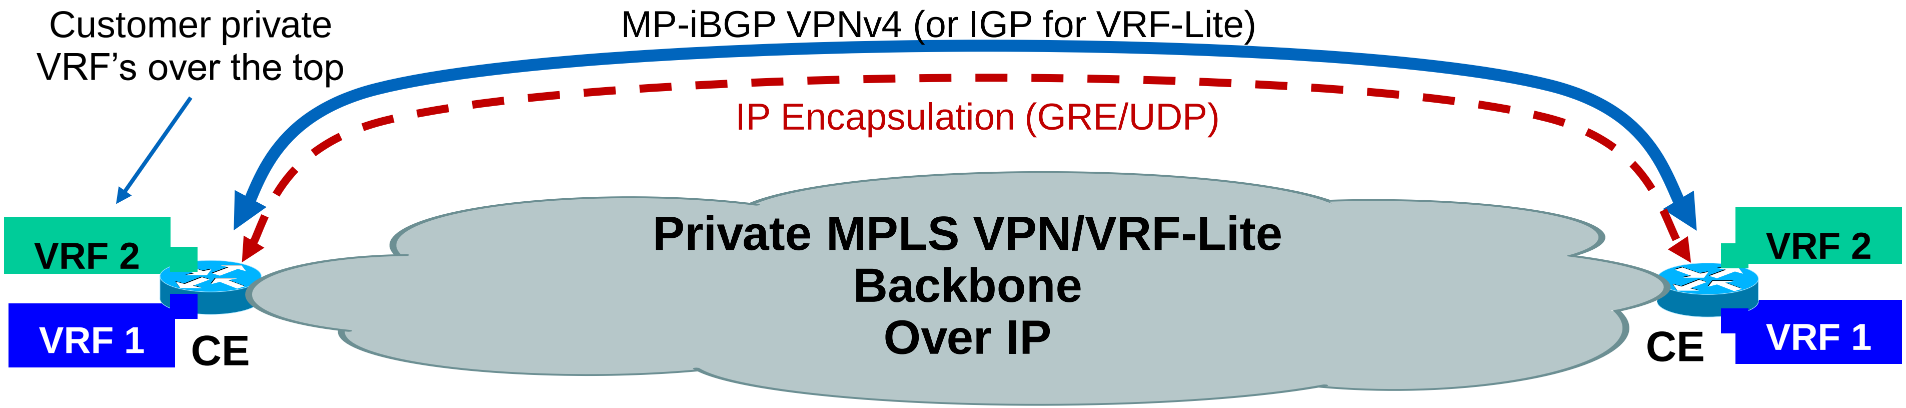


- MPLS VPN or VRF-Lite over IP Encapsulation
- Routing and data forwarding done "Over the Top" of the SP transport
- Enterprise routing - exchanged either inside IP tunnel, and/or over the top (BGP)
- Routing to SP – BGP/static and minimal (typically IP tunnel end-points)
- Multicast can be supported either (1) leveraging the SP service, or (2) inside the IP tunnel

# GRE Tunnel Encapsulation (RFC 2784)

## Applicable over Any IP WAN Transport

**Bit 0:** Check Sum
**Bit 1-12:** Reserved
**Bit 13-15:** Version Number
**Bit 16-31:** Protocol Type

***Original IP Datagram*** *(Before Forwarding)*

| Original IP Header | IP Payload |
|---|---|
| 20 Bytes | |

***GRE Packet with New IP Header: Protocol 47*** *(Forwarded Using New IP Dst)*

| New IP Header | GRE Header | Original IP Header | IP Payload |
|---|---|---|---|
| 20 Bytes | 4 Bytes | 20 Bytes | |


**Can Also Leverage IPSec When IP Encryption Is Required of an Untrusted WAN**

# Encapsulation for MPLS in GRE (RFC 4023)

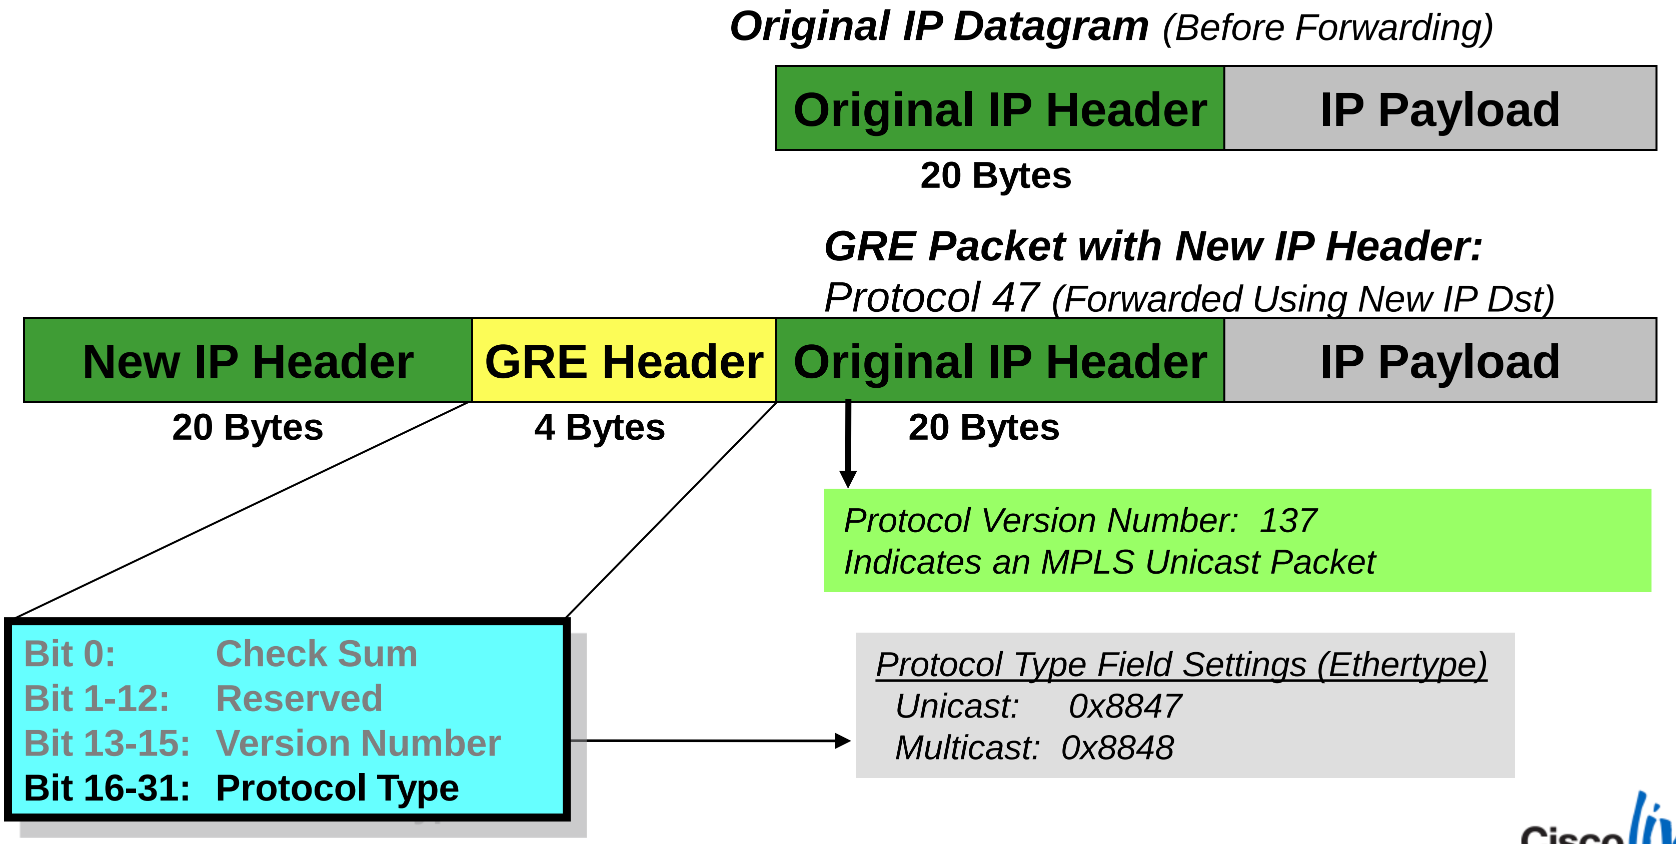

# GRE Tunnel Format with MPLS

(Reference: RFC 4023)

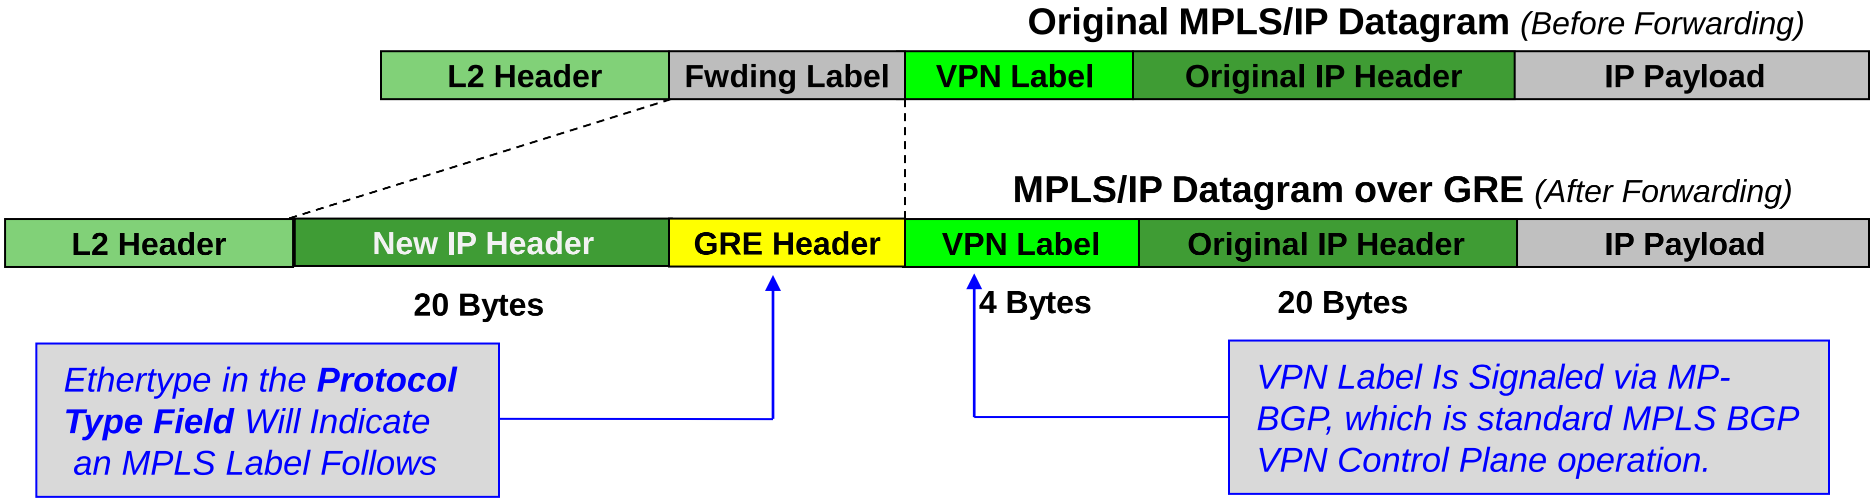


- MPLS Tunnel label (top) is replaced with destination PE's IP address
- Encapsulation defined in RFC 4023
- Most widely deployed form of MPLS over IP encapsulation

# MPLS VPN over IP/GRE

## Control Plane

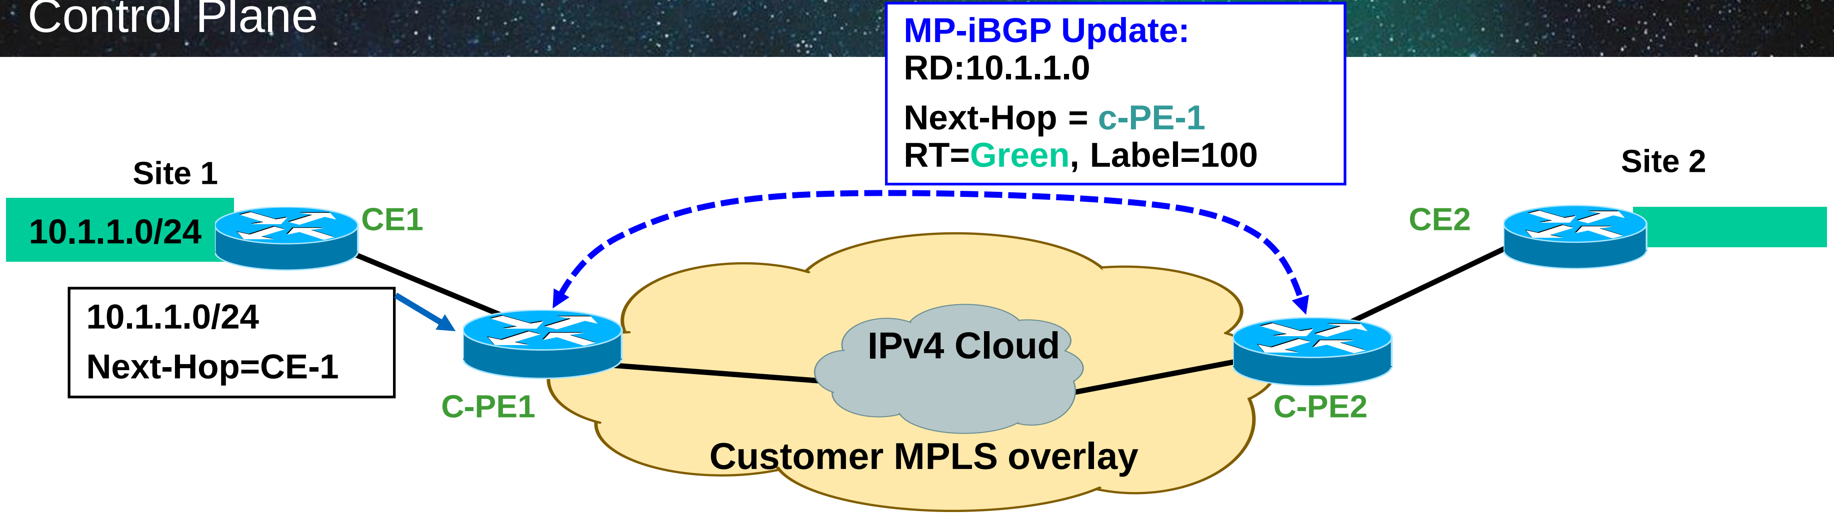


- C-PE1 receives an IPv4 update (eBGP/OSPF/ISIS/RIP/EIGRP)
- C-PE1 translates it into VPNv4 address, sends MP-iBGP update to other PE routers
- C-PE2 receives and checks whether the RT=green (40:103, say) is locally configured within any VRF, if yes, then
- C-PE2 translates VPNv4 prefix back into IPv4 prefix,
- All done over the GRE tunnel (point to point or DMVPN scenario)

# MPLS VPN over GRE

## Data Plane

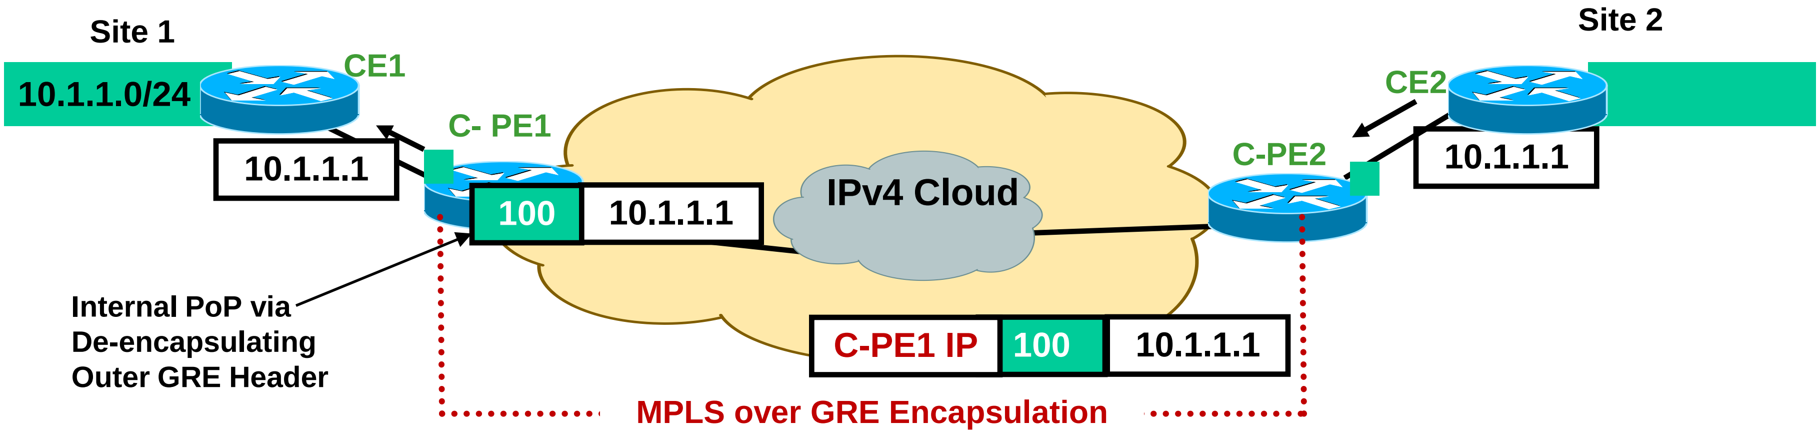


- c-PE2 normally imposes **two** labels for each packet going to the VPN destination 10.1.1.1, (1) top IGP derived label (2) VPN label

**For MPLS over GRE Encapsulation Case…**

- The top label is replaced with an IP Tunnel Header to the destination of c-PE1
- The 2nd label (inner) is the VPNv4 address learned via MP-BGP via GRE tunnel
- On c-PE1, the GRE header is removed, exposing the VPN label for forwarding
- From each c-PE view, the PE-PE connection is an implicit null (penultimate hop)

# MPLS VPN over Point-to-Point GRE

## Example is MPLS over Point-to-Point GRE Tunnel

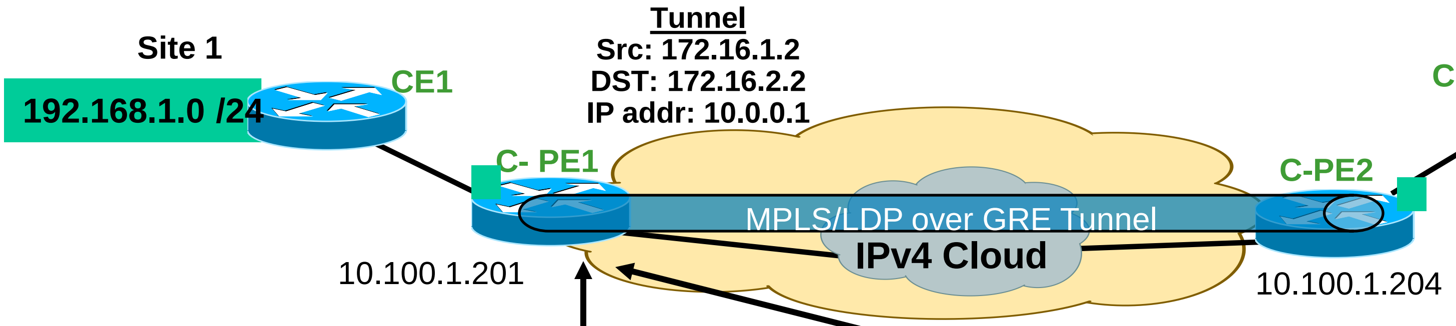


**Important:** You must force the forwarding of the BGP next-hop over the GRE tunnel

**Using 10.0.0.0/8 address space Forces Loopback 0 learning over GRE Tunnel**

**Enables MPLS/LDP over GRE**

```
ip vrf green
 rd 1:1
 route-target export 1:1
 route-target import 1:1
!
mpls label protocol ldp
!
interface Tunnel0
 ip address 10.0.0.1 255.255.255.0
 ip mtu 1400
 mpls ip
 tunnel source 172.16.1.2
 tunnel destination 172.16.2.2
 tunnel path-mtu-discovery
!
interface Loopback0
 ip address 10.100.1.201 255.255.255.255
!
```

```
router eigrp 1
 network 10.0.0.0
 no auto-summary
 !
router bgp 65000
 bgp router-id 10.100.1.201
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 10.100.1.204 remote-as 65000
 neighbor 10.100.1.204 update-source Loopback0
!
 address-family vpnv4
 neighbor 10.100.1.204 activate
 neighbor 10.100.1.204 send-community extended
 exit-address-family
 !
```

# GRE Tunnel Modes

"Stateful" vs. "Stateless"

## Point-to-Point GRE

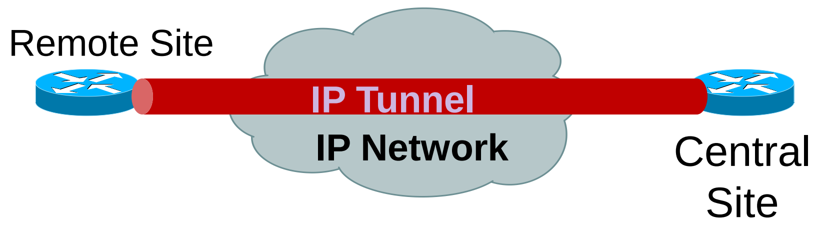


- Source and destination requires manual configuration
- Tunnel end-points are stateful neighbours
- Tunnel destination is explicitly configured
- Creates a logical point-to-point "Tunnel"

## Multipoint GRE

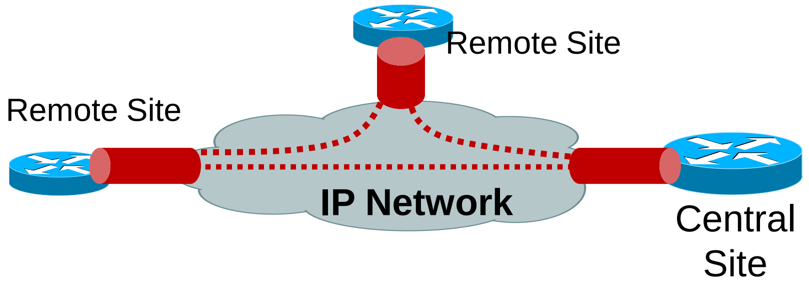


- Single multipoint tunnel interface is created per node
- Only the tunnel source is defined
- Tunnel destination is derived dynamically through some signalling mechanism (i.e. BGP, NHRP) or discovery end-point concept
- Creates an "encapsulation" using IP headers (GRE)

# MPLS VPN over Multipoint GRE (mGRE)

## MPLS VPNs over Multipoint GRE Using BGP for IP encapsulated Next-Hop

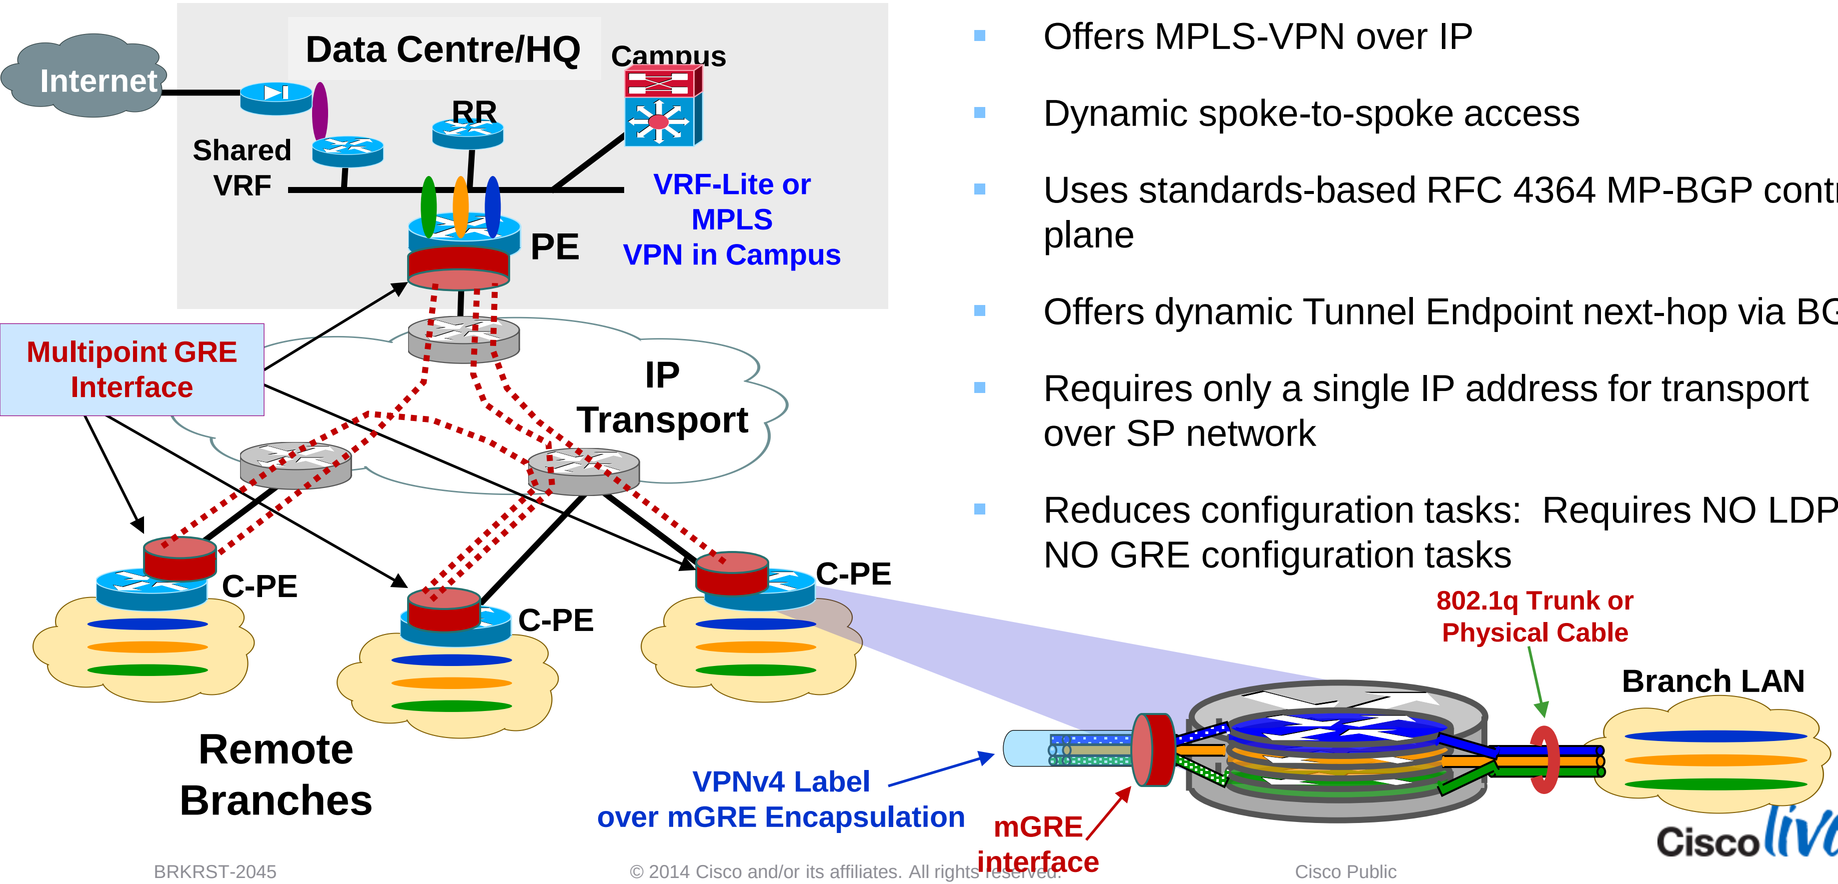


- Offers MPLS-VPN over IP
- Dynamic spoke-to-spoke access
- Uses standards-based RFC 4364 MP-BGP control plane
- Offers dynamic Tunnel Endpoint next-hop via BGP
- Requires only a single IP address for transport over SP network
- Reduces configuration tasks: Requires NO LDP, NO GRE configuration tasks

  

# MPLS VPN over Multipoint GRE (mGRE)

## Control/Data Plane Example over Service Provider Model

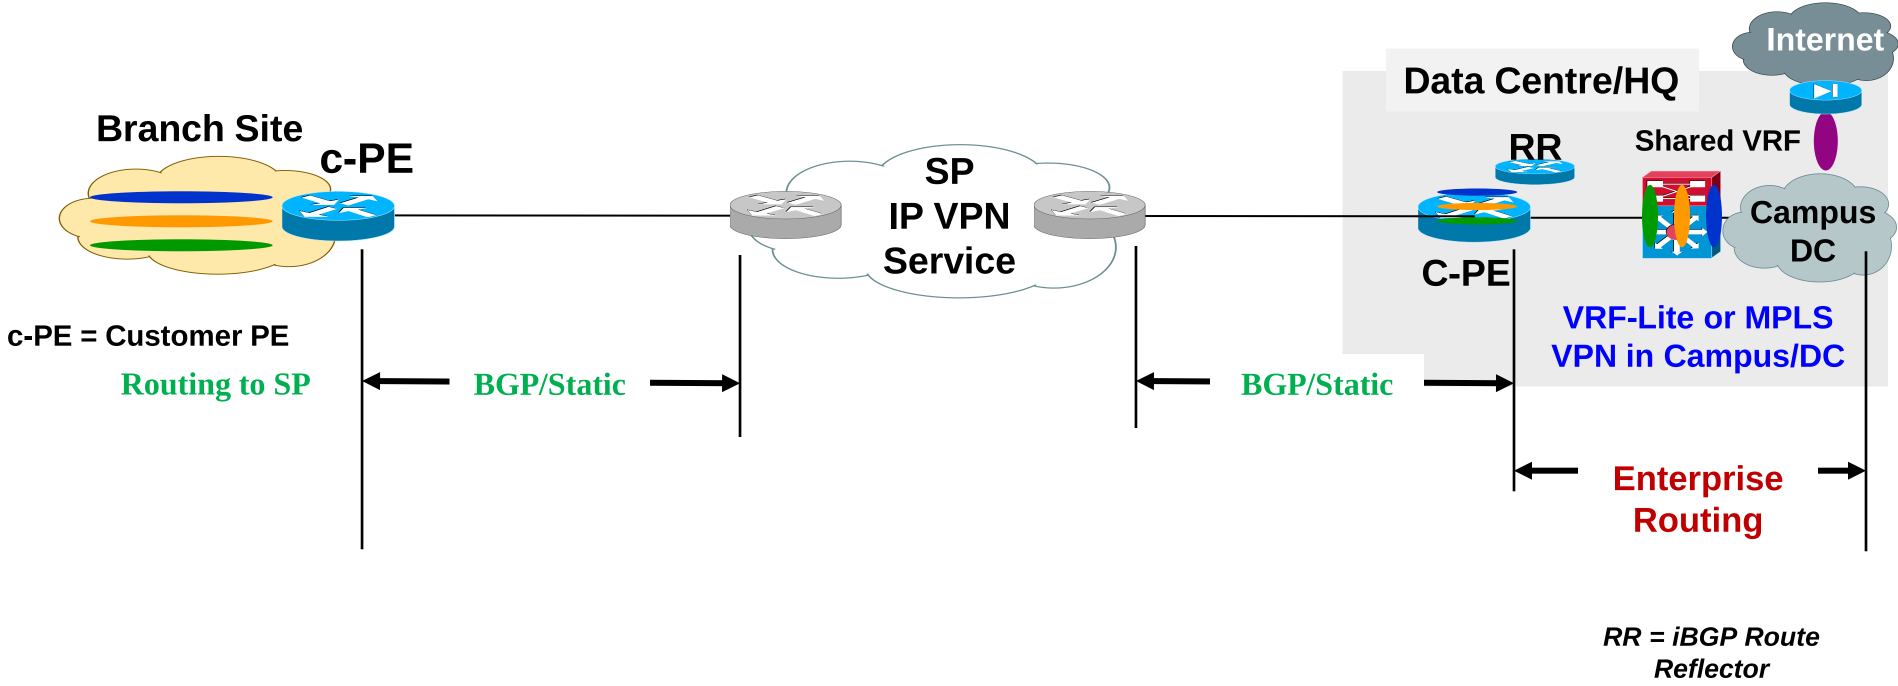

# MPLS VPN over Multipoint GRE (mGRE)

## Control/Data Plane Example over Service Provider Model

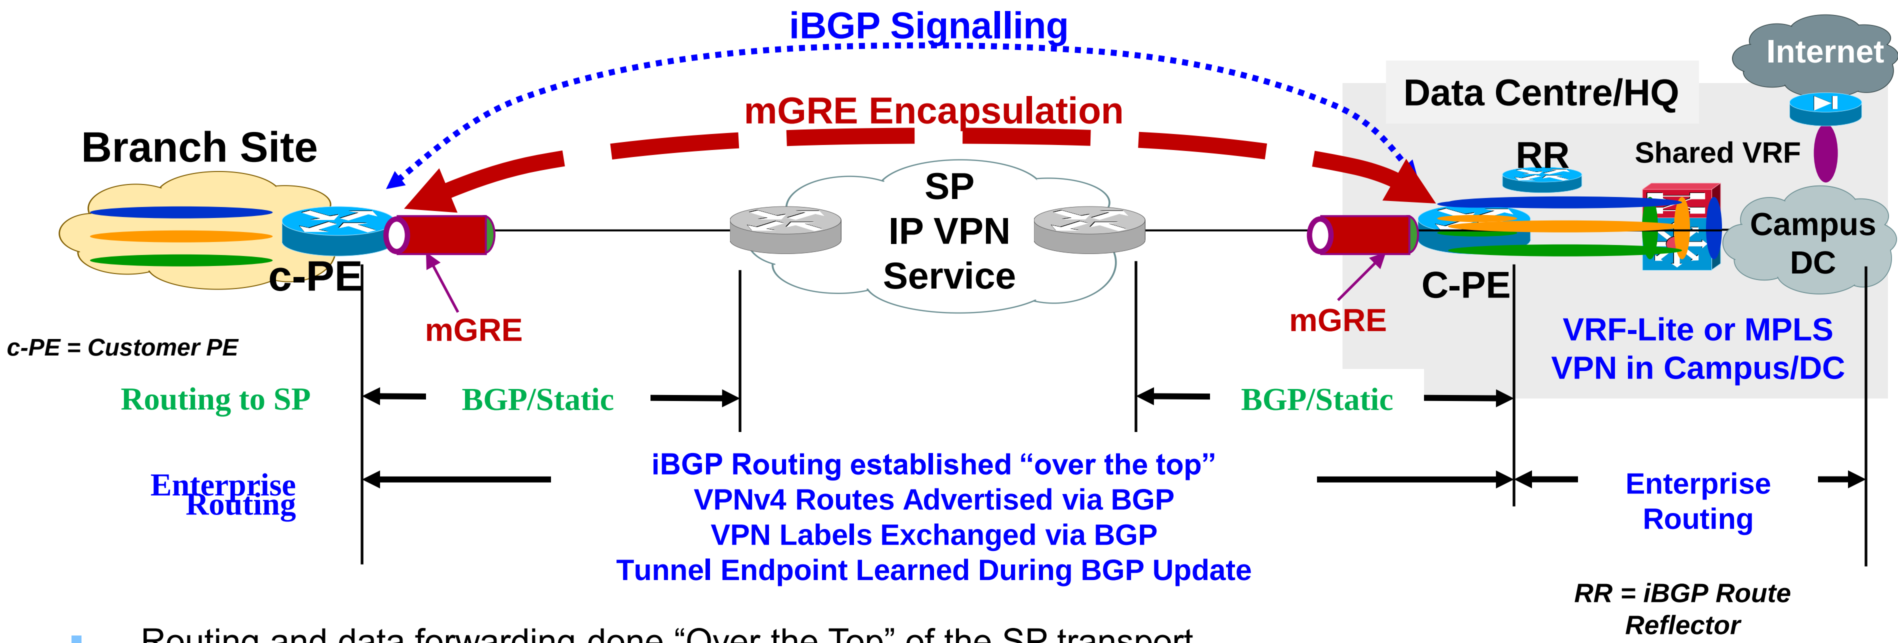


- Routing and data forwarding done "Over the Top" of the SP transport
- iBGP used to: (1) Advertise VPNv4 routes, (2) exchange VPN labels
- eBGP used to: (1) exchange tunnel end point routes with SP (optional static routes could be used)
- Only requires advertising ONE IP prefix to the SP network (e.g. IP tunnel "end points")

# MPLS VPN over Multipoint GRE (mGRE)

## Feature Components

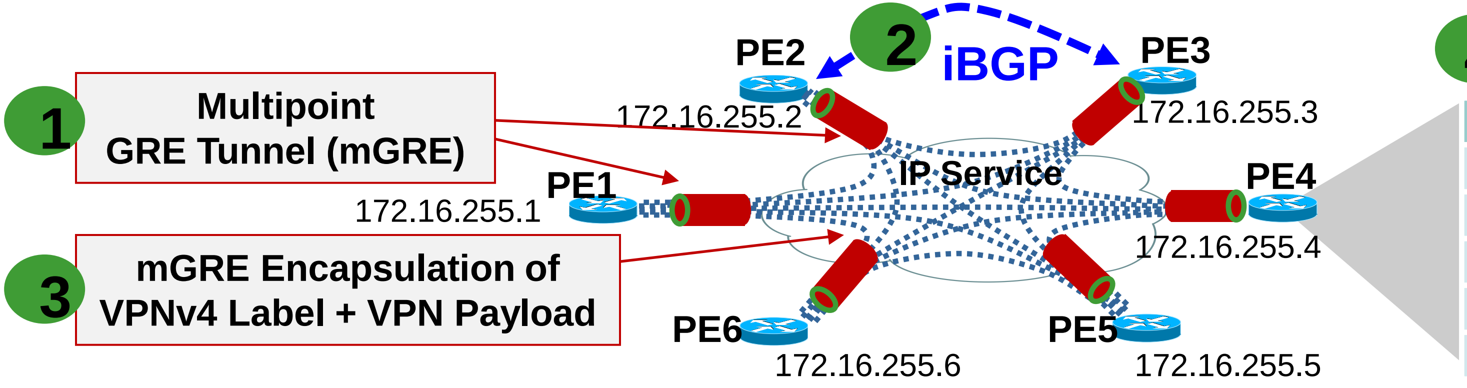


| Tunnel Endpoint |
|---|
| 172.16.255.1 |
| 172.16.255.2 |
| 172.16.255.3 |
| 172.16.255.5 |
| 172.16.255.6 |

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 Rcvd prefixes into tunnel database

- **Solution does NOT require manual configuration for GRE tunnels or ‘mpls ip’ on interface(s)**

**Multipoint GRE Interface**

# MPLS VPN over Multipoint GRE (mGRE)

## VPNv4 Configuration Example

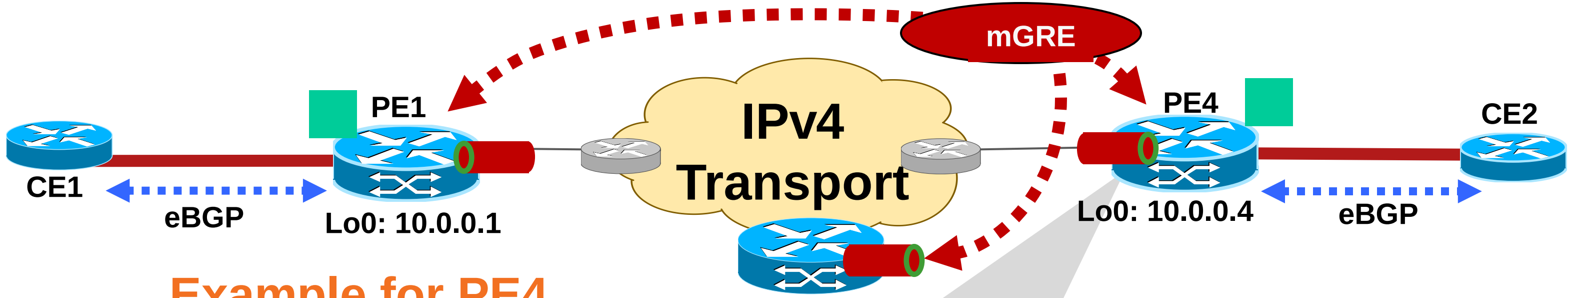


### Example for PE4

```
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

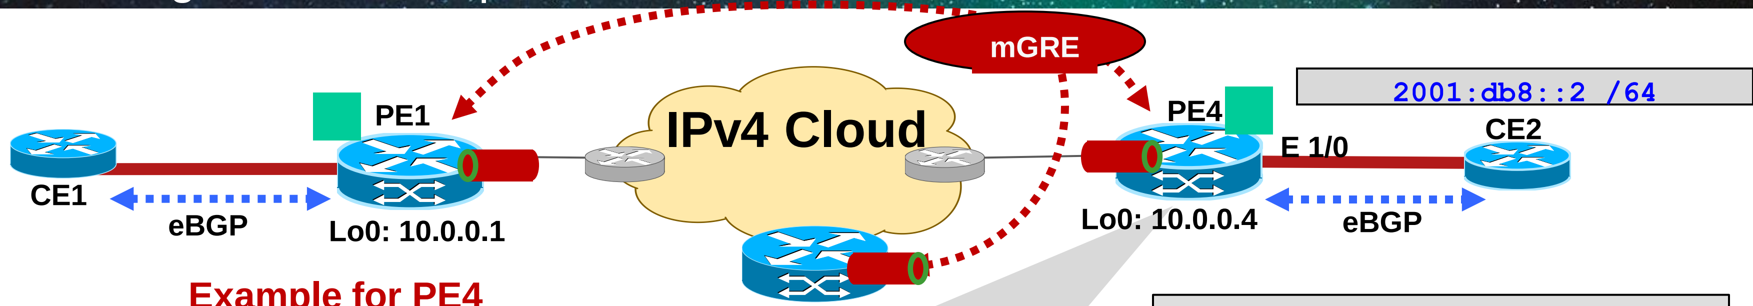


**Example for PE4**

```
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 over Multipoint GRE (mGRE)

## Summary and Configuration Notes

- Only requires advertising a single IP prefix to SP for mGRE operation
- Solution leverages standard MP-BGP control plane (RFC 2547/4364)
- Tunnel endpoint discovery is done via iBGP/route-map
- E-BGP can/is still used for route exchange with the SP
- Solution requires NO manual configuration of GRE tunnels or LDP (RFC 3036)
- 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)
- <u>Platform Support</u>

  <u>Today:</u> 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

  <u>Future:</u> IOS-XR Platforms (Future planning)

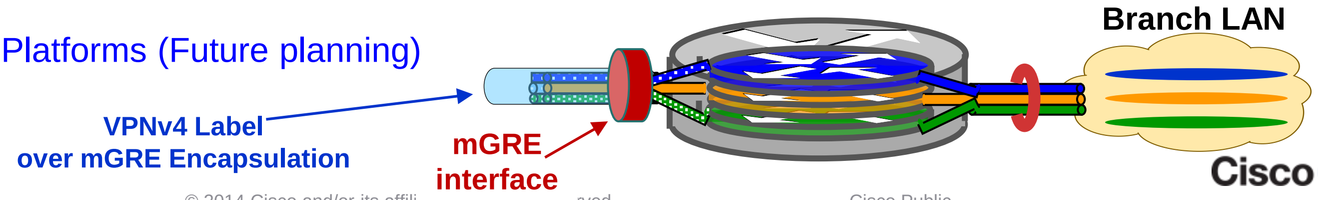

# MPLS VPN Deployment Considerations for WAN Designs (over IP)

## EXAMPLE: MPLS VPN over mGRE (BGP)

**Example: 50 – 1000 Sites**

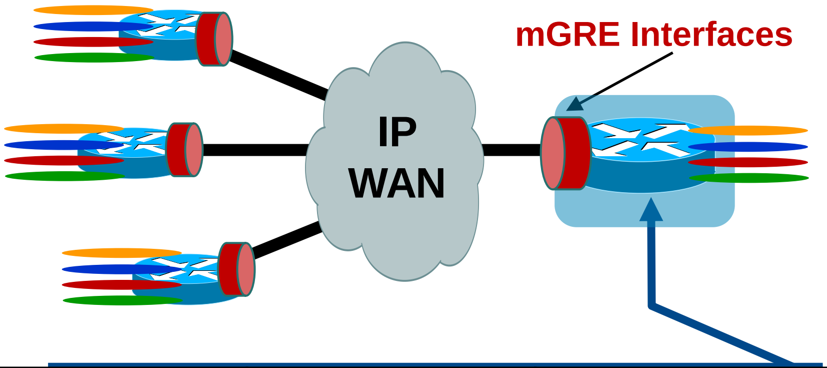


| VRFs | Neighbours | GRE Tunnel Interface |
|---|---|---|
| 50 | 50 | 1 |
| 100 | 100 | 1 |
| 250 | 200 | 1 |
| 500+ | 1000 | 1 |

**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 mGRE – “Config” and “Show” Examples

# MPLS VPN over Multipoint GRE (mGRE)

## Configuration Example – Router 218

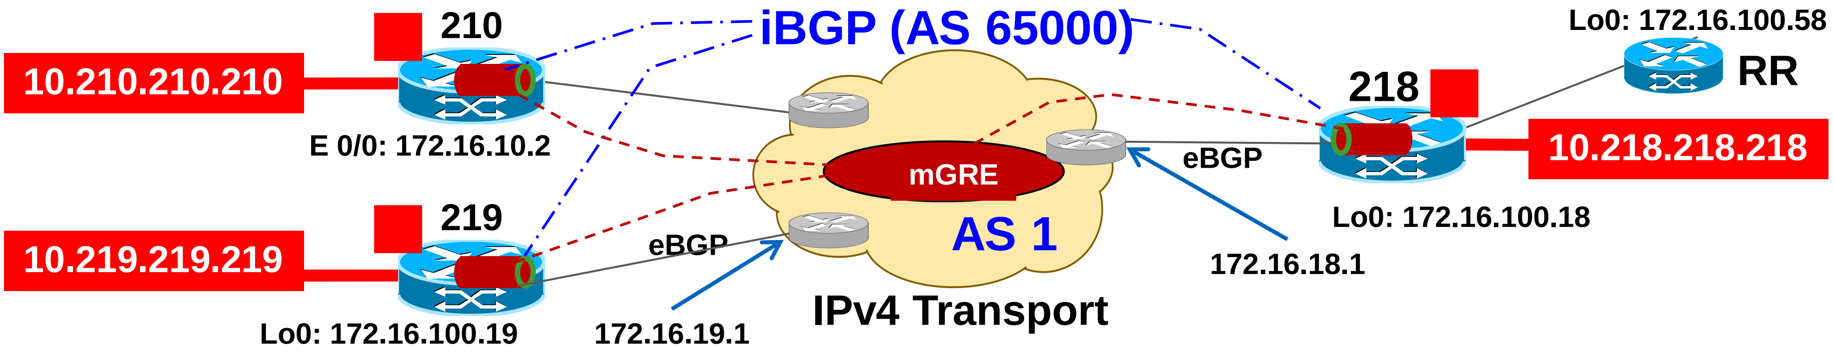


```
!
vrf definition red
 rd 1:1
 route-target export 1:1
 route-target import 1:1
 !
 address-family ipv4
!
interface Loopback0
 ip address 172.16.100.18 255.255.255.255
!
interface Ethernet0/0
 ip address 172.16.18.2 255.255.255.0
 service-policy output parent
!
```

```
!
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)

## Configuration Example – Router 218

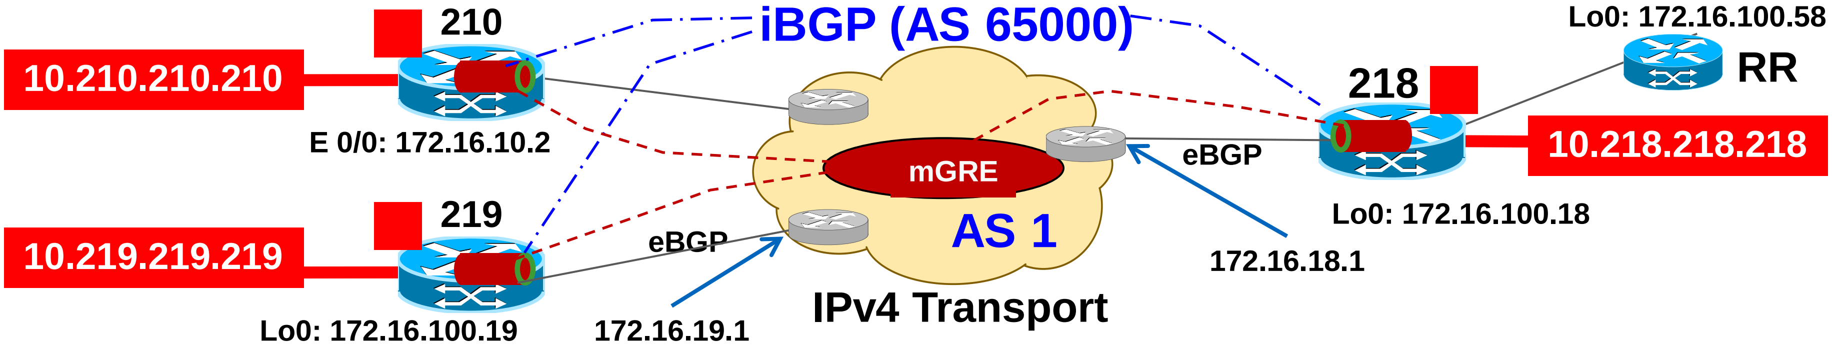


```
!
router bgp 65000
 neighbor 172.16.18.1 remote-as 1
 neighbor 172.16.18.1 update-source Eth 0/0
 neighbor 172.16.100.58 remote-as 65000
 neighbor 172.16.100.58 update-source Loop 0
 !
 address-family ipv4
  network 172.16.100.18 mask 255.255.255.255
  neighbor 172.16.18.1 activate
  neighbor 172.16.18.1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

## Configuration Example – Router 218

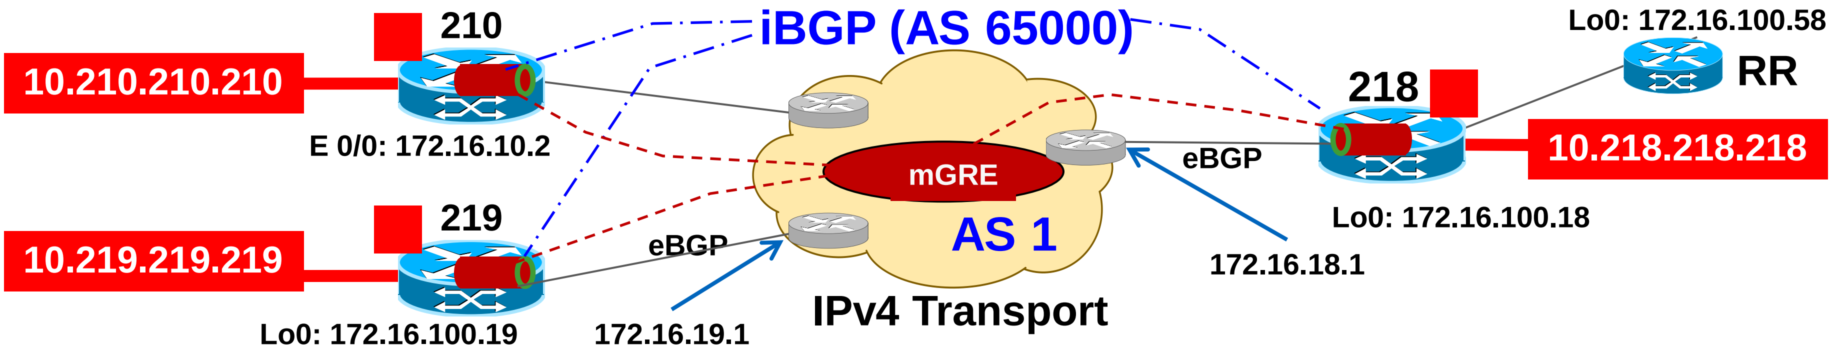


```
218#conf t
Enter configuration commands, one per line.  End with CNTL/Z.

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)

## Configuration Example – Router 218

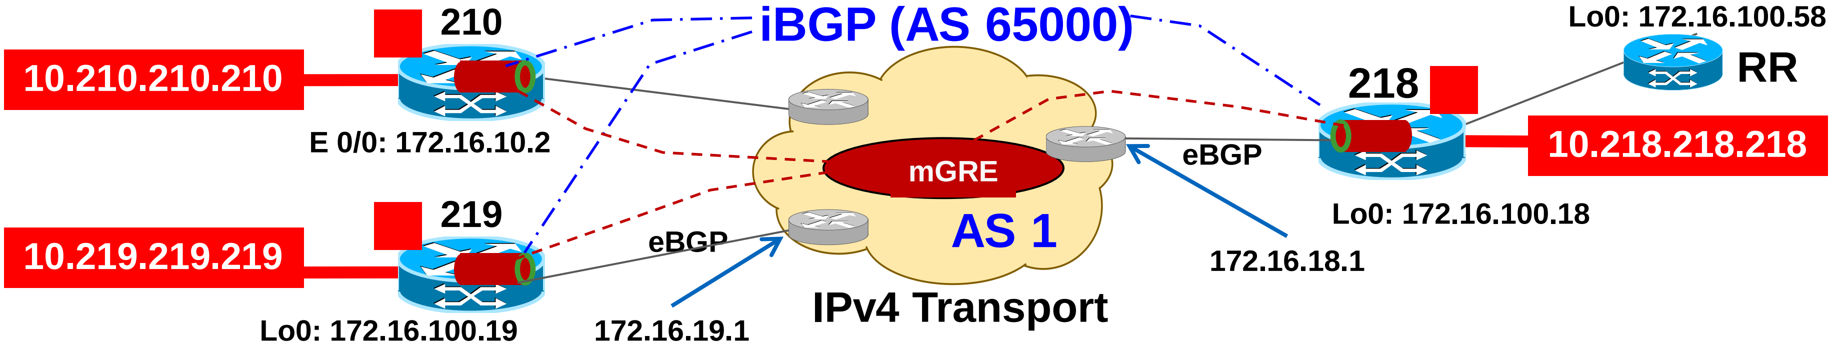


```
218#sh adjacency tunnel 0
Protocol Interface                 Address
IP       Tunnel0                   172.16.10.2(3)
TAG      Tunnel0                   172.16.10.2(3)
IP       Tunnel0                   172.16.100.19(3)
TAG      Tunnel0                   172.16.100.19(3)
```

```
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]
```

# MPLS VPN over Multipoint GRE (mGRE)

## Configuration Example – Router 218

```
218#sh adjacency tunnel 0 encapsulation

IP        Tunnel0                    172.16.100.19(3)
  Encap length 24
  4500000000000000FF2F9B88AC106412
  AC1064130000800
  Provider: TUNNEL

TAG       Tunnel0                    172.16.100.19(3)
  Encap length 24
  4500000000000000FF2F9B88AC106412
  AC10641300008847
  Provider: TUNNEL
  Protocol header count in macstring: 2
    HDR 0: ipv4
       dst: static, 172.16.100.19
       src: static, 172.16.100.18
      prot: static, 47
       ToS: static, 0
       ttl: static, 255
        df: static, cleared
      per packet fields: ident tl chksm
    HDR 1: gre
      prot: static, 0x8847
      per packet fields: none
```

AC106412 → 172.16.100.18

AC106413 → 172.16.100.19

# MPLS VPN over Multipoint GRE (mGRE)

## Configuration Example – Router 218

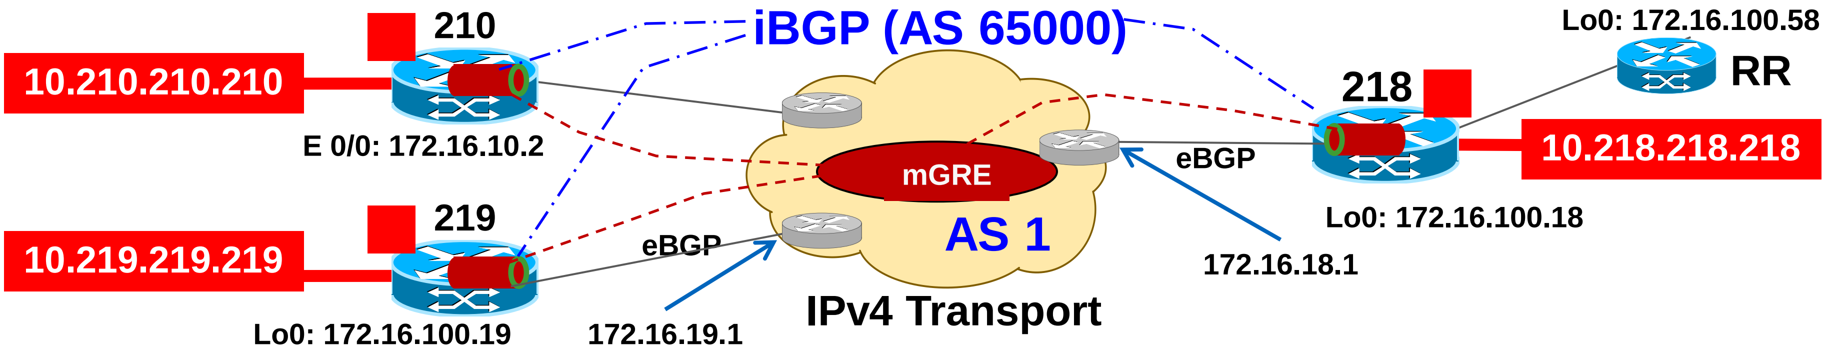


```
218#sh ip route

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 11 subnets, 2 masks
B        172.16.10.0/24 [20/0] via 172.16.18.1, 5d15h
L        172.16.18.2/32 is directly connected, Ethernet0/0
B        172.16.19.0/24 [20/0] via 172.16.18.1, 5d15h
C        172.16.58.0/24 is directly connected, Ethernet1/0
C        172.16.100.18/32 is directly connected, Loopback0
B        172.16.100.19/32 [20/0] via 172.16.18.1, 02:18:31
O        172.16.100.58/32 [110/11] via 172.16.58.2, 5d15h, Ethernet1/0
```

Router 210

Route Reflector

Router 219

# MPLS VPN over Multipoint GRE (mGRE)

## Configuration Example – Router 218

```
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
```

# MPLS VPN over Multipoint GRE (mGRE)

## Configuration Example – Router 218

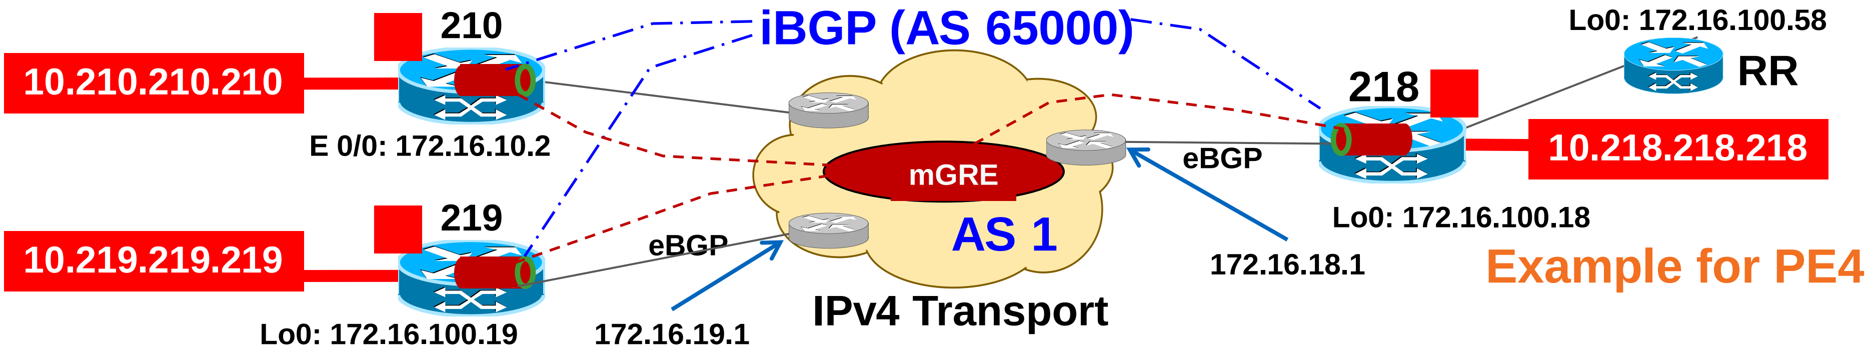


```
218#sh ip cef vrf red

Prefix                Next Hop              Interface
10.210.210.210/32     172.16.10.2           Tunnel0
10.218.218.218/32     receive               Loopback218
10.219.219.219/32     172.16.100.19         Tunnel0
```

```
218#sh ip cef vrf red 10.219.219.219

10.219.219.219/32
  nexthop 172.16.100.19 Tunnel0 label 16
```

# VRF-Lite over Dynamic Multipoint VPN (DMVPN)

## L3 Virtualisation Extension over DMVPN

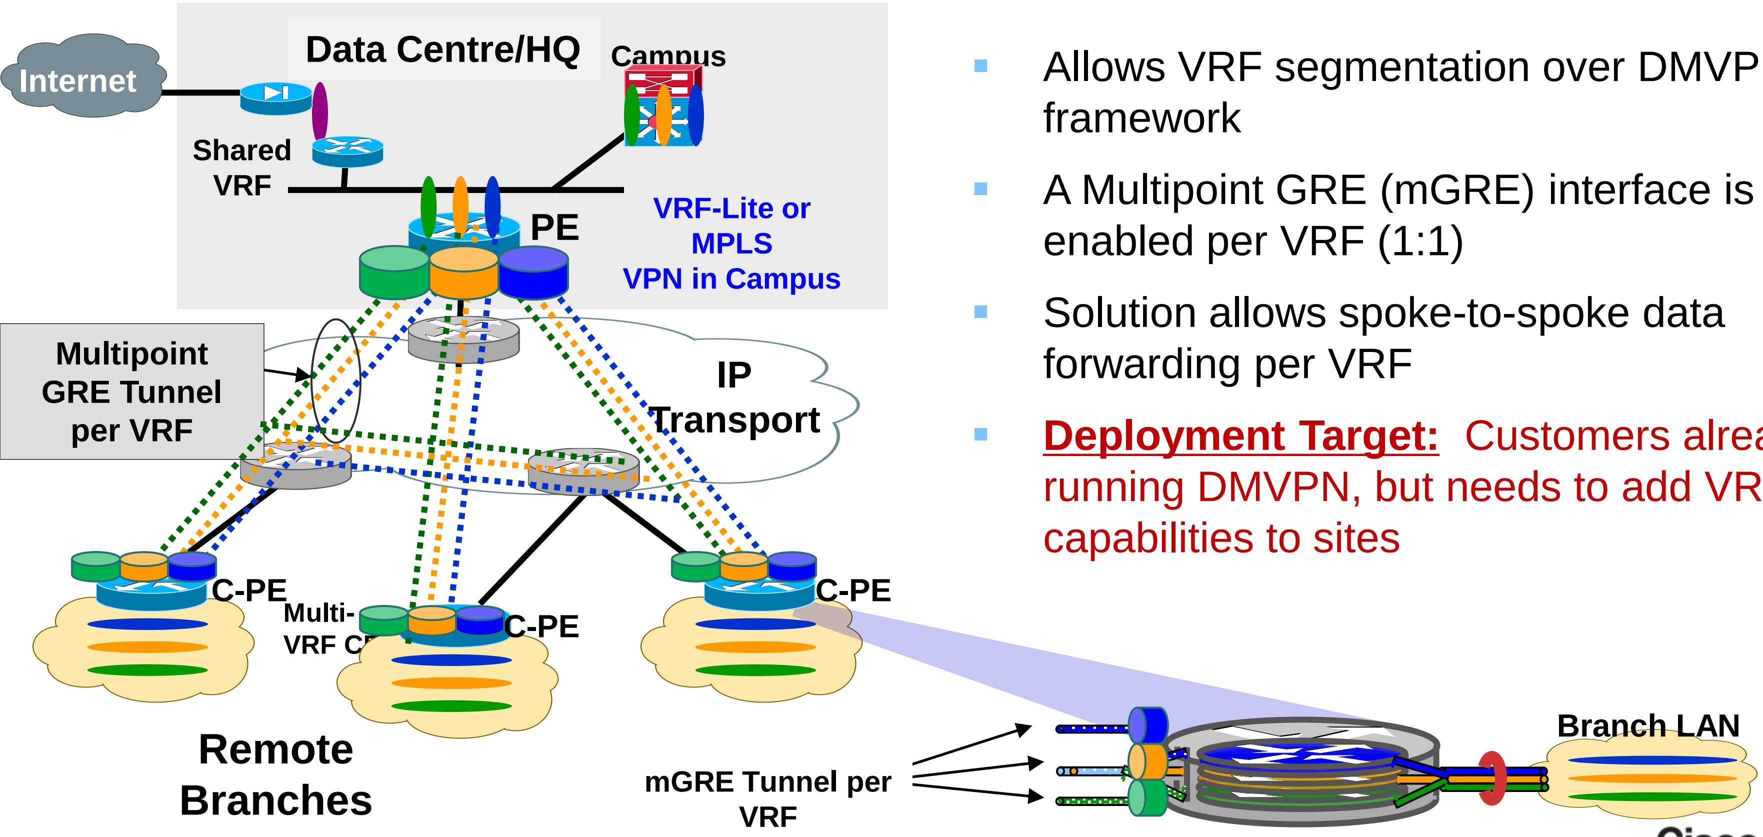


- Allows VRF segmentation over DMVPN framework
- A Multipoint GRE (mGRE) interface is enabled per VRF (1:1)
- Solution allows spoke-to-spoke data forwarding per VRF
- **Deployment Target:** Customers already running DMVPN, but needs to add VRF capabilities to sites

# MPLS VPN over Dynamic Multipoint VPN (DMVPN)

## MPLS VPN over a DMVPN Framework

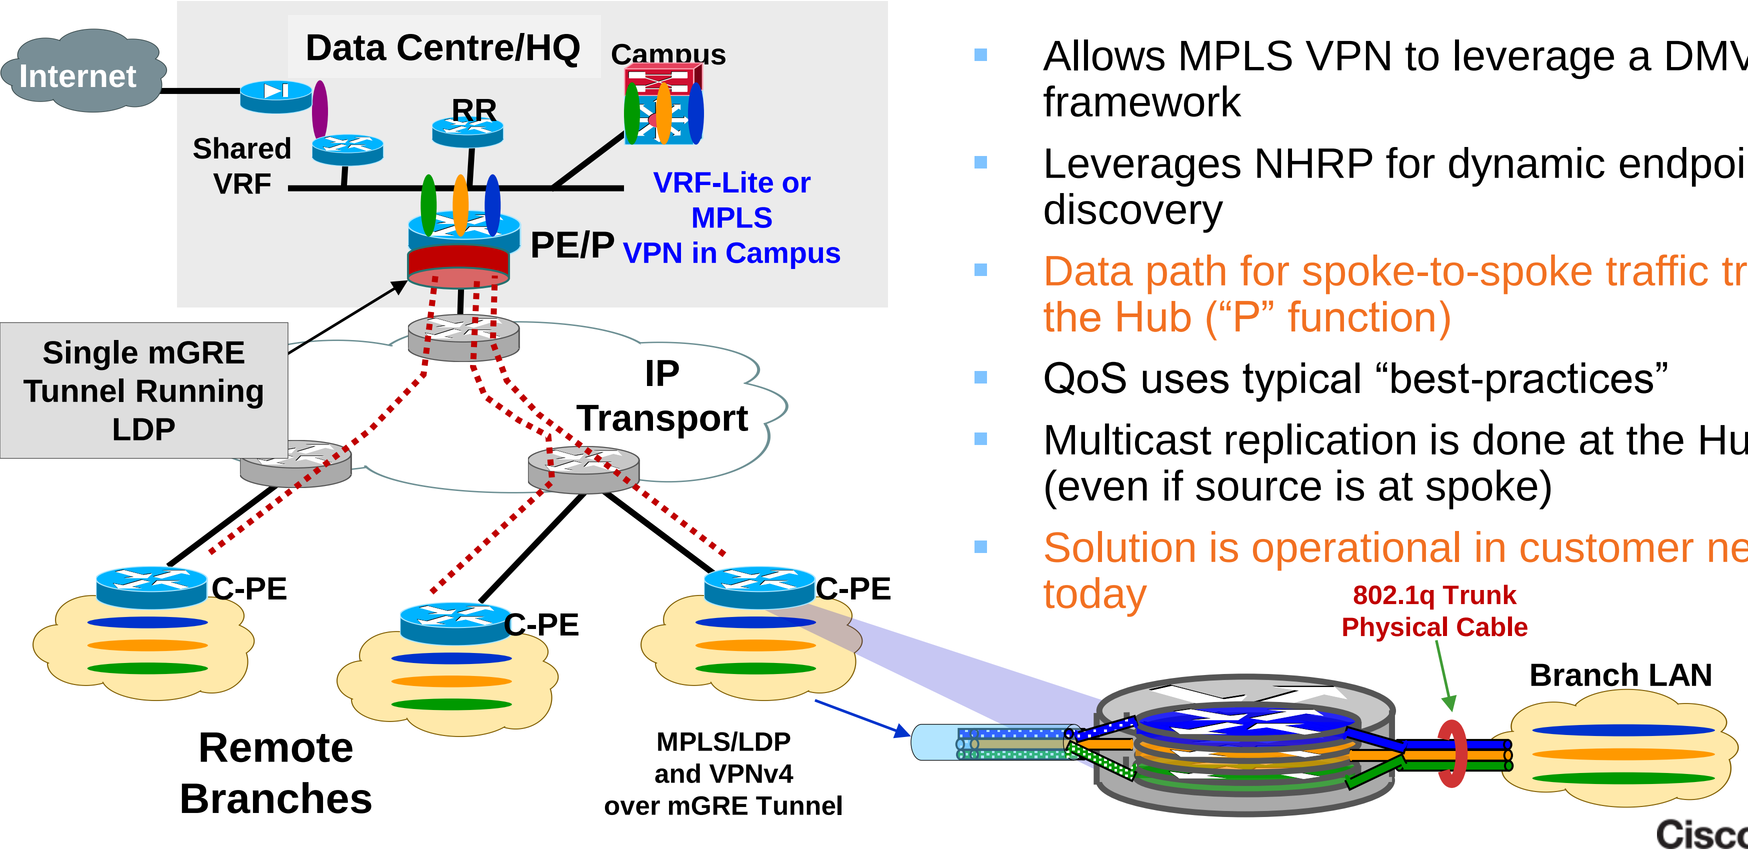


- Allows MPLS VPN to leverage a DMVPN framework
- Leverages NHRP for dynamic endpoint discovery
- Data path for spoke-to-spoke traffic transits the Hub (“P” function)
- QoS uses typical “best-practices”
- Multicast replication is done at the Hub (even if source is at spoke)
- Solution is operational in customer networks today

# MPLS VPN over GRE Solutions

## Comparison Matrix

| | MPLS VPN over mGRE | MPLS VPN over DMVPN | MPLS VPN over P2P GRE |
|---|---|---|---|
| Target Deployment | **Campus/WAN** | WAN | **Campus/WAN** |
| MPLS VPN Target VRFs | Yes (> 8 VRFs) | Yes (> 8 VRFs) | Yes (> 8 VRFs) |
| Uses a Dynamic Endpoint Discovery Mechanism | Yes (BGP) | Yes (NHRP) | No |
| Avoids Manual Full-Mesh GRE Configurations (mGRE) | Yes | Yes | No |
| Requires LDP over the Tunnel for Virtualisation with MPLS VPNs | No | Yes | Yes |
| Current Scaling of End Nodes (Tested) | 1000+ (Recommend RRs) | EIGRP – 1000 (ASR 1K)<br>OSPF – 600 (7200)<br>BGP – 1800 (ASR 1K) | 1000+ (Manually Intensive) |
| Supports IPSec Encryption | Yes (GET, SA) | Yes | Yes |
| Supports MVPN Multicast * | Yes (Platform Specific) | * Yes | Yes |
| Supports IPv6 VPN (6vPE) | Yes | No (Future) | Yes |

*** Platform Specific for support.  Also, DMVPN requires traffic be sent spoke-hub-spoke, if source is located at spoke site**

# Agenda

- Introduction - Network Virtualisation Drivers and Concepts
- SP WAN Transport Service Impact on L3 Virtualisation Solution Choices
- Technology and Deployment Deep-Dive for L3 Virtualised WAN
- **Integrating Encryption into L3 Virtualisation Solutions for mGRE**
- Recent “Innovations” Evolving in L3 Virtualisation
- QoS/H-QoS and MTU Considerations for L3 Virtualisation Deployments
- Summary and Wrap Up

# Group Encrypted Transport (GET) VPN

## Any to Any Encryption for “Stateless” IP Tunnels (mGRE, LISP…)

### Public/Private WAN

**Before: IPSec P2P Tunnels**

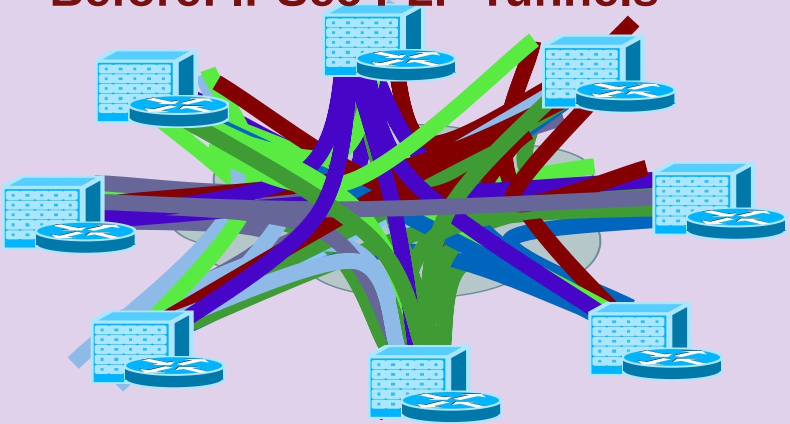

- Scalability—an issue  (N^2 problem)
- Overlay routing
- Any-to-any instant connectivity can’t be done to scale
- Limited QoS
- Inefficient Multicast replication

### Private WAN

**After: Tunnel-Less VPN**

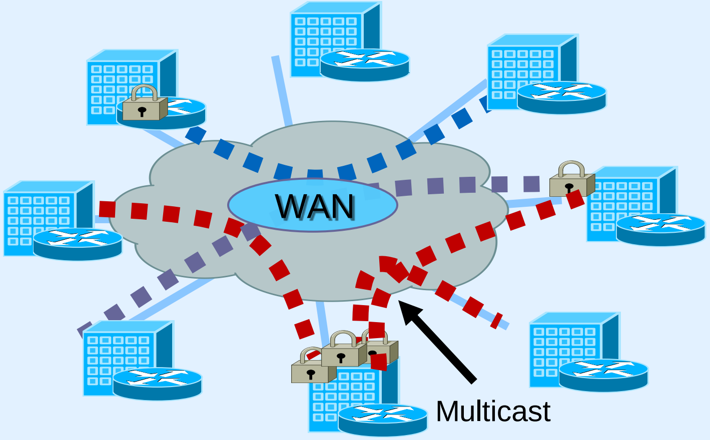


- Scalable architecture for any-to-any connectivity and encryption
- No overlays—native routing
- Any-to-any instant connectivity
- Enhanced QoS
- Efficient Multicast replication

# GETVPN Security Devices

**Key Server**
- Validate Group Members
- Manage Security Policy
- Create Group Keys
- Distribute Policy / Keys

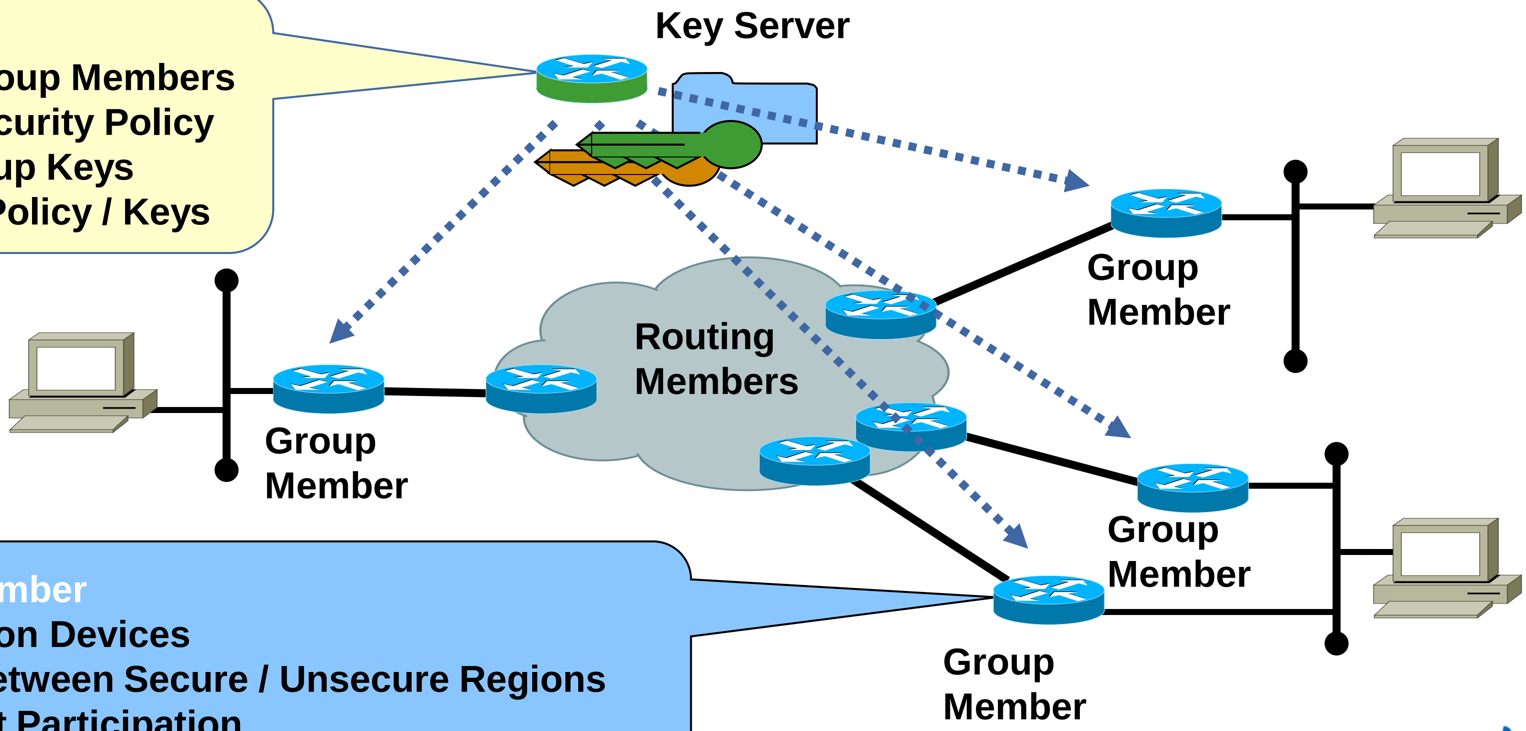


**Group Member**
- Encryption Devices
- Route Between Secure / Unsecure Regions
- Multicast Participation

# Combining Technologies into Secure L3 Virtualisation Solution

Leverage MPLS VPN over mGRE + GET VPN Encryption

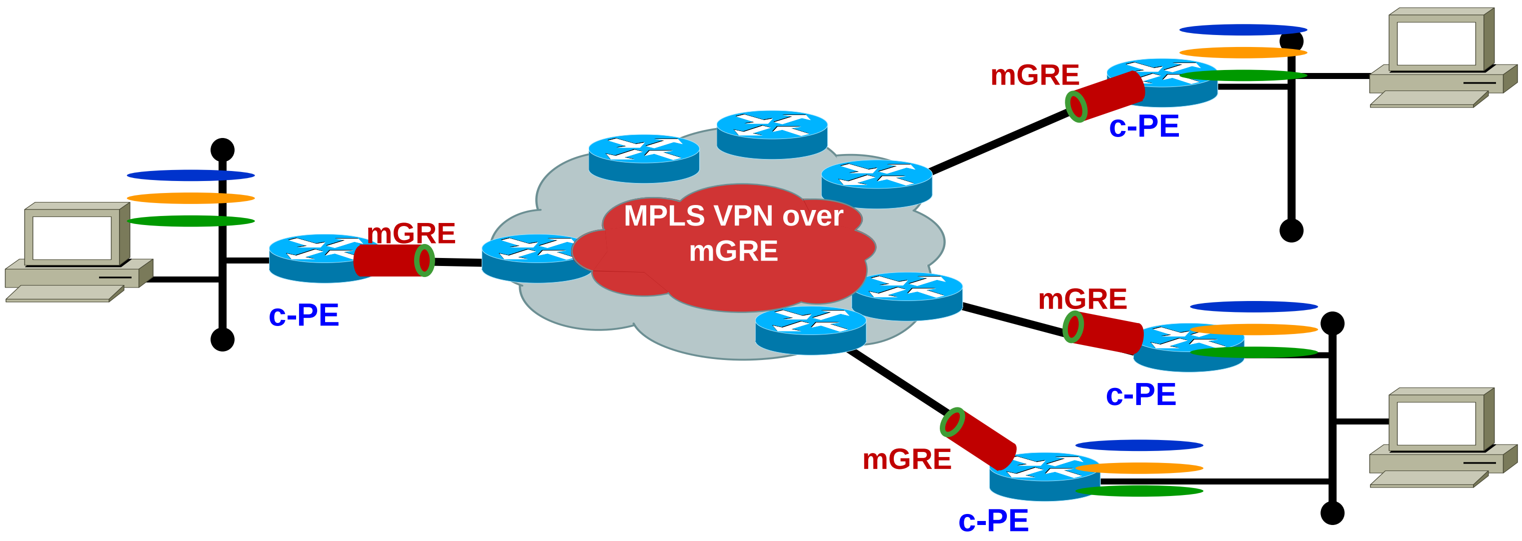

# Combining Technologies into Secure L3 Virtualisation Solution

## Leverage MPLS VPN over mGRE + GET VPN Encryption

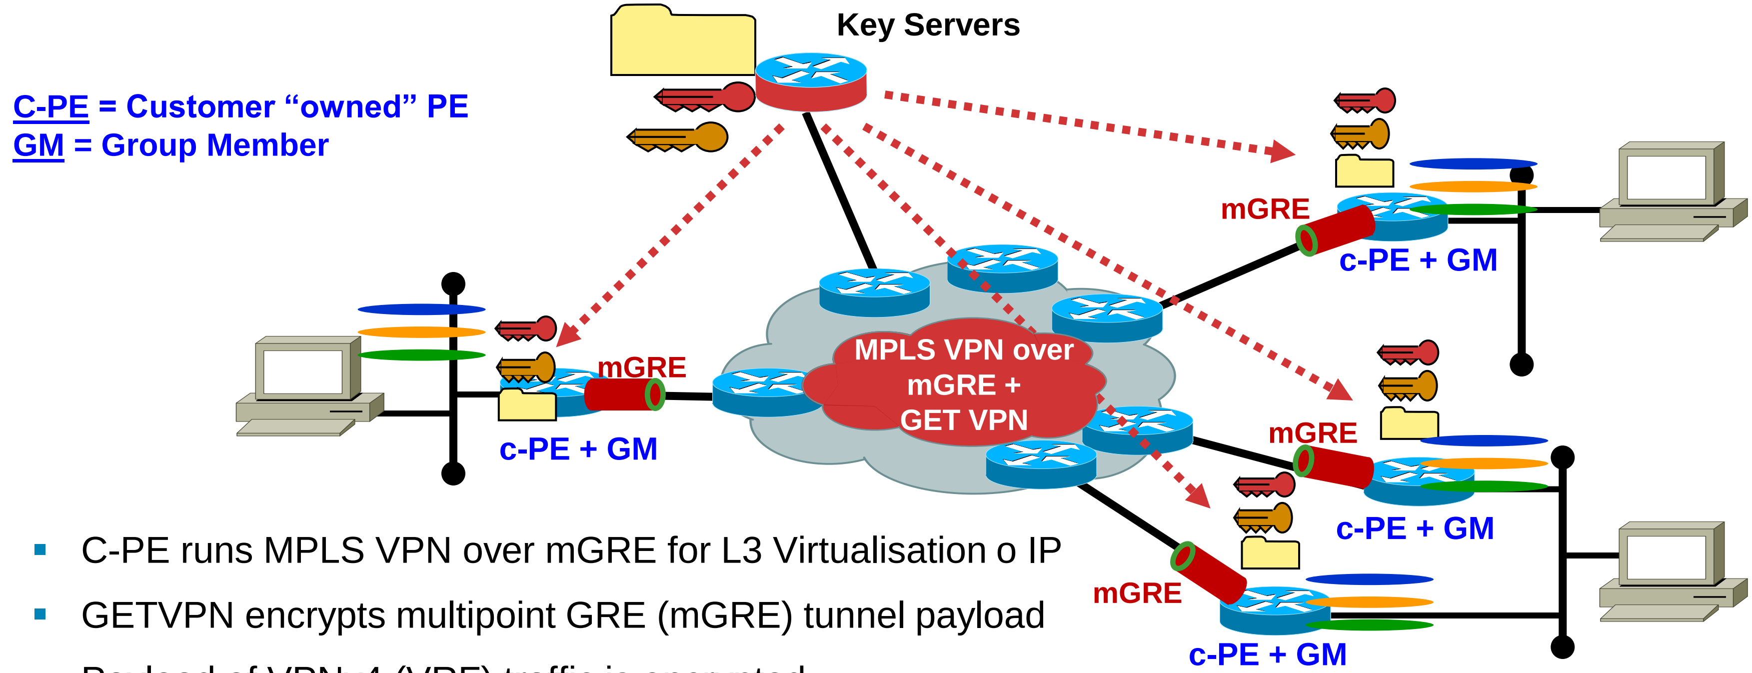


- C-PE runs MPLS VPN over mGRE for L3 Virtualisation o IP
- GETVPN encrypts multipoint GRE (mGRE) tunnel payload
- Payload of VPNv4 (VRF) traffic is encrypted
- Leverage simplicity of MPLS VPN over mGRE + GETVPN

**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

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

# Agenda

- Introduction - Network Virtualisation Drivers and Concepts
- SP WAN Transport Service Impact on L3 Virtualisation Solution Choices
- Technology and Deployment Deep-Dive for L3 Virtualised WAN
- Integrating Encryption into L3 Virtualisation Solutions
- **Recent “Innovations” Evolving in L3 Virtualisation**
- QoS/H-QoS and MTU Considerations for L3 Virtualisation Deployments
- Summary and Wrap Up




# Innovations

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

# 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**

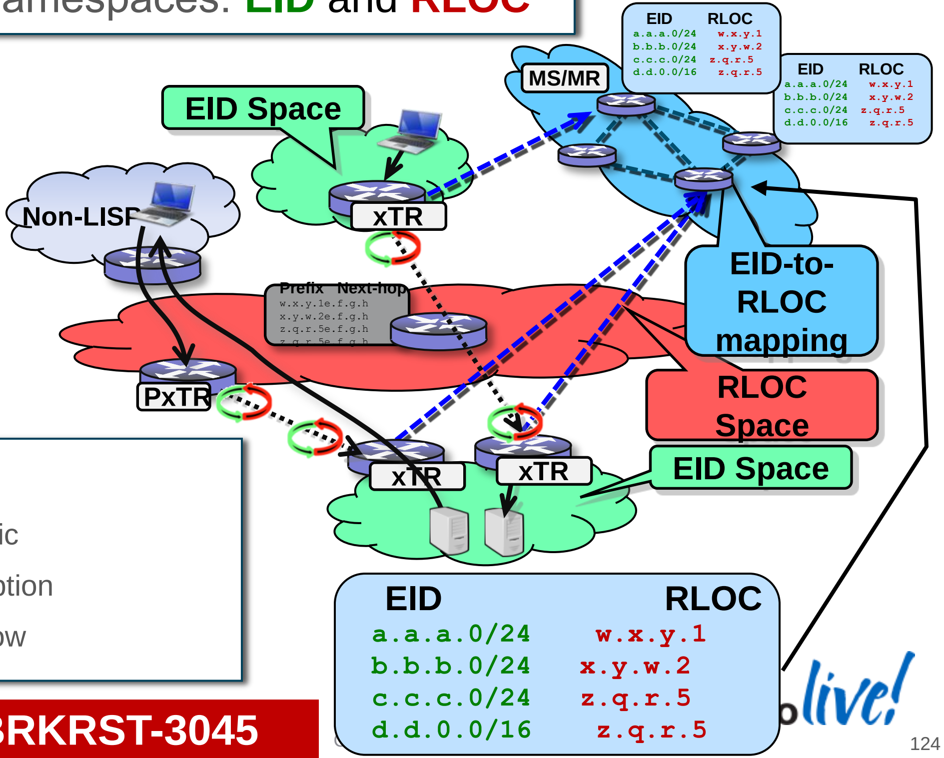


- Network-based solution
- No host changes
- Minimal configuration
- Incrementally deployable
- Support for mobility
- Address Family agnostic
- IPv4 to v6 Transition option
- In Cisco IOS/NX-OS now

**More Details on LISP Covered in Session BRKRST-3045**

# LISP Overview

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

## Today's Internet Behaviour
## Loc/ID "overloaded" semantic

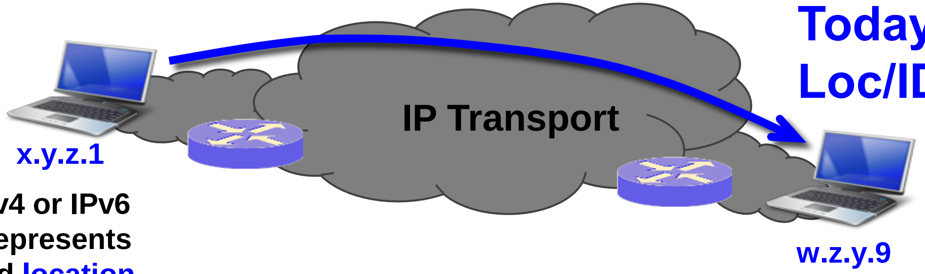


Device IPv4 or IPv6 address represents identityand location

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

## LISP Behaviour
## Loc/ID "split"

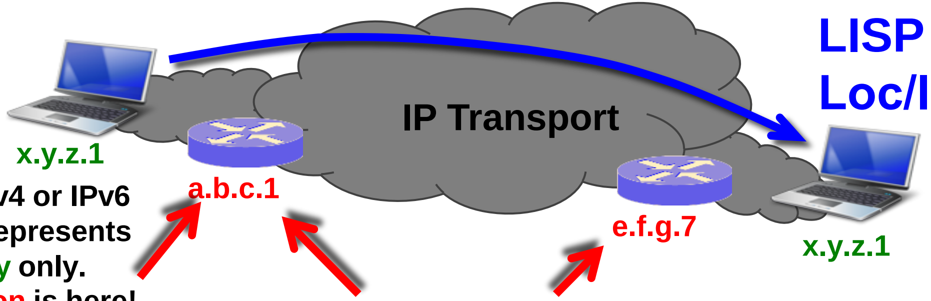


Device IPv4 or IPv6 address represents identity only.
Its location is here!

Only the location changes

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

# LISP Operations

## LISP IPv4 EID/IPv4 RLOC Header Example

draft-ietf-lisp-19

**IPv4 Outer Header: Router supplies RLOCs**

**UDP**

**LISP header**

**IPv4 Inner Header: Host supplies EIDs**

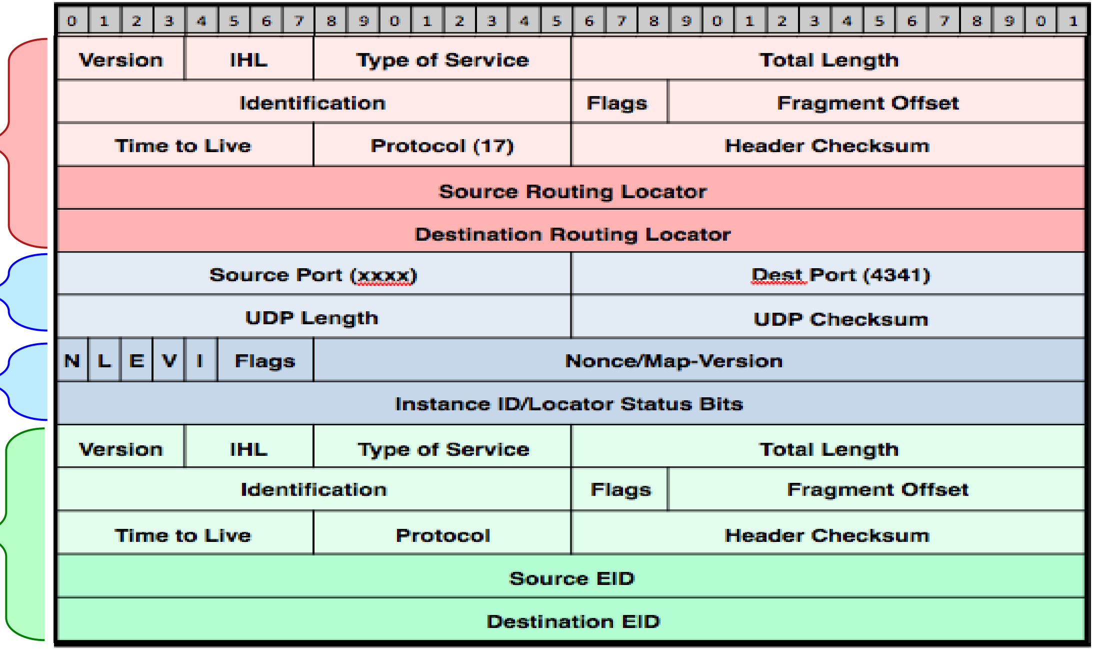

# LISP Operations

## 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**

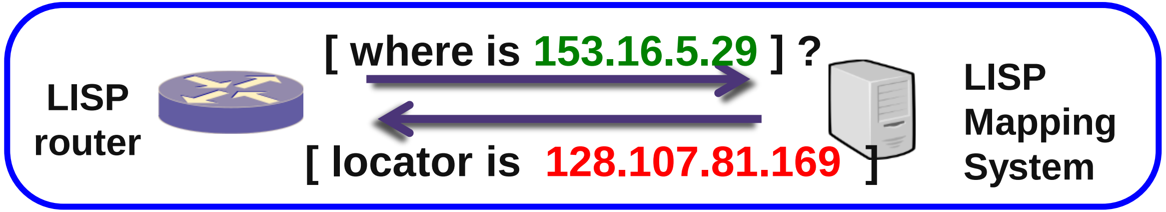


**LISP**
**Identity-to-locator**
**Mapping Resolution**

# LISP - Basic Routing Concept

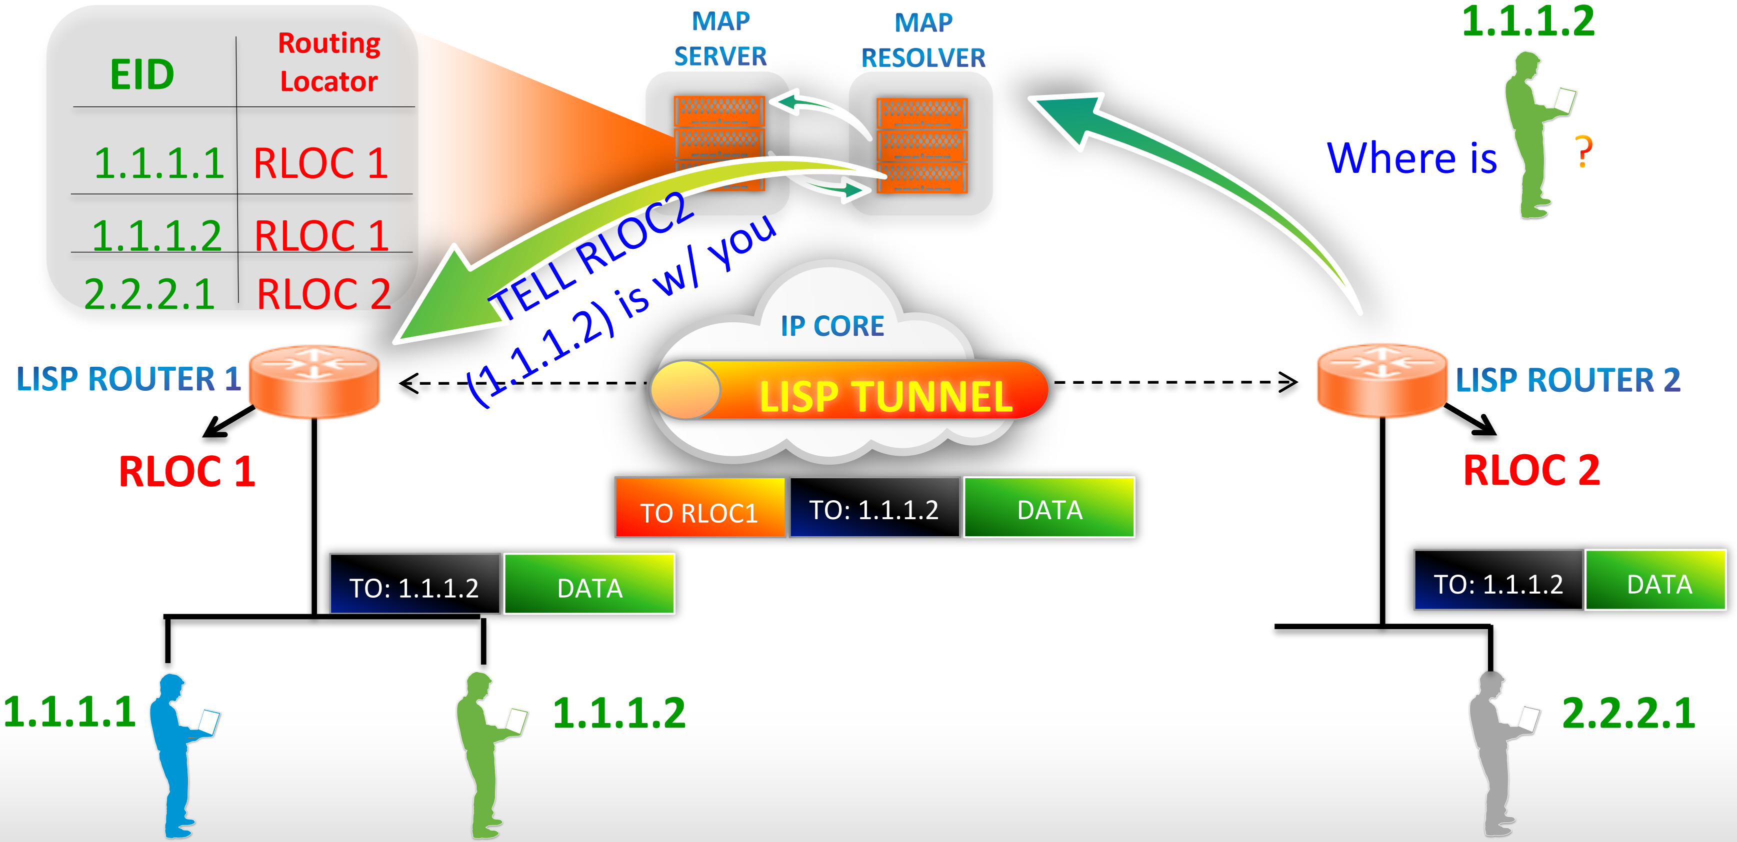

# LISP in Enterprise WAN/Branch

## Leverage LISP Framework for WAN Routing to/from Branch

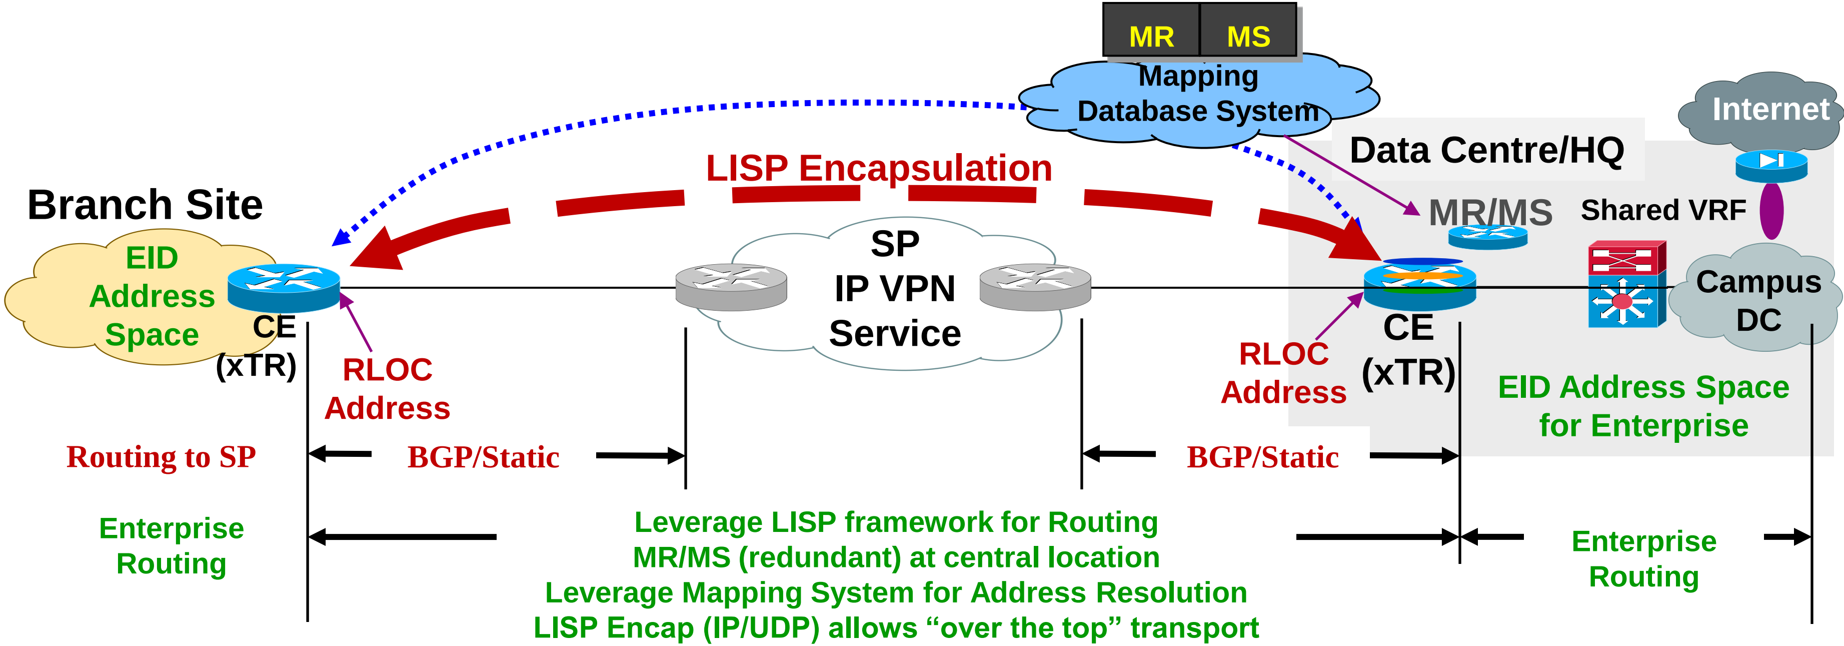


*MR = Map Resolver*
*MS = Map Server*

# LISP in Enterprise WAN/Branch

## Leverage LISP Framework for WAN Routing to/from Branch

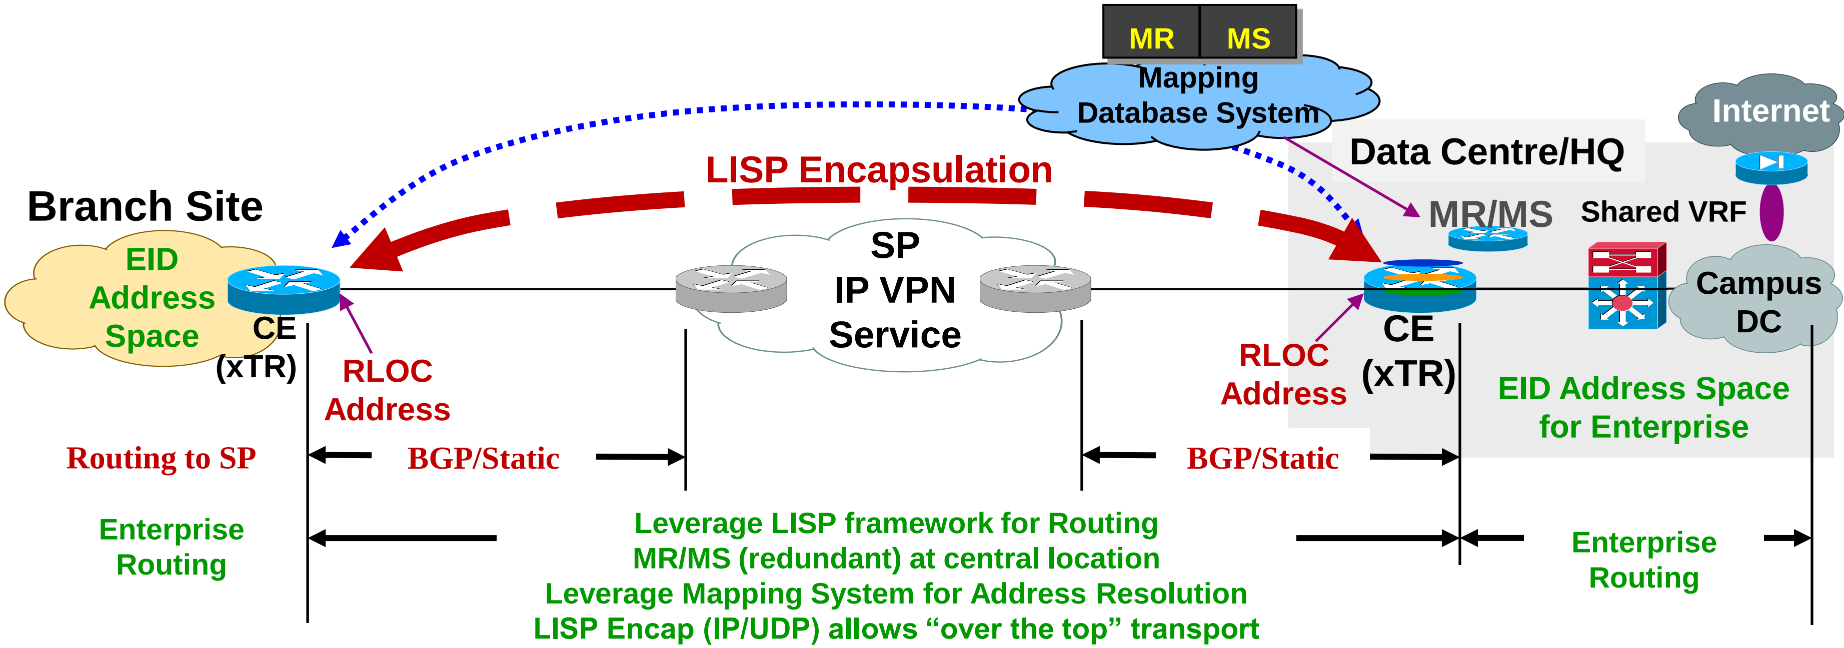


- Standard routing to SP for RLOC exchange (BGP/static)
- EID address space hidden from SP
- Uses Mapping System (MR/MS) to resolve EID location
- RLOC address carries EID to destination RLOC/EID

*MR = Map Resolver*
*MS = Map Server*

# LISP Use Cases

## The Five Core LISP Use-Cases

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

# LISP Operations

## LISP IPv4 EID/IPv4 RLOC Header Example

**IPv4 Outer Header: Router supplies RLOCs**

**UDP**

**LISP header**

**IPv4 Inner Header: Host supplies EIDs**

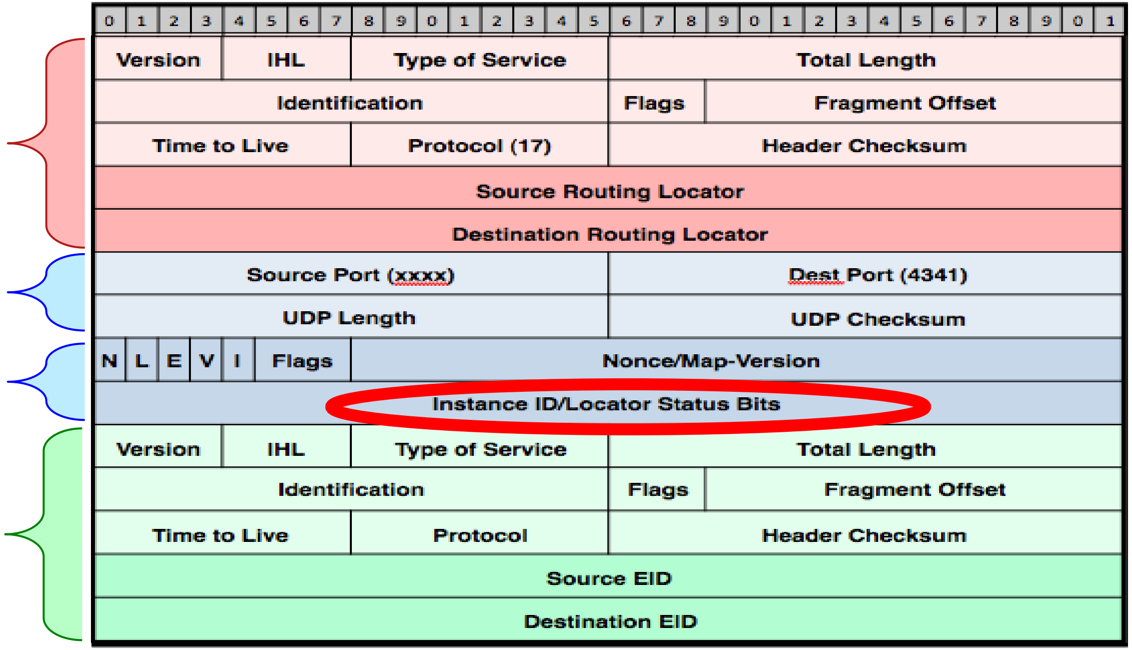

# LISP Virtualisation/VPN

## Efficient Virtualisation/Multi-Tenancy Support – Concepts…

- Because LISP considers virtualisation 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 Virtualisation/VPN

## Efficient Virtualisation/Multi-Tenancy Support – Shared Model…

- Shared Model – at the device level
  - 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

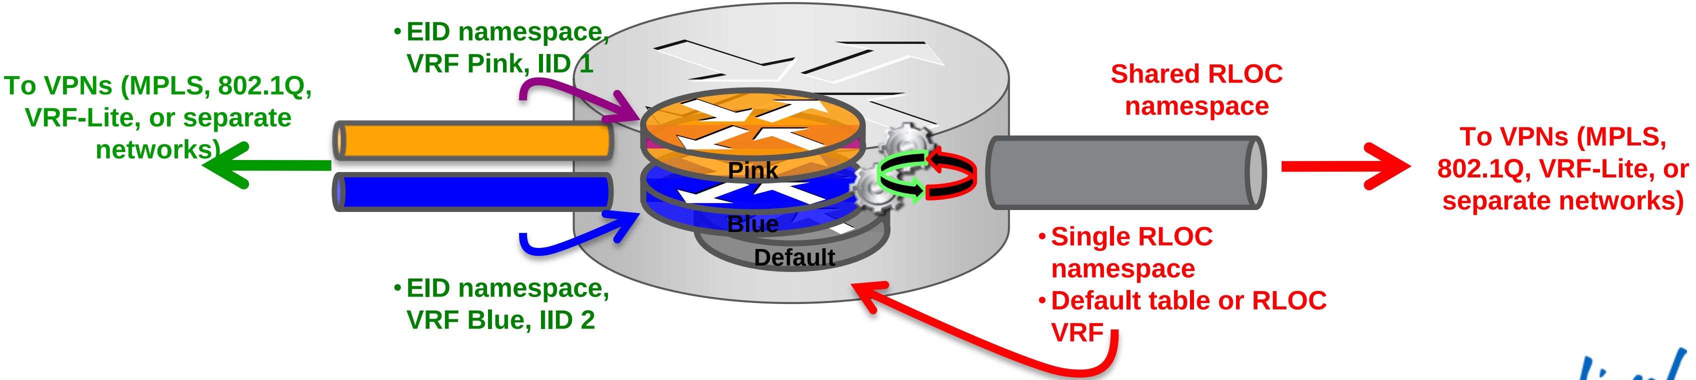

# LISP Virtualisation/VPN

Efficient Virtualisation/Multi-Tenancy Support – Parallel Model…

- Parallel Model – at the device level
  - 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

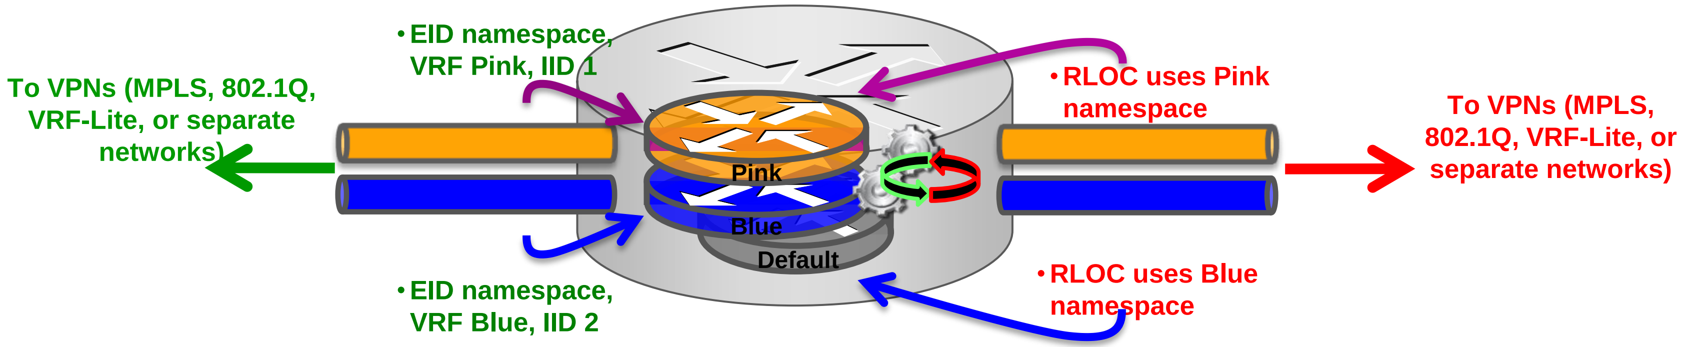

# LISP in Enterprise WAN/Branch

## Leverage LISP Framework for WAN Routing to/from Branch

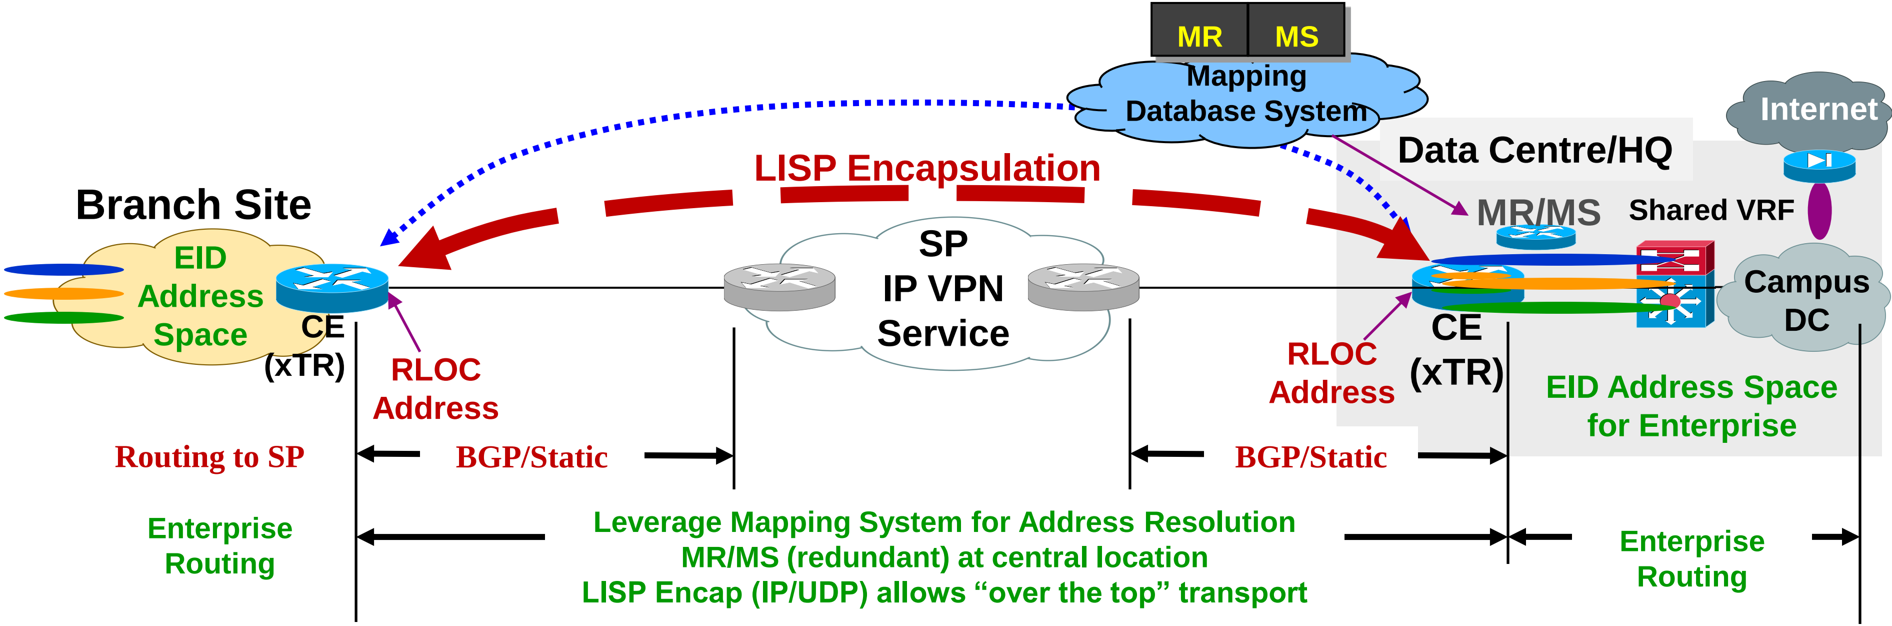


- Allows network segmentation on xTR (viewed as CE in L3 VPN model)
- PE routers require minimal routes (RLOC address only, which only SP knows)
- VRF Segmentation is applied to CE/xTR
- Offers another "over the top" Virtualisation solution (VRF capabilities
- Can leverage GET VPN for additional data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

*MR = Map Resolver*
*MS = Map Server*

# LISP References

## LISP Information and Mailing Lists

- LISP Information
  - IETF LISP Working Group — http://tools.ietf.org/wg/lisp/
  - LISP Beta Network Site — http://www.lisp4.net or http://www.lisp6.net
  - Cisco LISP Site — http://lisp.cisco.com (ipv4 and IPv6)
  - Cisco LISP Marketing Site — http://www.cisco.com/go/lisp/
- LISP Mailing Lists
  - IETF LISP Working Group — http://tools.ietf.org/wg/lisp/
  - LISP Interest (public) — lisp-interest@puck.nether.net
  - Cisco LISP Questions — lisp-support@cisco.com
  - LISPmob Questions — users@lispmob.org

# Agenda

- Introduction - Network Virtualisation Drivers and Concepts
- SP WAN Transport Service Impact on L3 Virtualisation Solution Choices
- Technology and Deployment Deep-Dive for L3 Virtualised WAN
- Integrating Encryption into L3 Virtualisation Solutions
- Recent “Innovations” Evolving in L3 Virtualisation
- **QoS/H-QoS and MTU Considerations for L3 Virtualisation Deployments**
- Summary and Wrap Up

# QoS with GRE, MPLS over GRE

## ToS/EXP Reflection Behaviour for “transit traffic” Through the Router

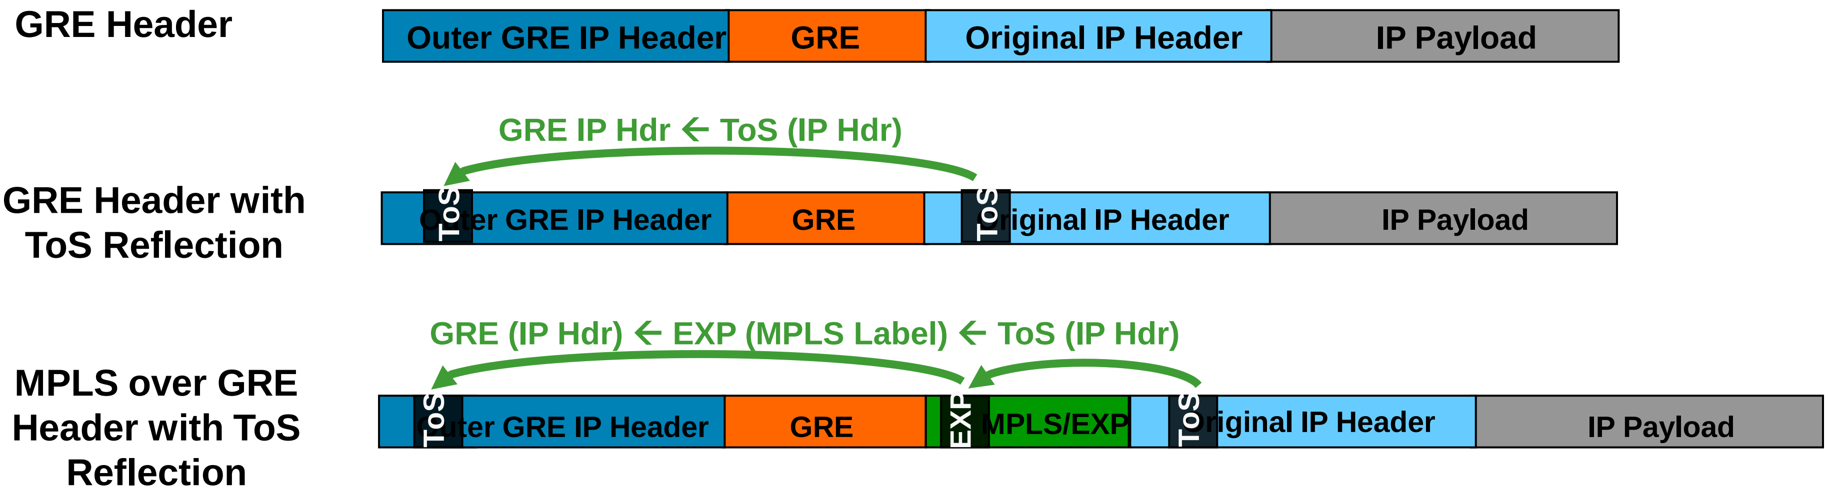


- Router will copy original ToS marking to outer GRE header
- For MPLS over GRE, the EXP marking is copied to the outer header of the GRE tunnel
- This allows the IPv4 “transport” to perform QoS on the multi-encapsulated packet

**Caveats:**

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

Cisco live!

# 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 Virtualisation

- **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. Prioritsed Model**

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

# QoS Deployment with Network Virtualisation

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

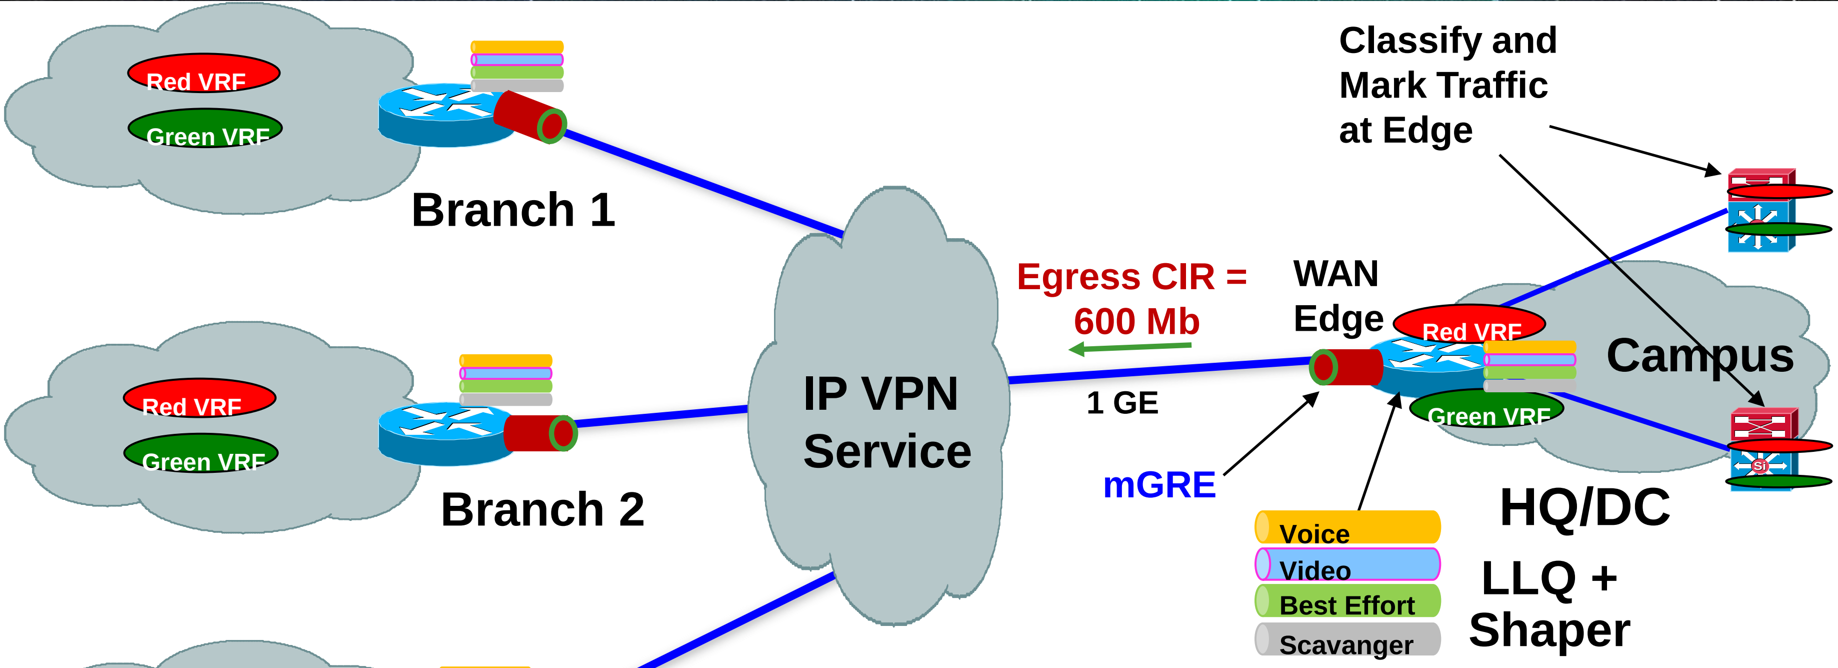


- **1st Layer – GRE Tunnel (Parent)**
  Shaper per GRE
- **2nd Layer - Service Queuing per GRE (child)**
  Queuing determines order of packets sent to shaper
- **H-QoS policy applies to main interface (not mGRE)**

# Hierarchical QoS Example

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

### Two MQC Levels

```
Policy-map PARENT
  class class-default
   shape average 600000000
   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
```

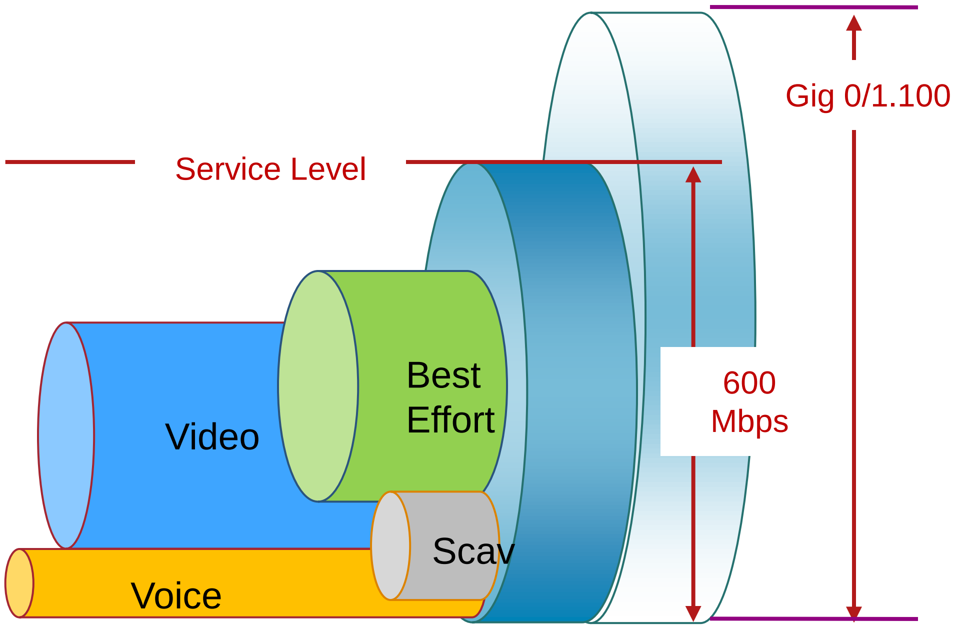

# QoS for Virtualisation – Summary

- **Aggregate** QoS model is the simplest and most straight forward approach **(Recommended)**
- Simplification using the **Aggregate** model recommends:
  - Traffic class marking identical to non Virtualisation scheme
  - Traffic class marking identical between VRF's
  - Leverage H-QoS on virtualised interfaces (GRE, .1Q)
  - Router dynamically copies ToS→EXP→ToS (GRE)
- **Prioritised** VRF model can be used to prefer traffic originating in one VRF over another (Becomes more complex, through techniques such as Policy-Based Routing, MPLS-TE, or a combination of both)
- **Summary: Consider implementing the same QoS approach that is used for non-virtualised deployments, when enabling QoS in virtualised enterprise network designs**

# MTU Considerations with GRE Tunnels

## Challenges


- 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

## Point to Point GRE

- ✓ 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

```
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

- 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

# Innovations Worth Investigating Further

- Easy Virtual Networking (EVN)
  - (in backup slides)
- VRF Aware Services Interface (VASI)
  - (in backup slides)
- EIGRP Over The Top
- Software Defined Networking (SDN)
  - Network Virtualisation Use cases
- Flex VPN in Virtualised Networking Environments
- Using MP-BGP control plane, with VxLAN data plane

# Agenda

- Introduction - Network Virtualisation Drivers and Concepts
- SP WAN Transport Service Impact on L3 Virtualisation Solution Choices
- Technology and Deployment Deep-Dive for L3 Virtualised WAN
- Integrating Encryption into L3 Virtualisation Solutions
- Recent “Innovations” Evolving in L3 Virtualisation
- QoS/H-QoS and MTU Considerations for L3 Virtualisation Deployments
- **Summary and Wrap Up**

# WAN Virtualisation - Key Takeaways

- The ability for an enterprise to extend Layer 3 (L3) Virtualisation 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 Virtualisation 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 Virtualisation) and how they can help simplify network Virtualisation in the WAN for future designs
- **Leverage the technology, but "Keep it Simple" when possible ☺**

# Recommended Reading

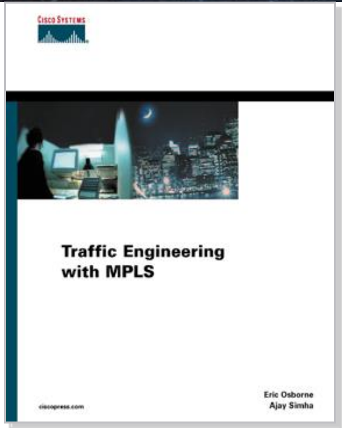


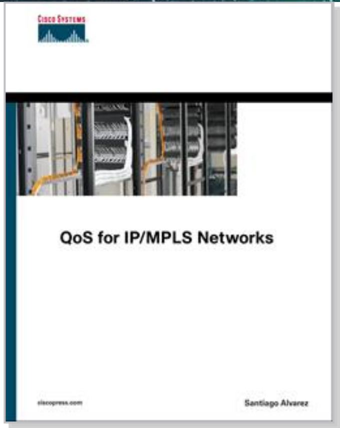


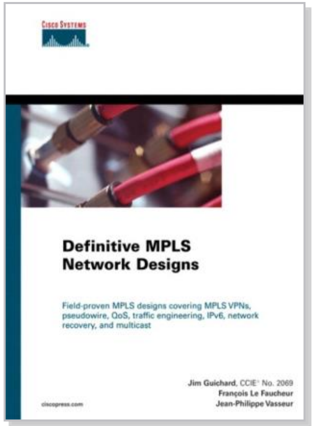


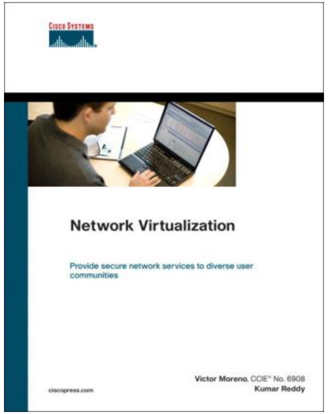


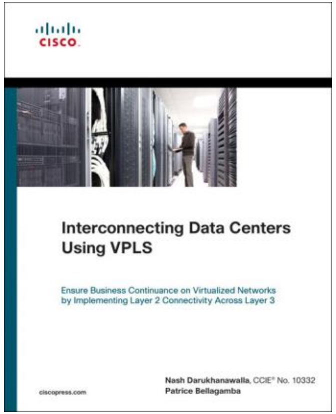




# Q & A

# Complete Your Online Session Evaluation

## Give us your feedback and receive a Cisco Live 2014 Polo 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 www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 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™