What You Make Possible

TOMORROW
starts here.

Network Virtualisation Design Concepts Over the WAN

BRKRST-2045

TOMORROW starts here.

Session Assumptions & Disclaimers

- Participants should have a:
  - Intermediate knowledge of IP routing, IP/GRE tunnels, VRF's, and WAN design fundamentals and technologies
  - Basic knowledge of MP-BGP, MPLS VPNs, 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

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

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

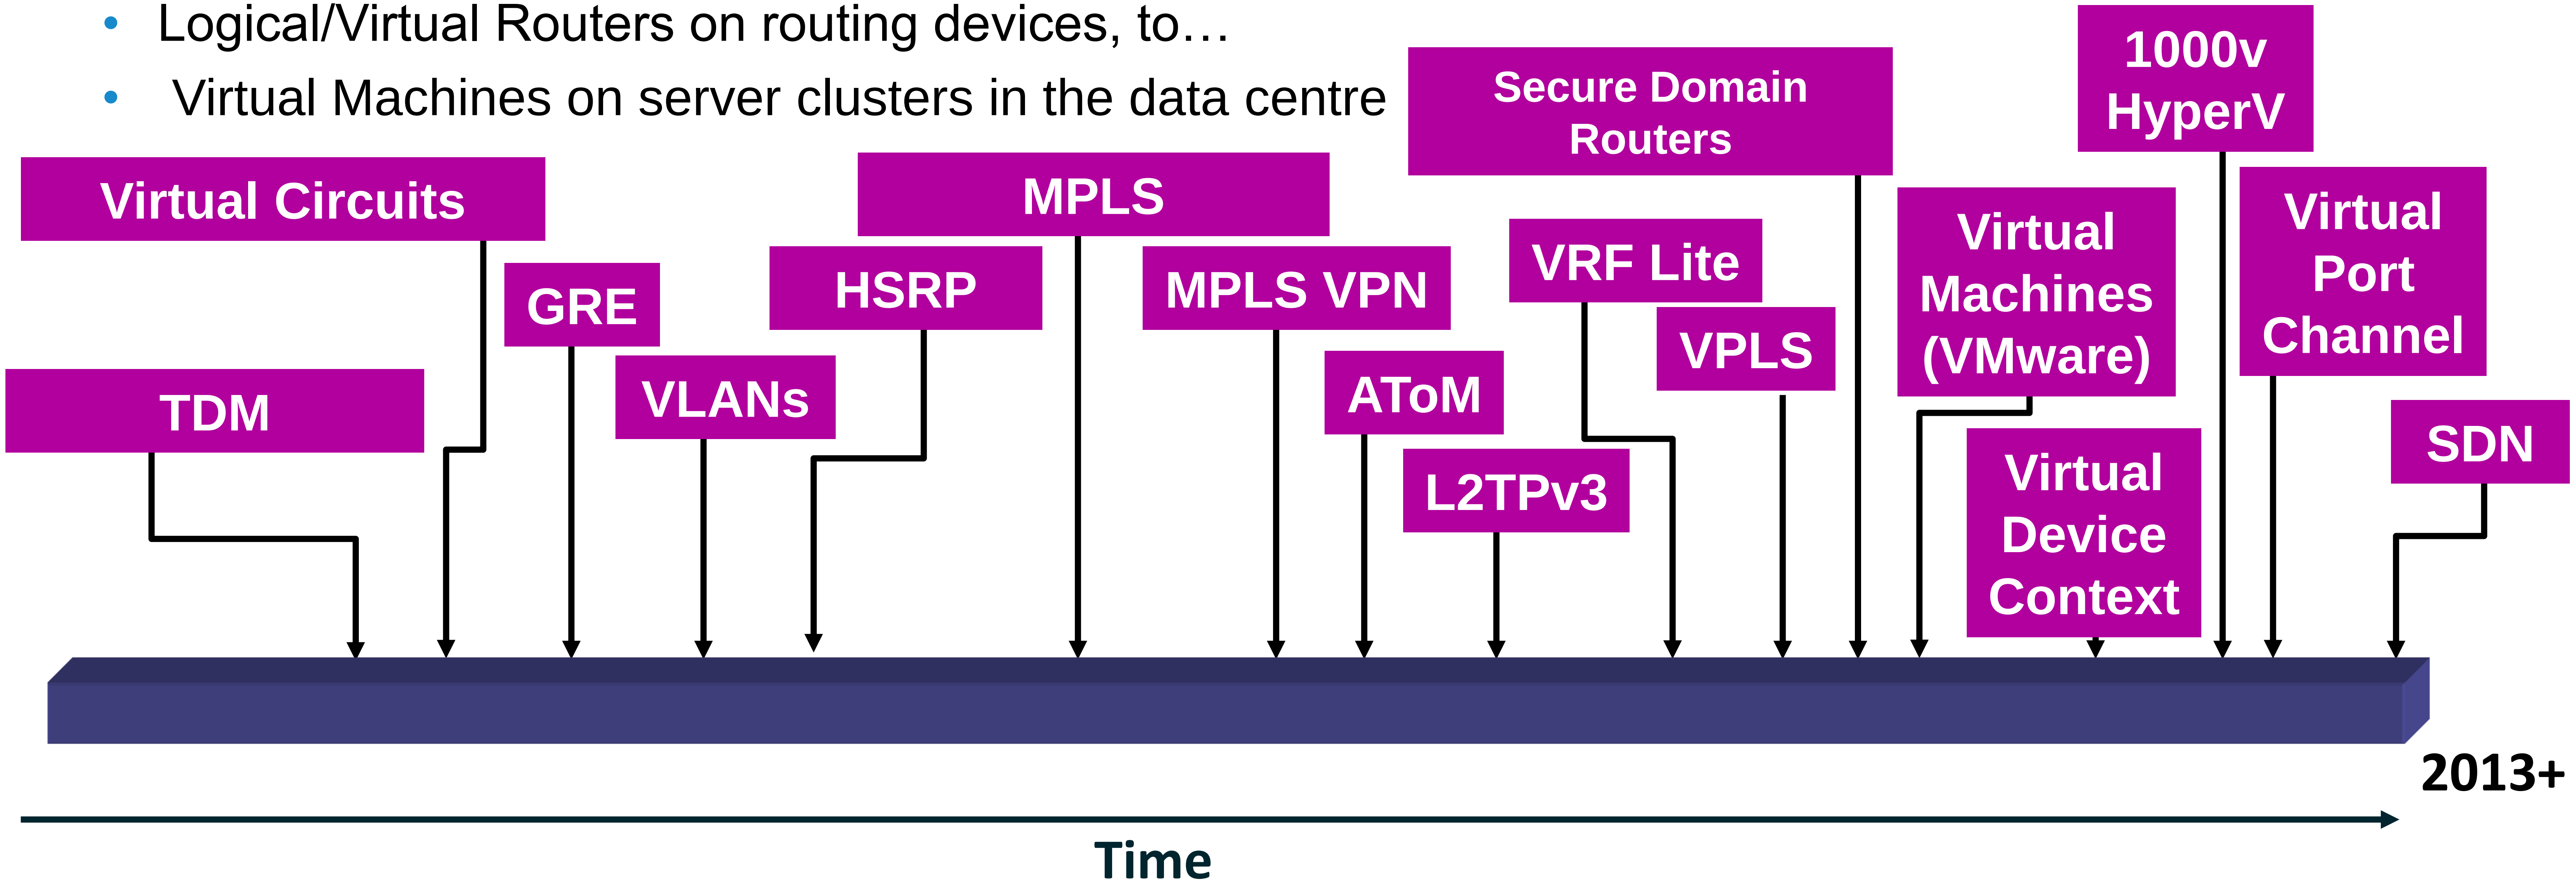

What Is Enterprise "Network" Virtualisation?

- Giving One physical network the ability to support multiple 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
- Allows for better utilisation of network resources

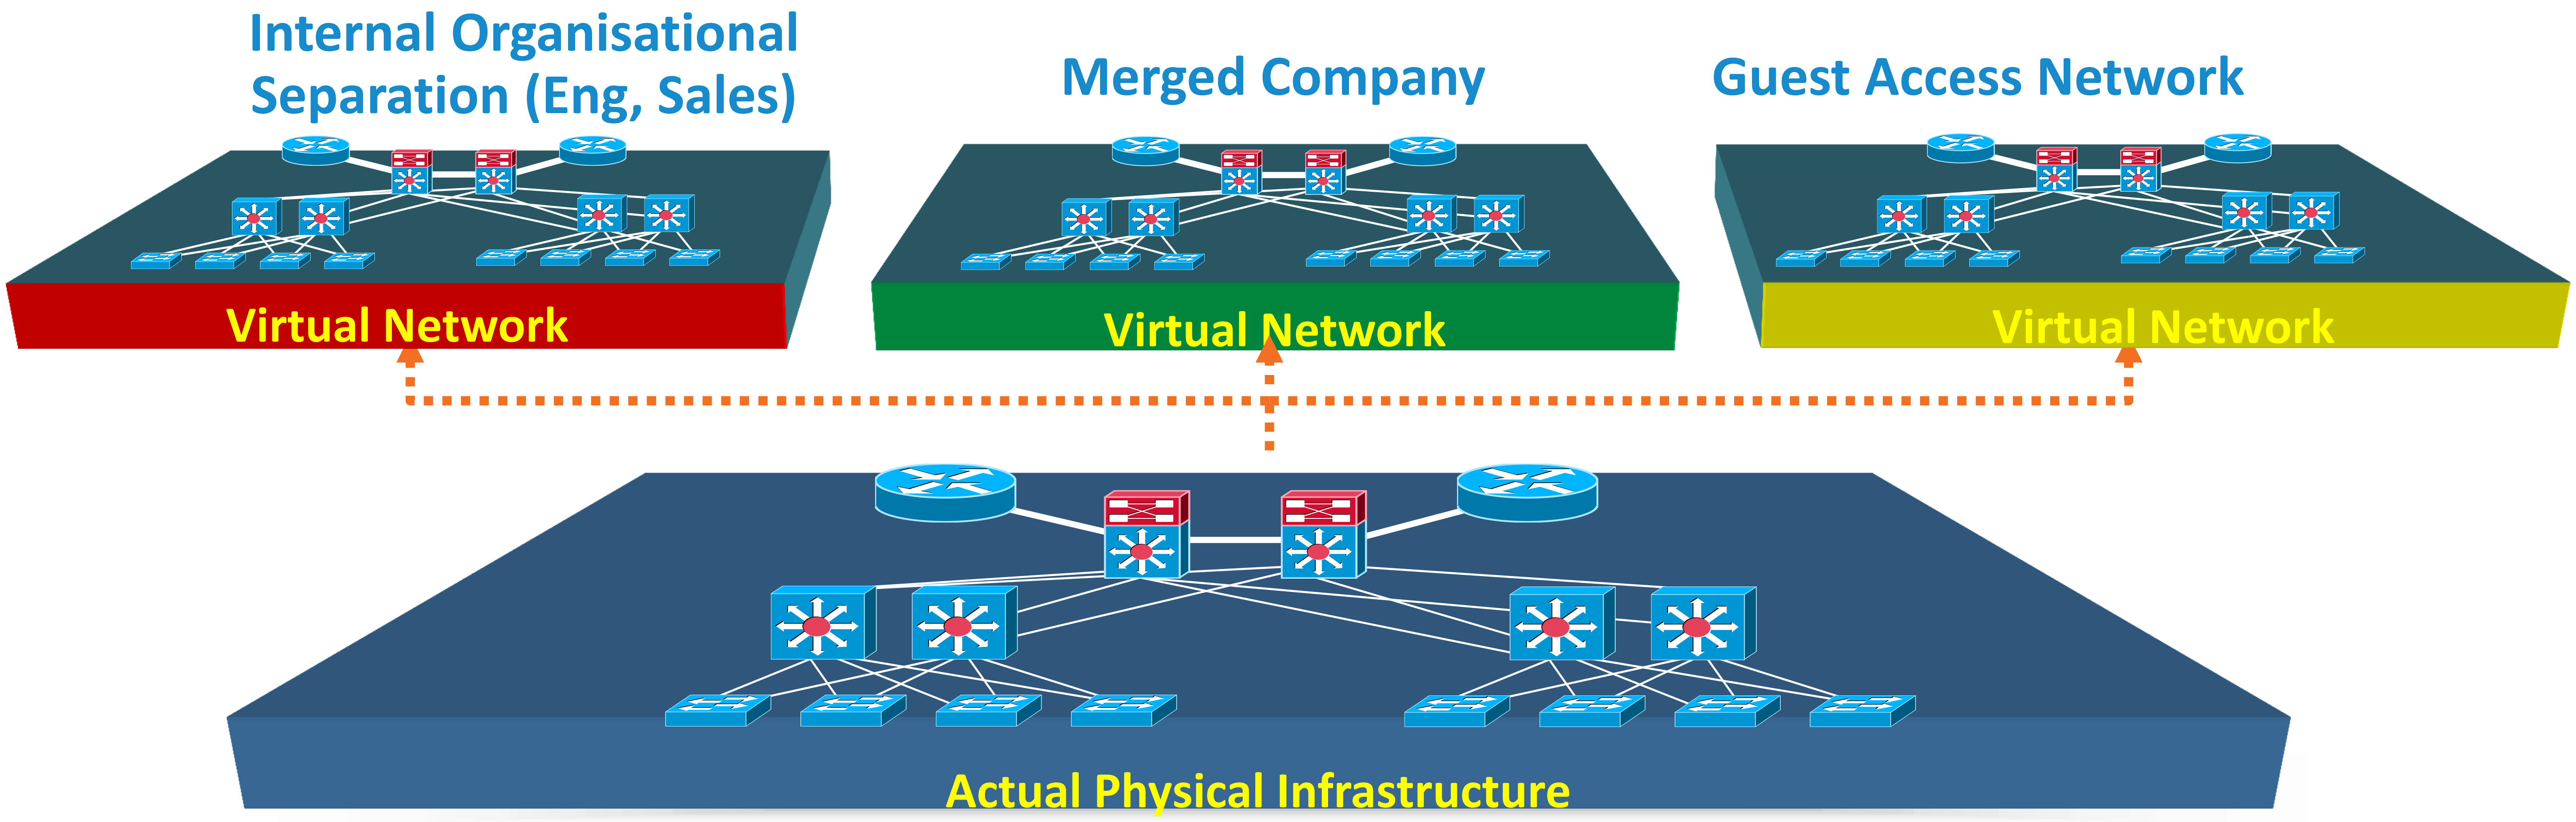

Why Network Virtualisation?

Key Benefits

- **Cost Reduction**—allowing a single physical network the ability to support 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 Multi-tenant DC's for Cloud Computing
- **Common Use Cases**
  - Guest Access, Airports, Cloud Computing IaaS, Physical Security Separation, Company Mergers
  - Regulation/Compliance – Health Care (HIPPA), Credit Card (PCI)

Enterprise Network Virtualisation

Key Building Blocks

Device Partitioning

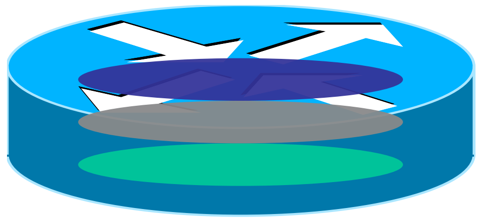

"Virtualising" the Routing and Forwarding of the Device

Virtualised Interconnect

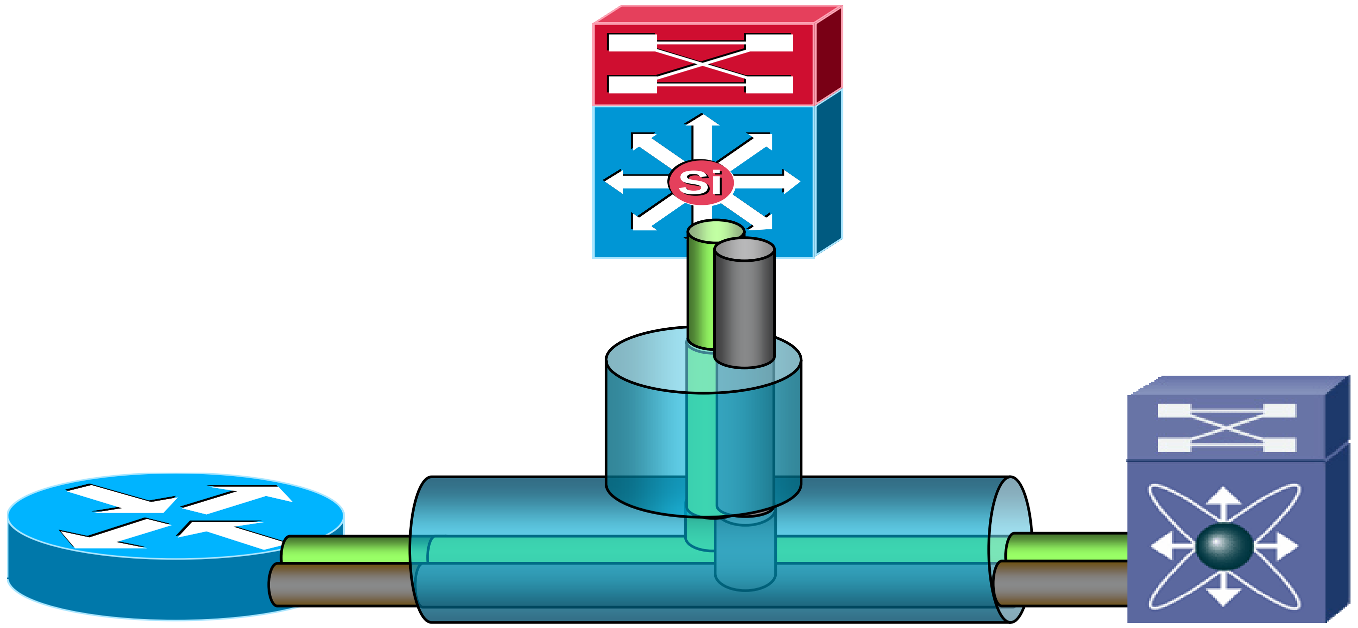

Extending and Maintaining the "Virtualised" Devices/Pools over Any Media

Device Pooling

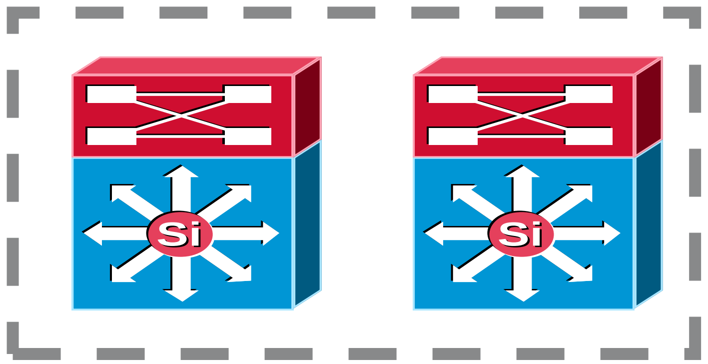

"Virtualising" Multiple Devices to Function as a Single Device

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)

Virtualised 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

- VSS
- Stackwise
- Virtual Port Channel (vPC)
- HSRP/GLBP

Device Partitioning – Layer 2 & 3

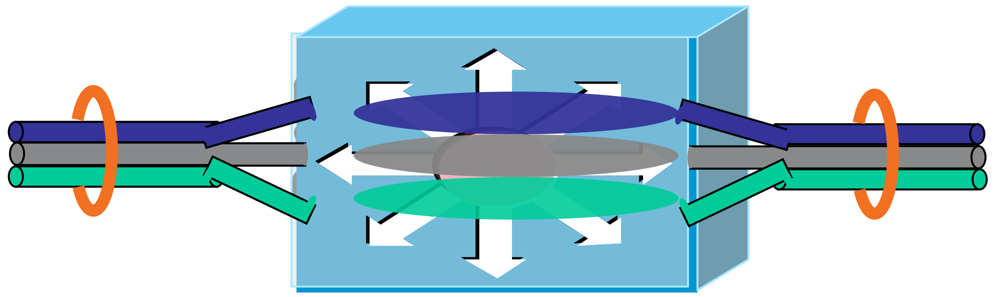

VLAN—Virtual LAN

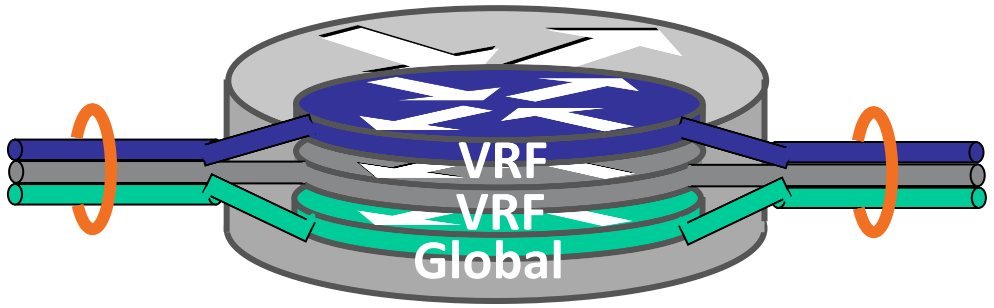

VRF—Virtual Routing and Forwarding

- Virtualise at Layer 2 forwarding
- Associates to one or more L2 interfaces on switch
- Has its own MAC forwarding table and spanning-tree instance per VLAN
- Interconnect options?
  - VLANs are extended via phy cable or virtual 802.1q trunk

- Virtualise at Layer 3 forwarding
- Associates to one or more Layer 3 interfaces on router/switch
- Each VRF has its own
  - Forwarding table (CEF)
  - Routing process (RIP, OSPF, BGP)
- Interconnect options (VRF-Lite)?
  - 802.1q, GRE, sub-interfaces, physical cables, signalling

Enterprise Network Virtualisation

The Building Blocks – Example Technologies

Device Partitioning

VLANs

VRFs
VDCs

SDR (XR)

FW Contexts

Virtualised 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

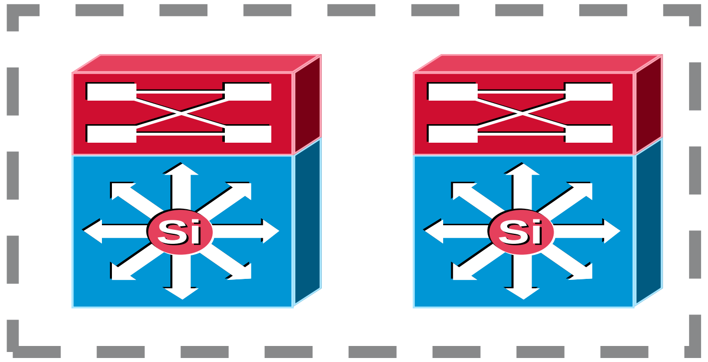

Virtual Sw System (VSS)

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

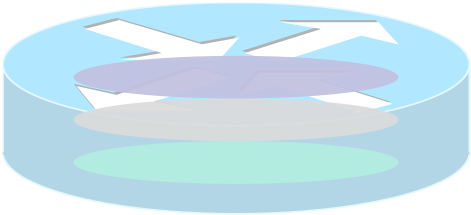

VLANs

VRFs

EVN

(Easy Virtual Network)

VDC (NX-OS)

(Virtual Device Context)

SDR (IOS-XR)

(Secure Domain Routers)

FW Contexts

Virtualised Interconnect

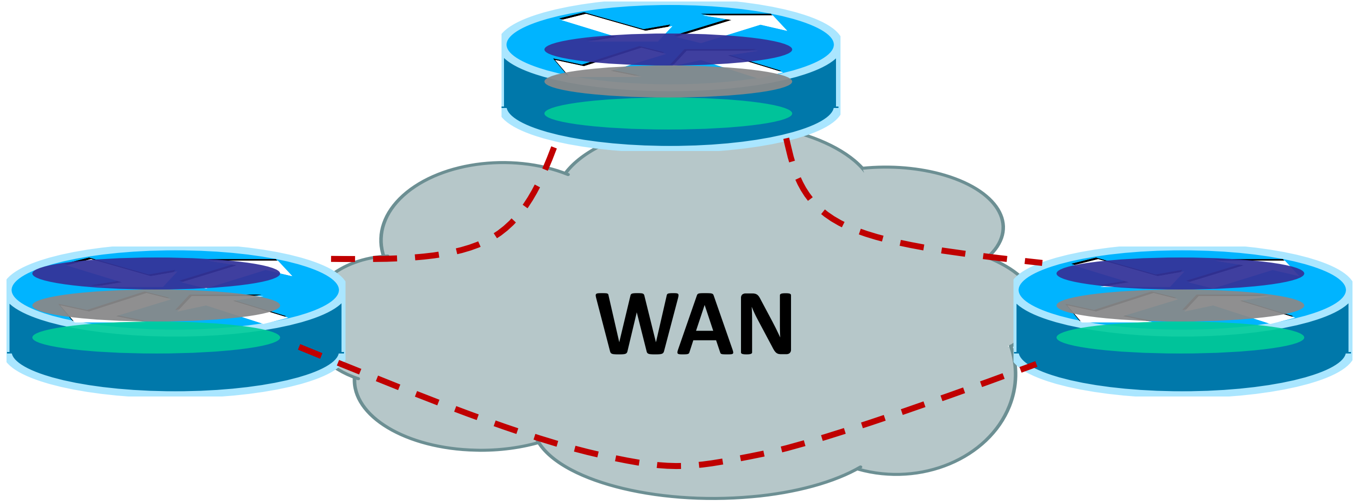

<u>L3 VPNs</u> – MPLS VPNs, VRF-Lite, MPLS VPN or VRF-Lite over IP

<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, LISP Virtualisation, VxLAN

MPLS-TP = MPLS Transport Profile

Device Pooling

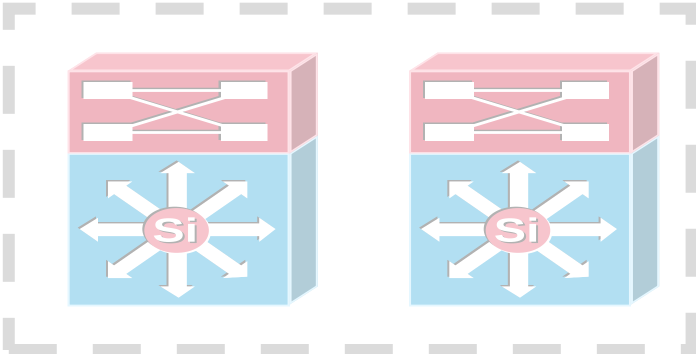

VSS

Stackwise

Virtual Port Channel (vPC)

HSRP/GLBP

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

Today's WAN Transport Options

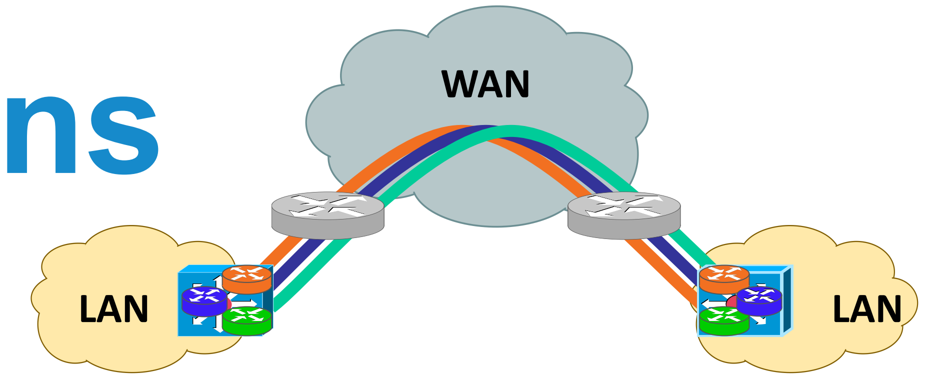

Topologies

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

VPN Offerings

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

Media

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

Public Transport

- L3 – Public (Internet)
- L3 – Broadband/WiFi

Network Deployment Options - Self Deployed vs. L3 Managed

Self Deployed IP Backbone

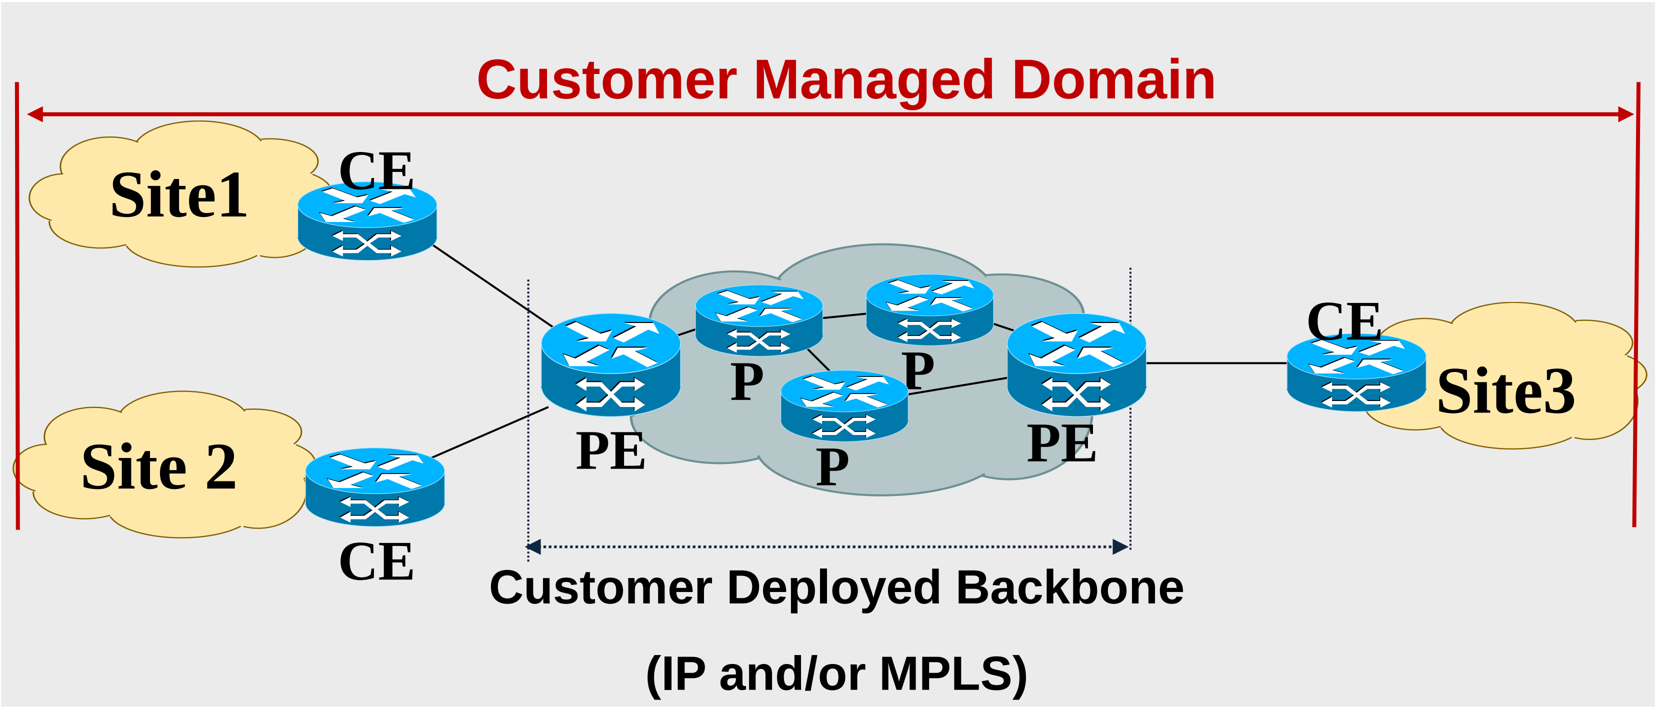

- MPLS Deployment Common for Service Richness
- Customer manages and owns:
  - IP routing, provisioning
  - Transport links for PE-P, P-P, PE-CE
  - SLA’s, to “end” customer, QoS
  - How rapidly services are turned up
- **Allows customer full control E2E**

SP Managed IP VPN Service

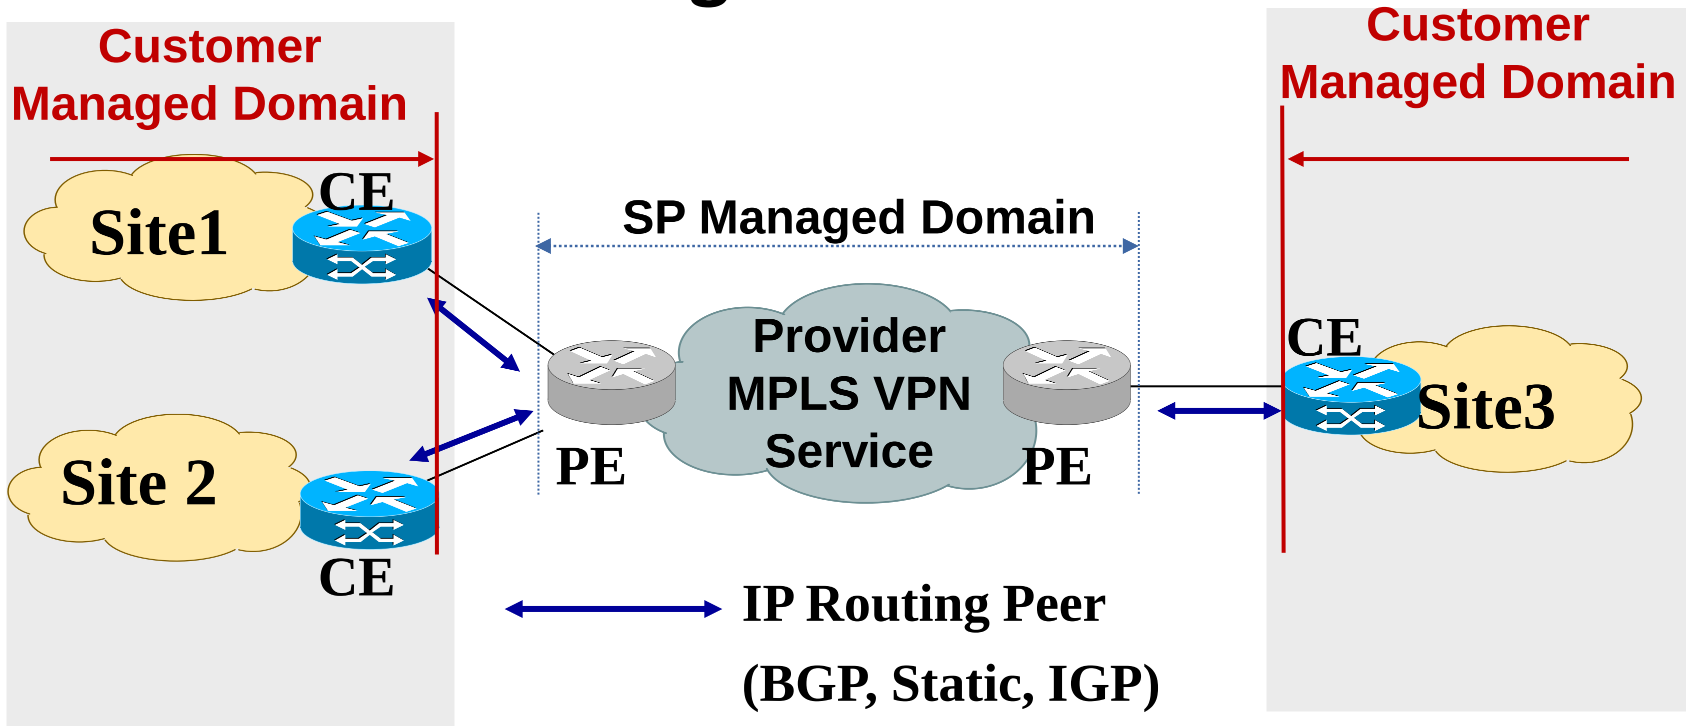

*** No Labels Are Exchanged with the SP**

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

Network Deployment Options - Self Deployed vs. L3 Managed

Service Provider (SP) Transport Options

Self Deployed IP Backbone

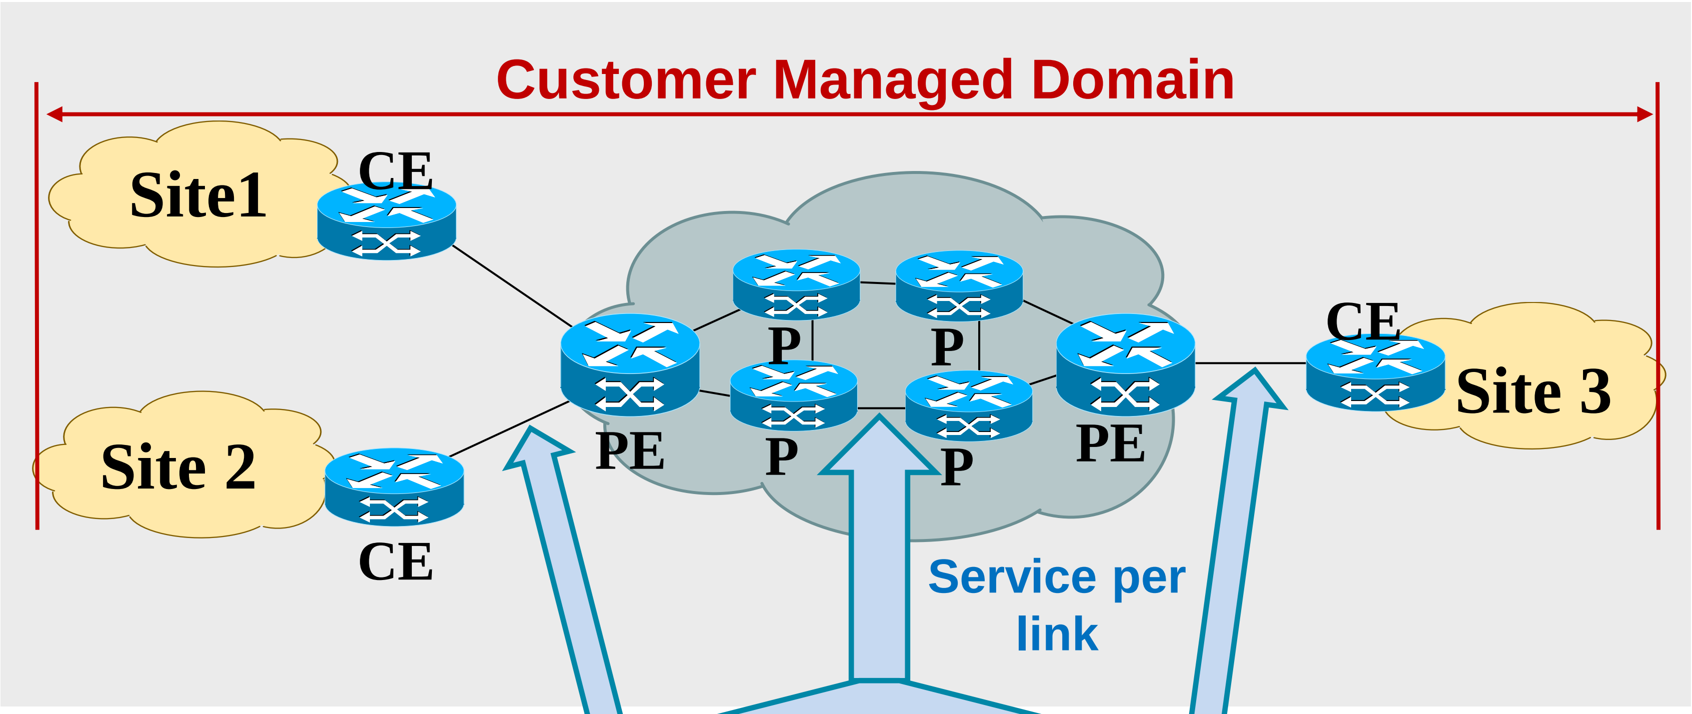

Point to Point Service Examples

- Ethernet Virtual Circuit (EVC/E-LINE)
- Ethernet Multipoint Service (E-LAN)
- T1/E1, T3/E3
- SONET/SDH
- Optical/Lambda

SP Managed IP VPN Service

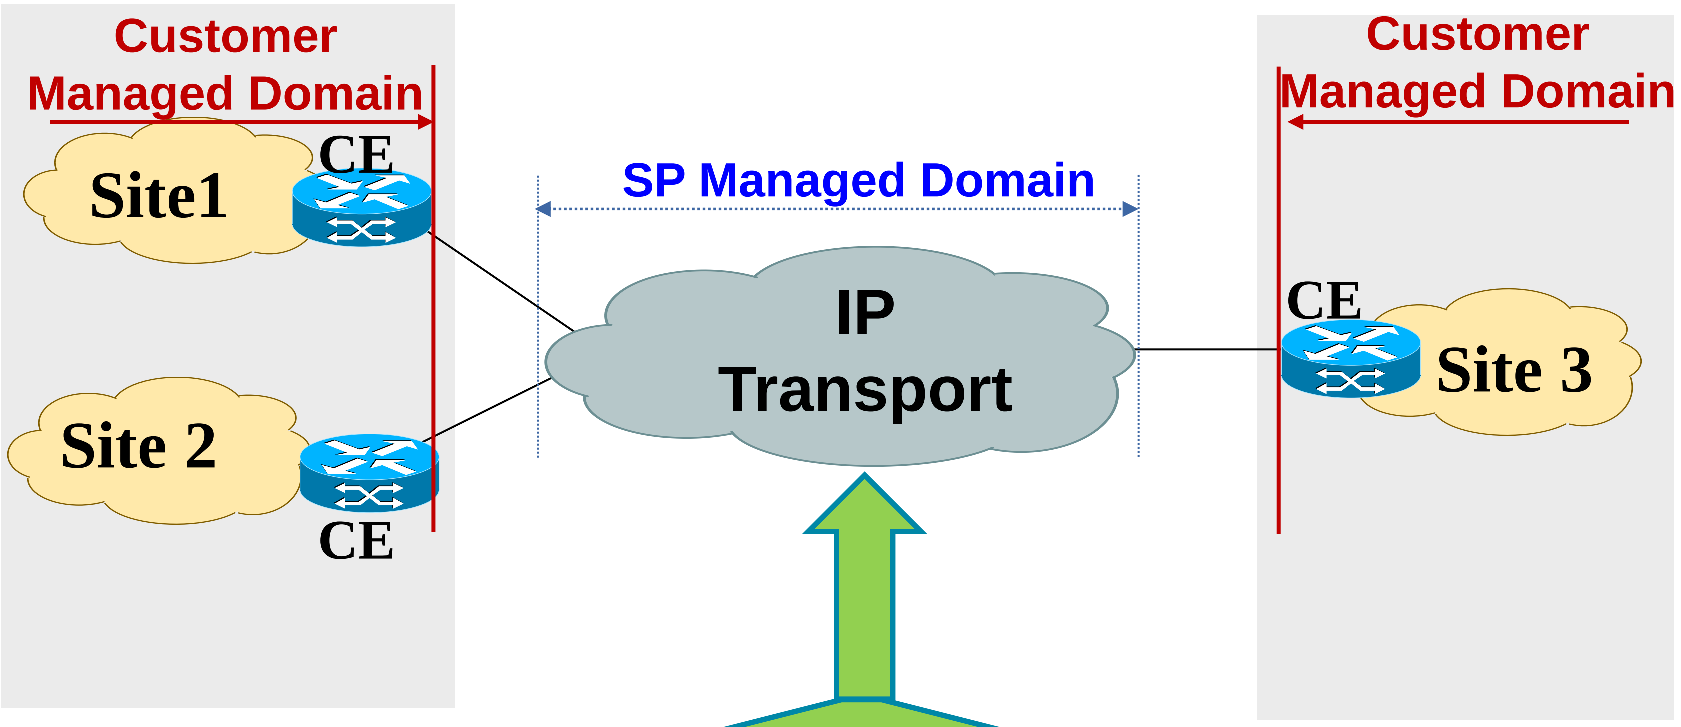

IP Overlay Transport Service Examples

- IP VPN Service (private)
- Internet Transport (public)
- Last-mile broadband

Note: Any of the “Point to Point Service” examples could also apply to transport options for the “Managed Transport” service as well.

Network Deployment Options - Self Deployed vs. L3 Managed

Point to Point Service Model

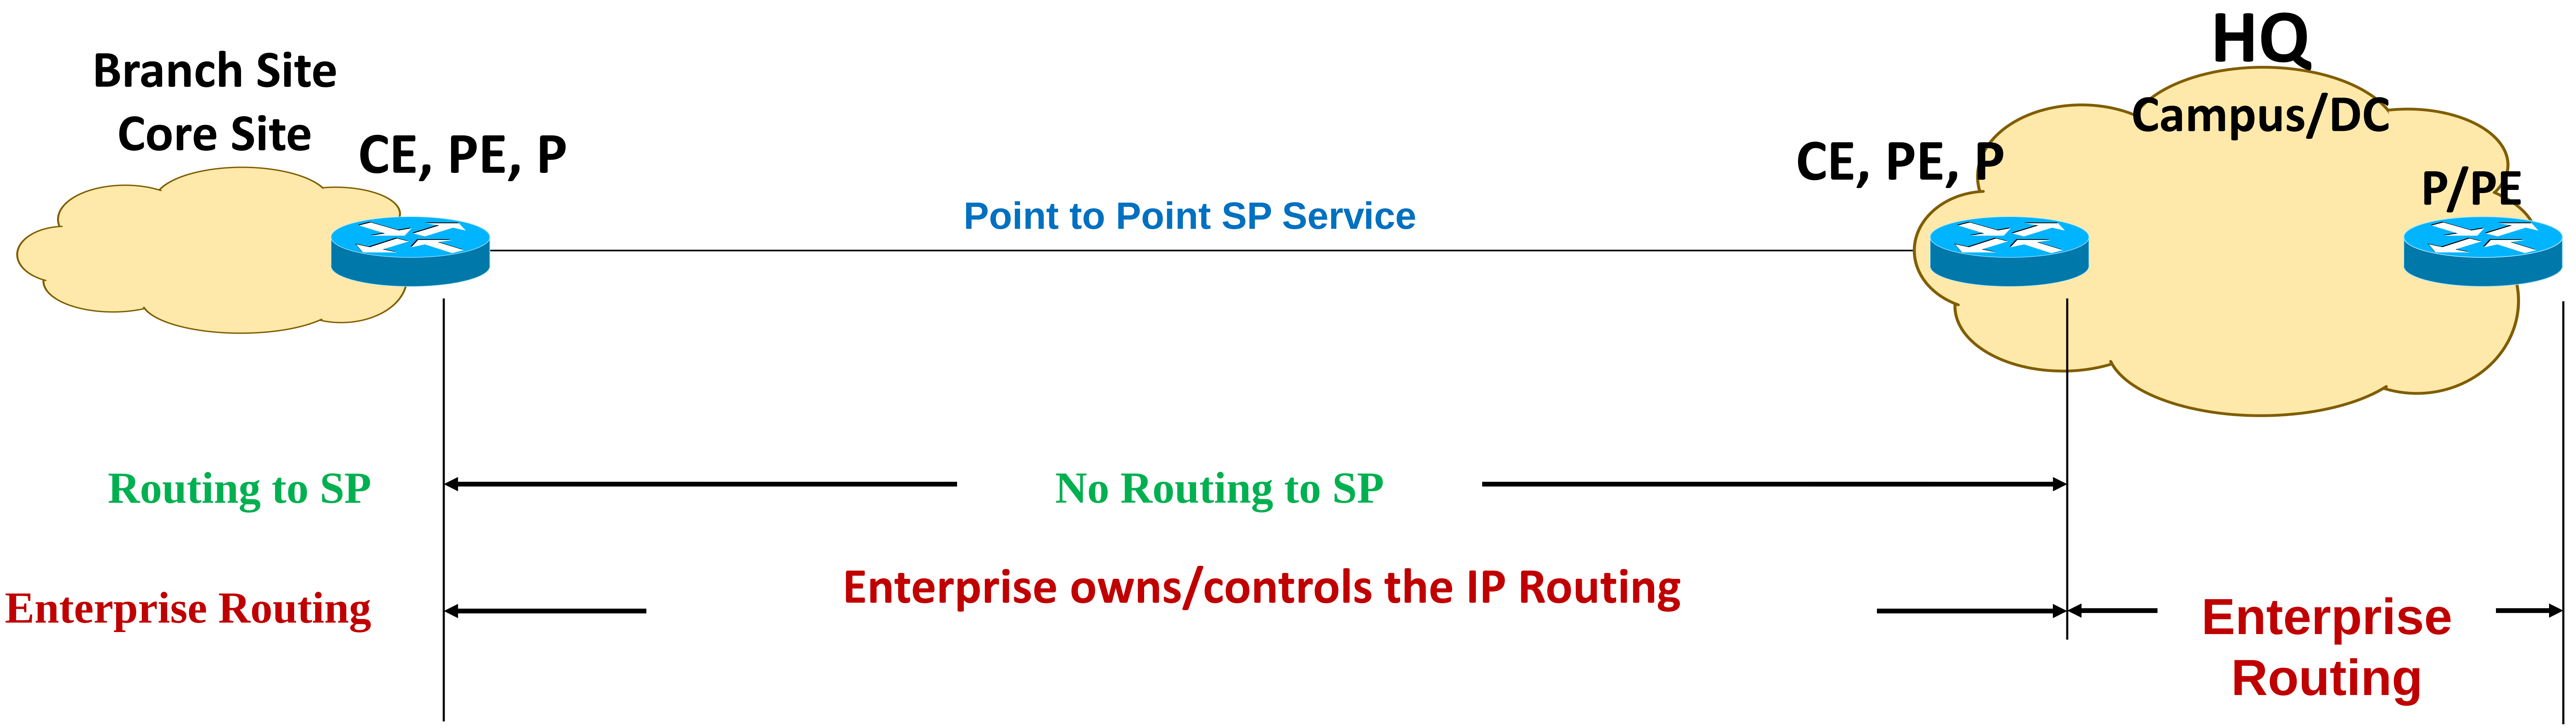

- No IP routing required to the SP
- Routing is controlled by the Enterprise
- Enterprise-to-SP interaction may require certain L2 parameters, dictated by the SP

  EXAMPLE: Ethernet – specific 802.1Q tag, ATM - VPI/VCI, Frame Relay - DLCI
- For Ethernet services, QoS markings will be dictated by the SP

Network Deployment Options - Self Deployed vs. L3 Managed

IP Overlay Transport Model

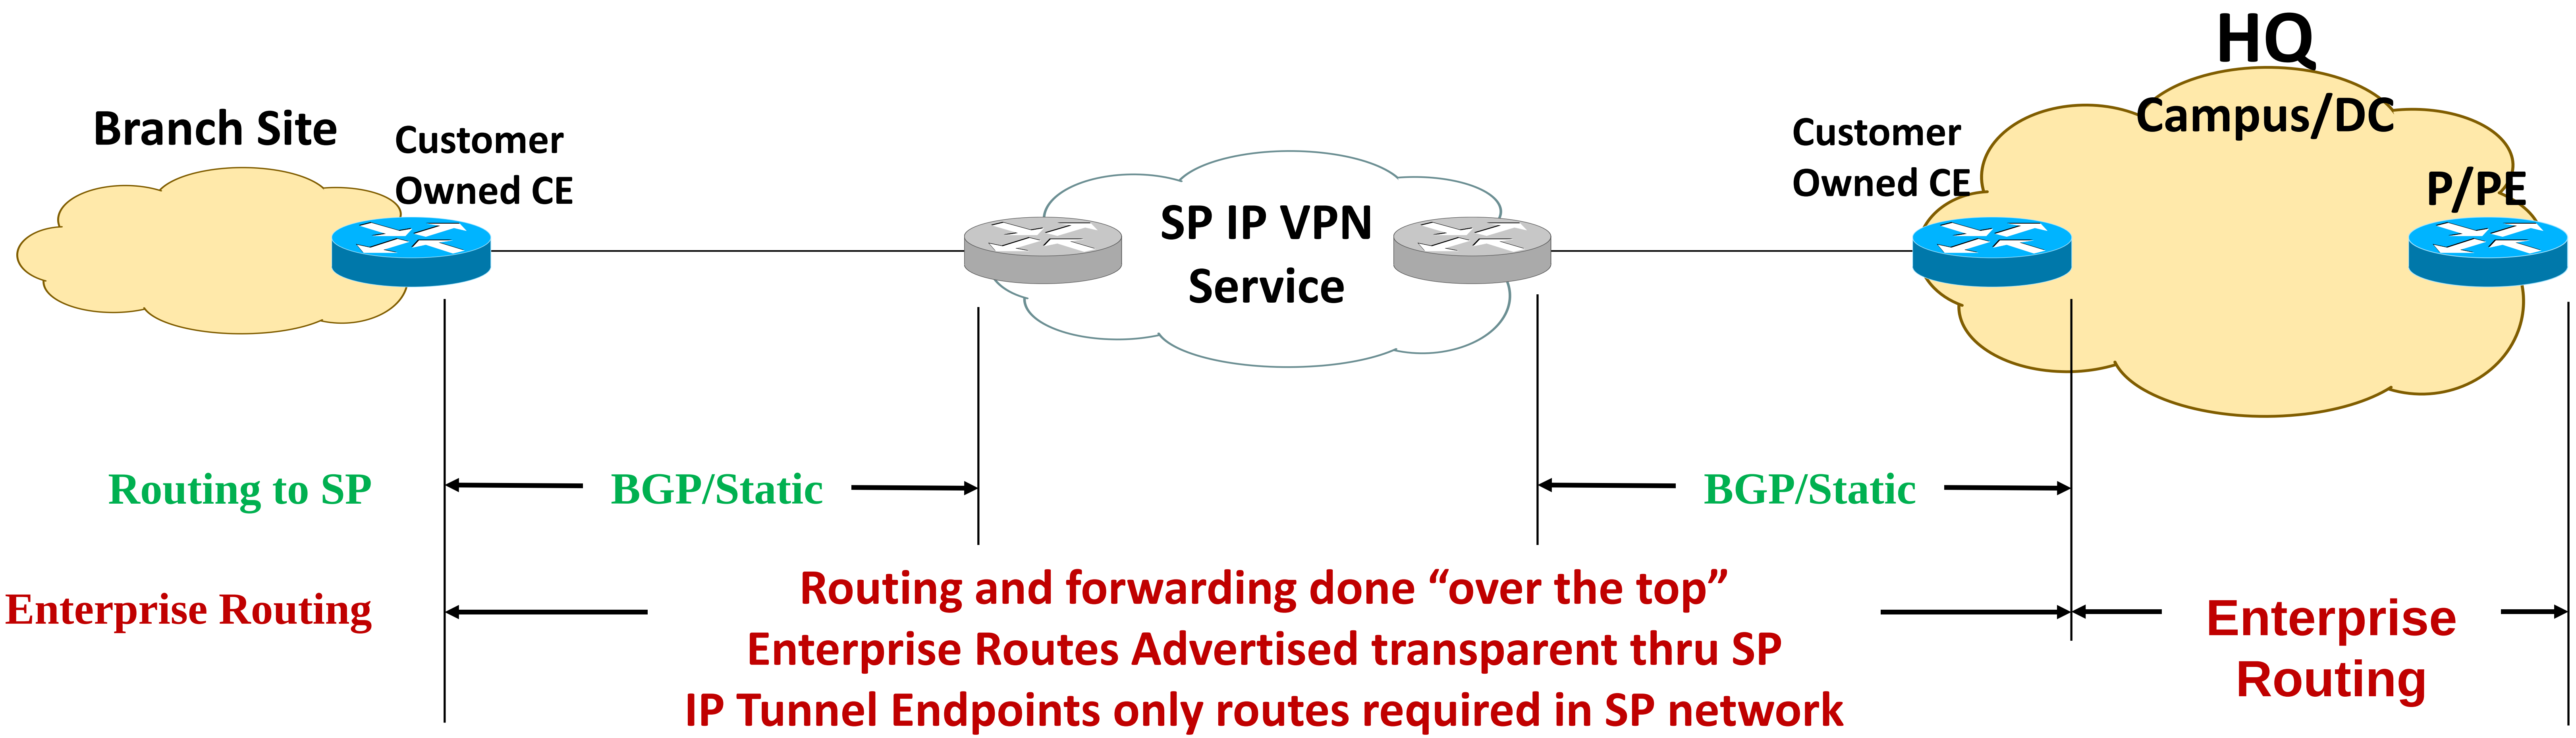

Network Deployment Options - Self Deployed vs. L3 Managed

IP Overlay Transport Model – MPLS/VRF "over the top" of SP Transport

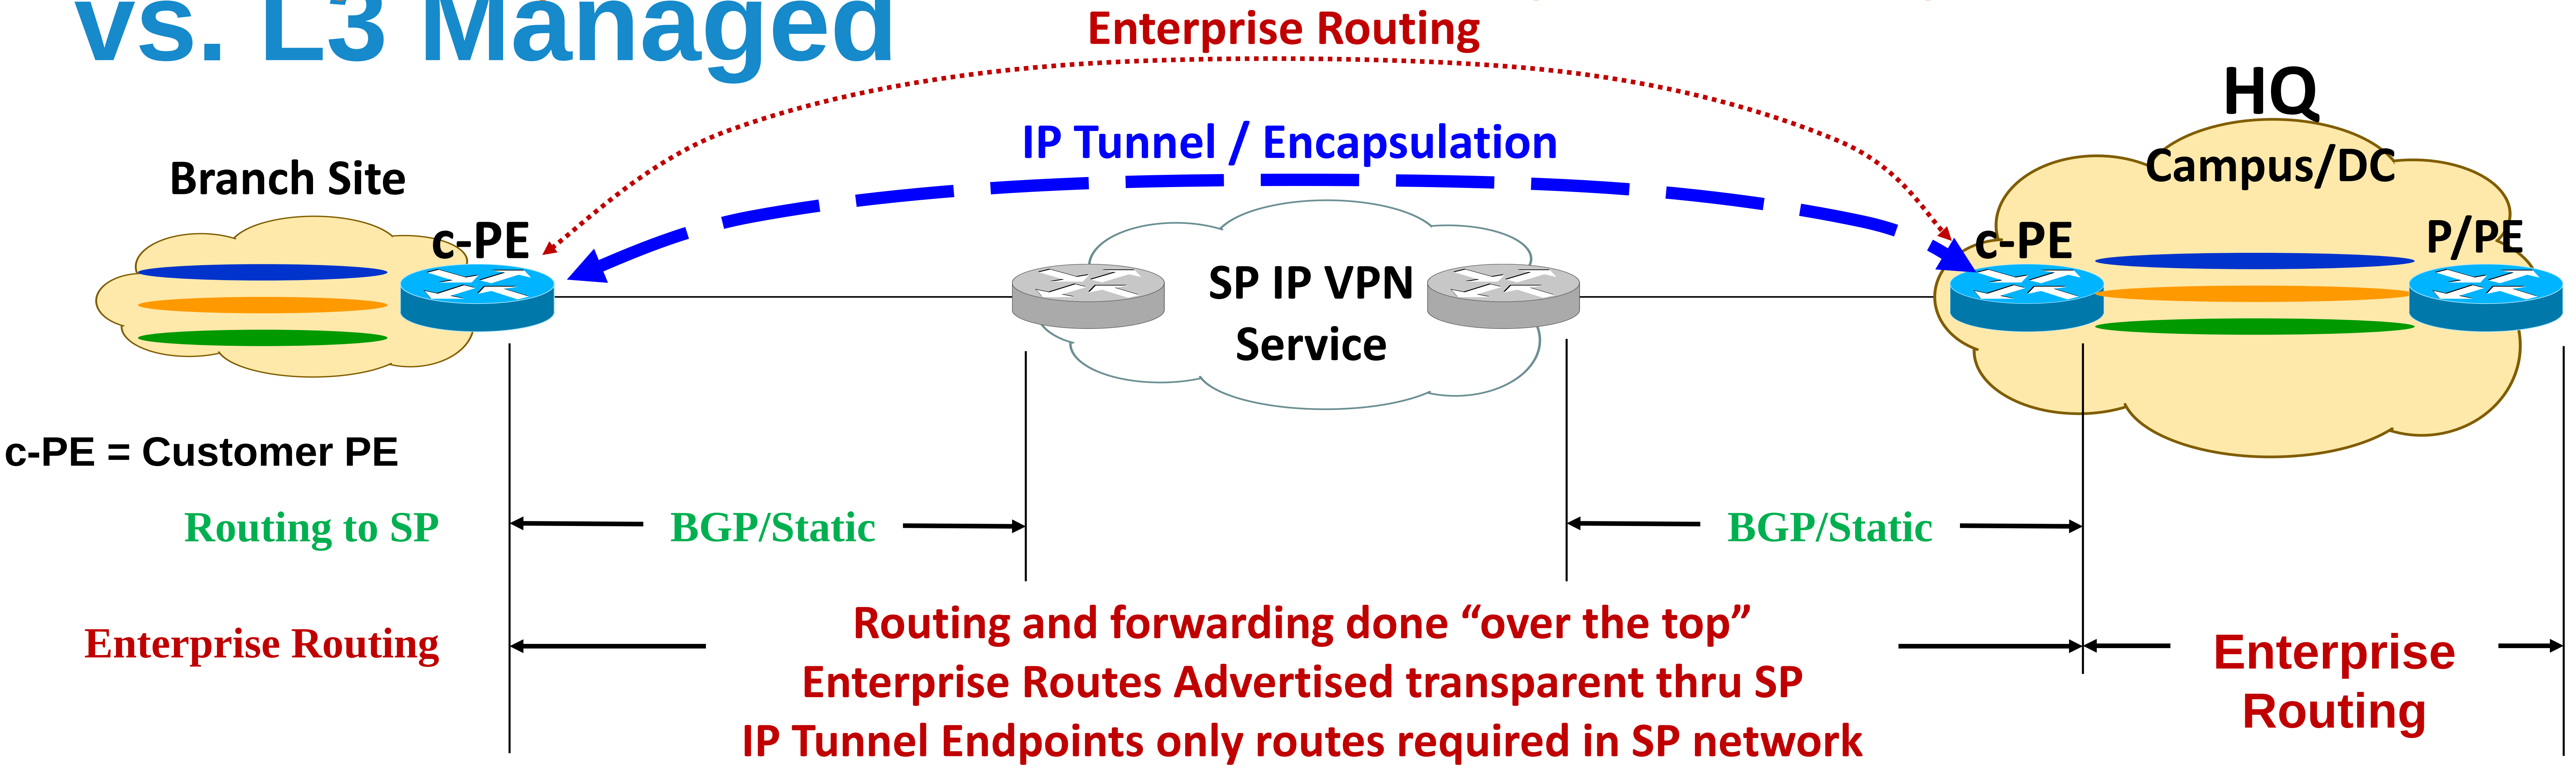

- Routing and data forwarding done "Over the Top" of the SP transport
- Enterprise routing - exchanged either inside of a IP tunnel, and/or over the top (i.e. BGP)
- Routing to SP – BGP/static, and minimal for "over the top" (normal only IP tunnel "end points")
- QoS is supported in-line with the SP offering and Service Level Agreements (SLA)
- Multicast can be supported either (1) leveraging the SP service, or (2) inside the IP tunnel

Self Deployed MPLS VPN “over the top” of SP L3 Managed Service

Creates a “Enterprise” version of “Carrier supporting Carrier” Model

SP Managed IP VPN Service

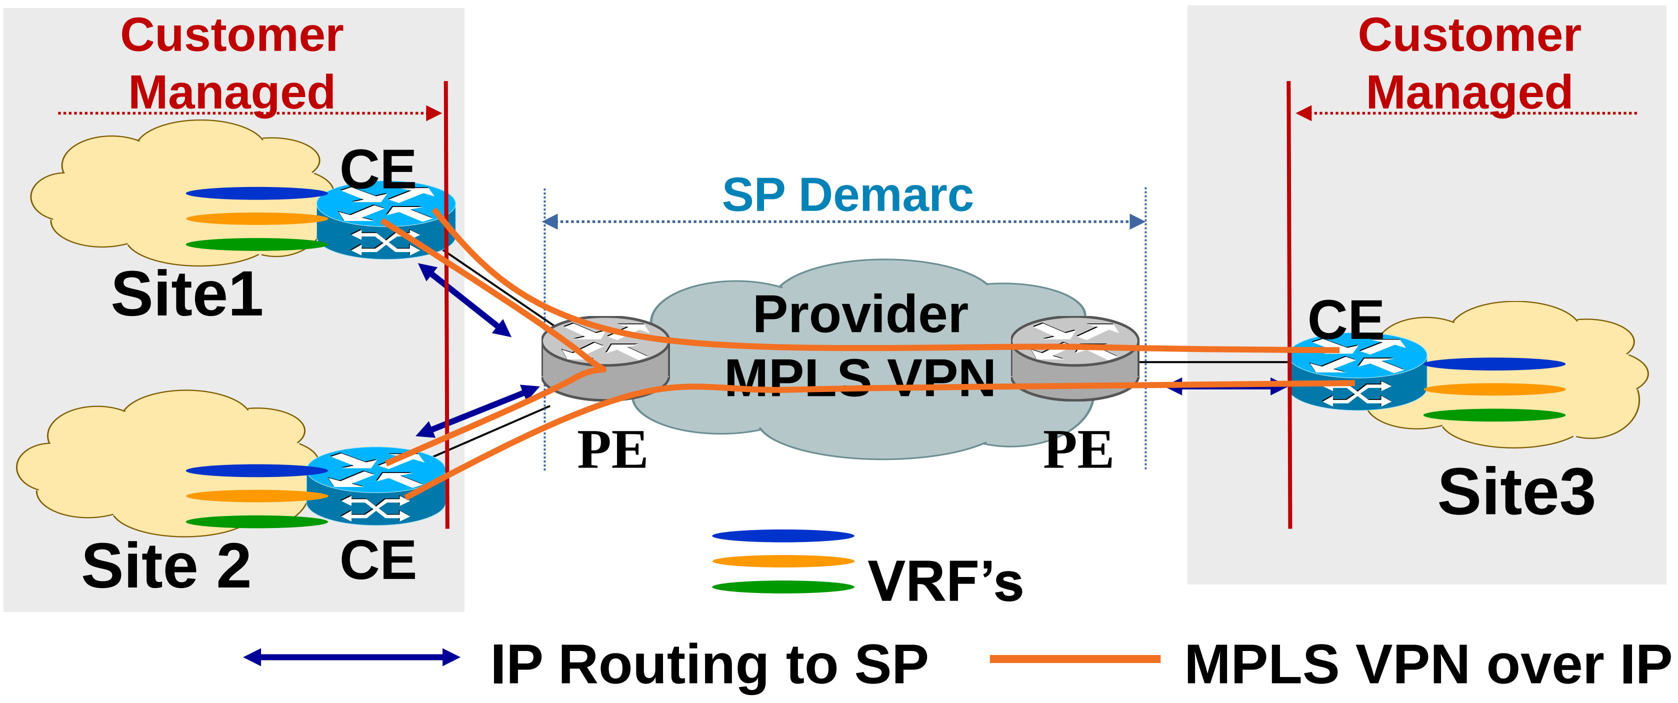

- CE Routers owned by customer
- PE Routers owned by SP
- Customer “peers” to “PE” via IP
- Exchanges routing with SP
- **Add overlay of IP that allows self-deployed MPLS over an IP Service**

Customer Controlled MPLS VPN over IP

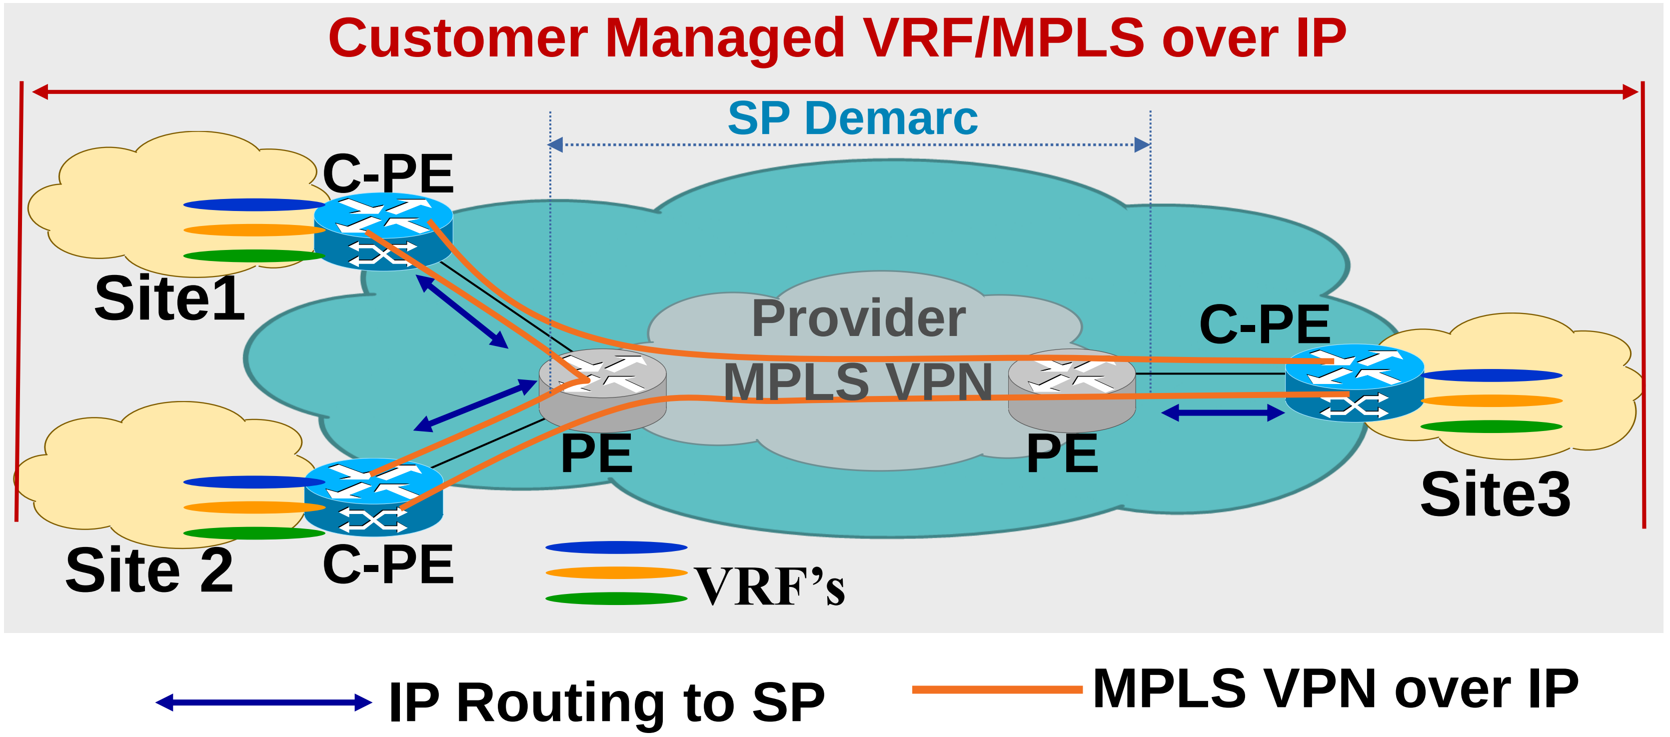

- CE routers become Customer-PE (c-PE)
- VRFs or MPLS labels are encapsulated in IP
- Other options not as scalable or more complex:
  - Carrier Supporting Carrier
  - Back to Back VRFs/Inter-AS Option “A”
  - Layer 2 Service (e.g. VPLS)

Enterprise Virtualisation End to End

WAN Virtualisation

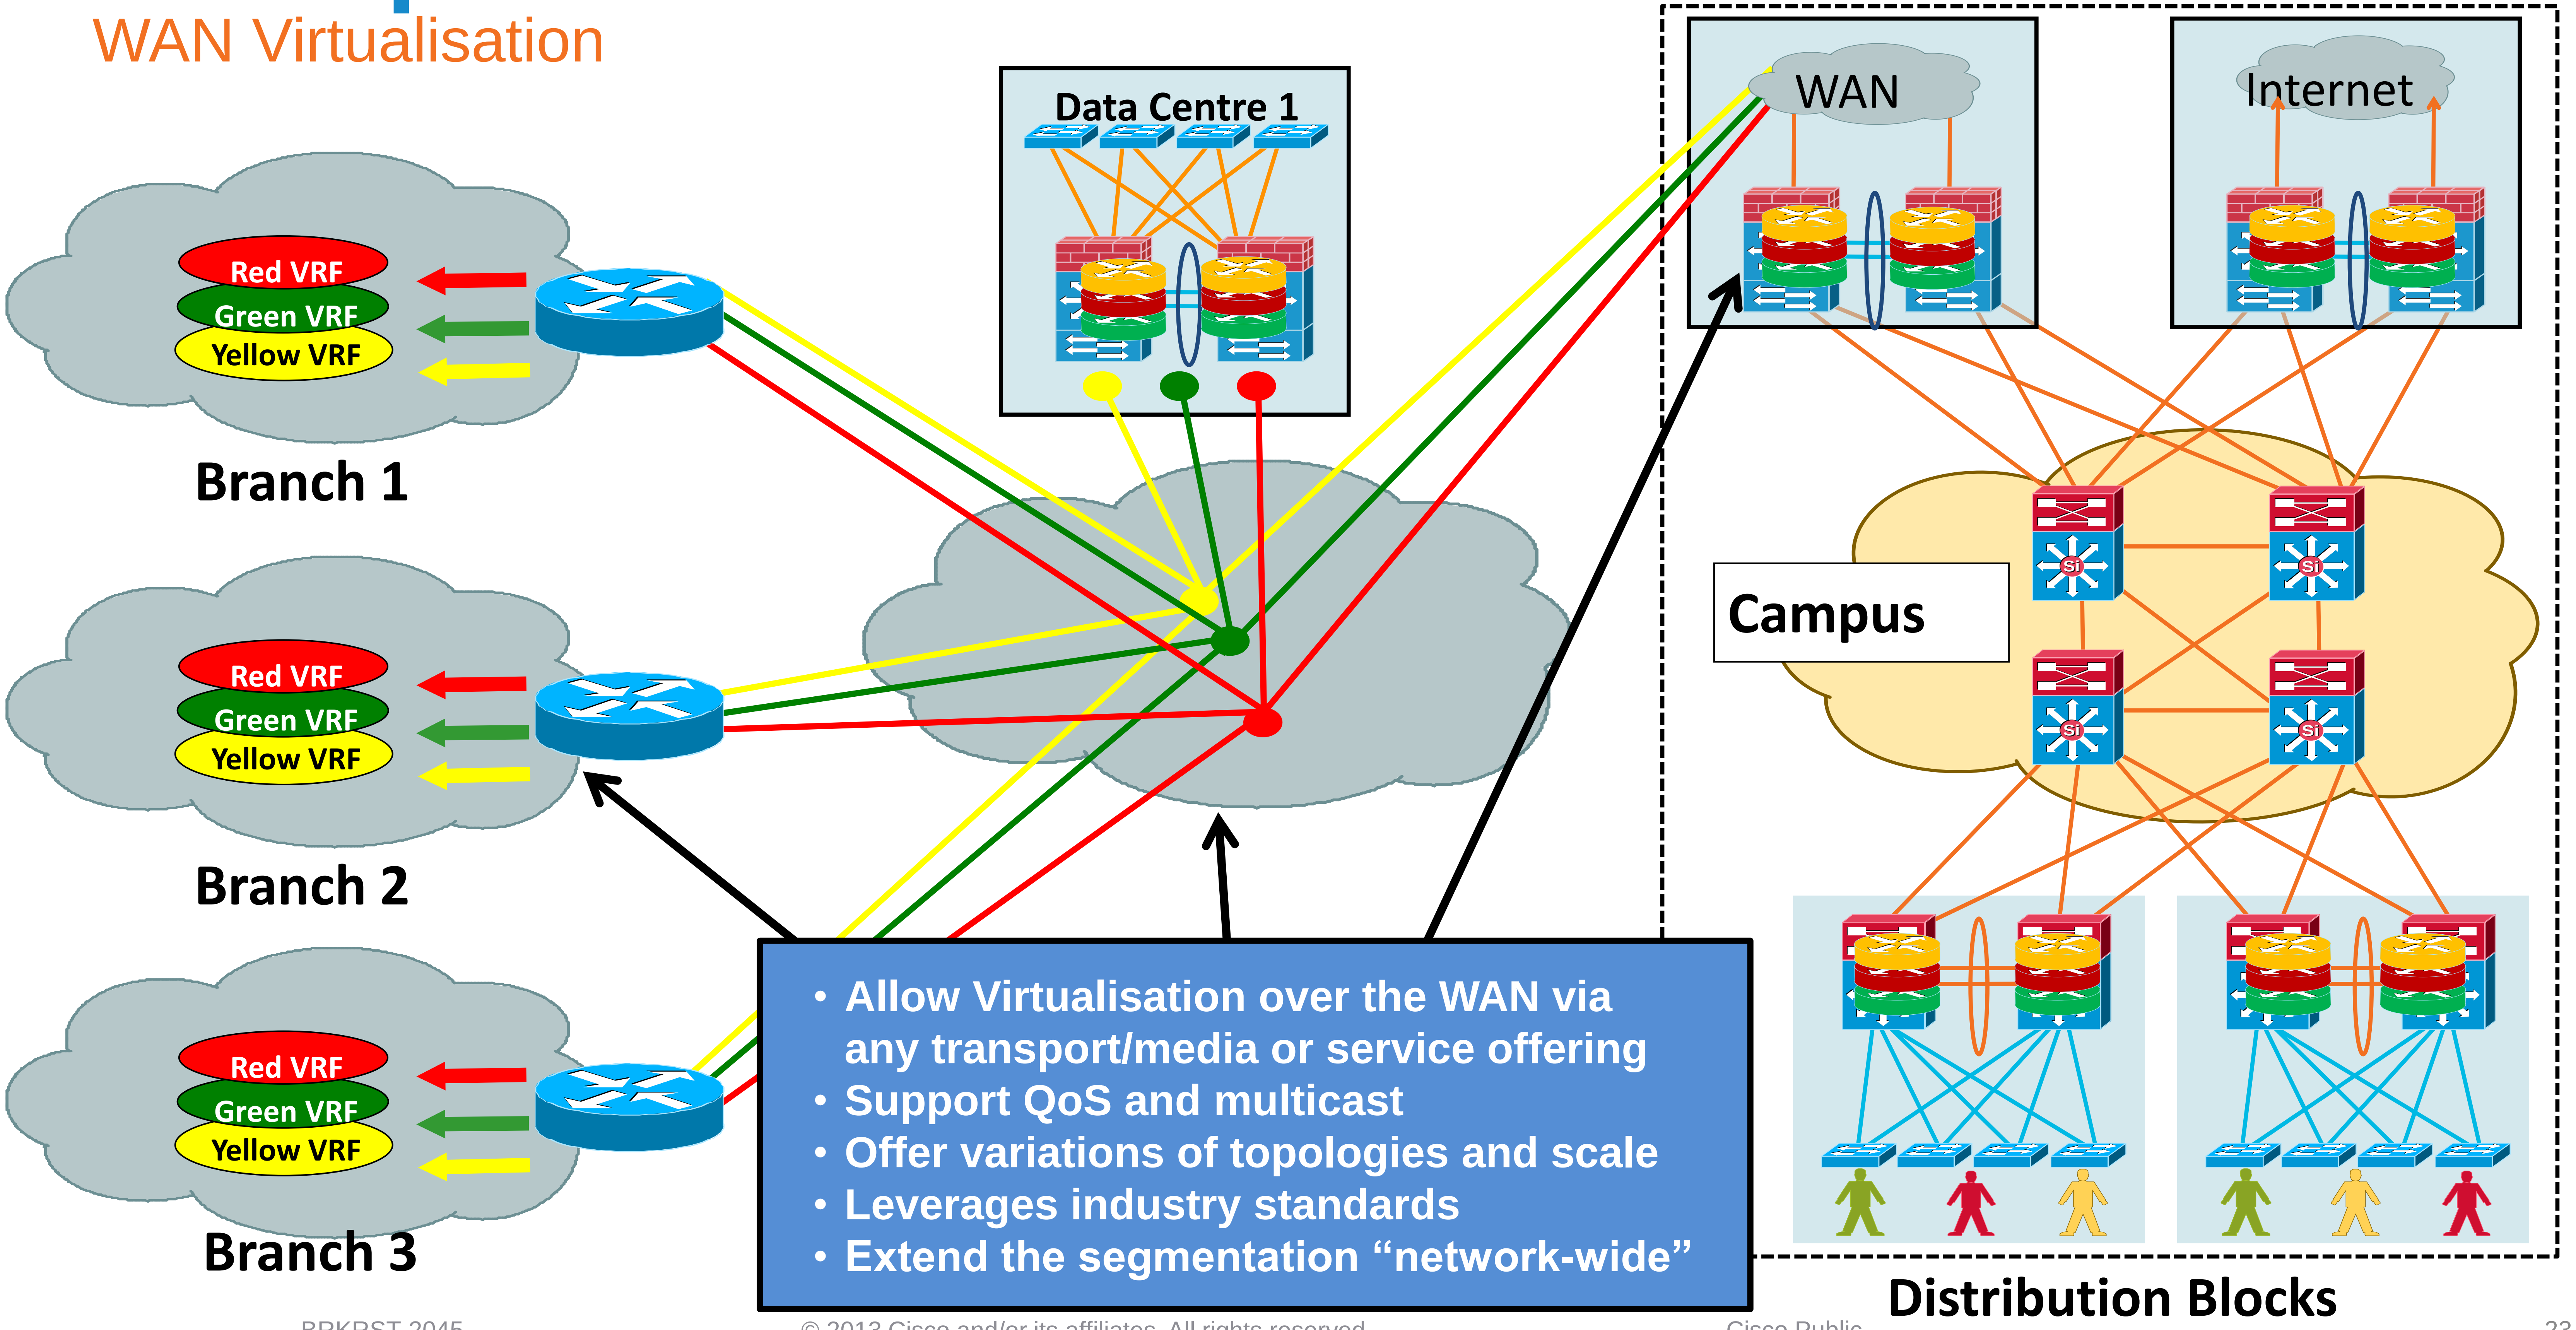

- Allow Virtualisation over the WAN via any transport/media or service offering
- Support QoS and multicast
- Offer variations of topologies and scale
- Leverages industry standards
- Extend the segmentation "network-wide"

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
  - VRF Lite over the WAN
  - MPLS VPN over L2 WAN Transport
  - L3 Virtualisation over IP
    - VRF-Lite over IP
    - MPLS VPN over IP
  - L3 Virtualisation over Multipoint GRE Tunnels (mGRE)
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
  - **VRF Lite over the WAN**
  - MPLS VPN over L2 WAN Transport
  - L3 Virtualisation over IP
    - VRF-Lite over IP
    - MPLS VPN over IP
  - L3 Virtualisation over Multipoint GRE Tunnels (mGRE)
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

What Is VRF Lite?

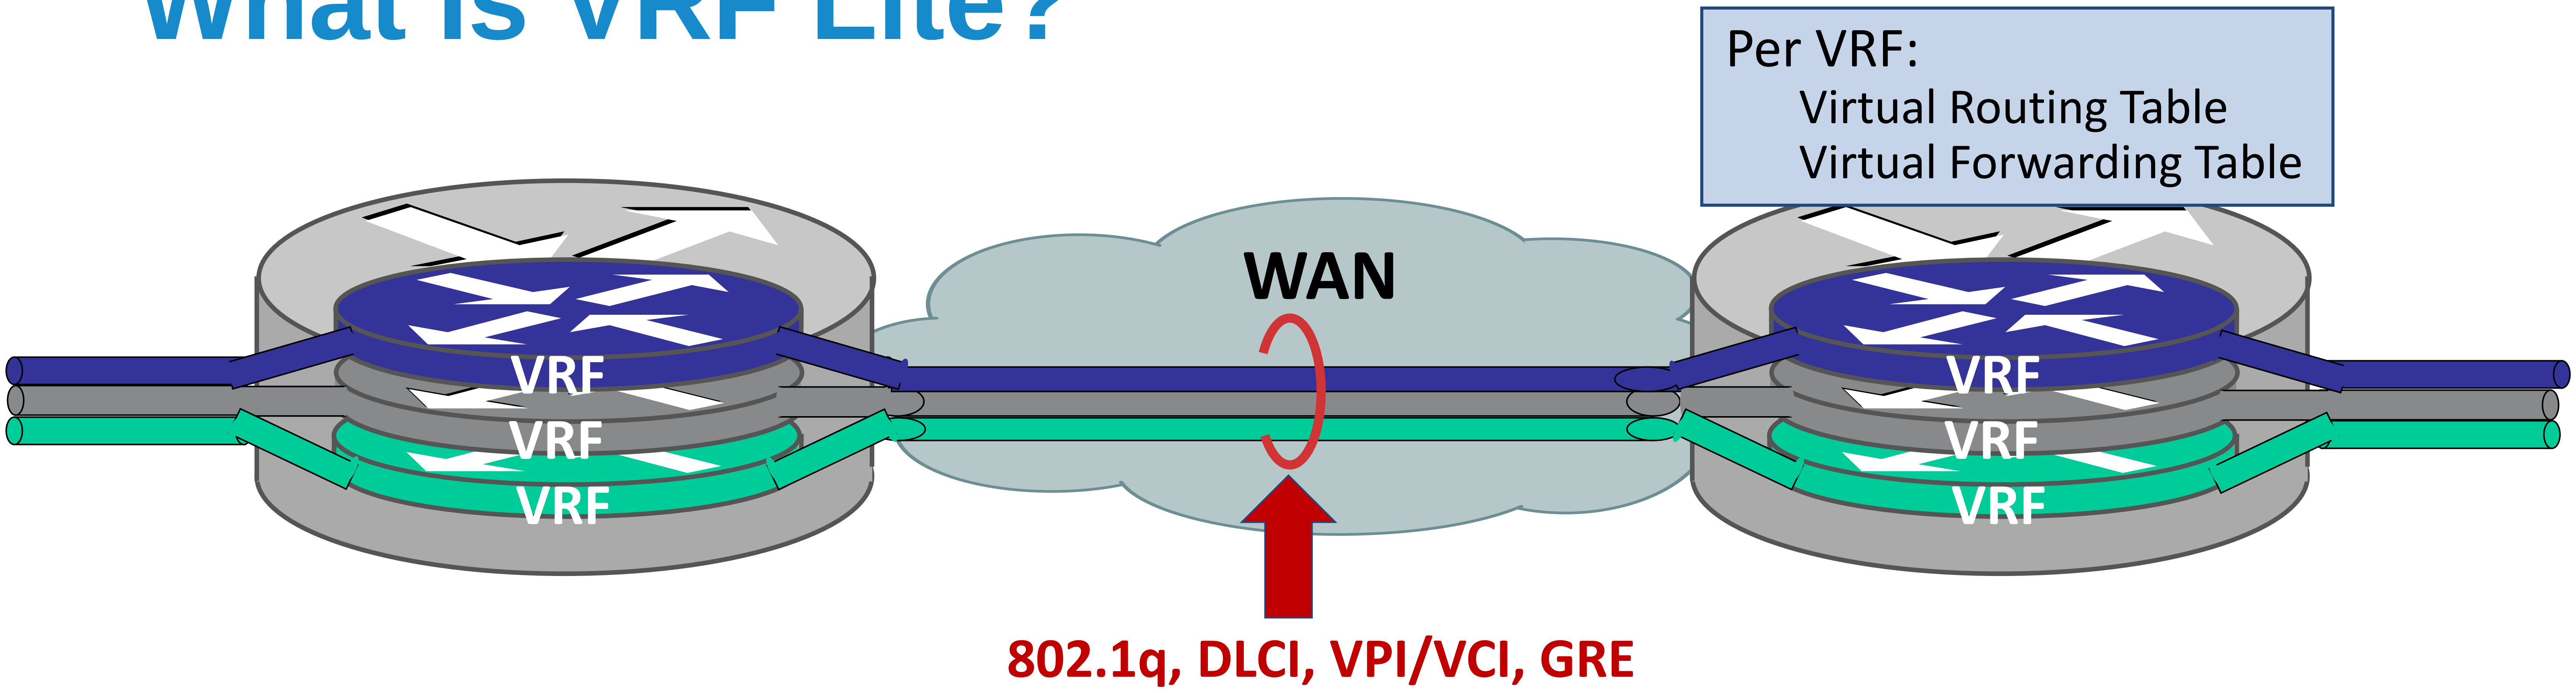

- Defines router supports **routing (RIB), forwarding (FIB), and interface per VRF !!**
- Leverages “Virtual” encapsulation for separation:
  - ATM VCs, Frame Relay, Ethernet/802.1Q
- The routing protocol is also “VRF aware”
  - EIGRP, OSPF, BGP, RIP/v2, static (per VFR)
- Layer 3 VRF interfaces cannot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

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

Extend Virtualisation over WAN L2 Service

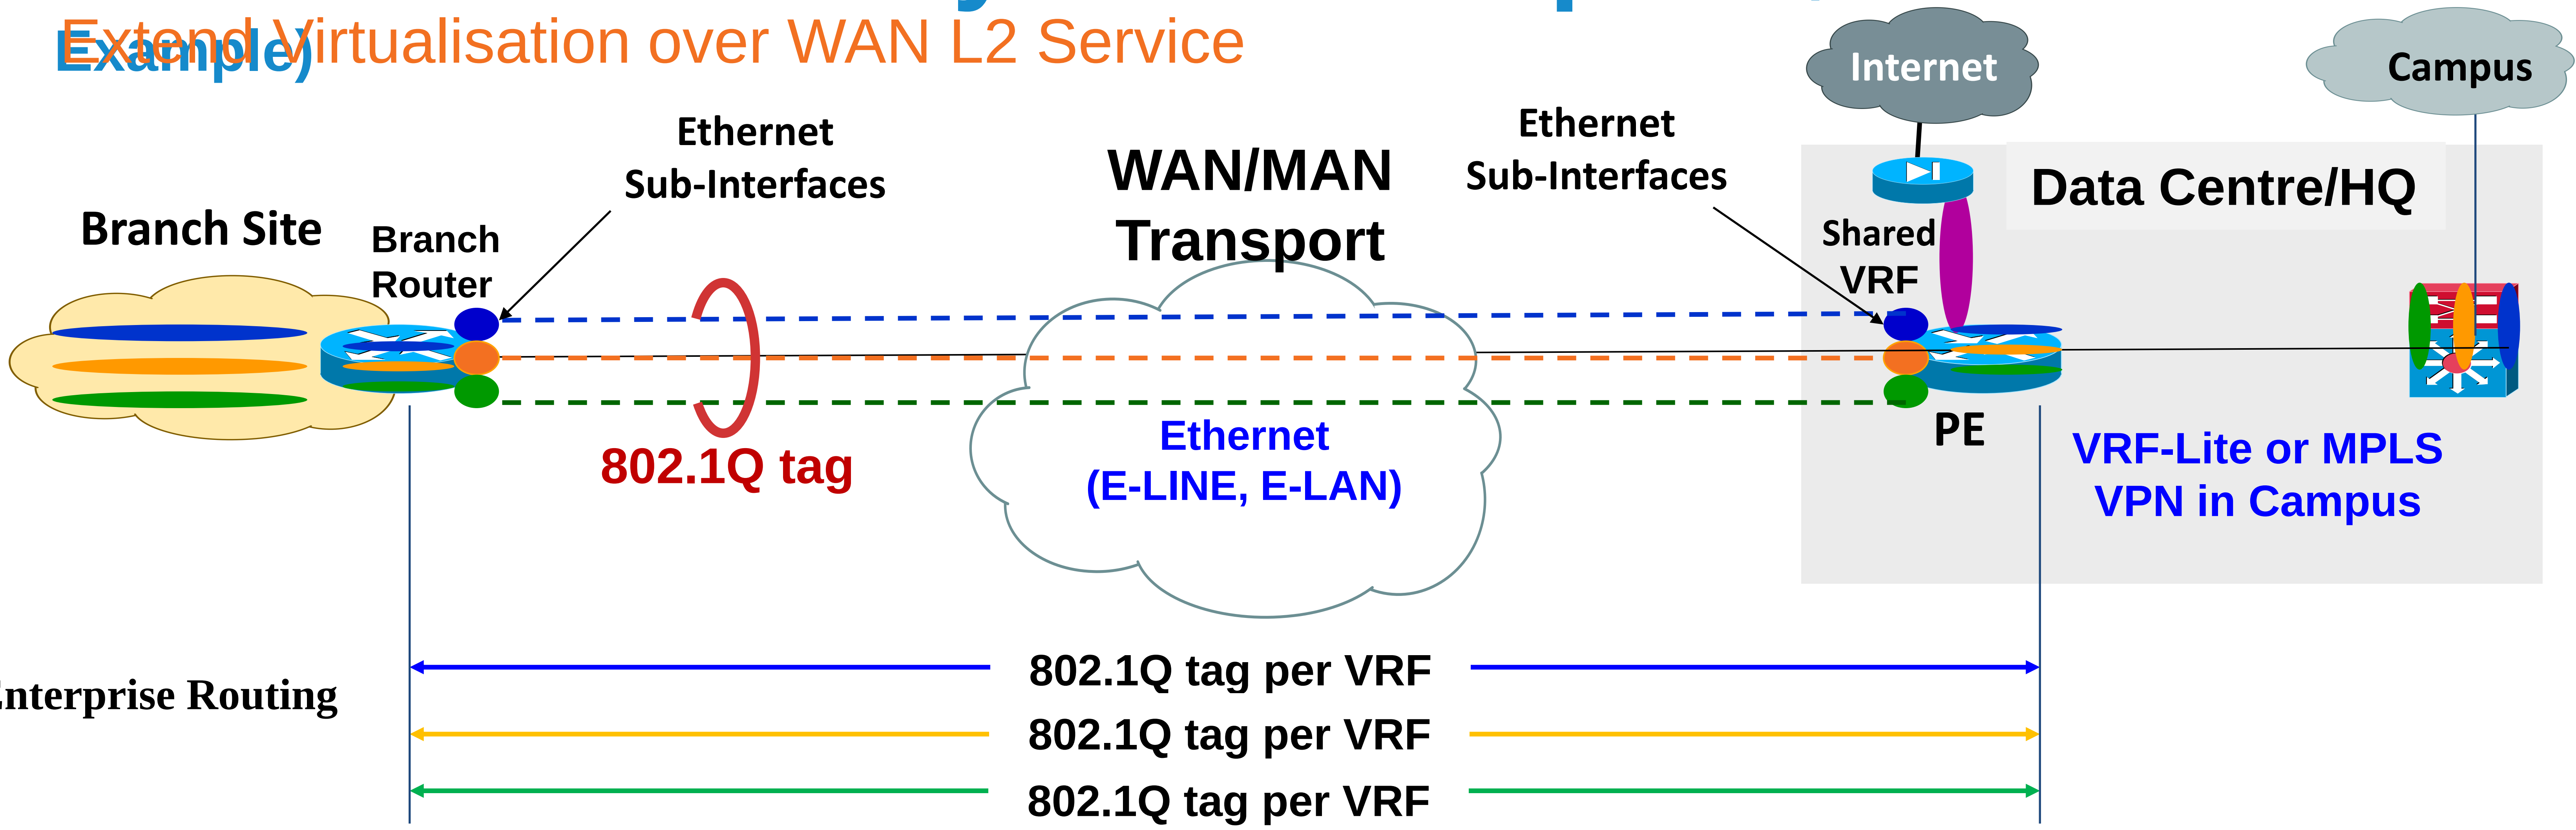

- Each Ethernet interface (or serial) leverages a sub-interface
- Unique DLCI (frame relay) or 802.1Q tag (Ethernet) per VRF
- IGP process created per VRF in both Branch/Campus
- Offers virtualised segmentation within a single “physical” interface

VRF-Lite over Layer 2 Transport

Example: Ethernet Service (Point to Point)

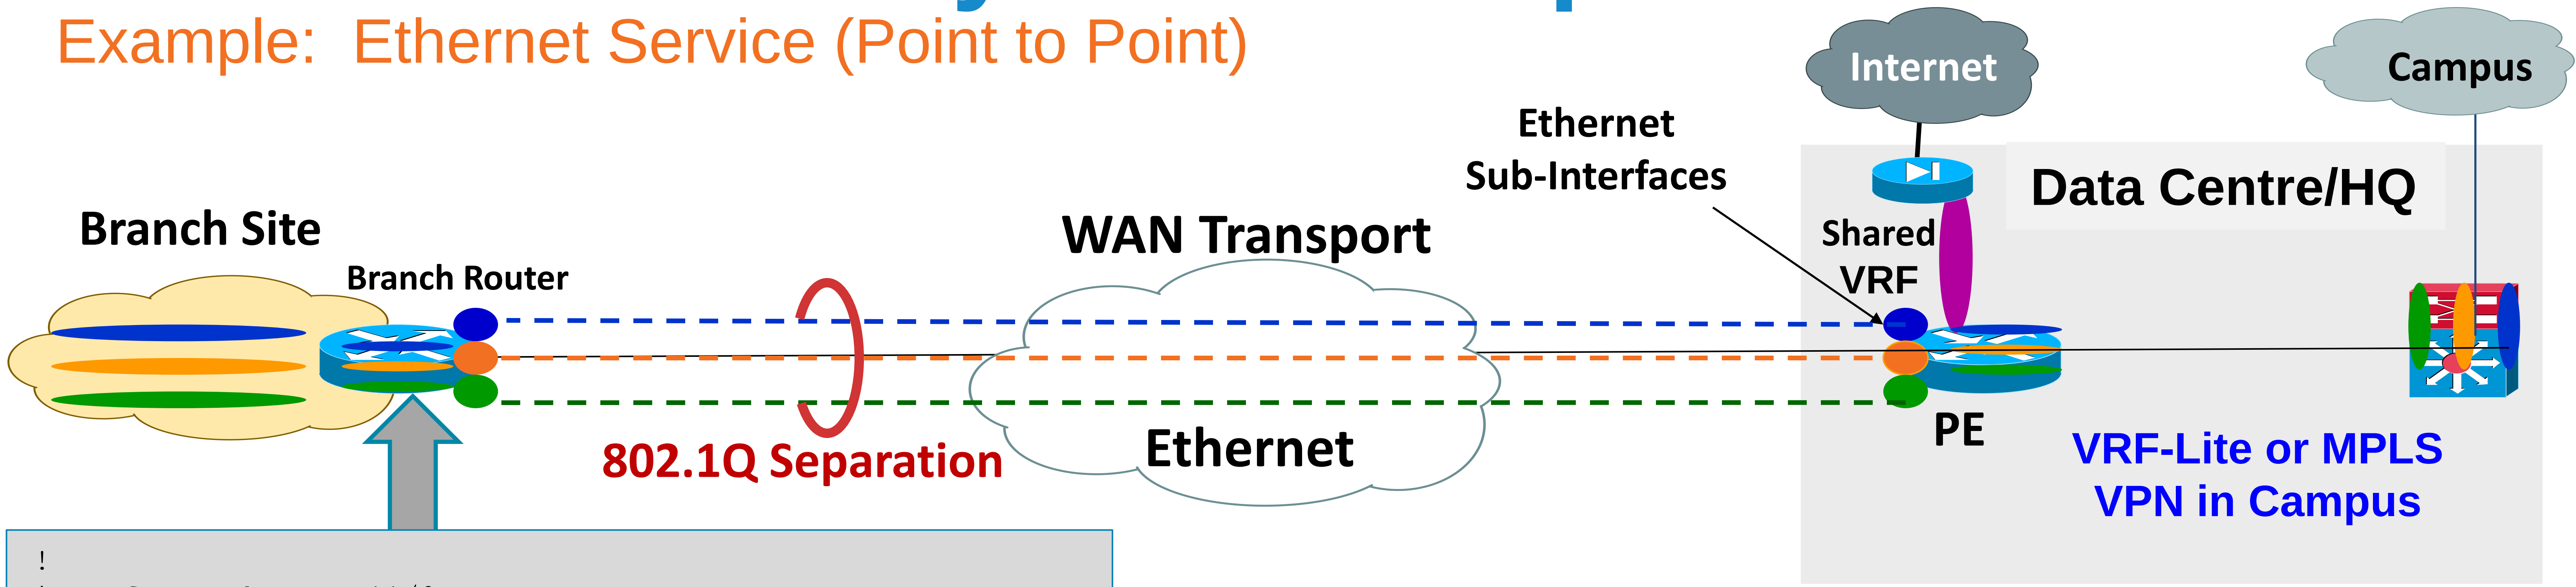


```
!
interface Ethernet 11/0
!
interface Serial 11/0.1
 encapsulation dot1Q 100
 ip vrf forwarding blue
 ip address 192.168.51.2 255.255.255.0

!
interface Ethernet 11/0.2
 encapsulation dot1Q 200
 ip vrf forwarding green
 ip address 192.168.61.2 255.255.255.0
!
interface Ethernet 11/0.3
 encapsulation dot1Q 300
 ip vrf forwarding yellow
 ip address 192.168.71.2 255.255.255.0
!
```

Configuration Notes:

- Frame Relay encapsulation can be used to virtualise a leased line
- 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?

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 have the expertise to manage an MPLS VPN network, if that is the best solution?

Example: 4 Sites with 4 VRFs

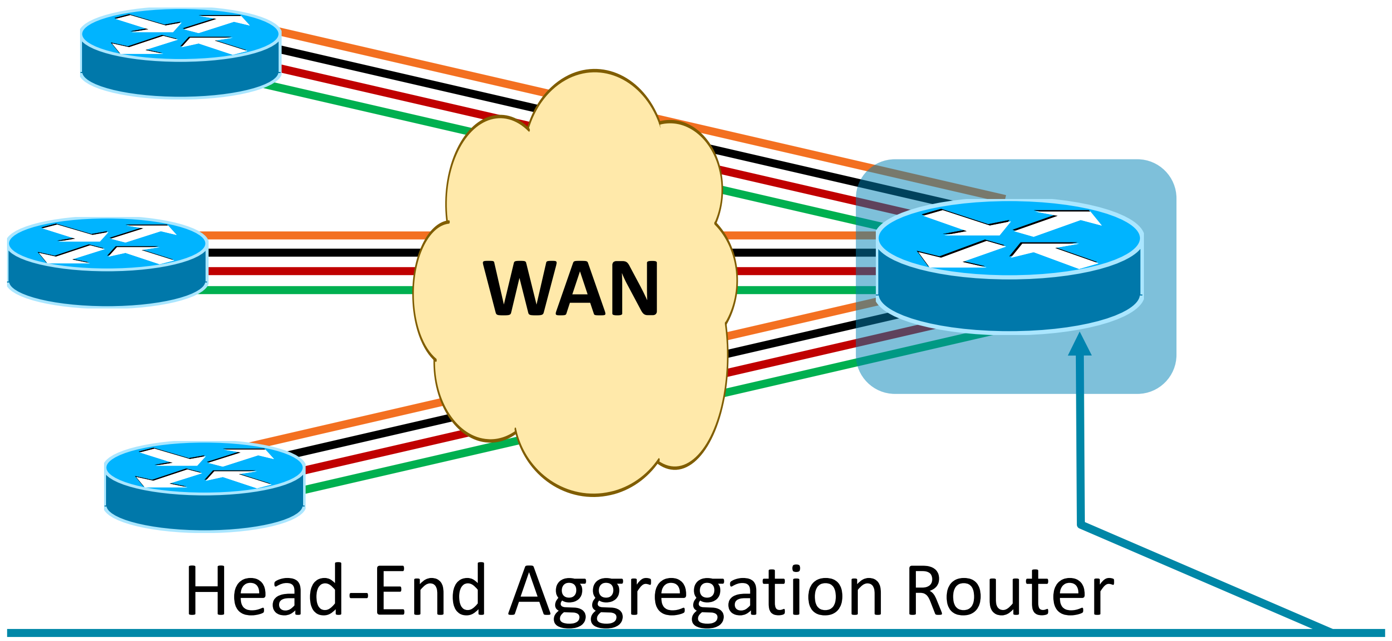

Virtual Networks	Neighbours	VRF Sub-interfaces
4	3	12
10	3	30
20	3	60
30	3	90

VRF-Lite for Branch Back-haul over the WAN

Summary

- Leverages VRF in router (RIB/FIB, interface) and interface for segmentation
- No MPLS, LDP, or BGP required
- Optimal solution when VRF count is small (~ <8)
- Scale usually dependent on routing protocol scale
- Supports multicast and QoS solutions
- Most common deployments?
  - Branch Back-haul to campus/DC, Branch Back-haul to aggregation PE running full MPLS VPN

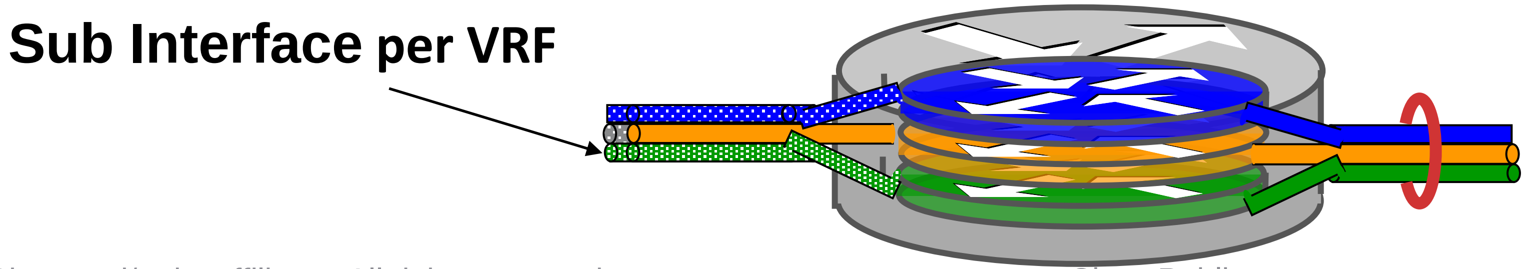

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
  - VRF Lite over the WAN
  - **MPLS VPN over L2 WAN Transport**
  - L3 Virtualisation over IP
    - VRF-Lite over IP
    - MPLS VPN over IP
  - L3 Virtualisation over Multipoint GRE Tunnels (mGRE)
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

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) Infrastructure
- Unified Control Plane (Generalised MPLS)

Key Virtualisation Mechanisms over an IP Infrastructure

MPLS Label Encapsulations

Applicable When Using MPLS over Layer 2 Transport

L2 Header Options

- PPP Header (Packet over SONET/SDH)
- Ethernet Header (LAN MAC)
- ATM Header (RFC 1483)

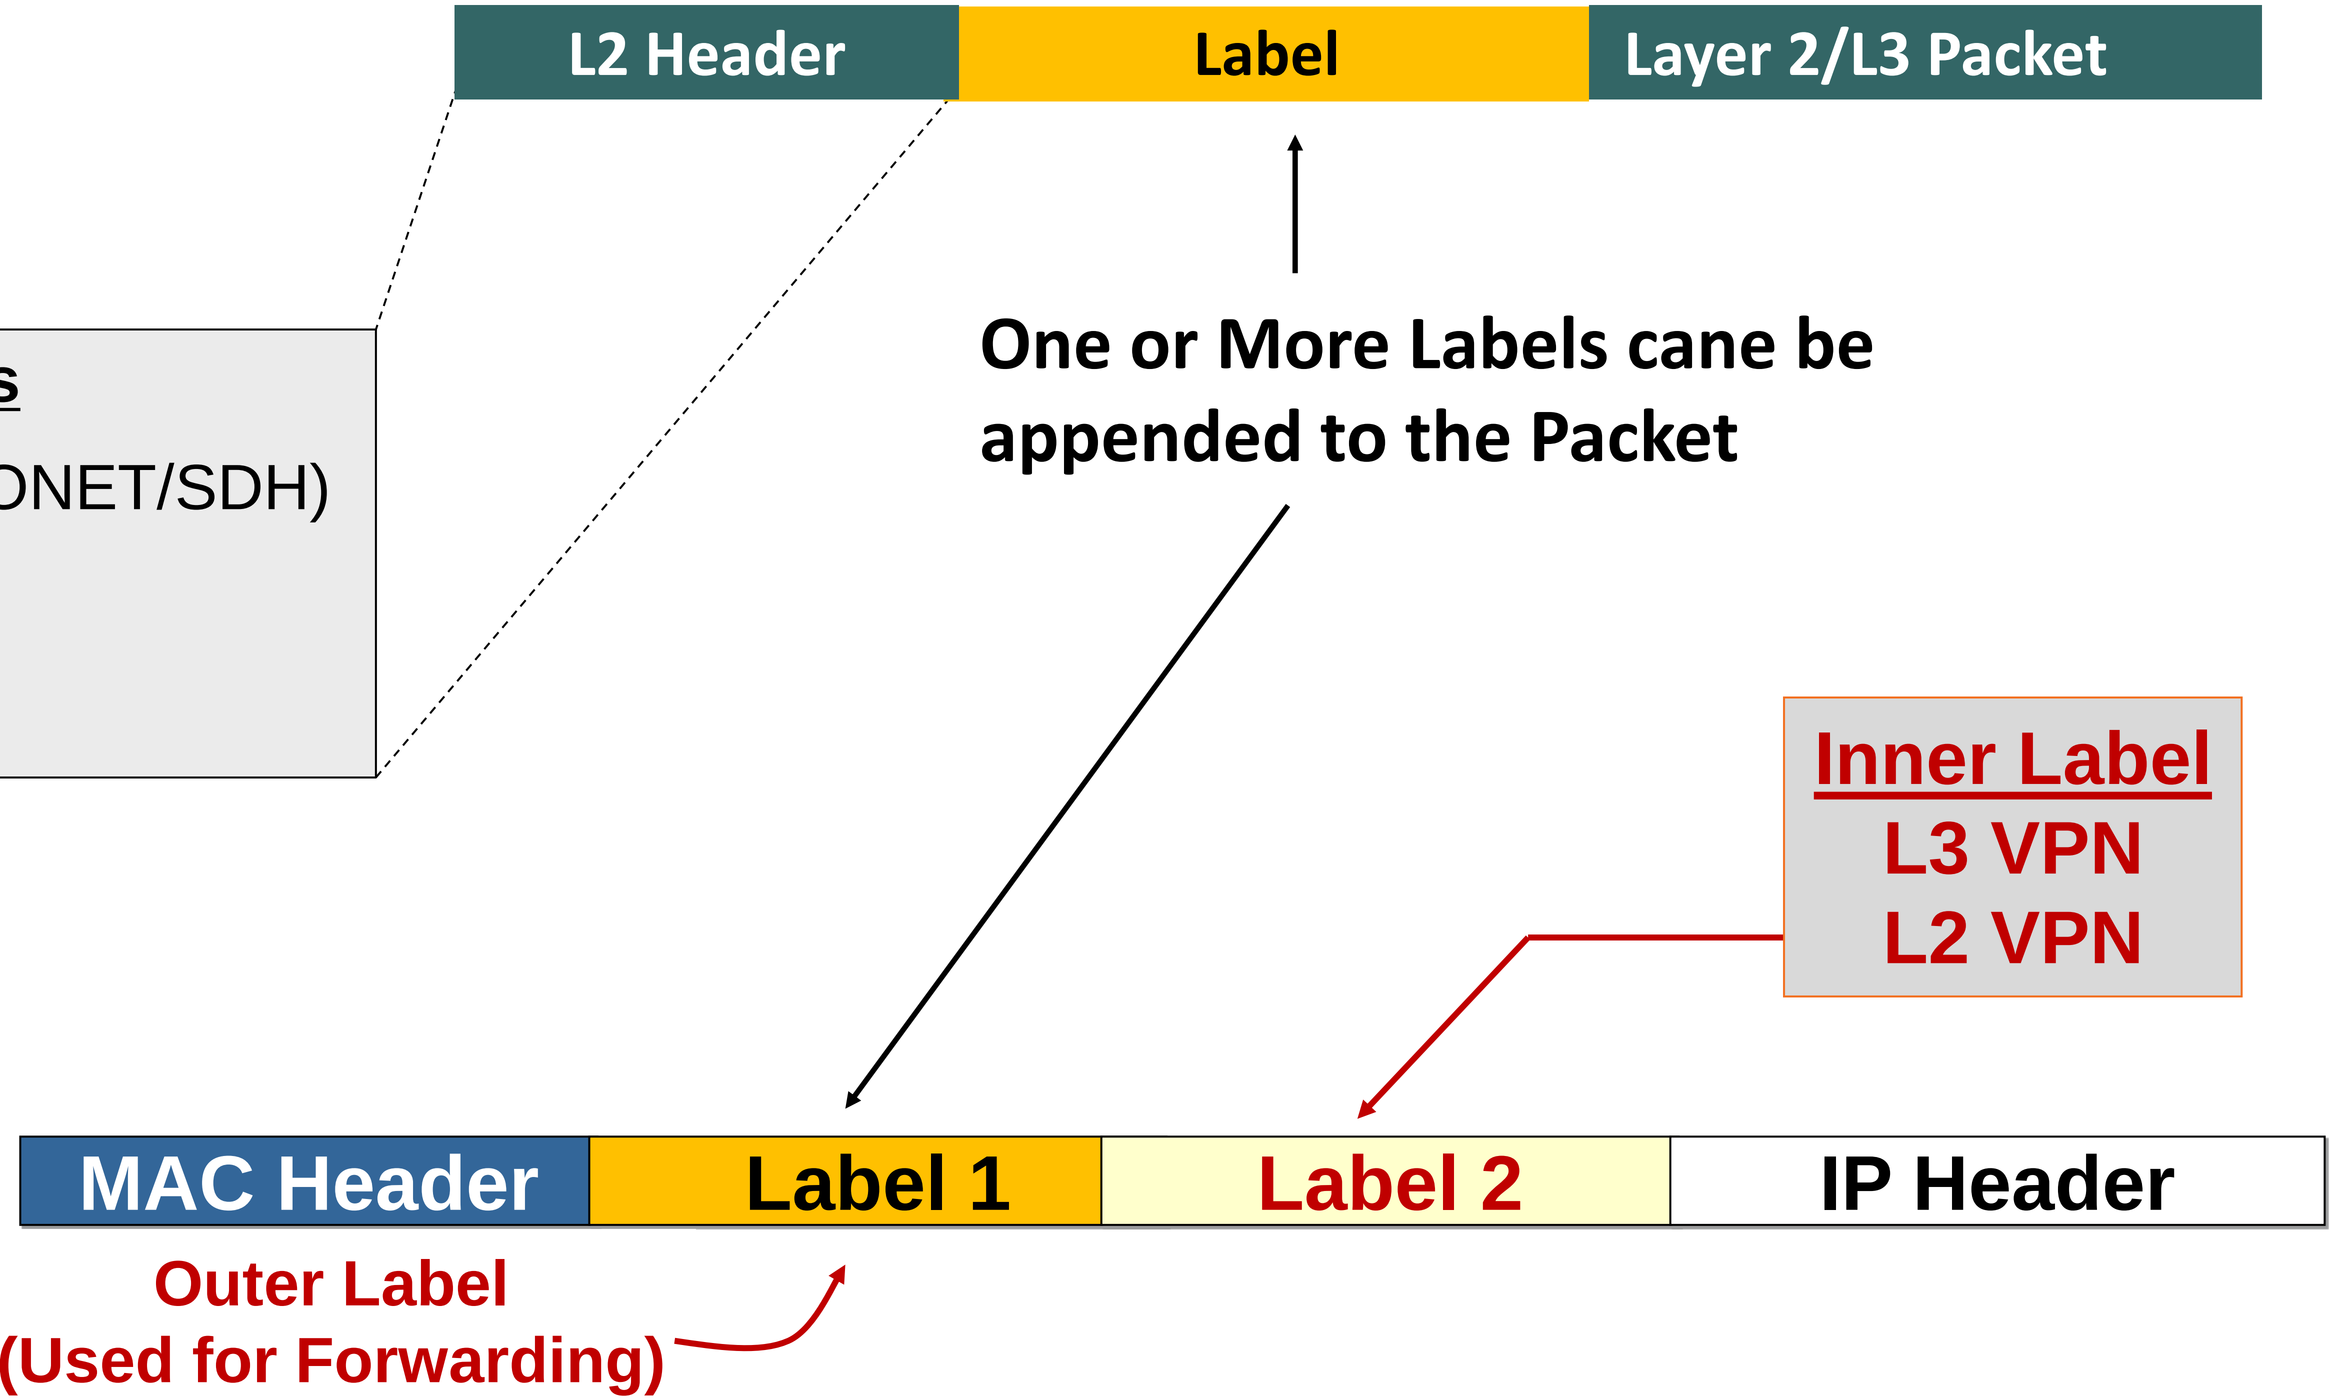

Label Stacking (LAN Example)

MPLS VPN Technology—Refresher

MPLS VPN Connection Model

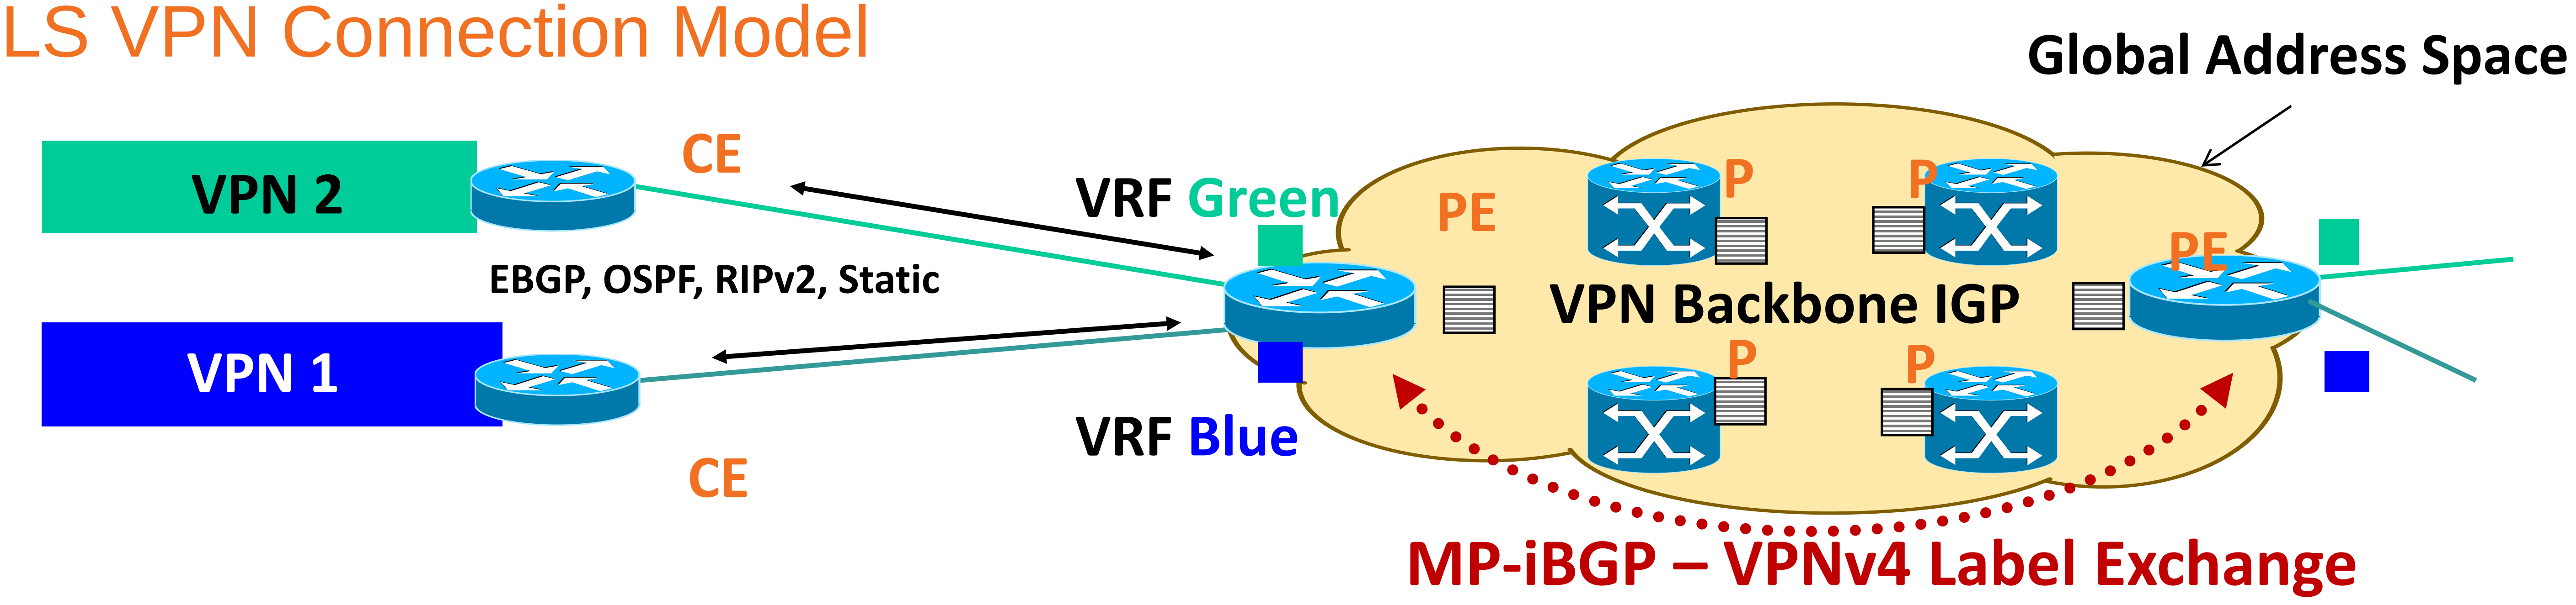

CE Routers

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

  (static, RIP, BGP, EIGRP, OSPF)

PE Routers

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

P Routers

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

MPLS VPN over L2

Configuration Example (IOS)

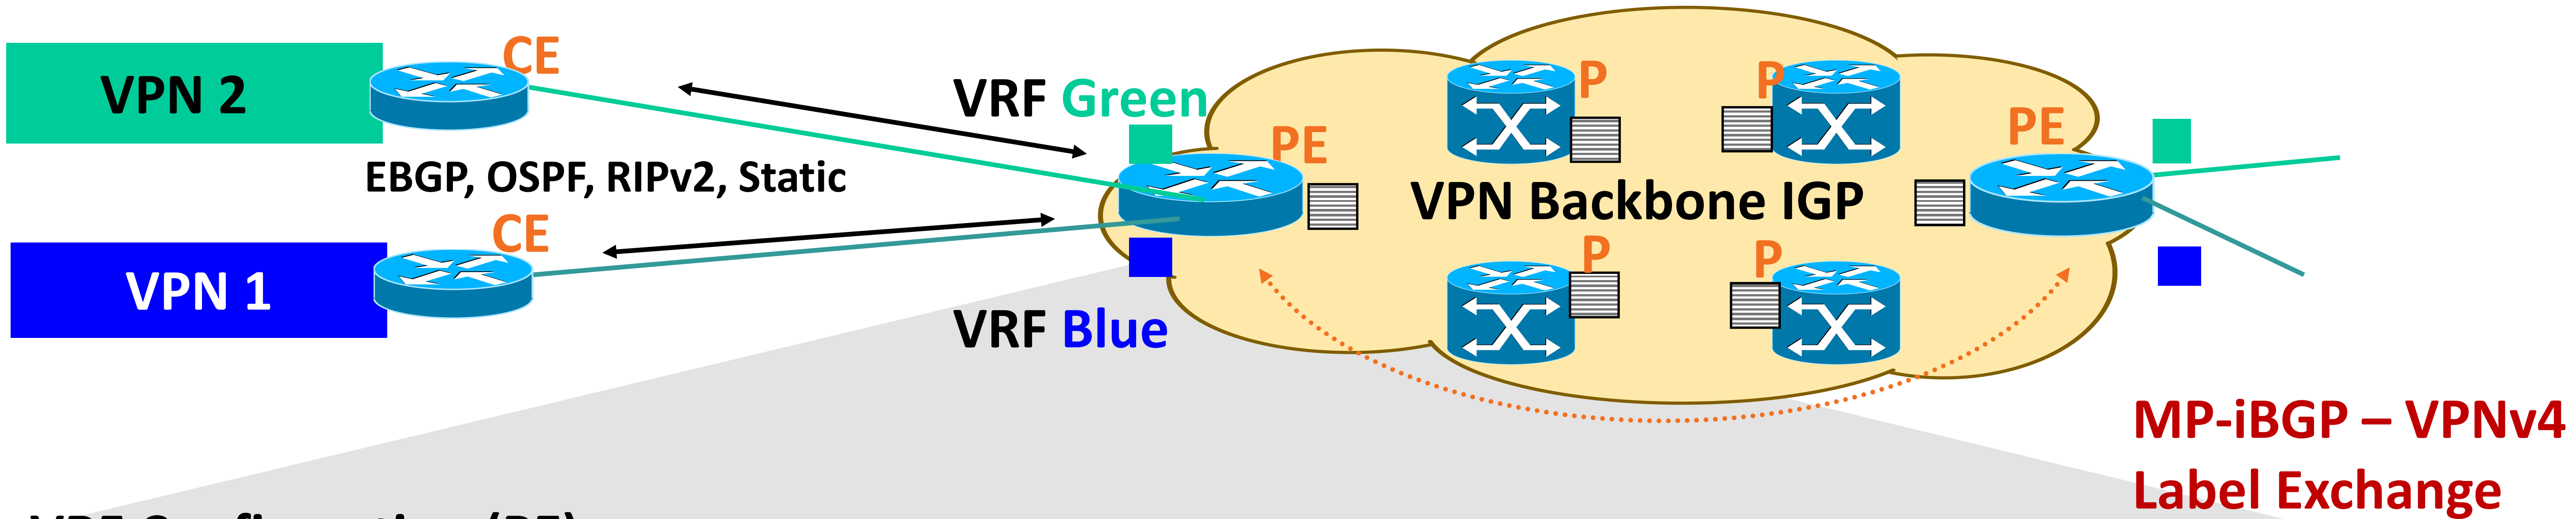

VRF Configuration (PE)

```
! 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 GigabitEthernet0/1.10
 ip vrf forwarding blue
interface GigabitEthernet0/1.20
 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
```

MPLS VPN Technology

MPLS VPN Using Separate “CE” router per VRF

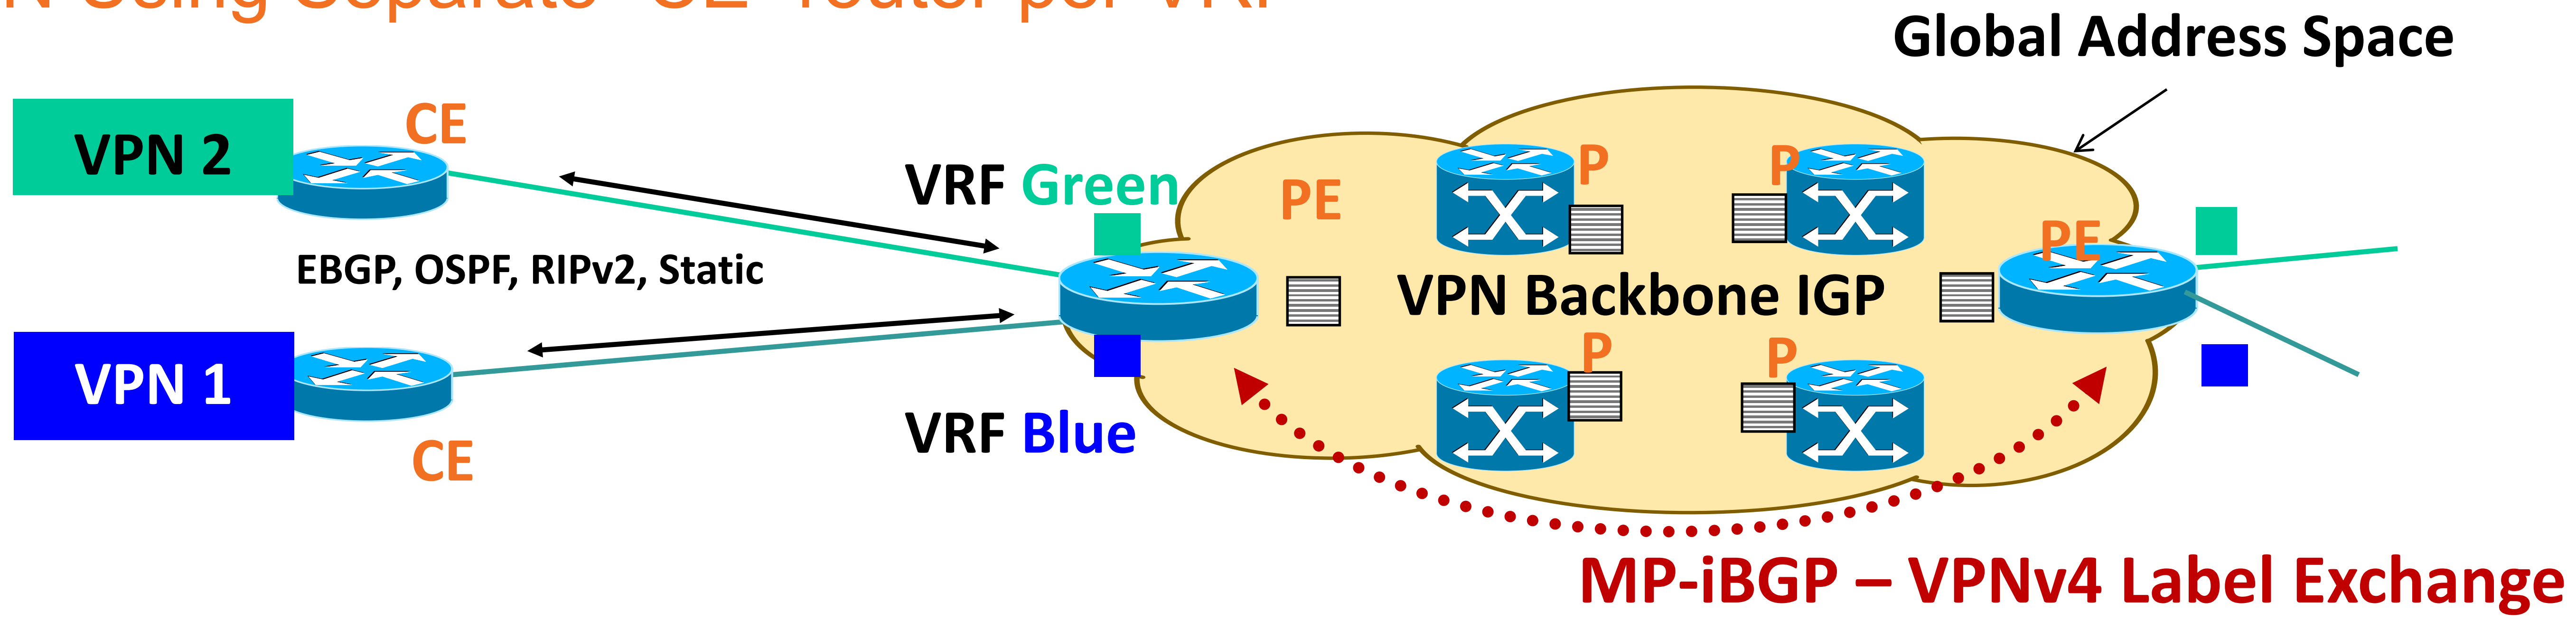

MPLS VPN Technology

MPLS VPN + VRF-Lite (PE-CE)

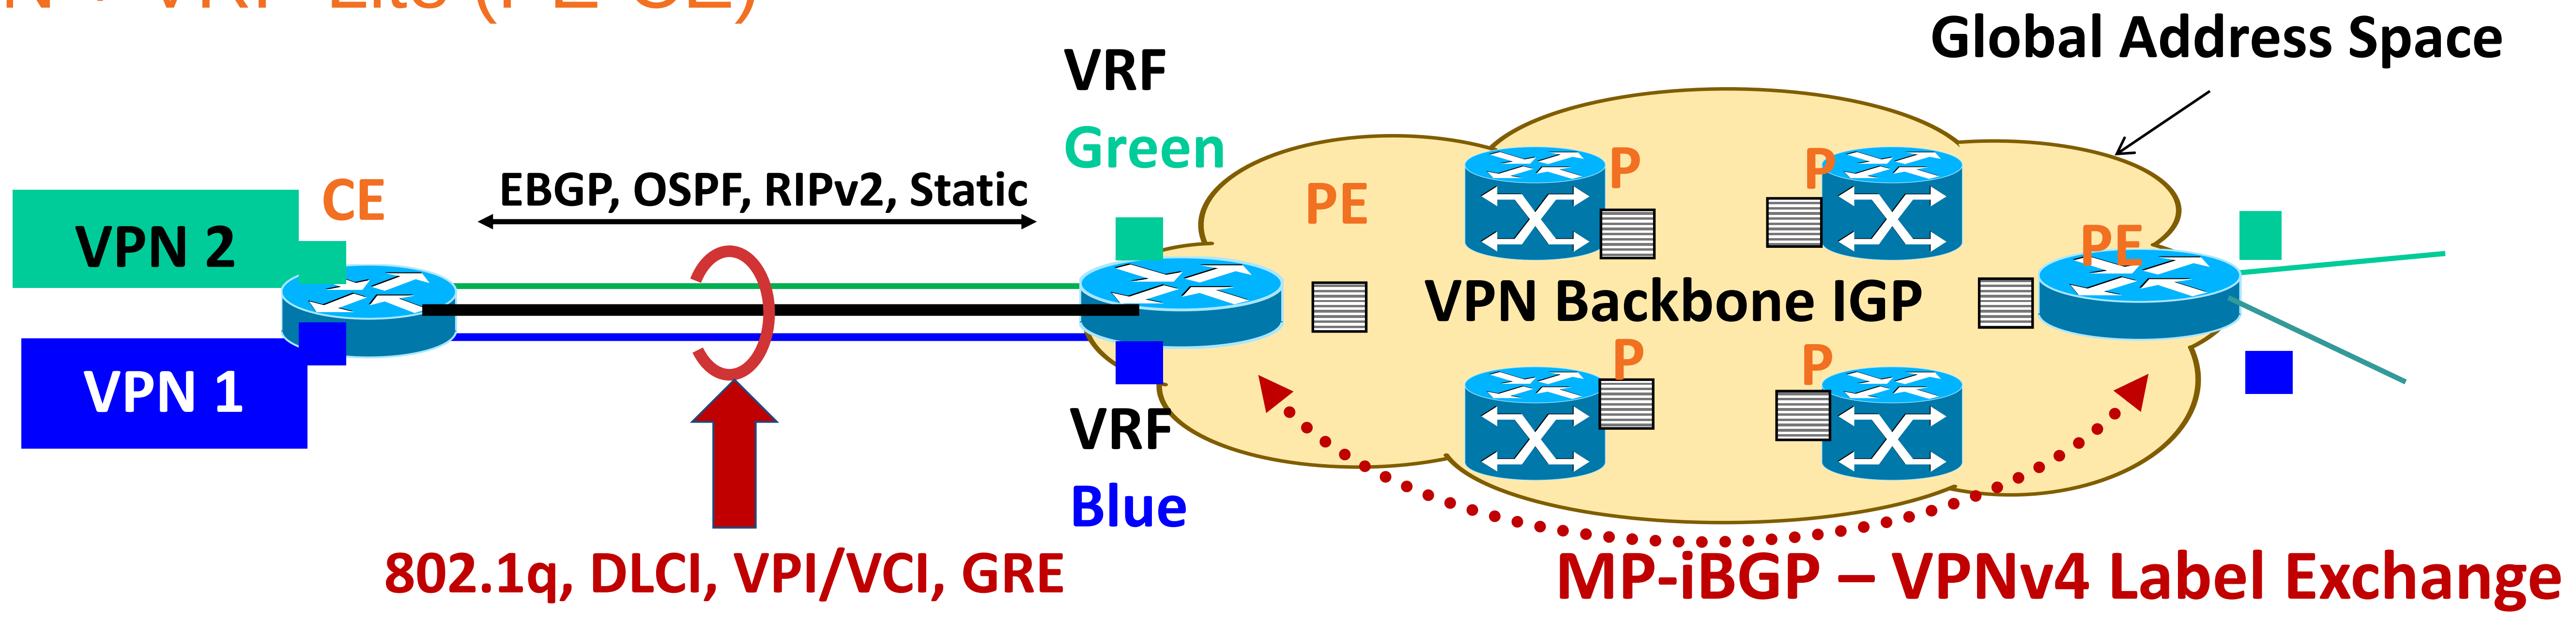

- MPLS VPN backbone remains the same
- Leverage VRF-Lite CE to PE
- CE to PE can be “local” (fibre, copper) or remote (WAN, Metro service)
- Transport will dictate technology chosen for CE – PE
  - Local or Ethernet service (802.1Q), WAN (DLCI), IP WAN (GRE)
- VRF has its own instance for the routing protocol

  (static, RIP, BGP, EIGRP, OSPF)

MPLS VPN over L2 WAN Links

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

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
  - VRF Lite over the WAN
  - MPLS VPN over L2 WAN Transport
  - **L3 Virtualisation over IP**
    - VRF-Lite over IP
    - MPLS VPN over IP
  - L3 Virtualisation over Multipoint GRE Tunnels (mGRE)
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

Why Do We Need IP/MPLS Virtualisation over IP?

- **VRF-Lite Requires Layer 2 for Separation**
  - – Need to leverage IP for broader reach, and more transport options
- **Not all "transport/transit" networks are MPLS**
  - – MPLS is not available for transport on every network
- **IP is the only Transit Option Between MPLS Islands (i.e. networks)**
  - – Core/transit network is IP and is not owned by Enterprise
  - – IP VPN Service from SP is only offering available (vs. L2 option)
  - – Customer uses "external" IP encryption units (i.e. device does not support MPLS)
- **MPLS packets require encryption** (no native MPLS encryption exists)
  - – Must encapsulate MPLS into IP, then leverage IPSec encryption technologies

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)

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	

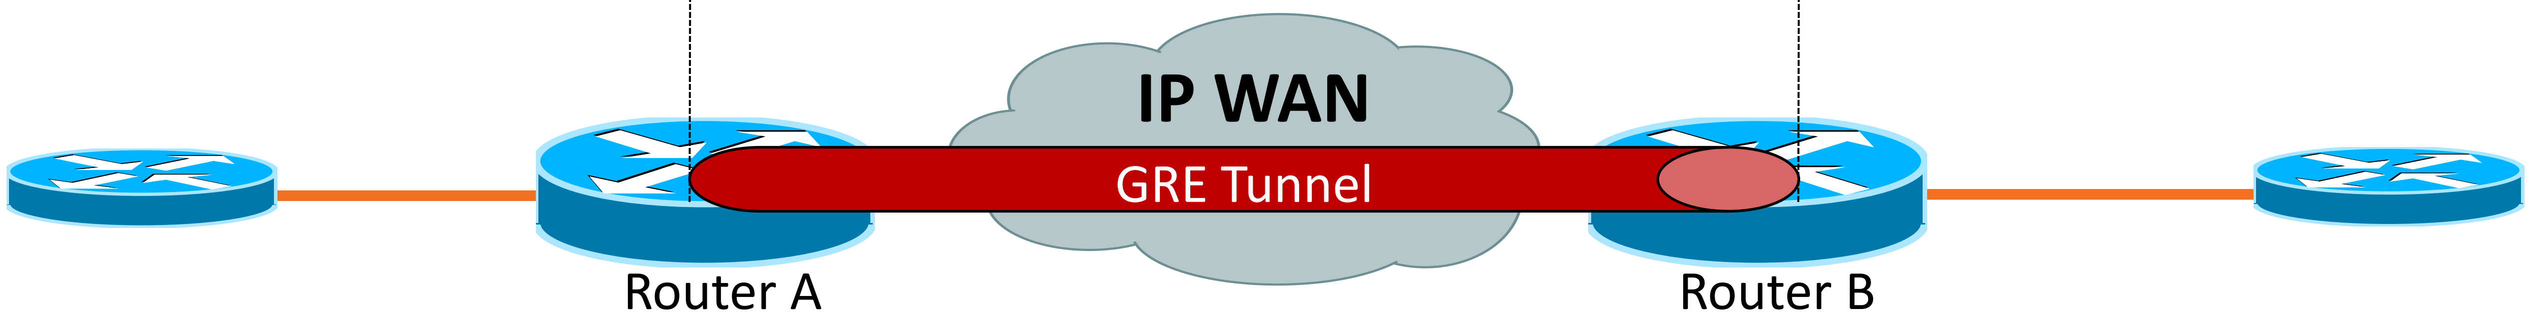

Can Also Leverage IPSec When IP Encryption Is Required of an Untrusted WAN

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
    - VRF Lite over the WAN
    - MPLS VPN over L2 WAN Transport
    - L3 Virtualisation over IP
        - **VRF-Lite over IP**
        - MPLS VPN over IP
    - L3 Virtualisation over Multipoint GRE Tunnels (mGRE)
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

VRF-Lite Over IP Transport

VRF-Lite over GRE

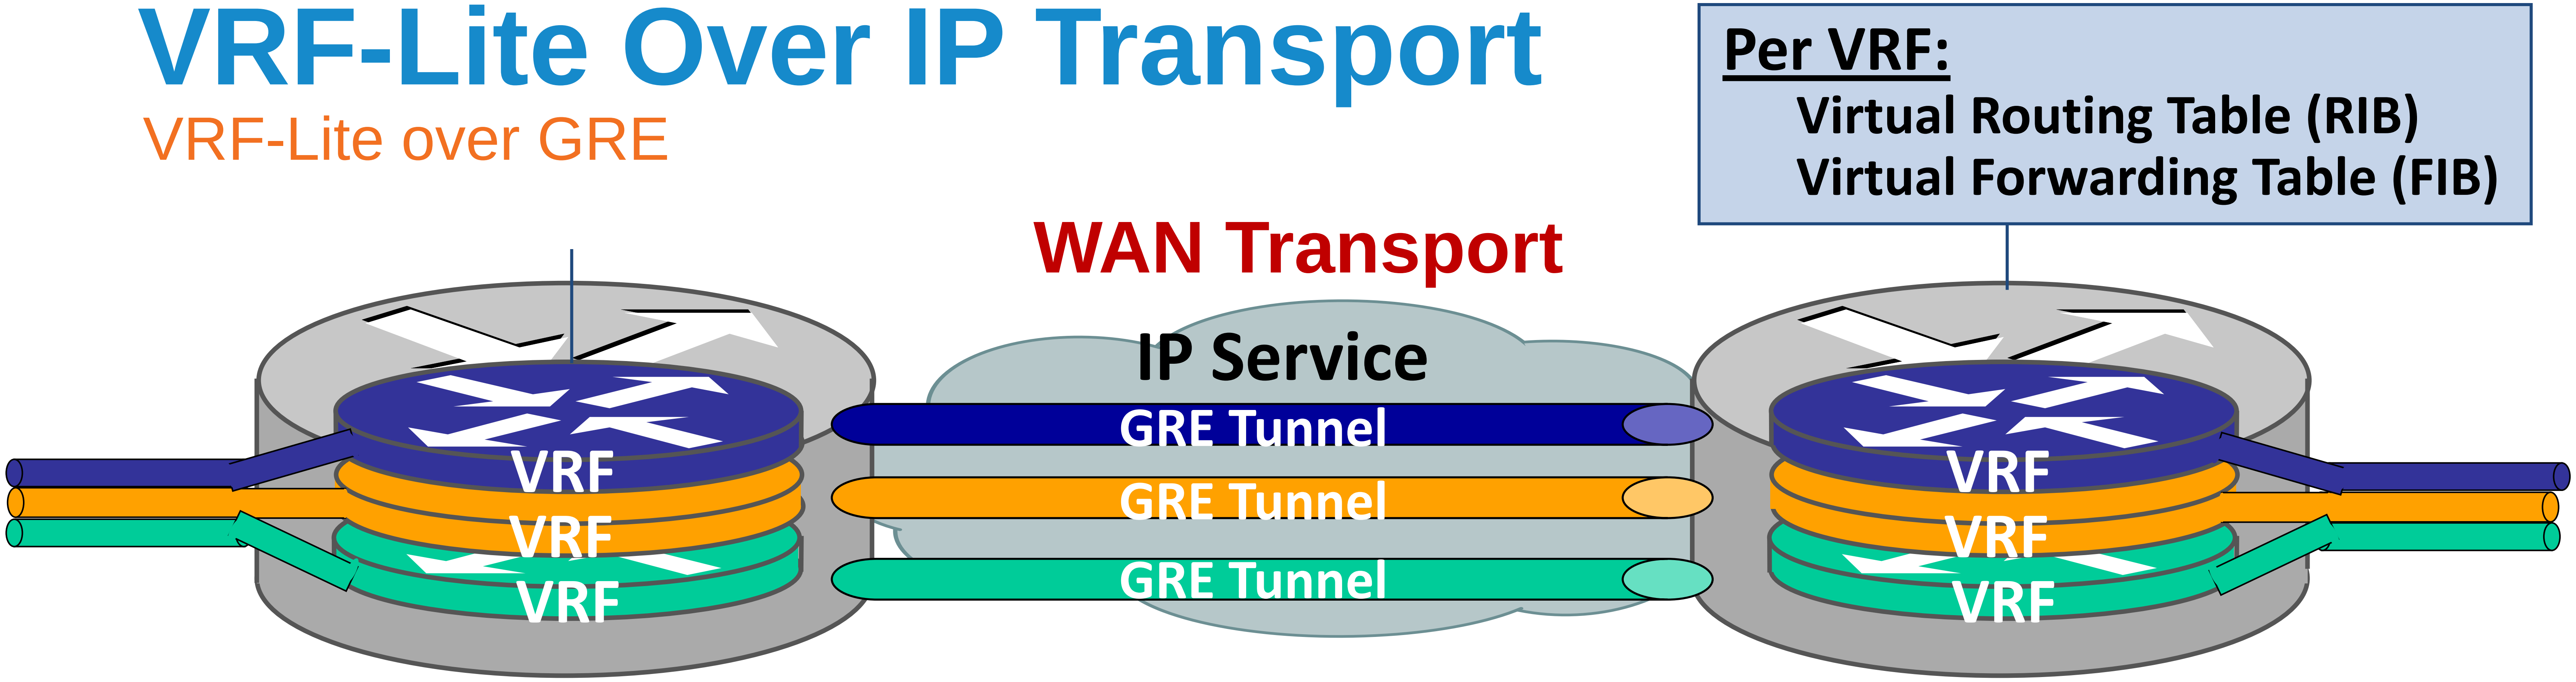

- VRF Lite can also leverage GRE tunnels as a segmentation technology
- Each VRF uses a unique GRE tunnel
- GRE tunnel interface is “VRF aware”

VRF-Lite Over the WAN

VRF-Lite per GRE Tunnels

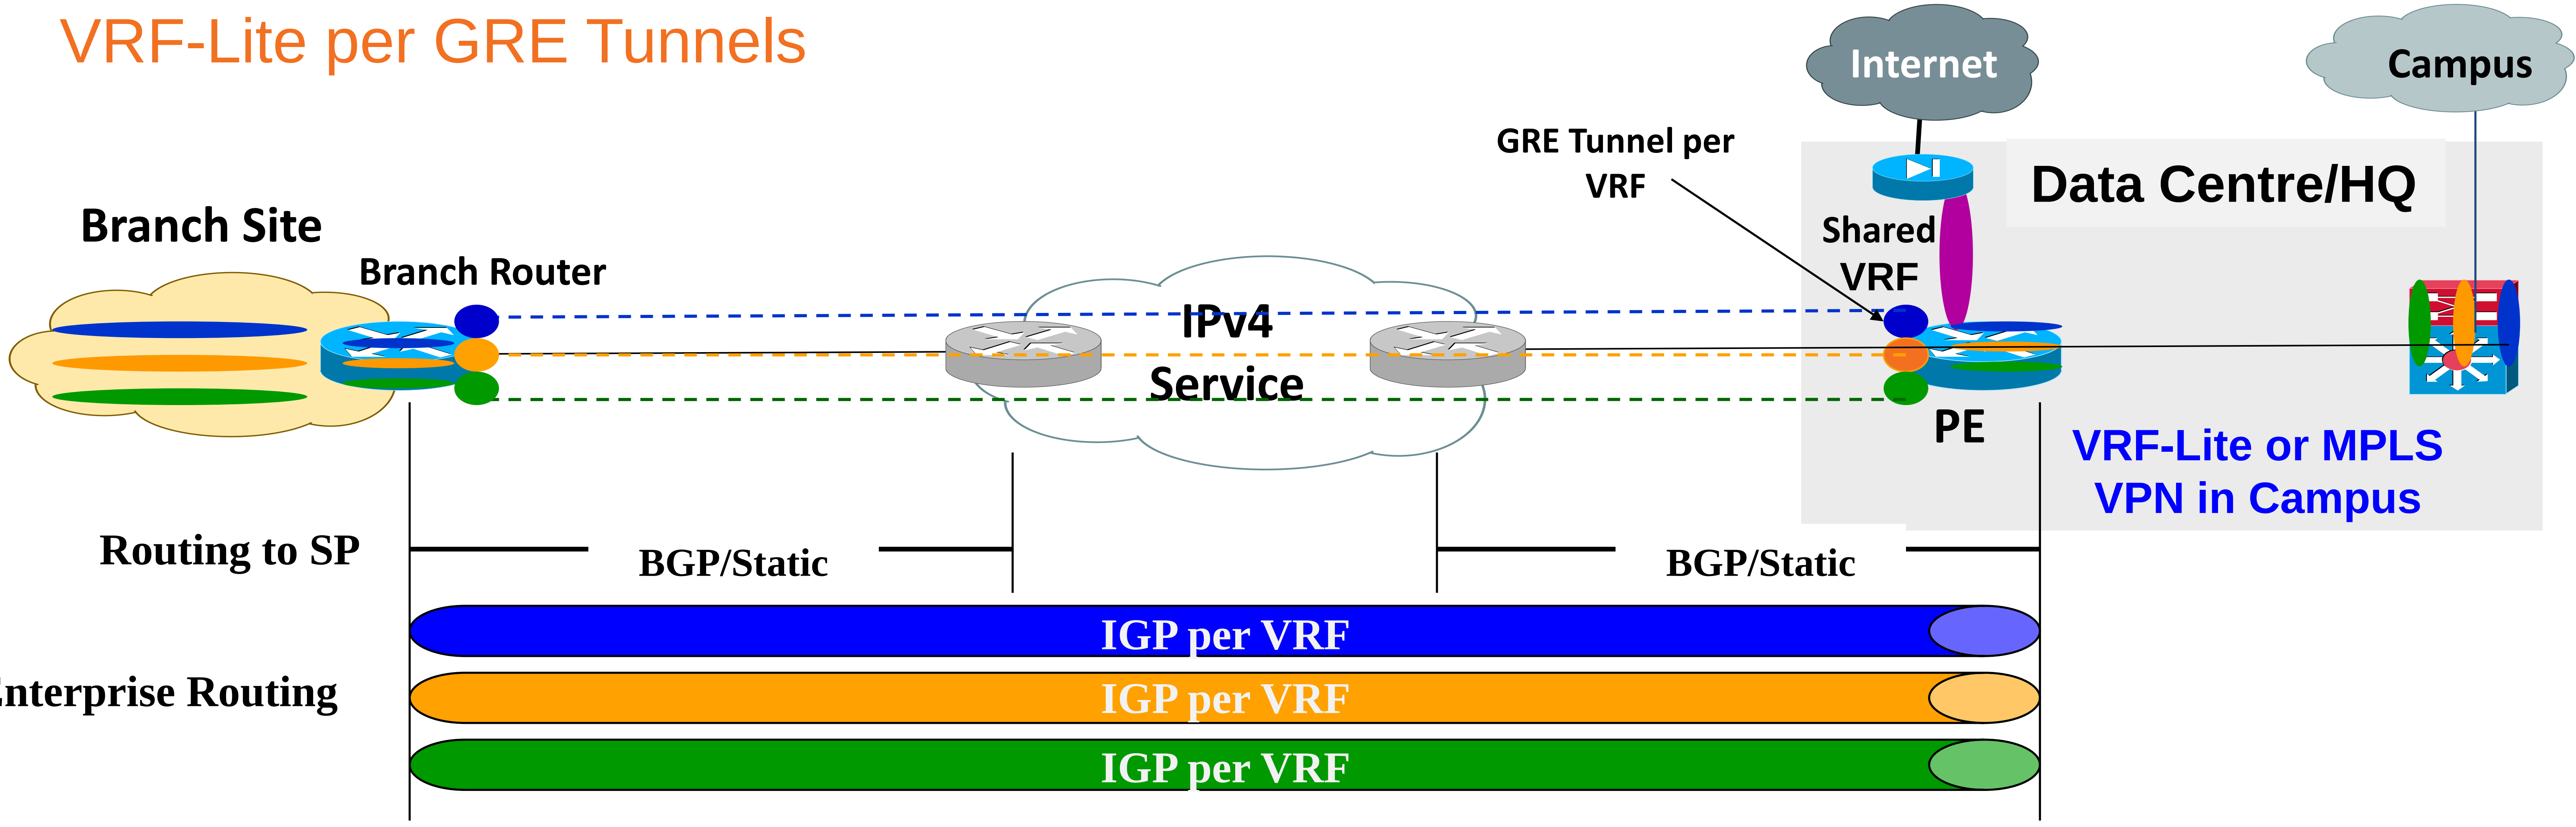

- Each GRE tunnel contains a VRF for extension
- Routing protocol process created per VRF (each end)
- Common Deployment: Branch→Aggregation Backhaul, low number of VRF's are required

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

VRF-Lite over Point-to-Point GRE

Example for “Blue” VRF (IOS)

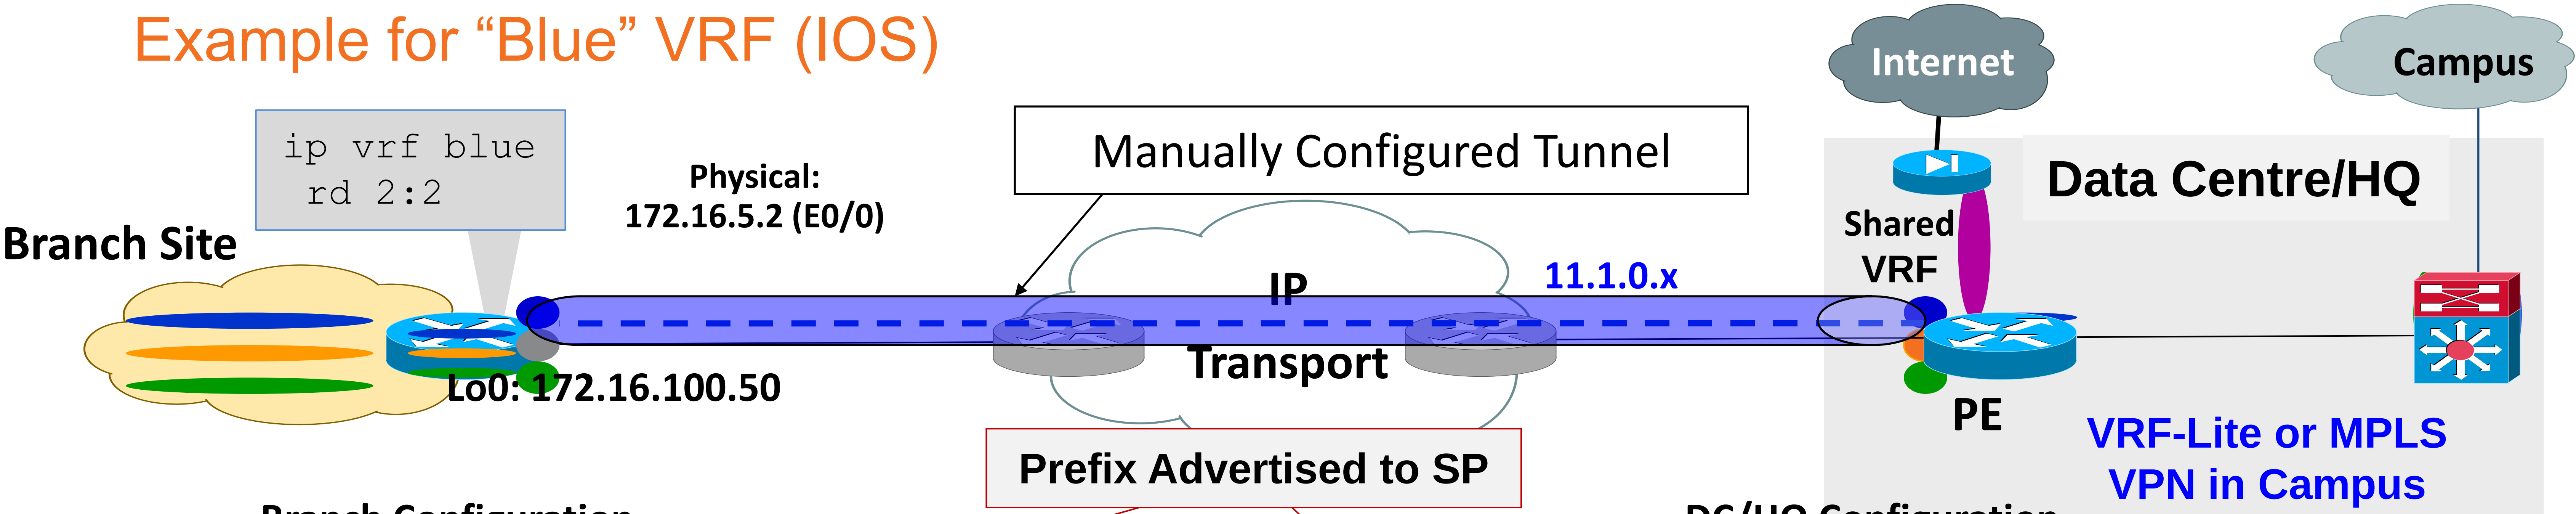

Branch Configuration

```
interface Loopback100
 ip address 172.16.100.50 255.255.255.255
!
interface Tunnel100
 Description GRE to PE router 201
 ip vrf forwarding blue
 ip address 11.1.0.2 255.255.255.0
 tunnel source Loopback100
 tunnel destination 172.16.100.10
!
interface Ethernet0/0
 ip address 172.16.5.2 255.255.255.0
!
router eigrp 1
 !
 address-family ipv4 vrf blue autonomous-system 1
  network 11.0.0.0
  no auto-summary
 exit-address-family
 no auto-summary
```

VRF Command Applied per GRE Tunnel

DC/HQ Configuration

```
interface Loopback100
 ip address 172.16.100.10 255.255.255.255
!
interface Tunnel100
 Description GRE to PE router 201
 ip vrf forwarding blue
 ip address 11.1.0.1 255.255.255.0
 tunnel source Loopback100
 tunnel destination 172.16.100.50
!
interface Ethernet0/0
 ip address 172.16.6.2 255.255.255.0
!
router eigrp 1
 !
 address-family ipv4 vrf blue autonomous-system 1
  network 11.0.0.0
  no auto-summary
 exit-address-family
 no auto-summary
```

VRF-Lite Considerations in WAN Deployments (VRF-Lite over GRE)

Is VRF-Lite the Best Fit for My Network?

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 too?
- What is the transport? (i.e. is VRF-Lite over GRE required?)
- Do I have the expertise to manage an MPLS VPN network?

Example: 4 Sites with 4 VRFs

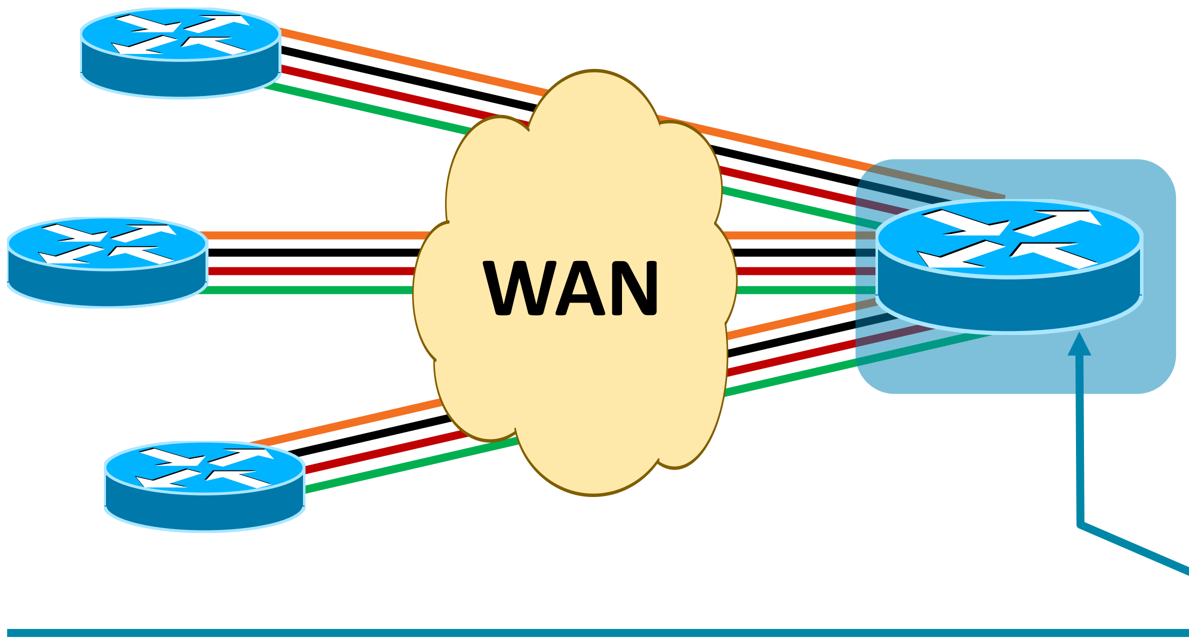

Virtual Networks	Neighbours	GRE Tunnels (1 per VRF)
4	3	12
10	3	30
20	3	60
30	3	90

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
  - VRF Lite over the WAN
  - MPLS VPN over L2 WAN Transport
  - L3 Virtualisation over IP
    - VRF-Lite over IP
    - **MPLS VPN over IP**
  - L3 Virtualisation over Multipoint GRE Tunnels (mGRE)
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

GRE (RFC 2784) with GRE+MPLS (RFC 4023)

Packet Format

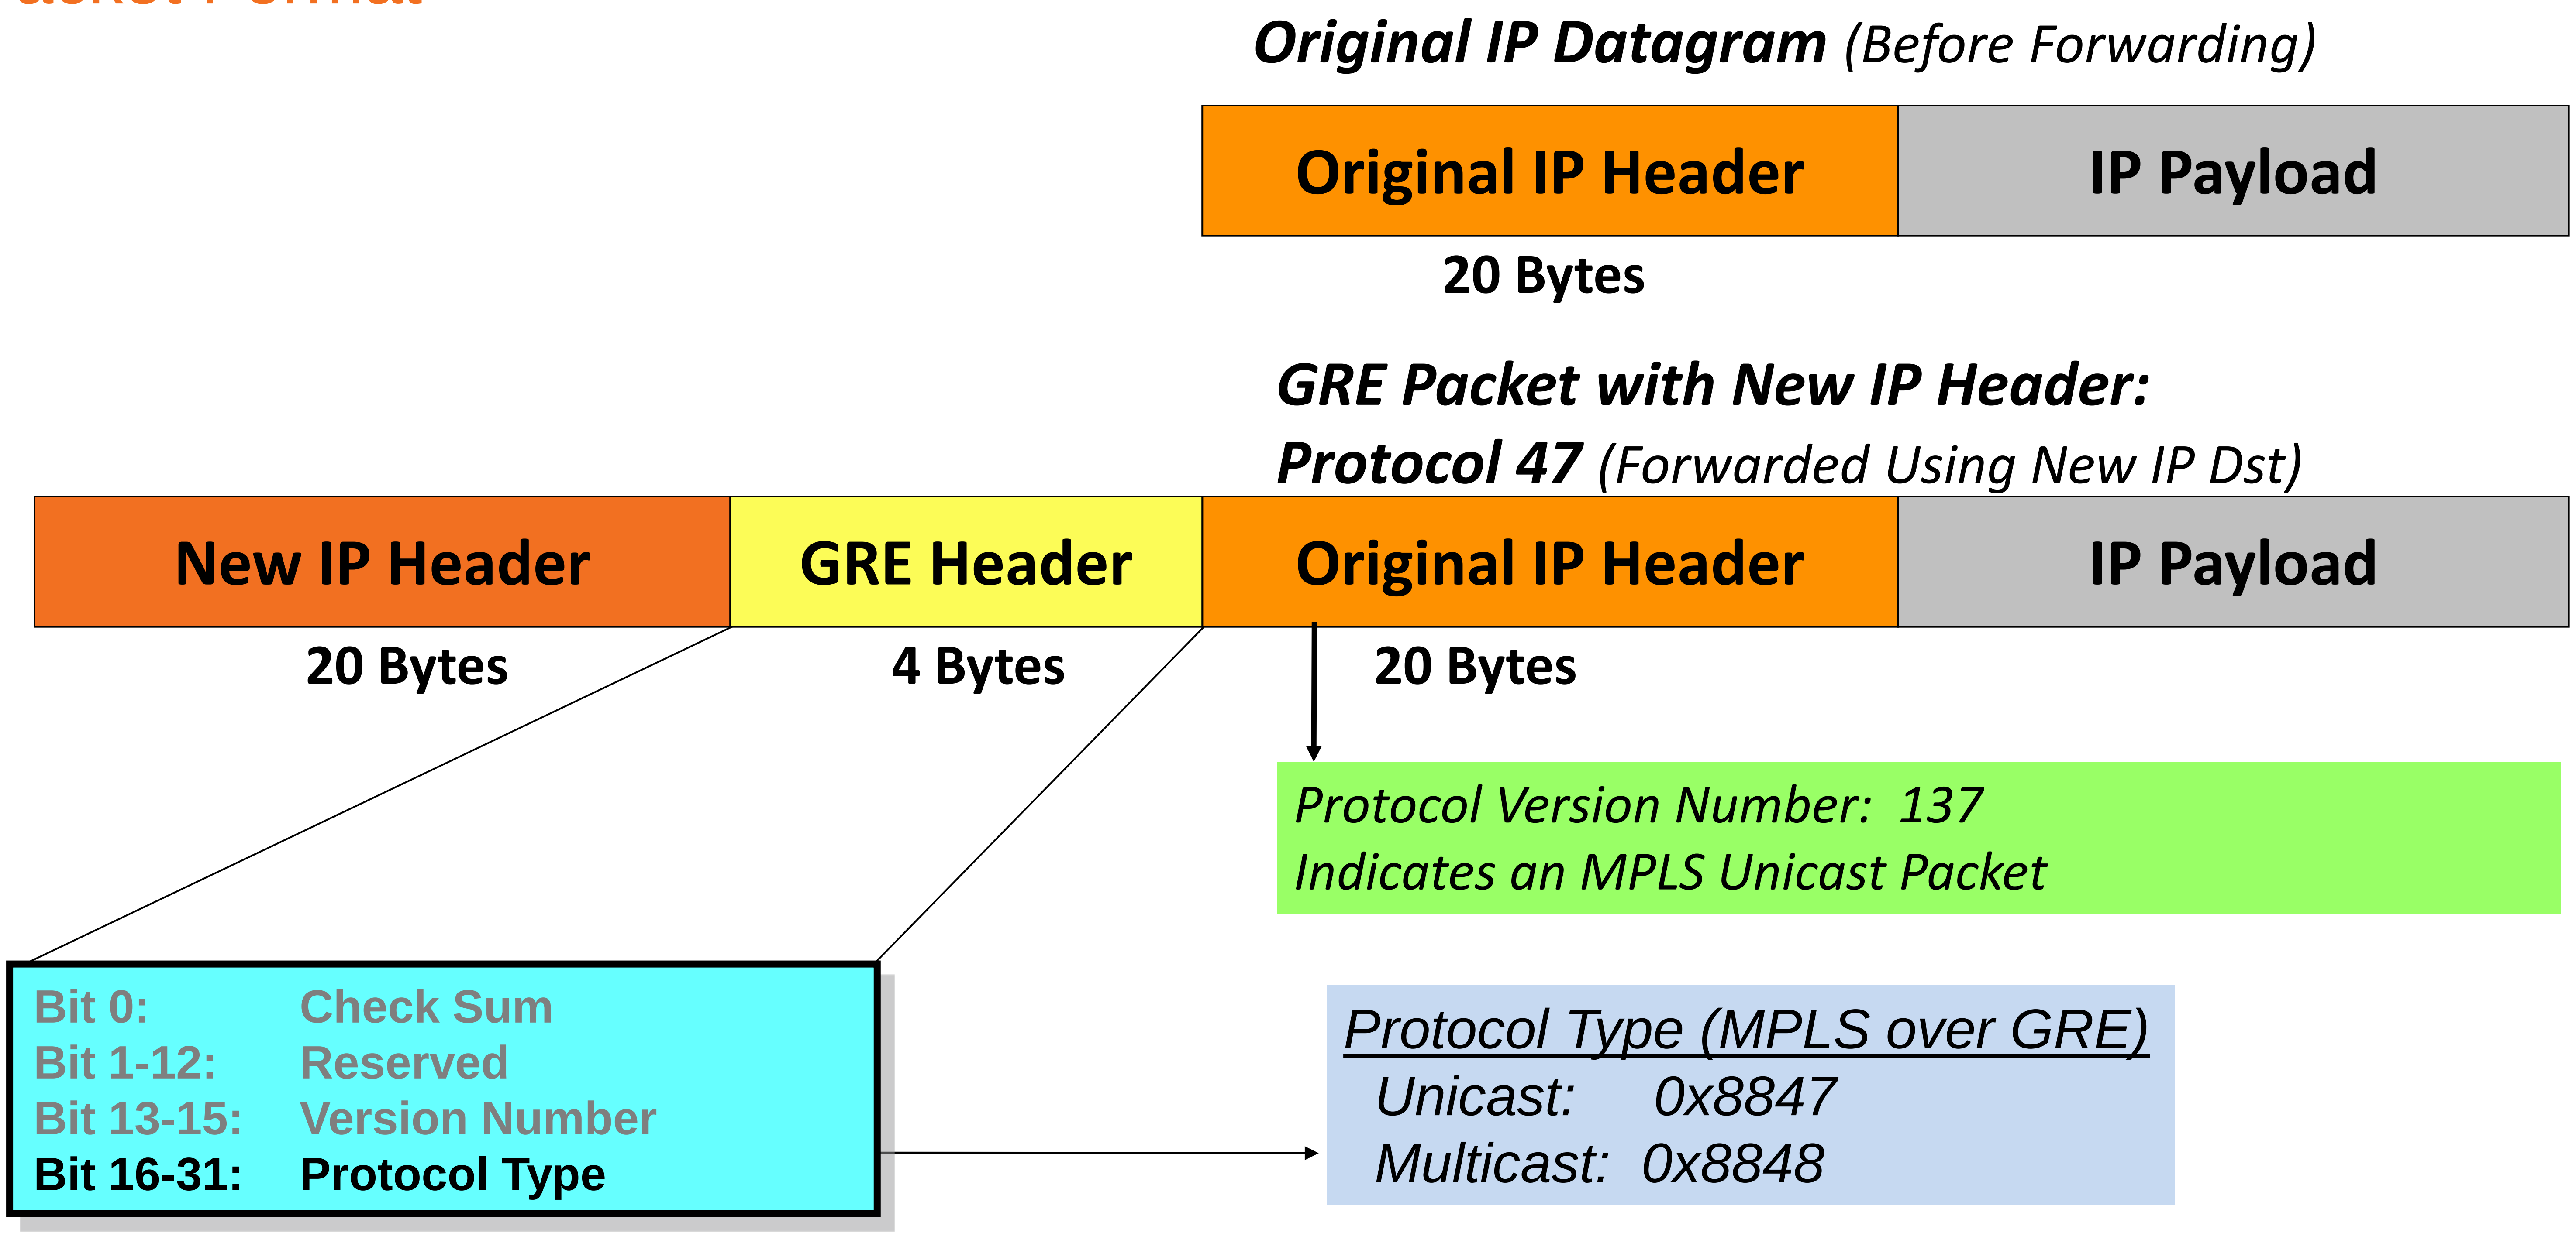

GRE Tunnel Format with MPLS

(Reference: RFC 4023)

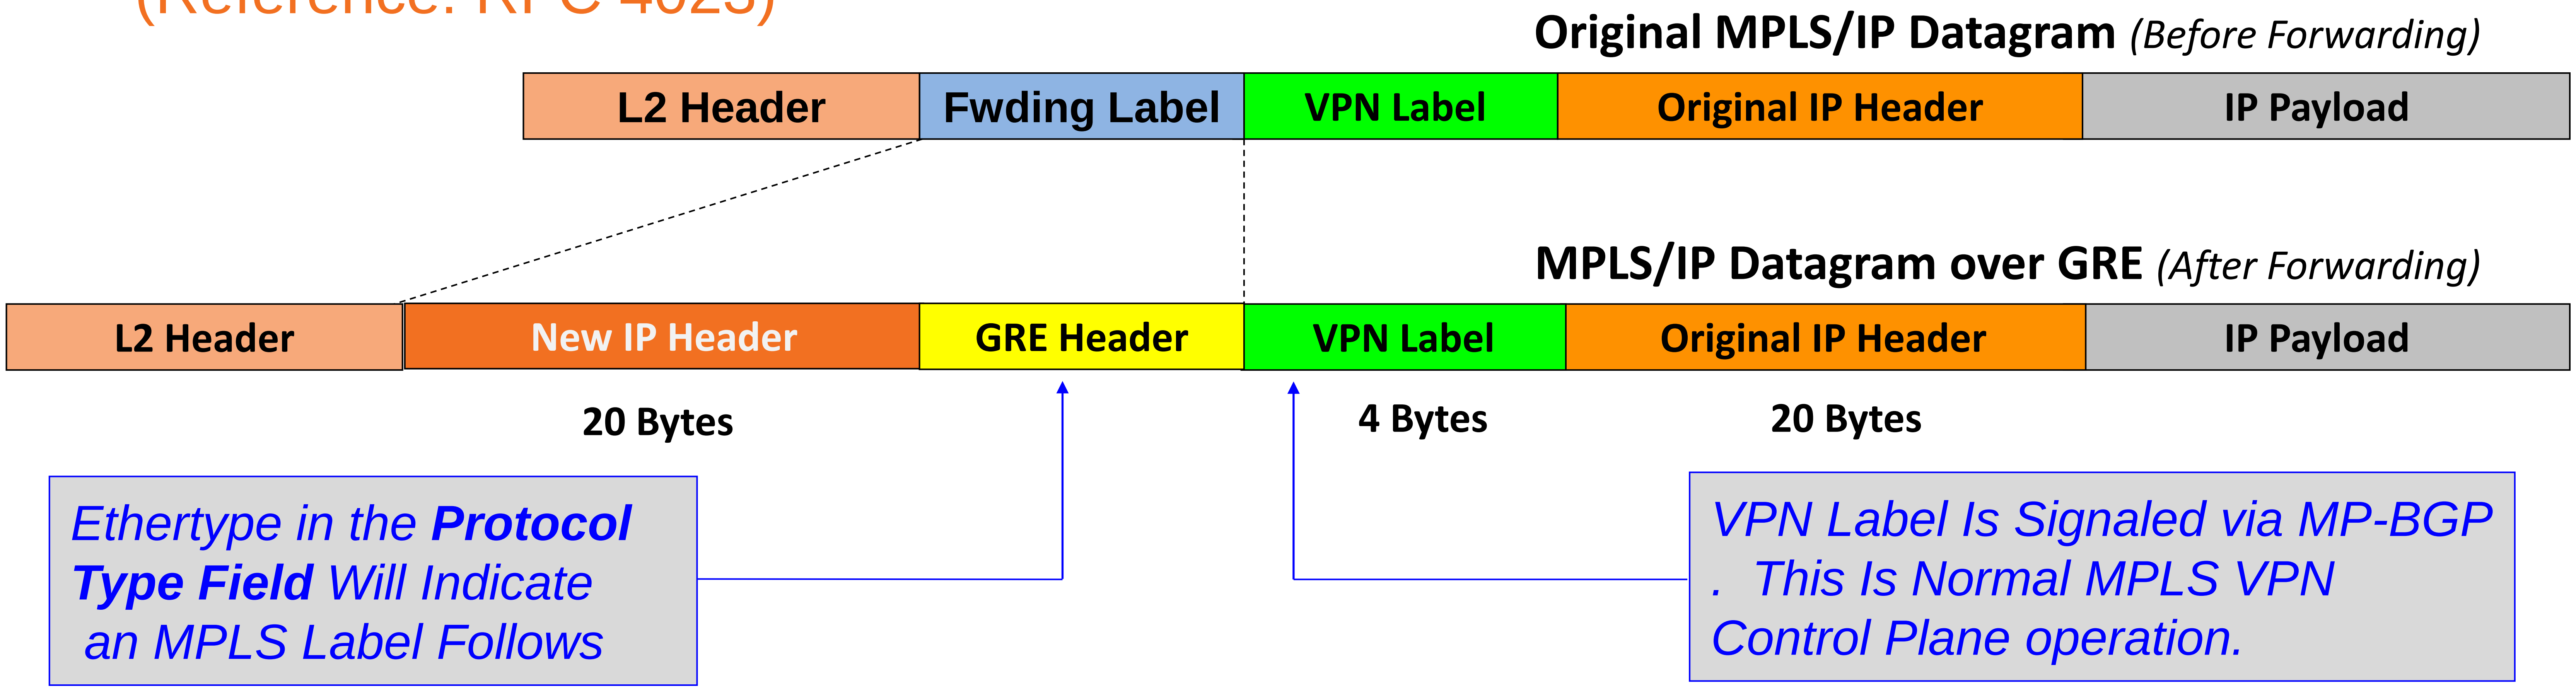

- MPLS Tunnel label (top) is replaced with destination PE's IP address
- Encapsulation defined in RFC 4023
- Most widely deployed form of MPLS over IP encapsulation

MPLS VPN over Point-to-Point GRE

Example is MPLS over Point-to-Point GRE Tunnel

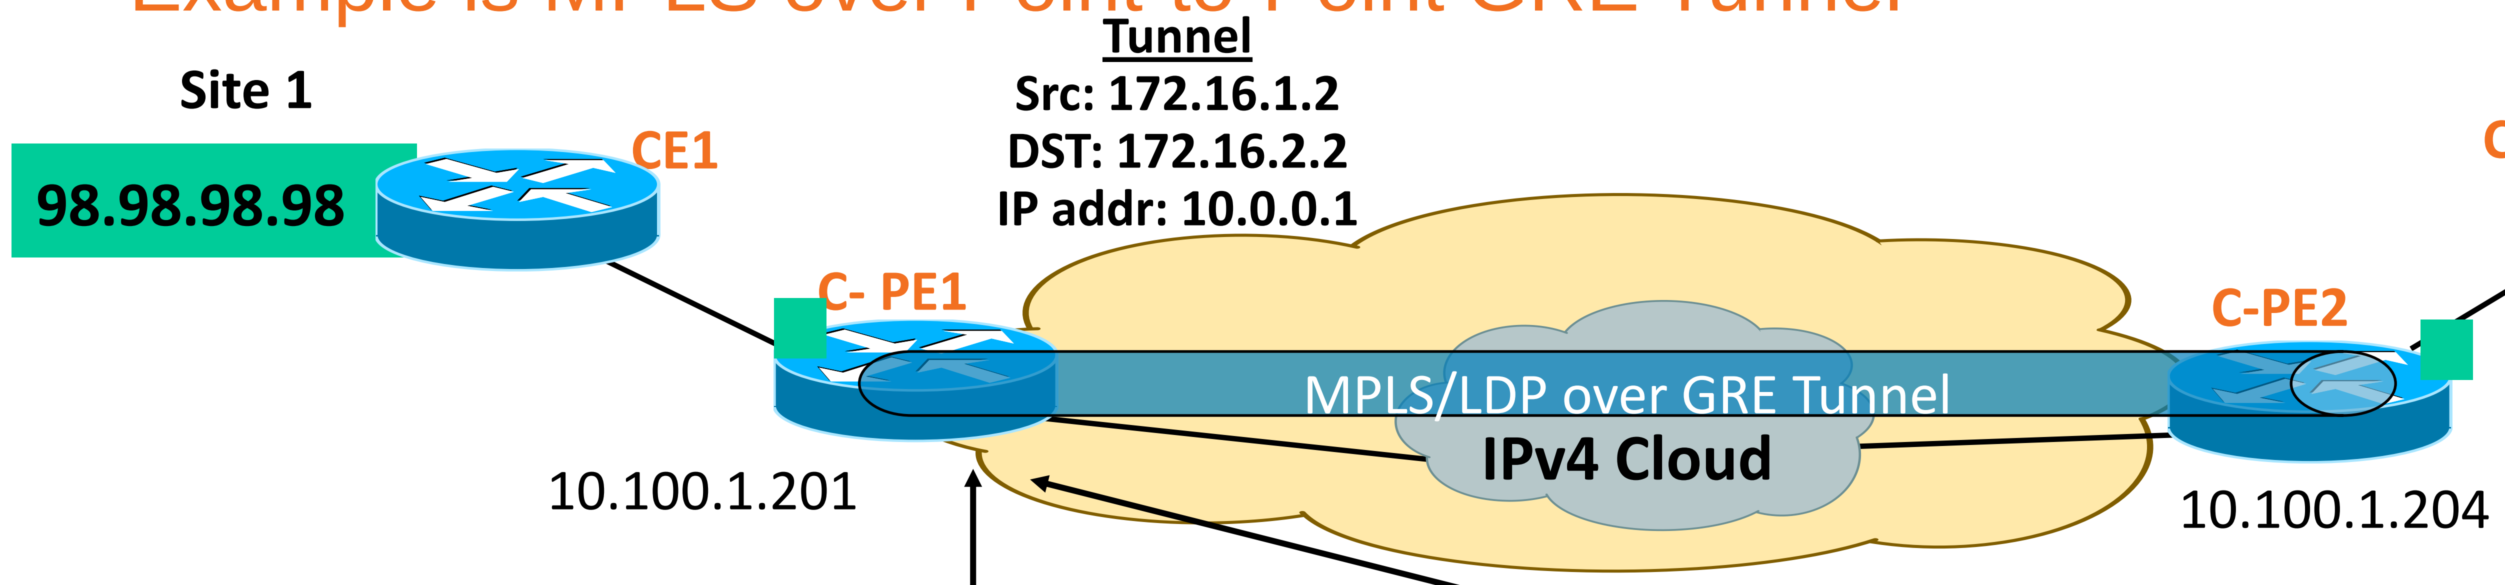

Important: You must force the forwarding of the BGP next-hop over the GRE tunnel

```
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
!
```

Enables MPLS/LDP over GRE (→ mpls ip)

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

```
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
 !
```

VRF-Lite Considerations in WAN Deployments

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

Example: 4 Sites with 4 VRFs

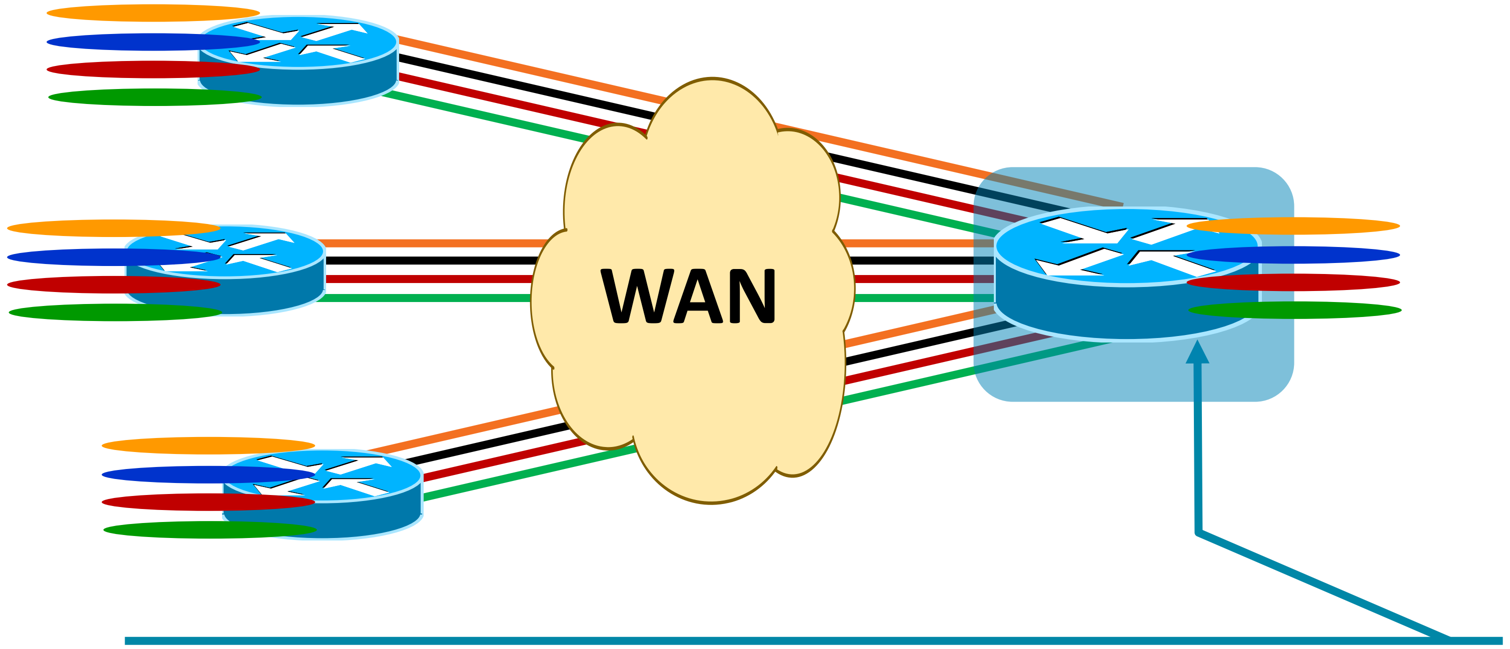

VRFs	Neighbours	GRE Tunnels (1 per VRF)
4	3	12
10	3	30
20	3	60
30	3	90

Example: 4 Sites with 4 VRFs

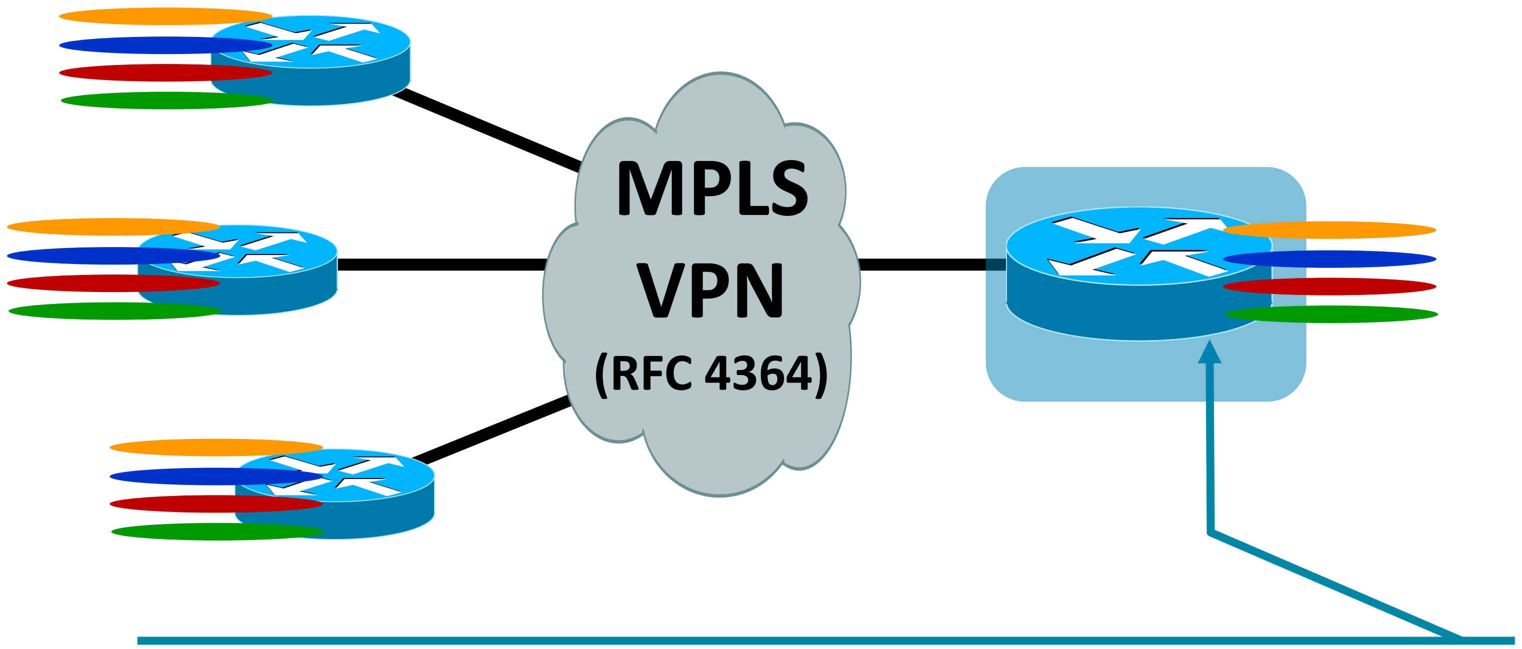

VRFs	Neighbours	Interfaces to the WAN
4	3	1
10	3	1
20	3	1
30	3	1

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
  - VRF Lite over the WAN
  - MPLS VPN over L2 WAN Transport
  - L3 Virtualisation over IP
    - VRF-Lite over IP
    - MPLS VPN over IP
  - **L3 Virtualisation over Multipoint GRE Tunnels (mGRE)**
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

GRE Tunnel Modes

“Stateful” vs. “Stateless”

Point-to-Point GRE

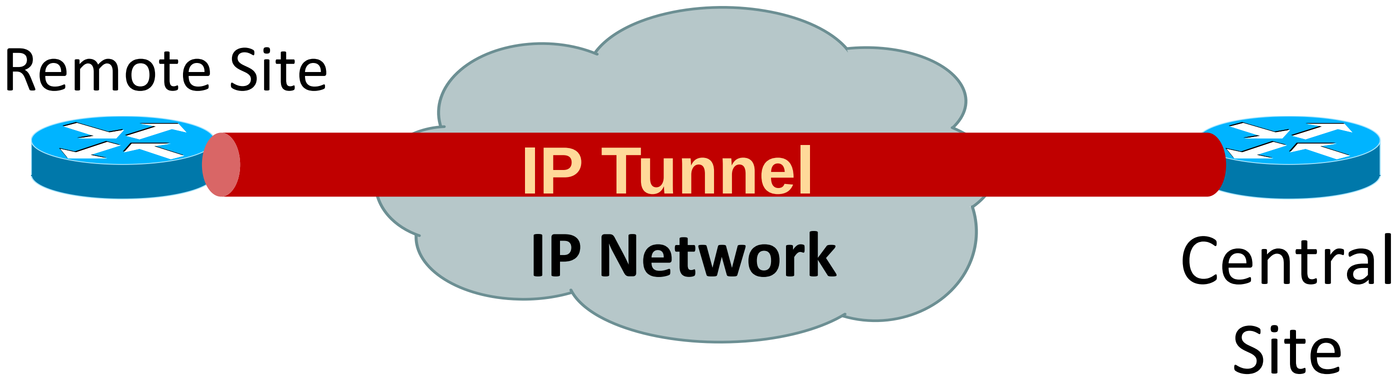

- Source and destination requires manual configuration
- Tunnel end-points are stateful neighbours
- Tunnel destination is explicitly configured
- Creates a logical point-to-point “Tunnel”

Multipoint GRE

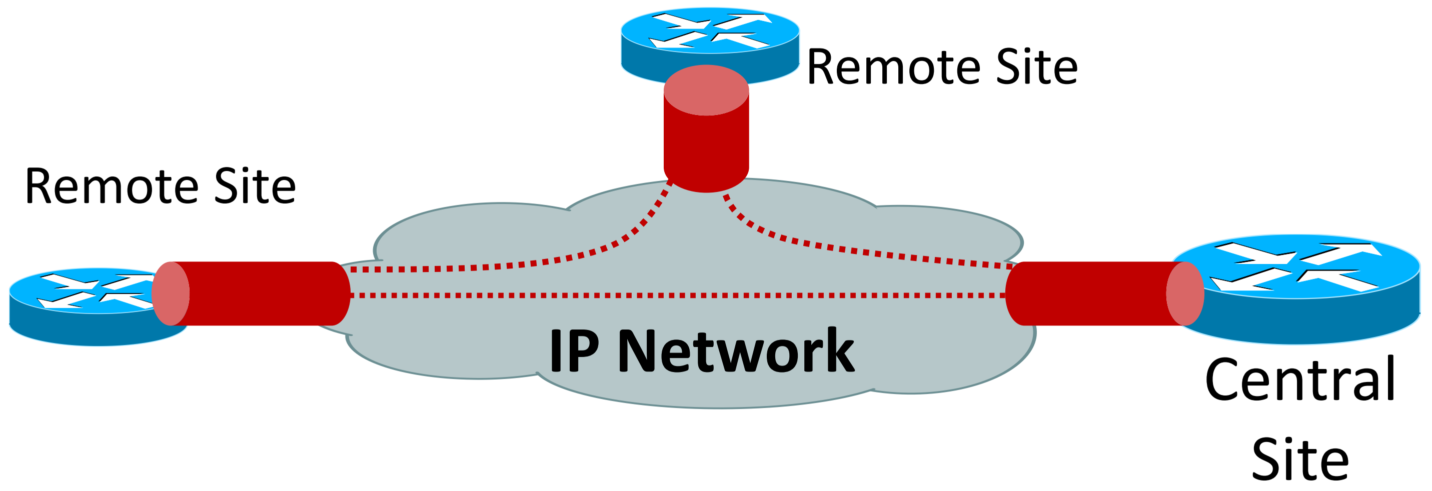

- Single multipoint tunnel interface is created per node
- Only the tunnel source is defined
- Tunnel destination is derived dynamically through some signalling mechanism (i.e. BGP, NHRP) or discovery end-point concept
- Creates an “encapsulation” using IP headers (GRE)

Dynamic Multipoint VPN

- Provides full meshed connectivity with simple configuration of hub and spoke
- Supports dynamically addressed spokes
- Facilitates zero-touch configuration for addition of new spokes
- Features automatic IPsec triggering for building an IPsec tunnel

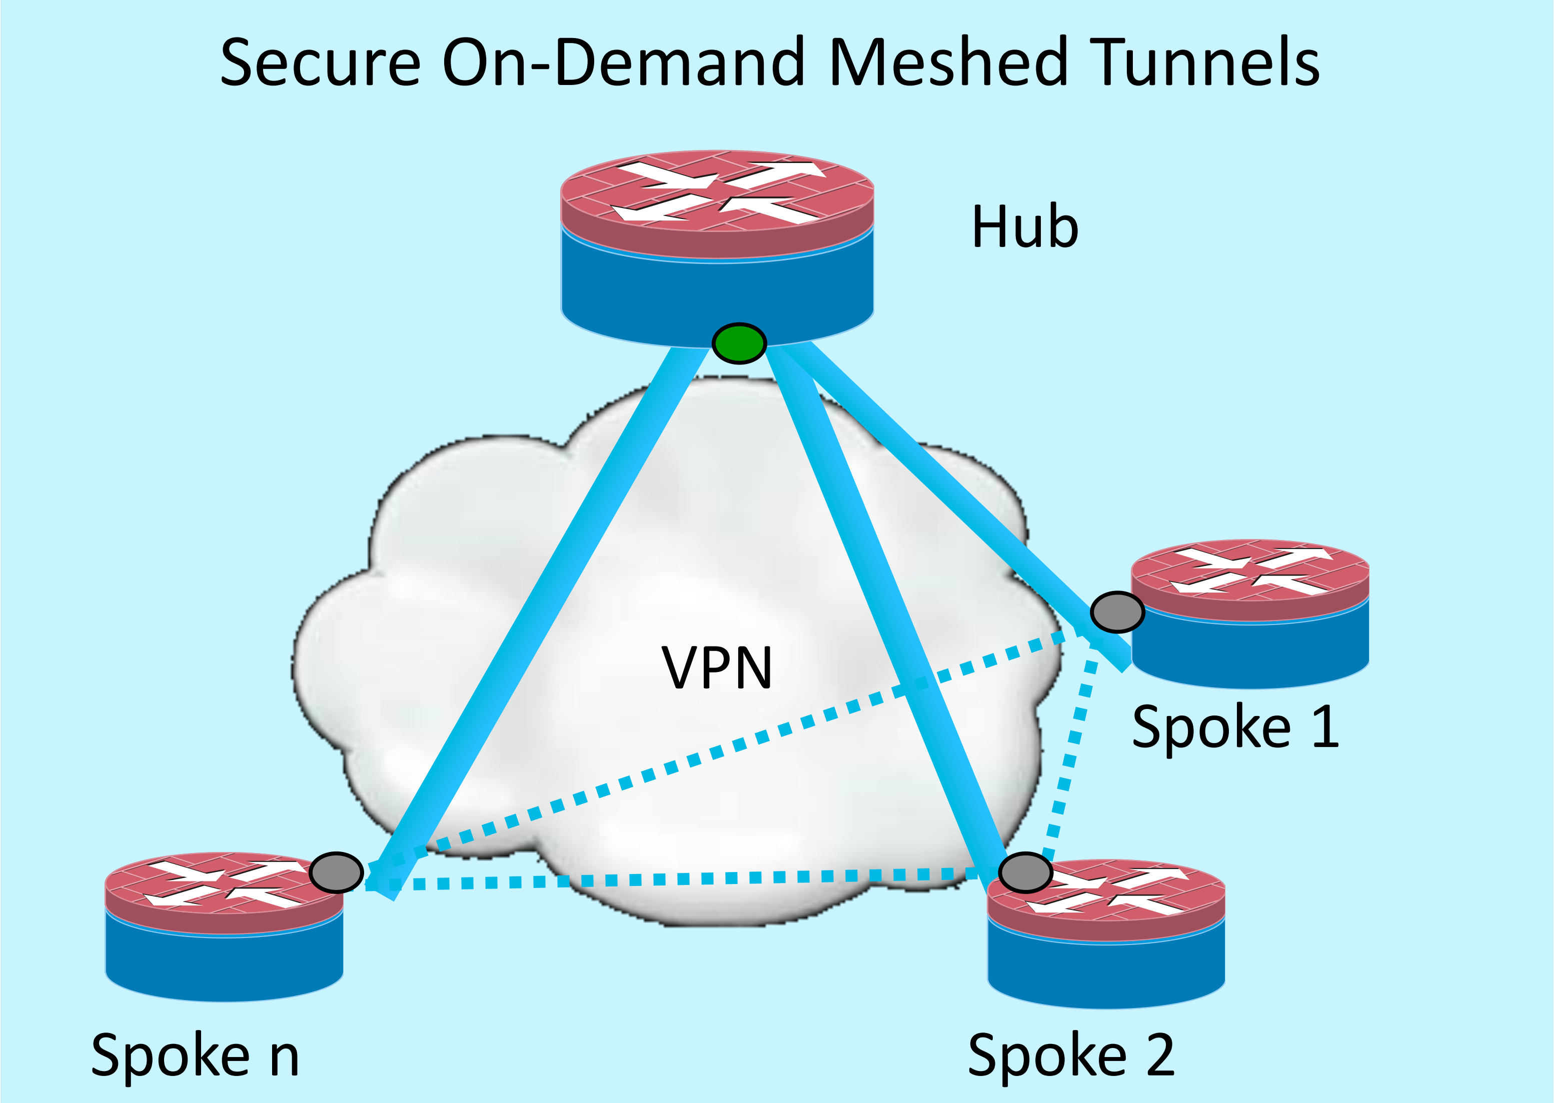

DMVPN Tunnels
Traditional Static Tunnels
Static Known IP Addresses
Dynamic Unknown IP Addresses

VRF-Lite Over Dynamic Multipoint VPN (DMVPN)

L3 Virtualisation Extension over DMVPN

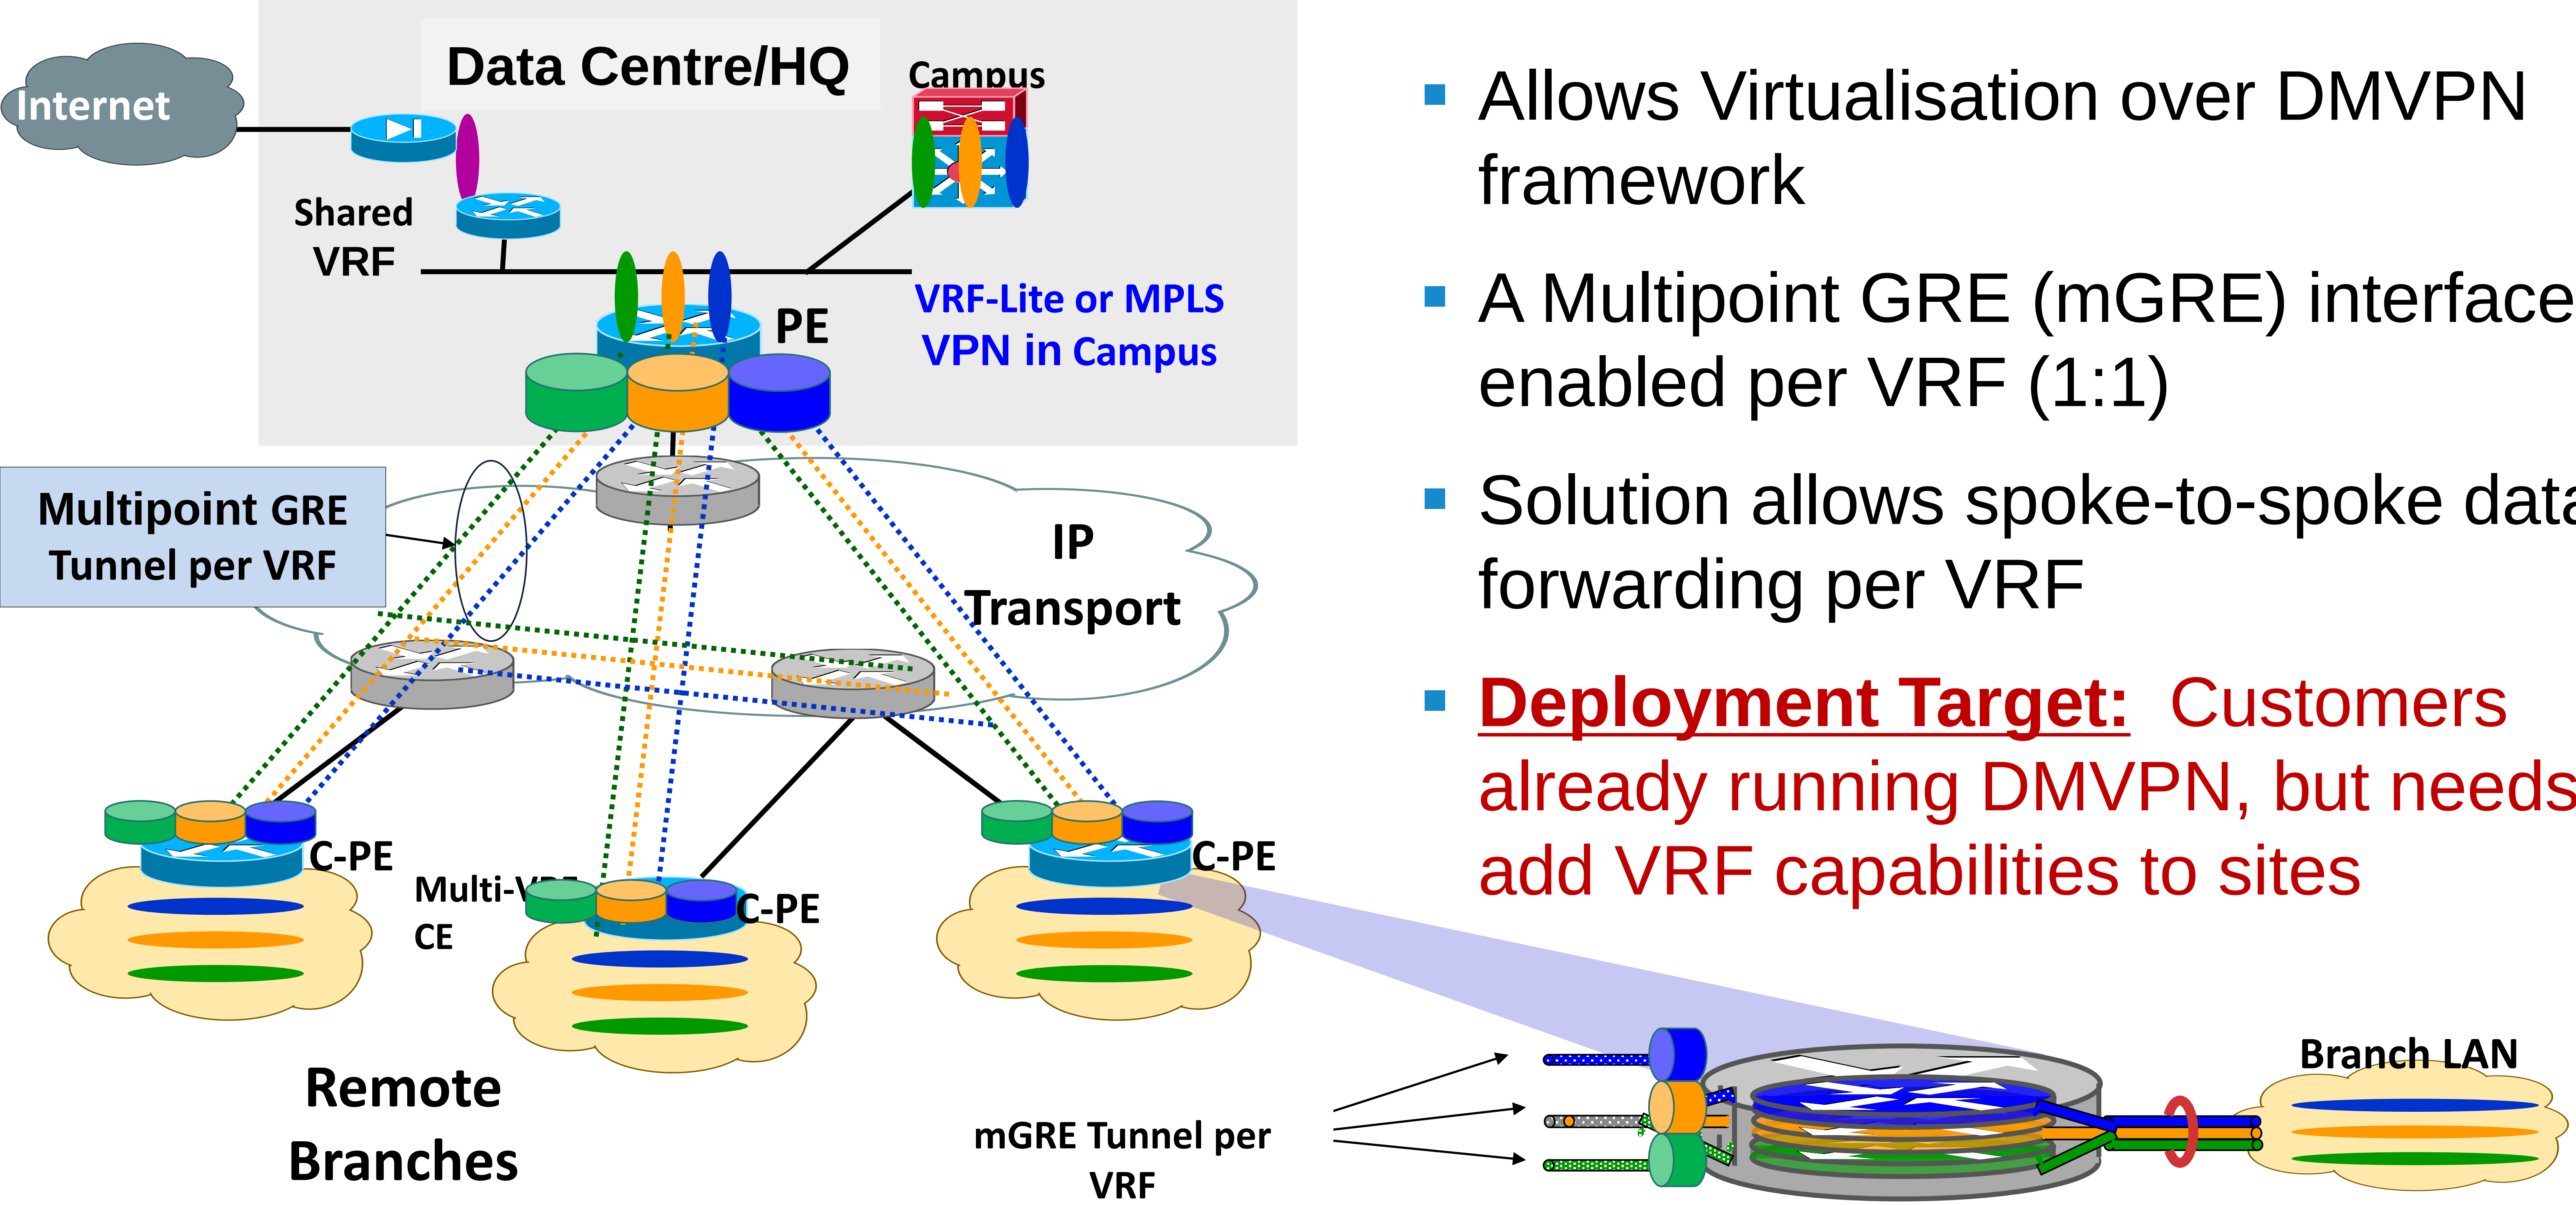

- Allows Virtualisation over DMVPN framework
- A Multipoint GRE (mGRE) interface is enabled per VRF (1:1)
- Solution allows spoke-to-spoke data forwarding per VRF
- **Deployment Target:** Customers already running DMVPN, but needs to add VRF capabilities to sites

VRF-Lite Over DMVPN

Multipoint GRE per VRF

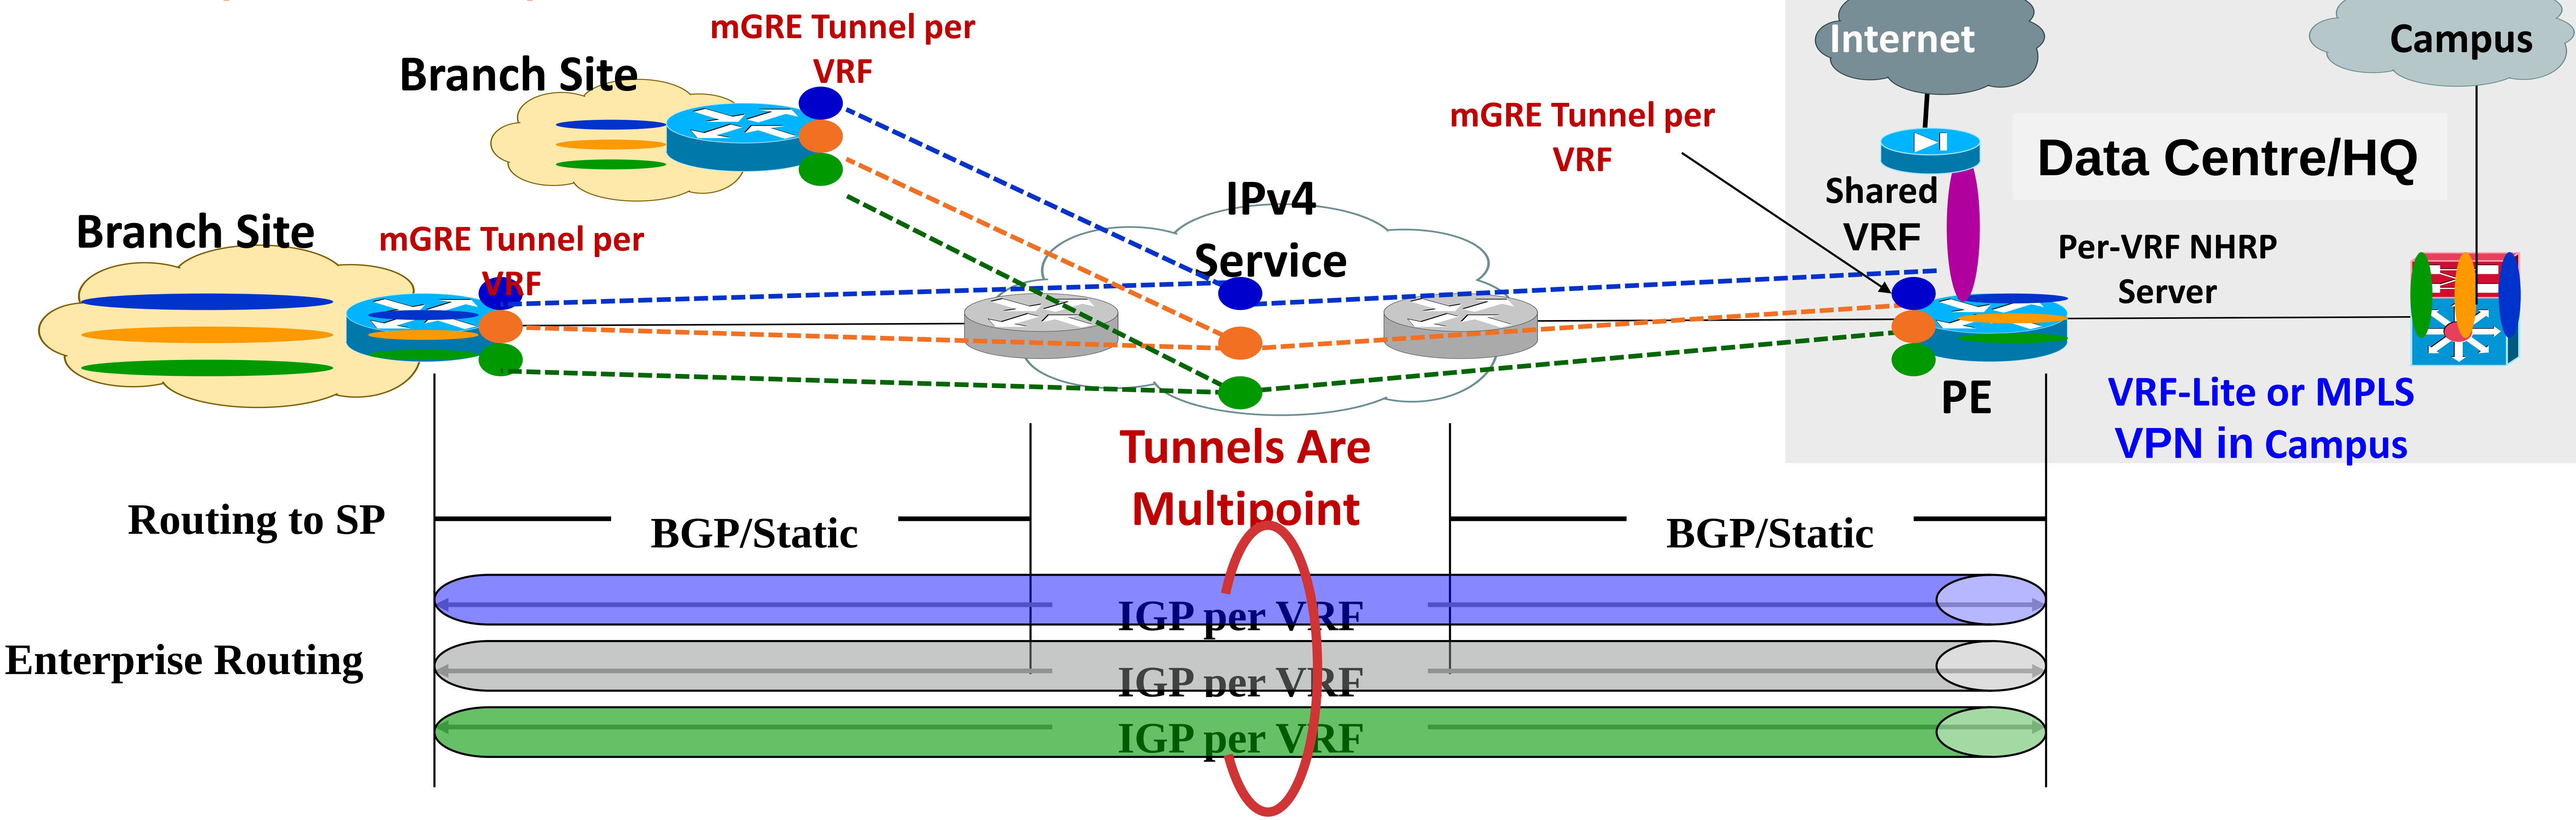

- **Unique RIB, FIB, and mGRE interface per VRF**
- Routing to the provider is based on the “global” address space
- Each VRF uses a unique network ID for each NHRP server

VRF-Lite Over DMVPN

Example (IOS)

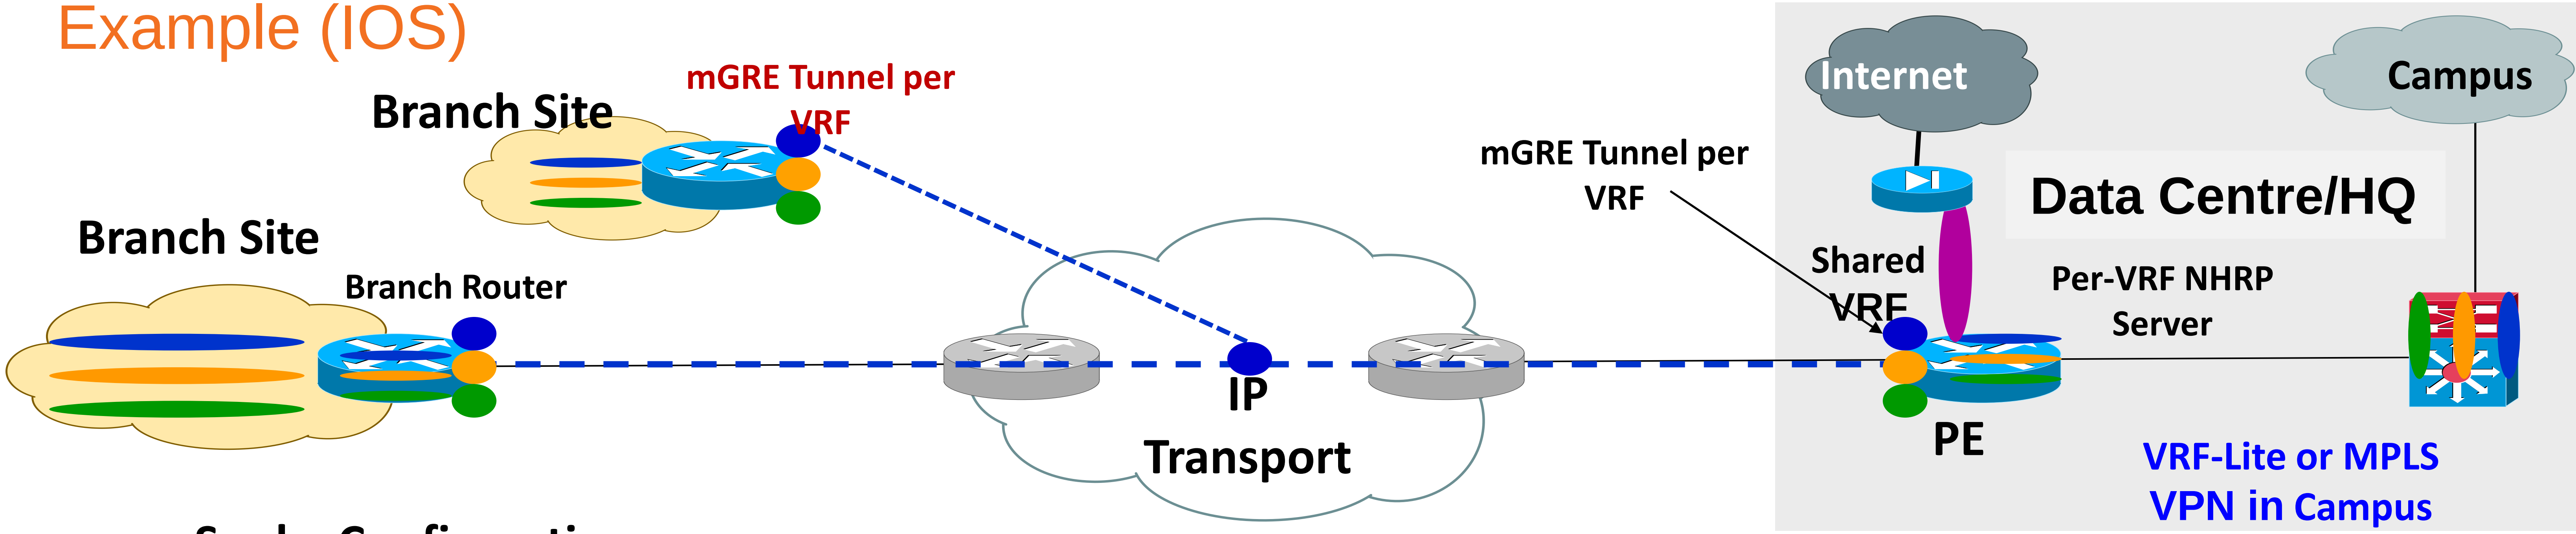

Spoke Configuration

```
ip vrf blue
!
interface Loopback0
 ip add 10.123.100.1 255.255.255.255
!
interface Tunnel0
 description GRE to hub
 ip vrf forwarding blue
 ip address 11.1.1.10 255.255.255.0
 ip nhrp network-id 100
 ip nhrp nhs 11.1.1.1
 tunnel source Loopback0
 tunnel destination 10.126.100.1
!
interface Vlan10
 description blue Subnet
 ip vrf forwarding blue
 ip address 11.1.100.1 255.255.255.0
```

Hub Configuration

```
ip vrf blue
!
interface Loopback0
 ip address 10.126.100.1 255.255.255.255
!
interface Tunnel0
 description mGRE for blue
 ip vrf forwarding blue
 ip address 11.1.1.1 255.255.255.0
 no ip redirects
 ip nhrp map multicast dynamic
 ip nhrp network-id 100
 tunnel source Loopback0
 tunnel mode gre multipoint
```

Unique “network-id” Parameter per VRF

VRF-Lite Over DMVPN

Example (IOS)

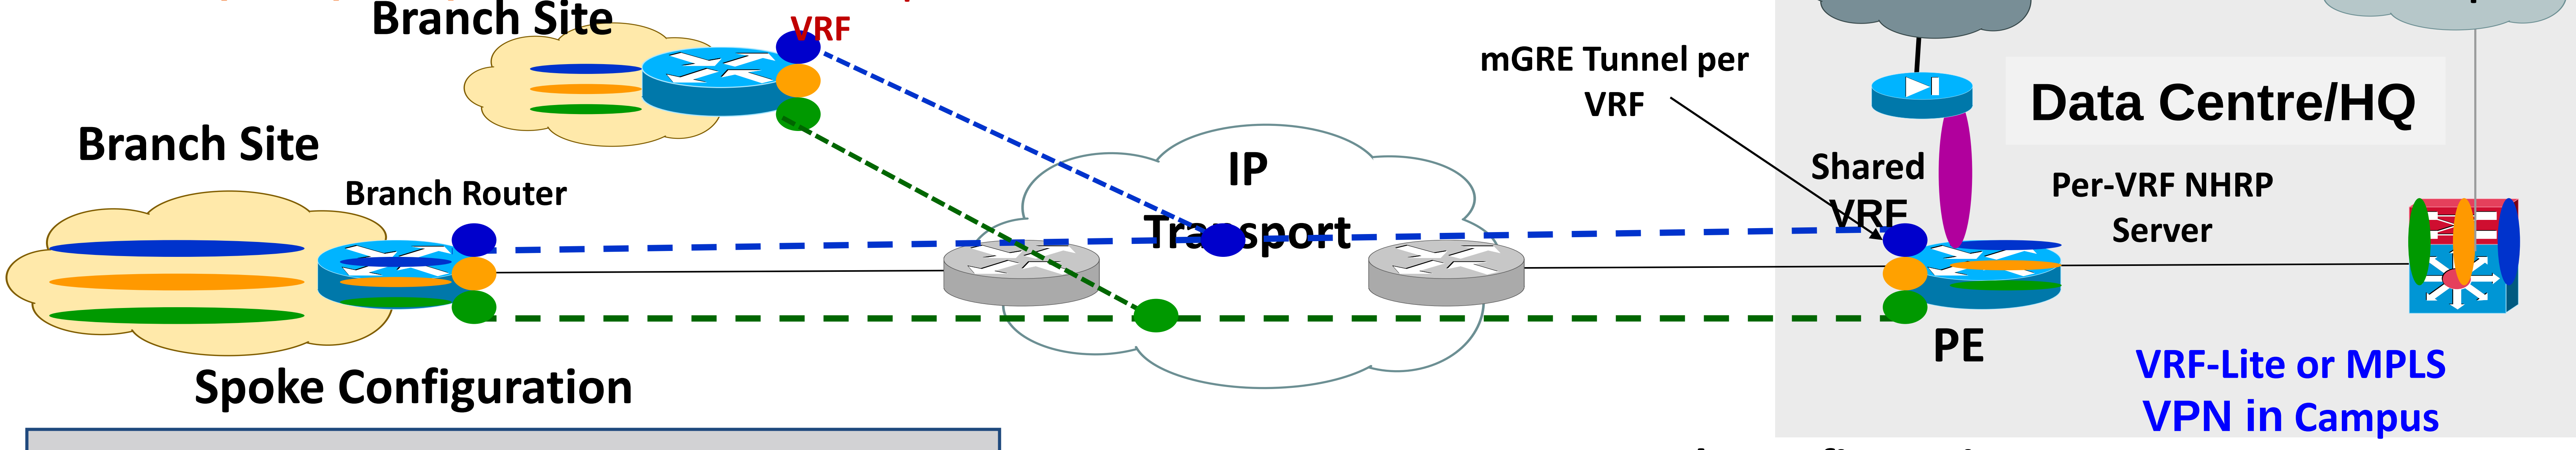

Spoke Configuration

```
ip vrf Green
!
interface Loopback0
 ip add 10.123.101.1 255.255.255.255
!
interface Tunnel0
 description GRE to hub
 ip vrf forwarding Green
 ip address 11.1.2.10 255.255.255.0
 ip nhrp network-id 101
 ip nhrp nhs 11.1.2.1
 tunnel source Loopback0
 tunnel destination 10.126.101.1
!
interface Vlan10
 description Green Subnet
 ip vrf forwarding Green
 ip address 11.1.101.1 255.255.255.0
```

Hub Configuration

```
ip vrf Green
!
interface Loopback0
 ip address 10.126.101.1 255.255.255.255
!
interface Tunnel0
 description mGRE for Green
 ip vrf forwarding Green
 ip address 11.1.2.1 255.255.255.0
 no ip redirects
 ip nhrp map multicast dynamic
 ip nhrp network-id 101
 tunnel source Loopback0
 tunnel mode gre multipoint
```

VRF-Lite Over DMVPN

Example (IOS)

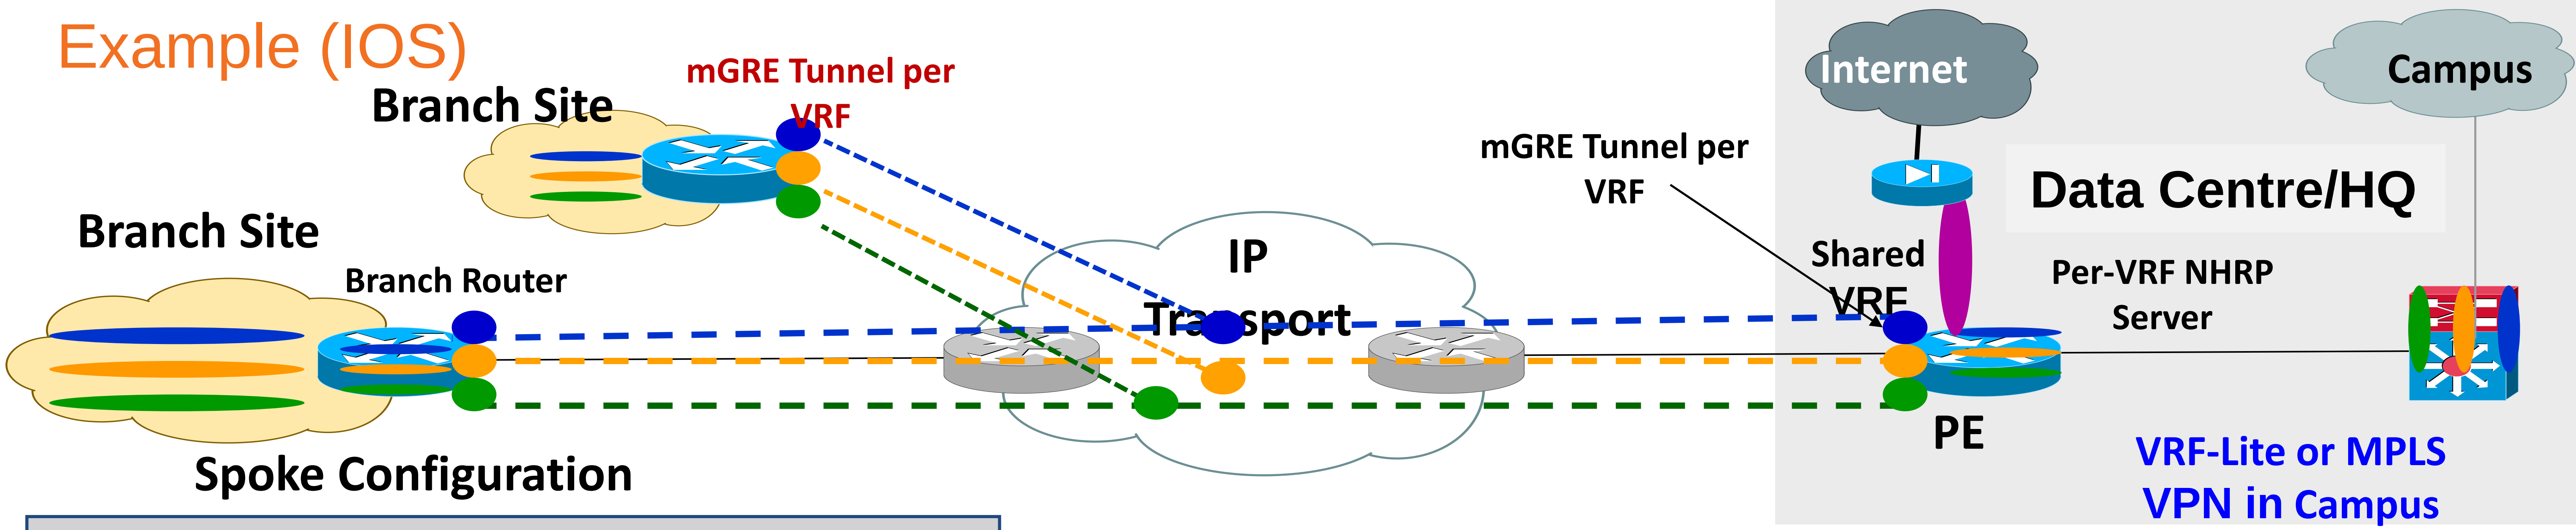

Spoke Configuration

```
ip vrf Yellow
!
interface Loopback2
 ip add 10.123.102.1 255.255.255.255
!
interface Tunnel2
 description GRE to hub
 ip vrf forwarding Yellow
 ip address 11.1.3.10 255.255.255.0
 ip nhrp network-id 103
 ip nhrp nhs 11.1.3.1
 tunnel source Loopback2
 tunnel destination 10.126.102.1
!
interface Vlan10
 description Green Subnet
 ip vrf forwarding Yellow
 ip address 11.1.102.1 255.255.255.0
```

Hub Configuration

```
ip vrf Yellow
!
interface Loopback2
 ip address 10.126.102.1 255.255.255.255
!
interface Tunnel2
 description mGRE for Yellow
 ip vrf forwarding Yellow
 ip address 11.1.3.1 255.255.255.0
 no ip redirects
 ip nhrp map multicast dynamic
 ip nhrp network-id 102
 tunnel source Loopback2
 tunnel mode gre multipoint
```

VRF-Lite Considerations in WAN Deployments (DMVPN)

Is VRF-Lite over DMVPN the Best Fit for My Network?

Key questions to ask yourself:

- What are the traffic patterns? (Hub-spoke? Spoke-spoke? Both?)
- 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?
- What is the transport? (i.e. is VRF-Lite over GRE required?)
- Do I have the expertise to manage an MPLS VPN network?

Example: 4 Sites with 4 VRFs

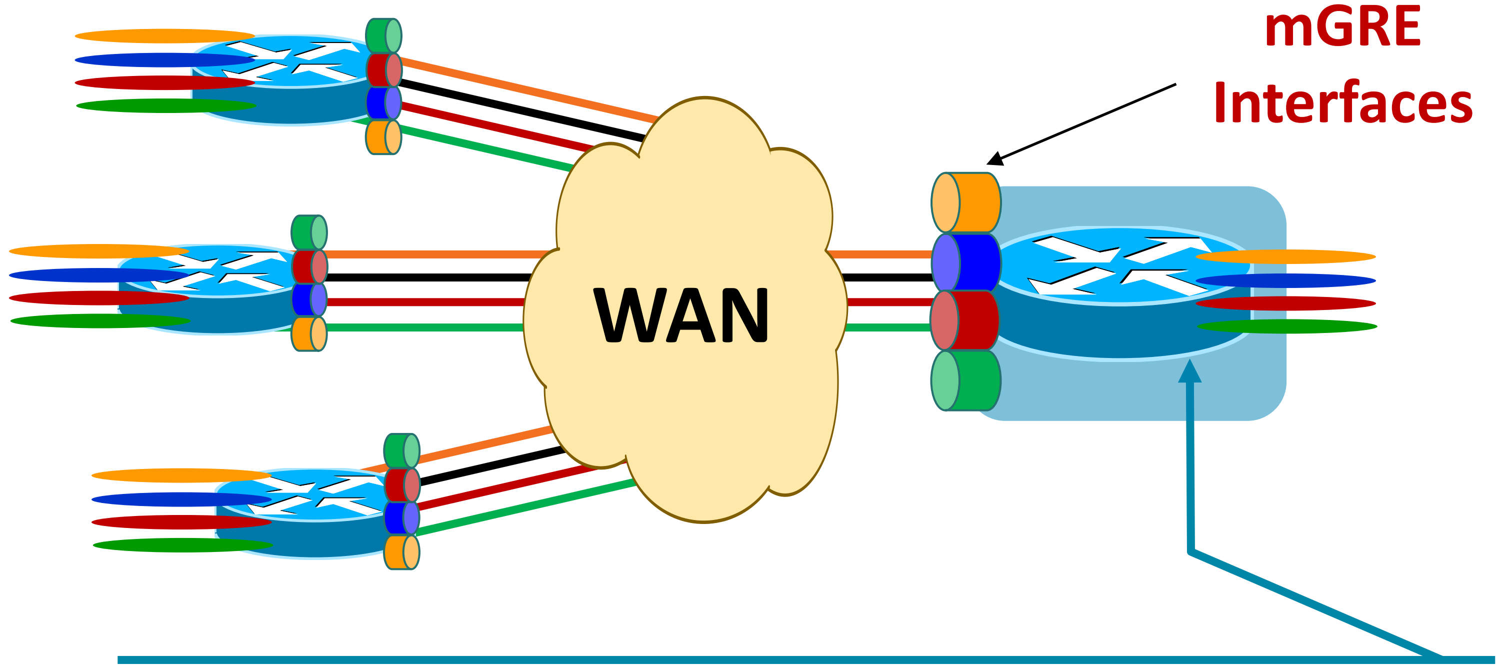

Virtual Networks	Neighbours	VRF Sub-interfaces
4	3	4
10	3	10
20	3	20
30	3	30

VRF-Lite over Dynamic Multipoint VPN (DMVPN)

Summary

- Allows Virtualisation over DMVPN framework
- Redundant Hub configurations can also be added for high availability
- Solution offers spoke-to-spoke traffic forwarding (bypass Hub), per VRF
- Multicast source at spoke is supported, but must traverse hub (traffic pattern is source → hub → spoke)

  RP <u>MUST</u> reside at the Hub Site
- Ideal solution when spoke-to-spoke traffic patterns are required
- Common QoS can be applied in VRF-Lite over DMVPN
- Tunnels in different VRF’s cannot share the same source address (unless tunnel key is used, which is not supported on the 7600/6500)

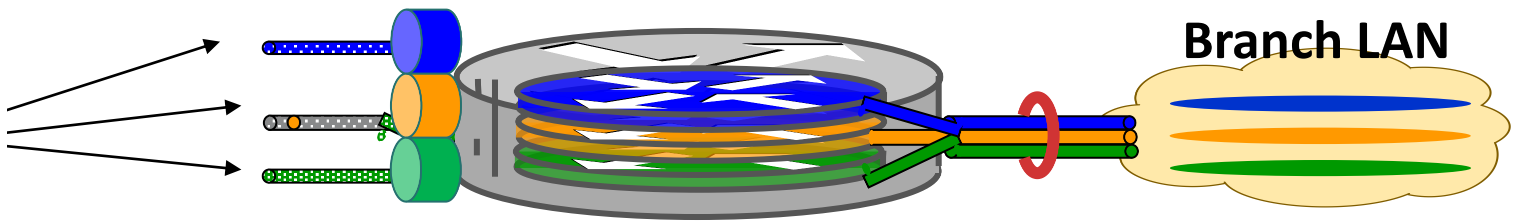

VRF-Lite Solutions over the WAN

Comparison Matrix

	VRF-Lite over Serial, FR/ATM	**VRF-Lite** over P2P GRE	**VRF-Lite** over DMVPN
Target Deployment	Campus/WAN (Ethernet in Campus)	Campus/WAN	WAN
Target Number of VRFs	< 8	< 8	< 8
Uses Dynamic Endpoint Discovery	No	No	Yes (NHRP)
Leverages Multipoint GRE Tunnels	No	No	Yes
Ability to Hide IP Addresses from SP	Yes	Yes	Yes
Supports VPN Multicast (per VRF)	Yes	Yes	Yes (Hub Sourced Only)
Support for IPv6 (Inside IPv4 Address Space)	Yes	Yes	Yes

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
    - VRF Lite over the WAN
    - MPLS VPN over L2 WAN Transport
    - L3 Virtualisation over IP
        - VRF-Lite over IP
        - MPLS VPN over IP
    - L3 Virtualisation over Multipoint GRE Tunnels (mGRE)
        - **L3 VPN over mGRE**
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

MPLS VPN Over Dynamic Multipoint VPN (DMVPN)

MPLS VPN over a DMVPN Framework

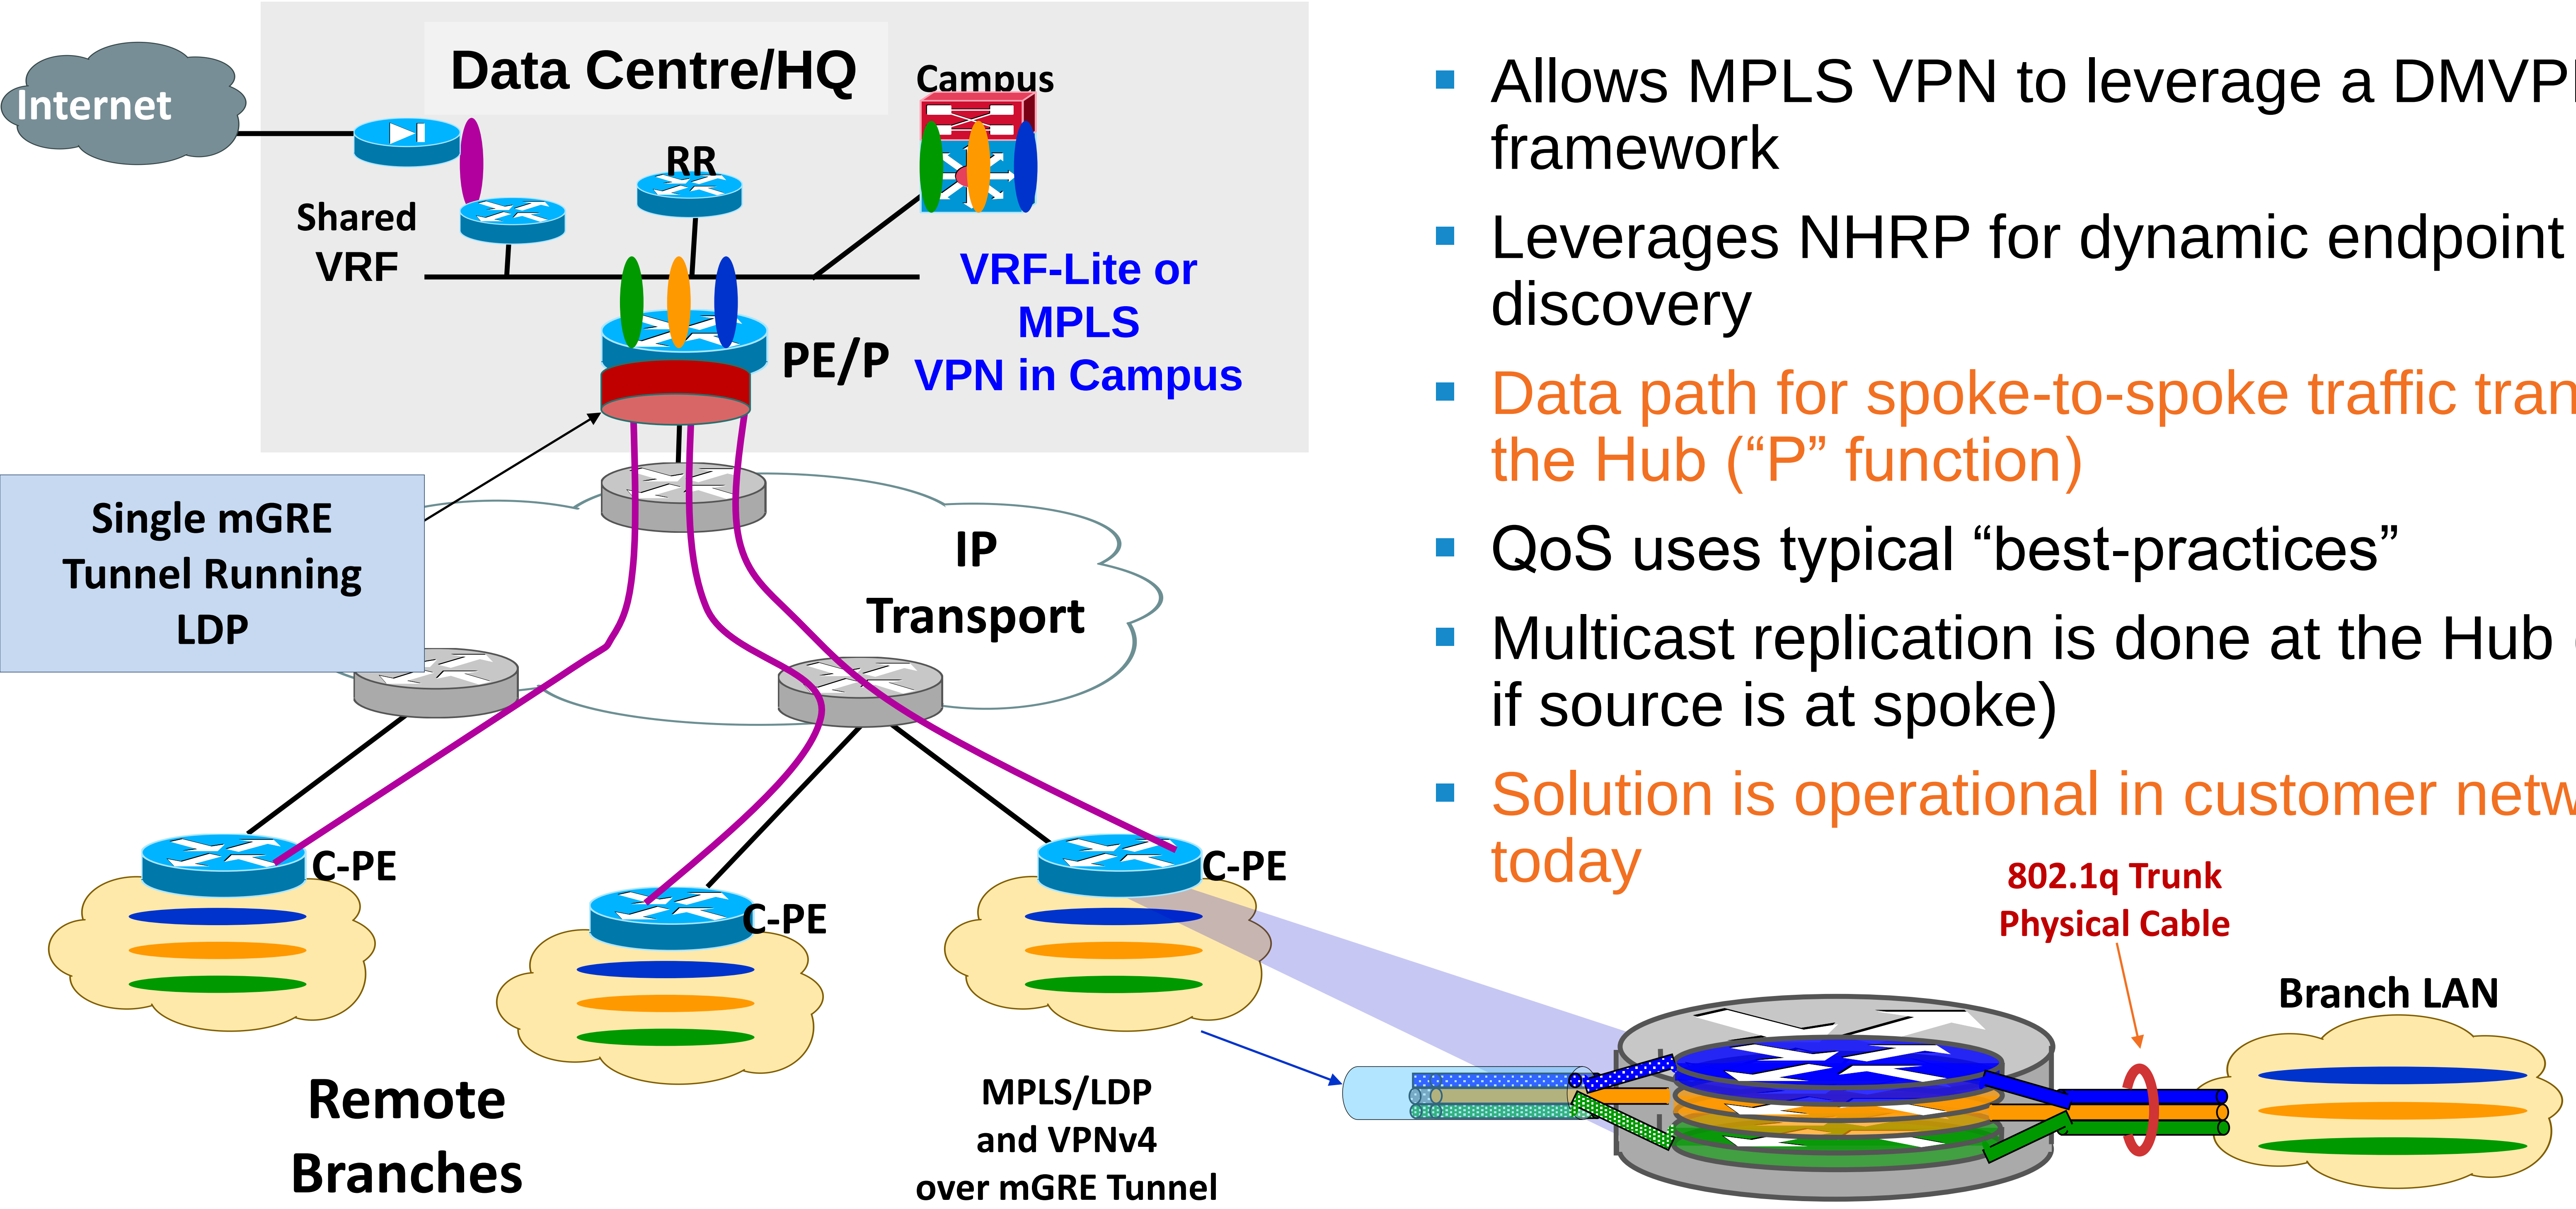

- Allows MPLS VPN to leverage a DMVPN framework
- Leverages NHRP for dynamic endpoint discovery
- Data path for spoke-to-spoke traffic transits the Hub (“P” function)
- QoS uses typical “best-practices”
- Multicast replication is done at the Hub (even if source is at spoke)
- Solution is operational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

MPLS VPN Over Multipoint GRE (mGRE)

MPLS VPNs over Multipoint GRE Using BGP for End Point Discovery

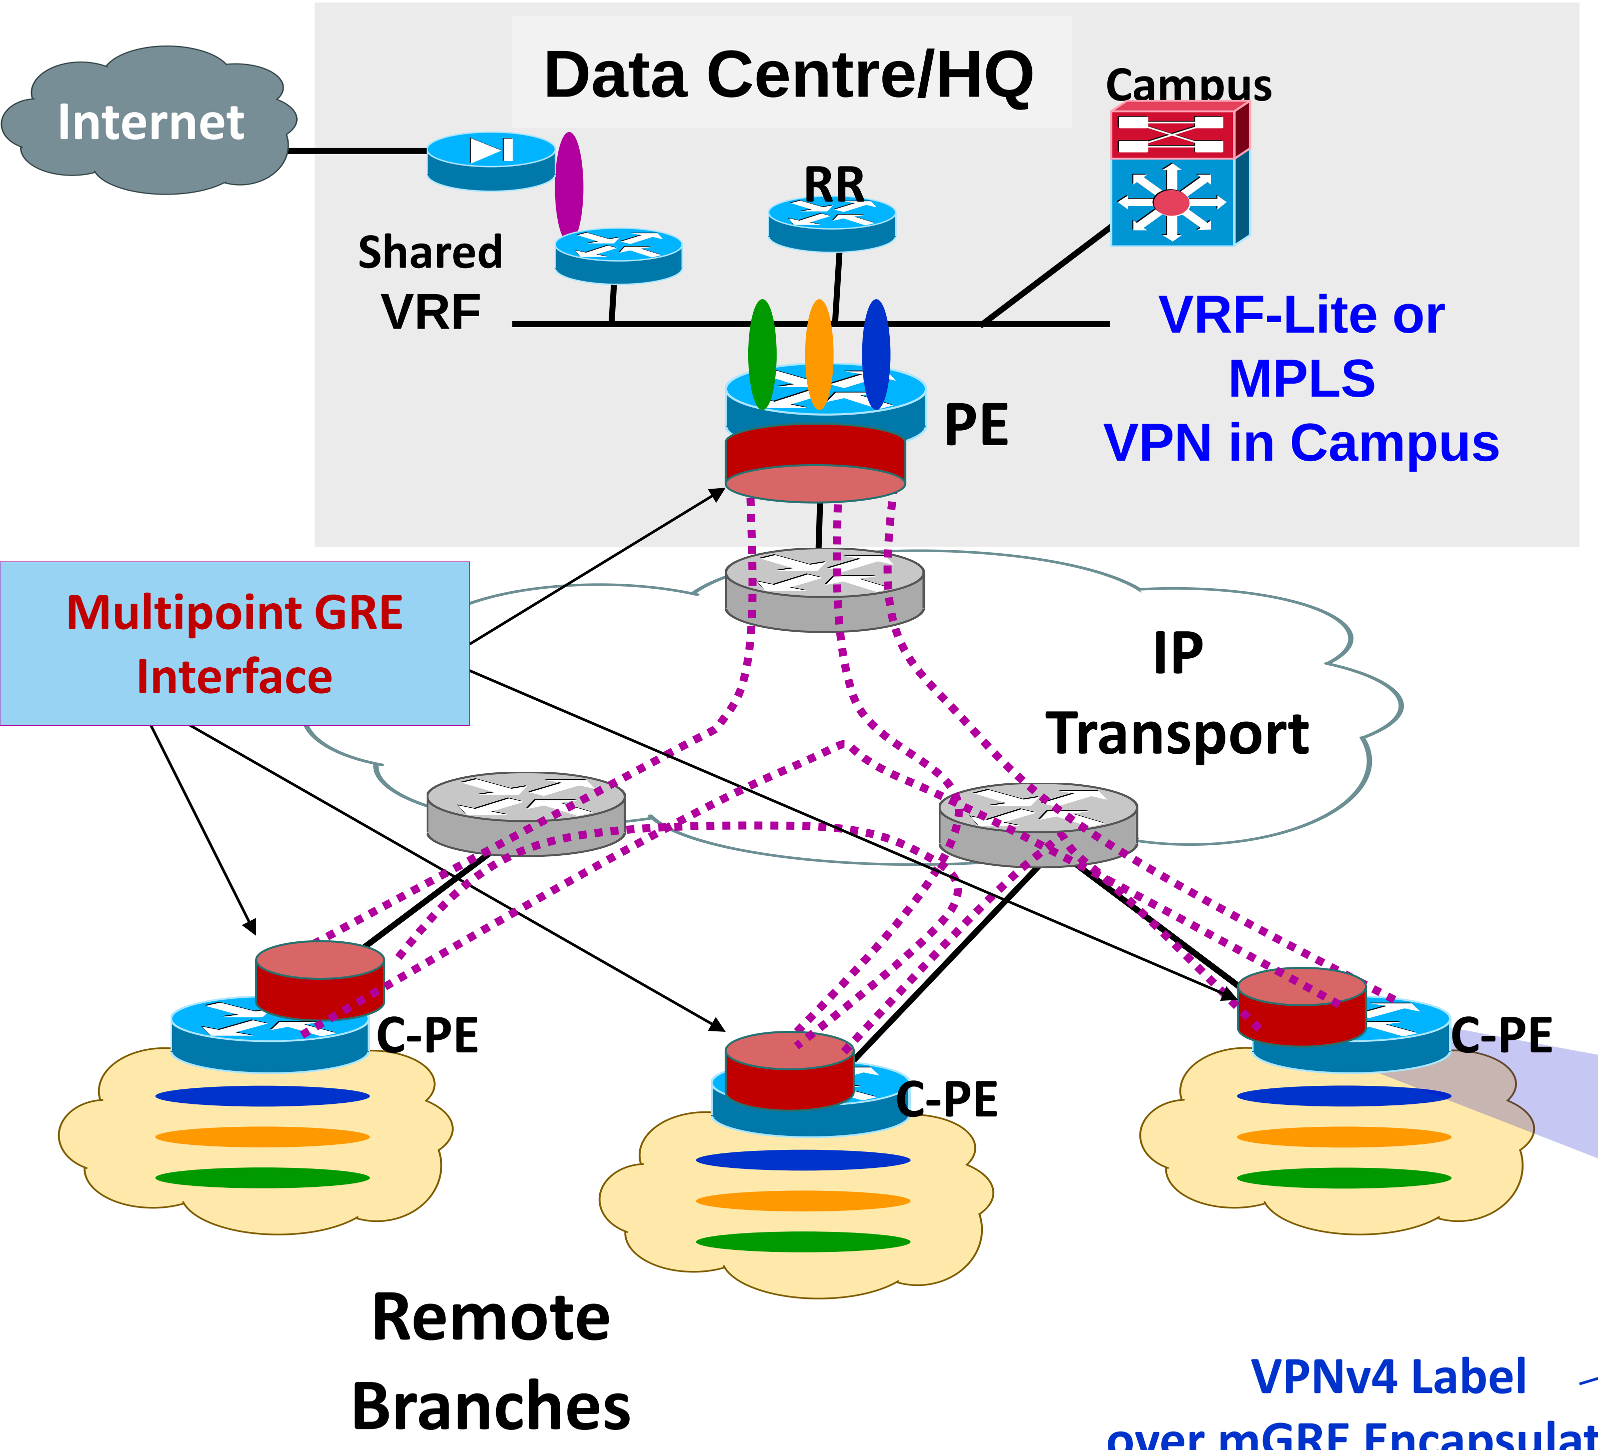

- Offers MPLS-VPN over IP
- Dynamic spoke-to-spoke access
- Uses standards-based RFC 2547 MP-BGP control plane
- Offers dynamic Tunnel Endpoint Discovery via BGP
- Requires only a single IP address for transport over SP network
- Reduces configuration tasks: Requires NO LDP, NO GRE configuration tasks

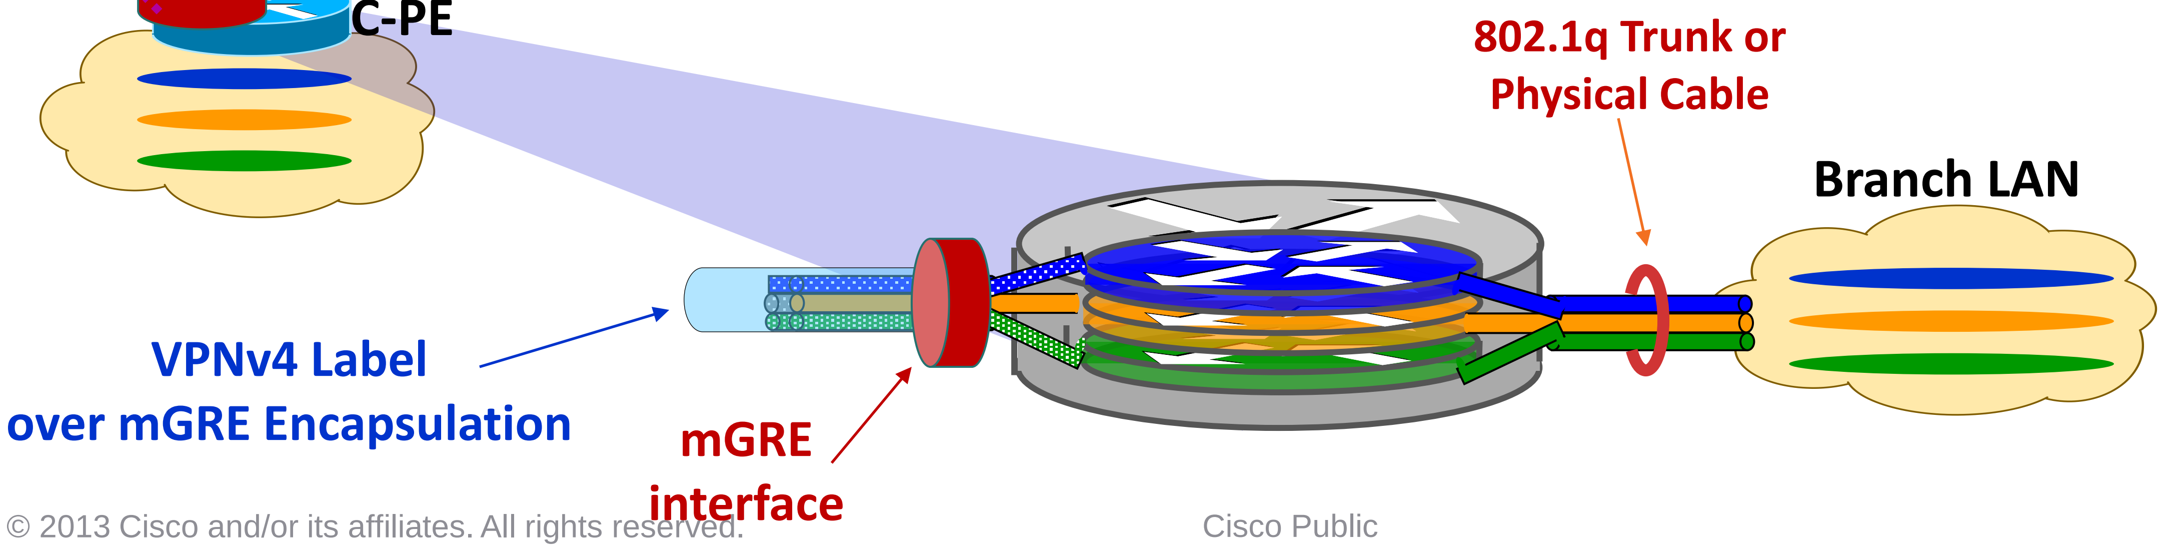

MPLS VPN Over Multipoint GRE (mGRE)

Control/Data Plane Example over Service Provider Model

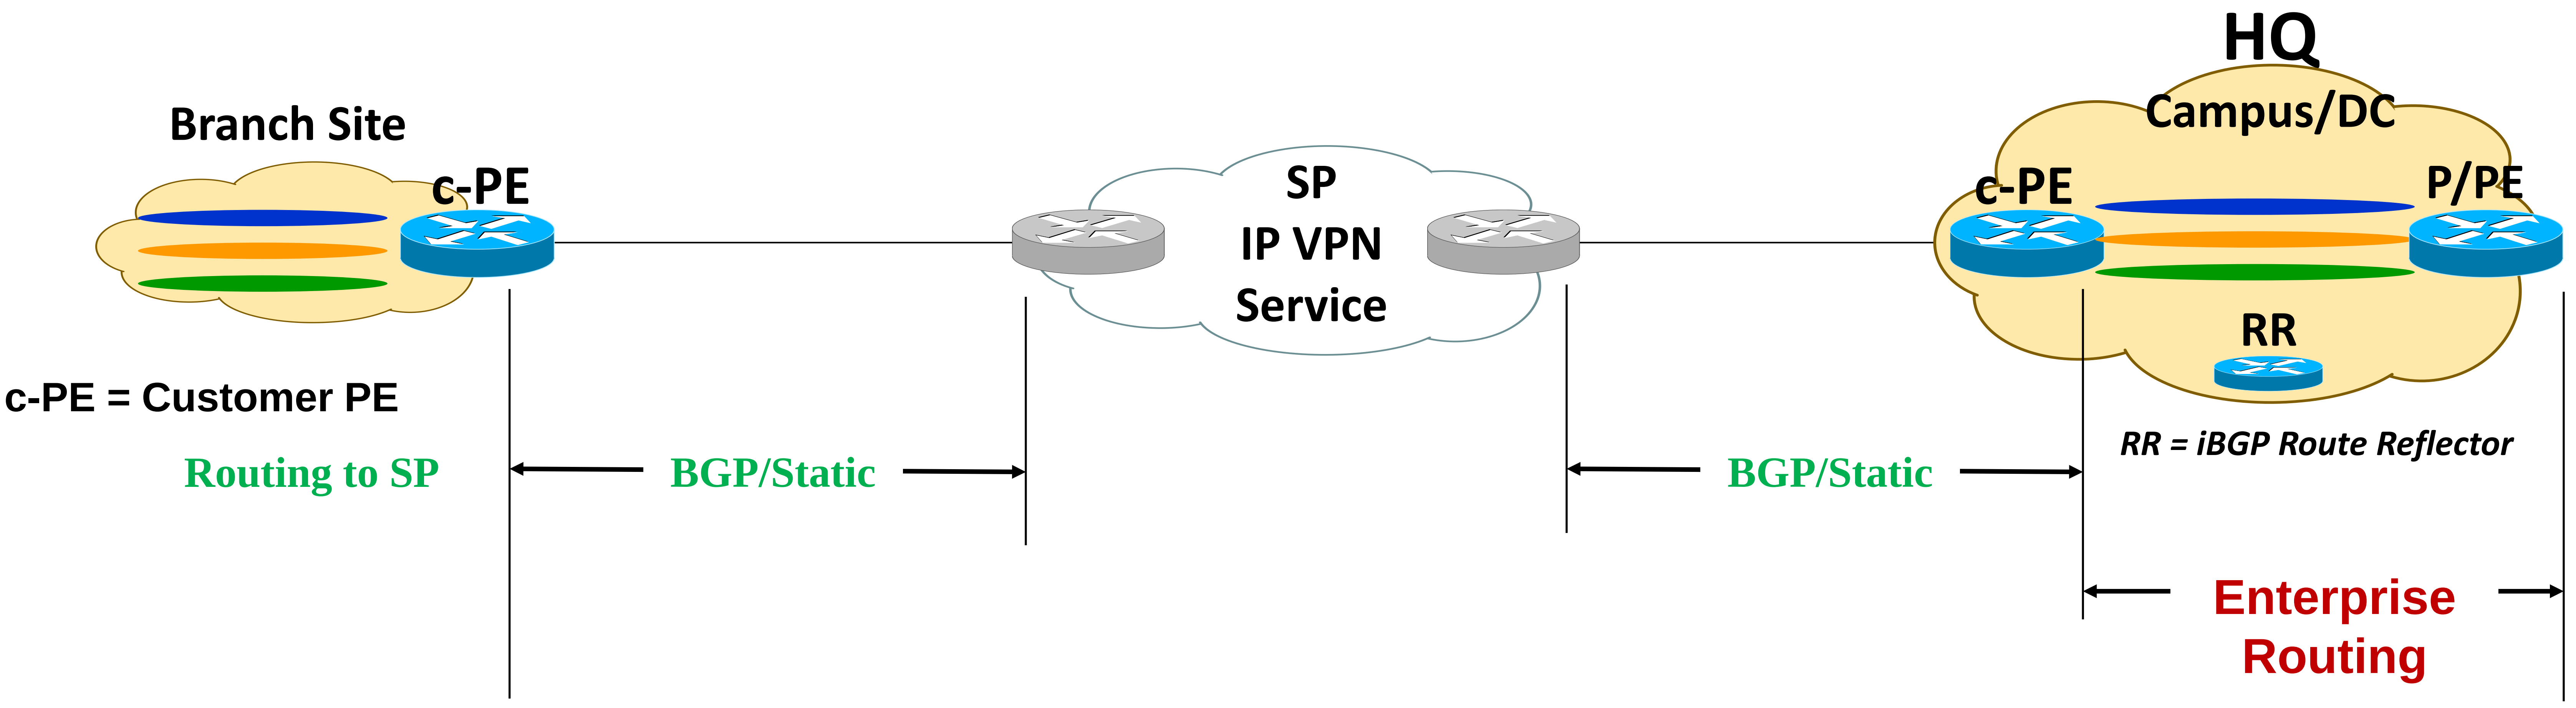

MPLS VPN Over Multipoint GRE (mGRE)

Control/Data Plane Example over Service Provider Model

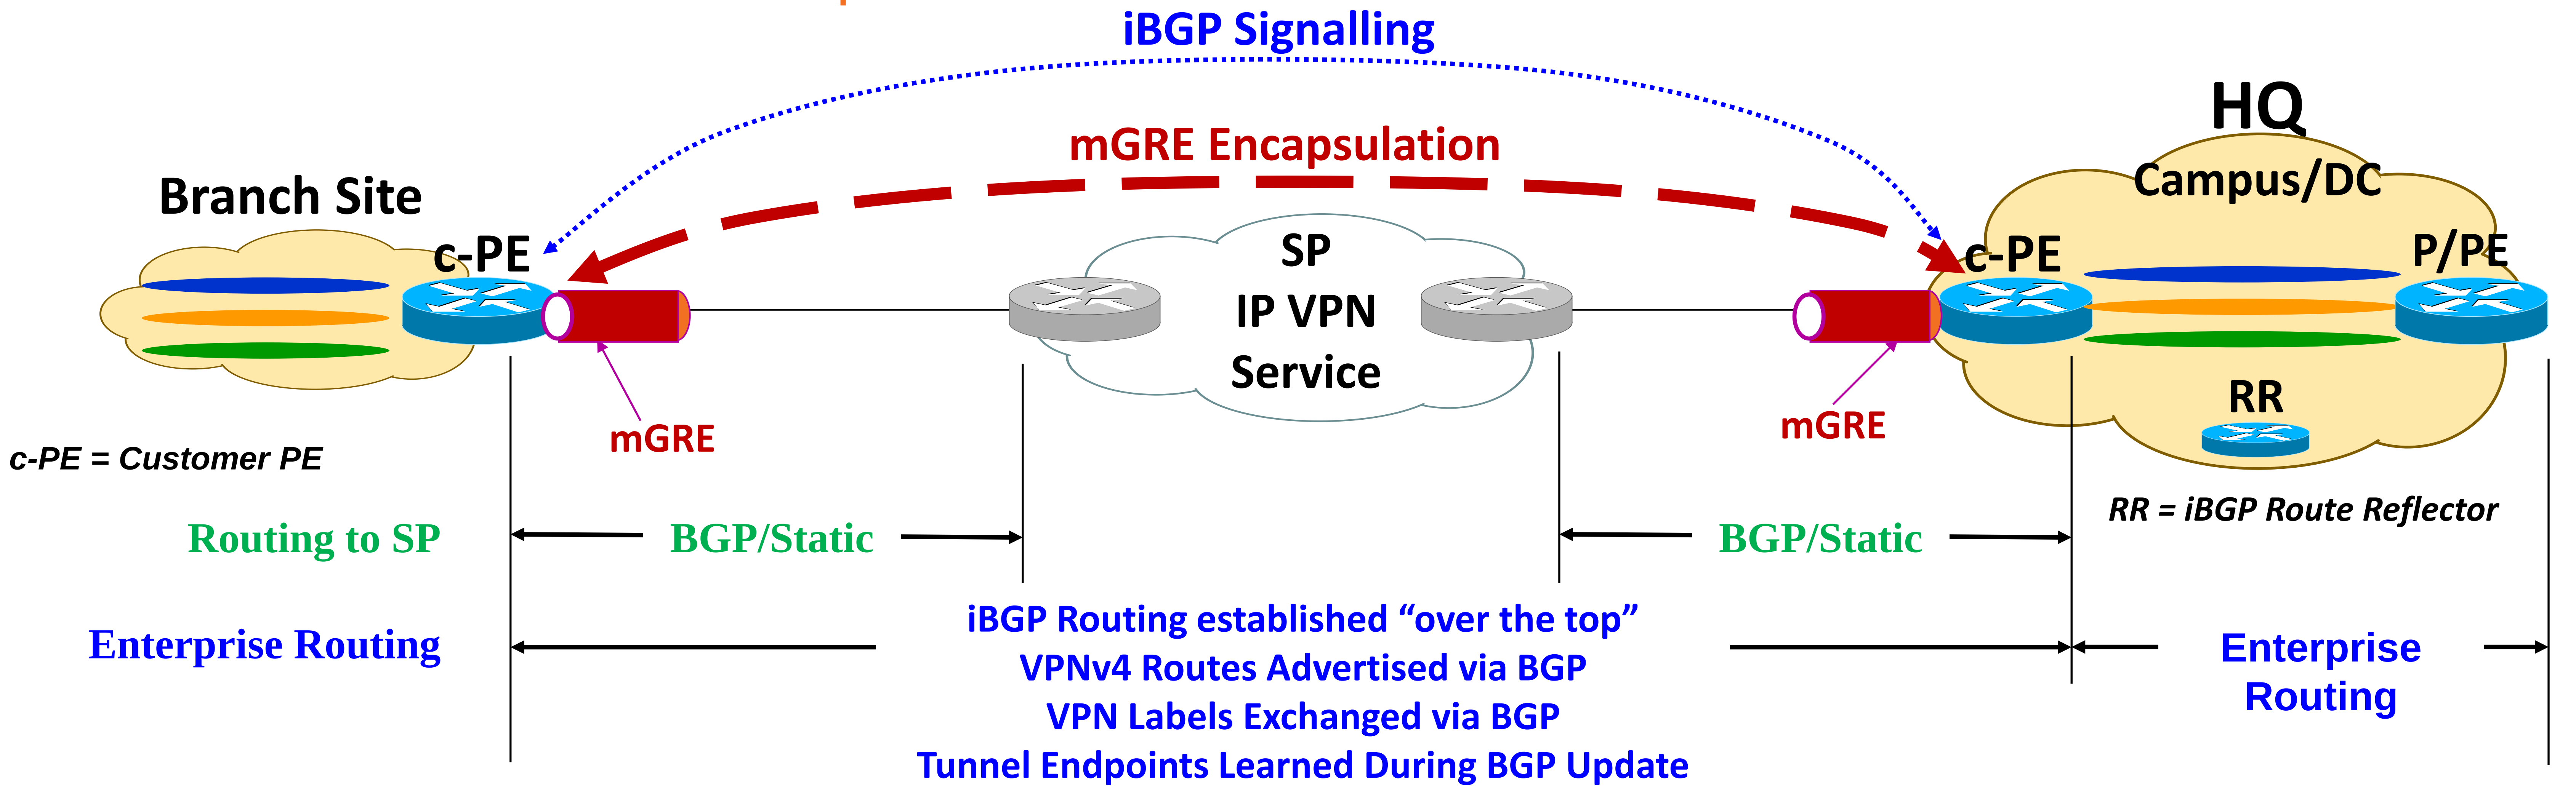

- Routing and data forwarding done "Over the Top" of the SP transport
- iBGP used to: (1) Advertise VPNv4 routes, (2) exchange VPN labels, (3) learn tunnel end-points
- eBGP used to: (1) exchange tunnel end point routes with SP (optional static routes could be used)
- Only requires advertising ONE IP prefix to the SP network (e.g. IP tunnel "end points")

MPLS VPN Over Multipoint GRE (mGRE)

Control Plane

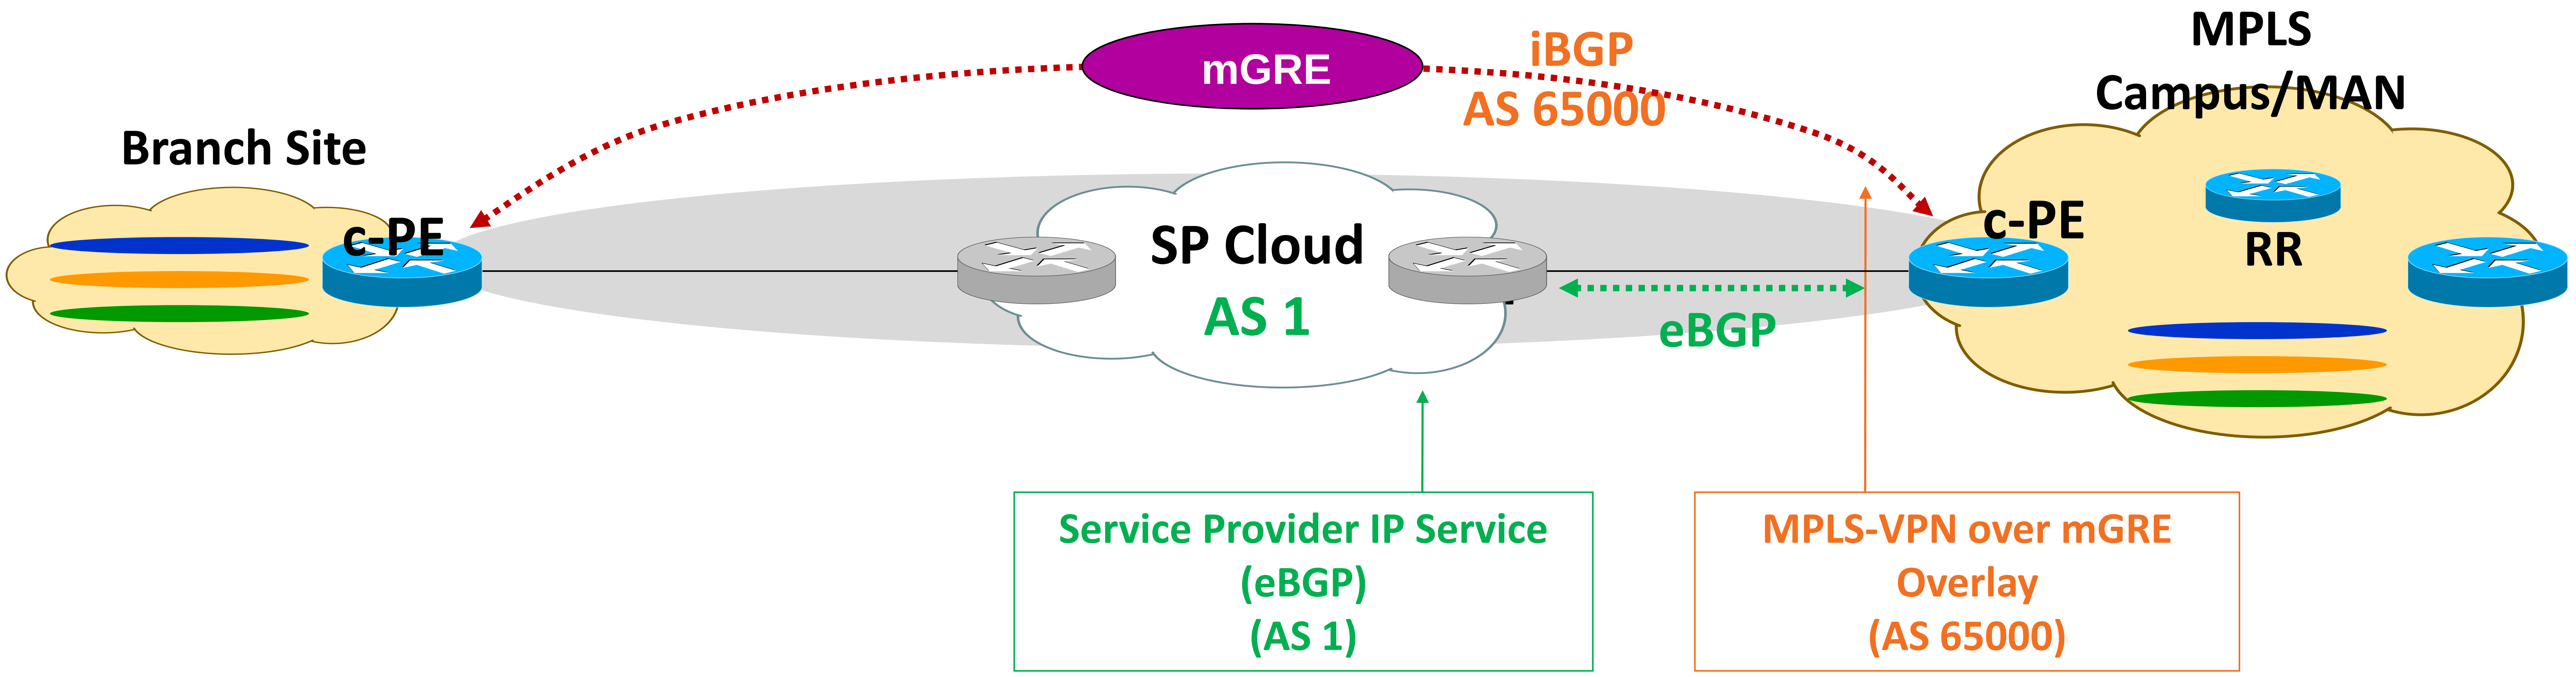

- eBGP (AS 1): used to peer to the SP PE router
- i-BGP (AS 65000): used for MP-BGP and VPNv4 prefix and label exchange
- C-PE for e-BGP appears as CE to the SP
- C-PE for i-BGP functions as a PE in supporting MPLS-VPN over mGRE

MPLS VPN Over Multipoint GRE (mGRE)

Feature Components

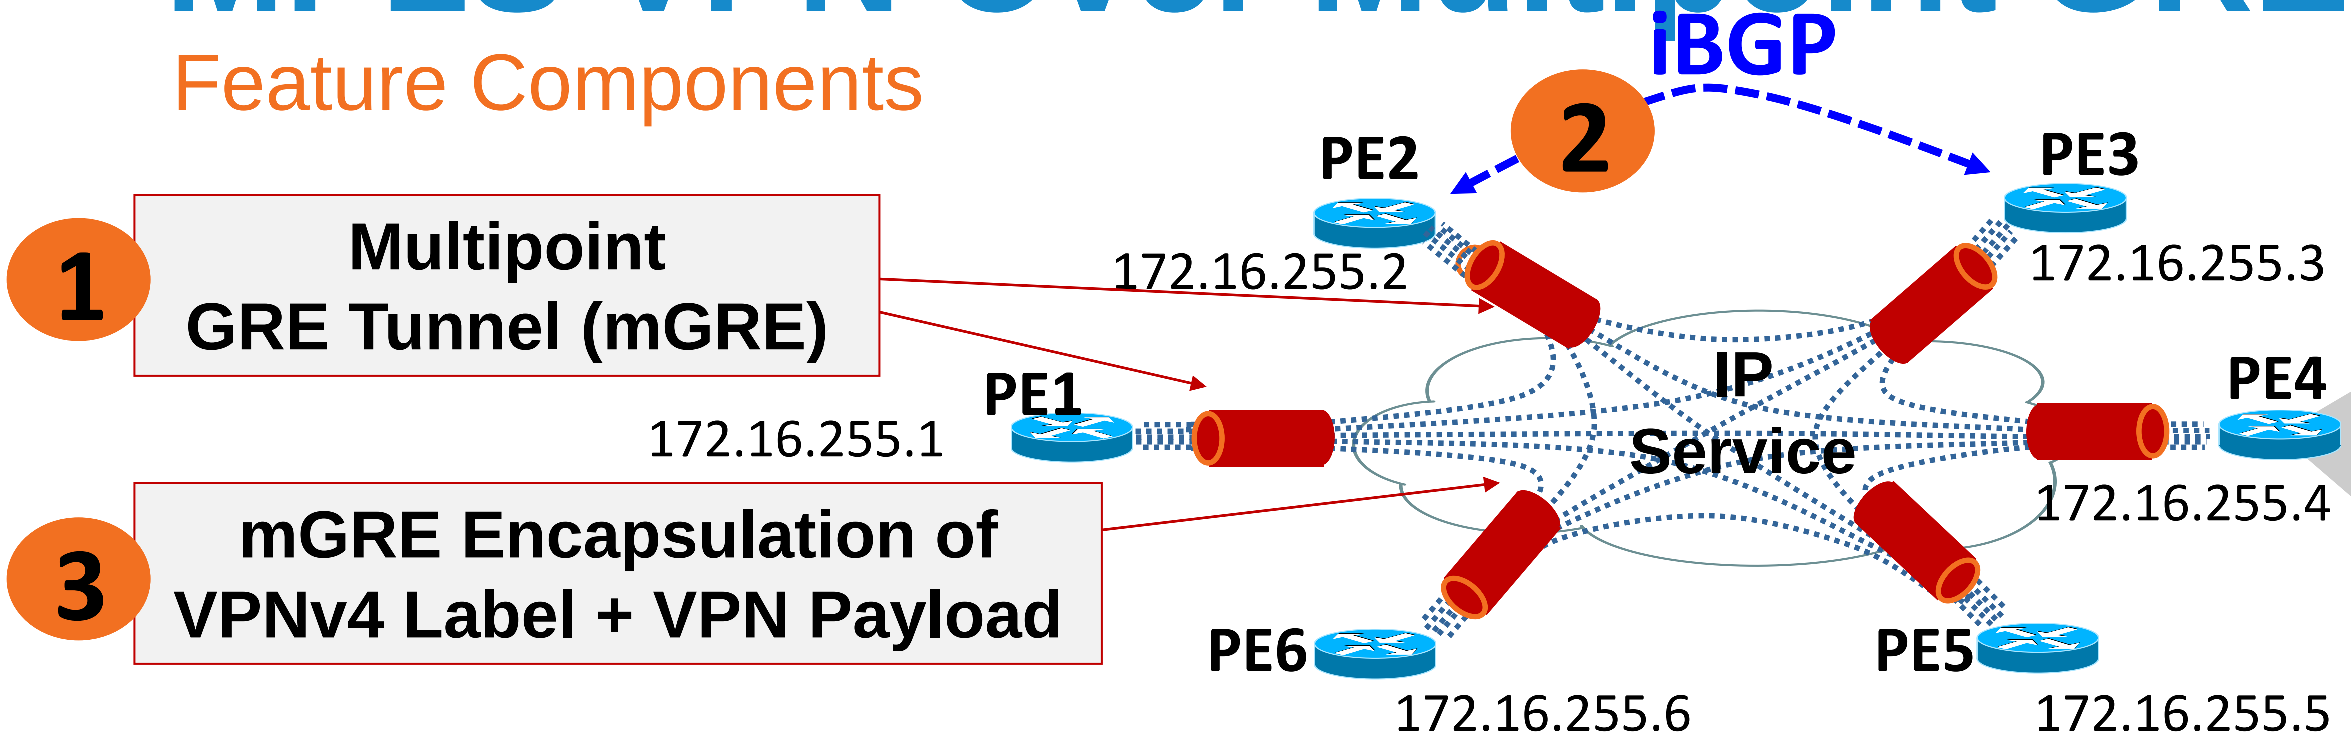

4 **View for PE 4**

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 2547 using MP-BGP

   Signalling VPNv4 routes, VPN labels, and tunnel endpoints
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

Multipoint GRE Interface

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

MPLS VPN Over Multipoint GRE (mGRE)

VPNv4 Configuration Example

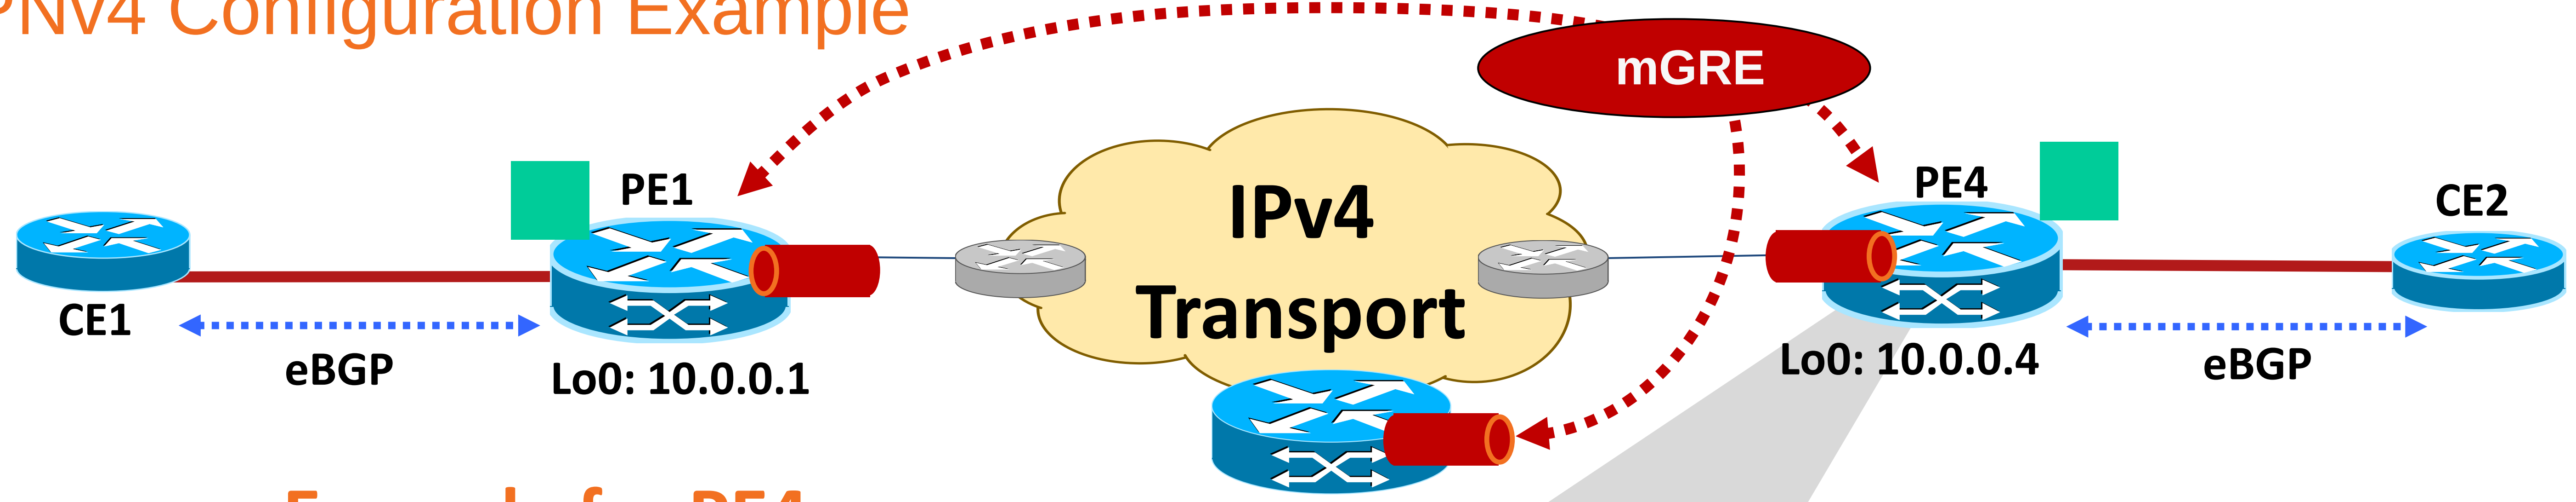

Example for PE4

```
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 Tunnel Interface

MPLS VPN Over Multipoint GRE (mGRE)

IPv6 Configuration Example

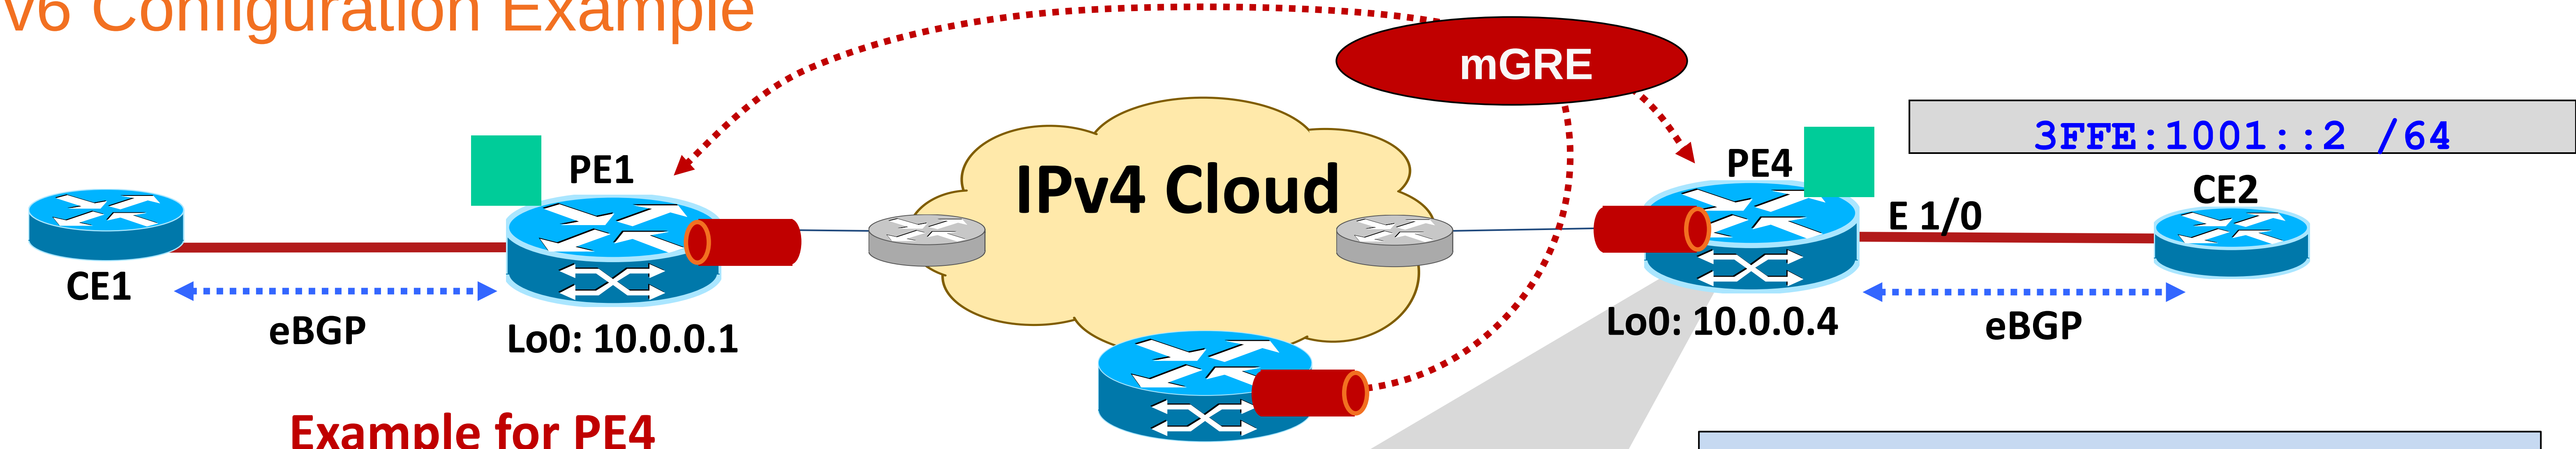

Example for PE4

```
interface Ethernet 1/0
 vrf forwarding green
 ip address 209.165.200.253 255.255.255.224
 ipv6 address 3FFE:1001::/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
```

<u>NOTE:</u> 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

- Solution 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: ASR 1000, ISR/G2, SUP-2T
- 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)

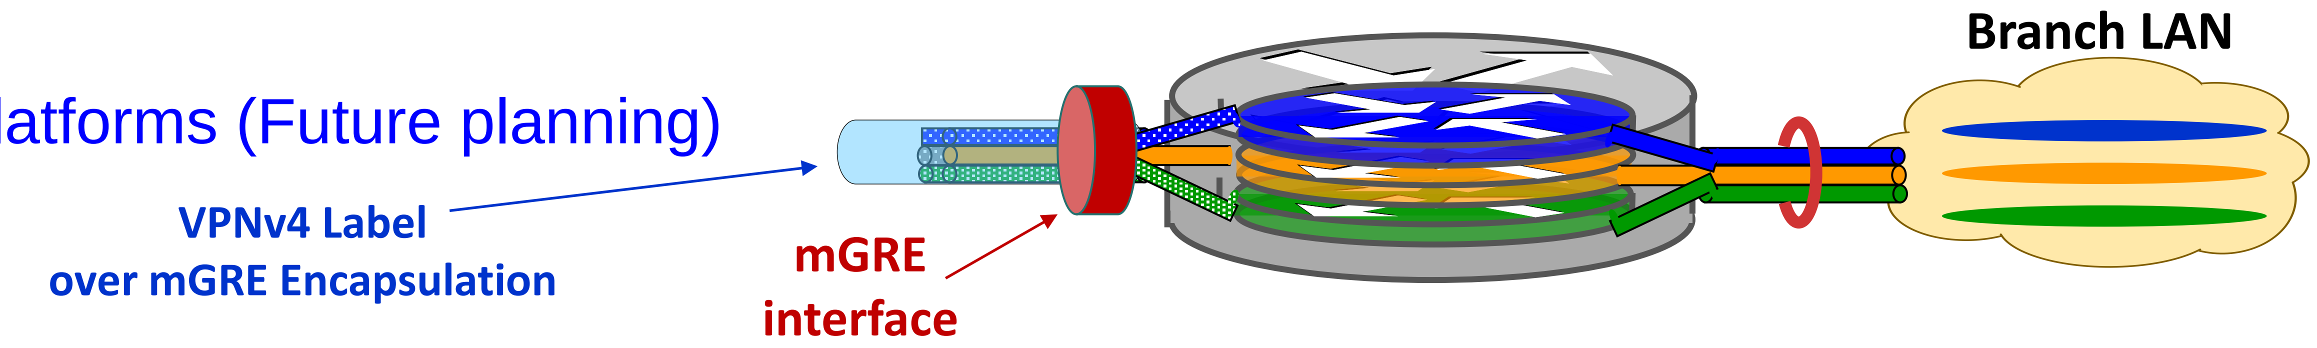

MPLS VPN Deployment Considerations for WAN Designs (over IP)

EXAMPLE: MPLS VPN over GRE (Point to Point Tunnels)

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?
- Is L2 VPN (PW, VPLS) required?
- How many locations will the network grow too?
- Do I require any-to-any traffic patterns?
- What is the transport? (i.e. is GRE required?)

Example: 50 – 1000 Sites (full-mesh)

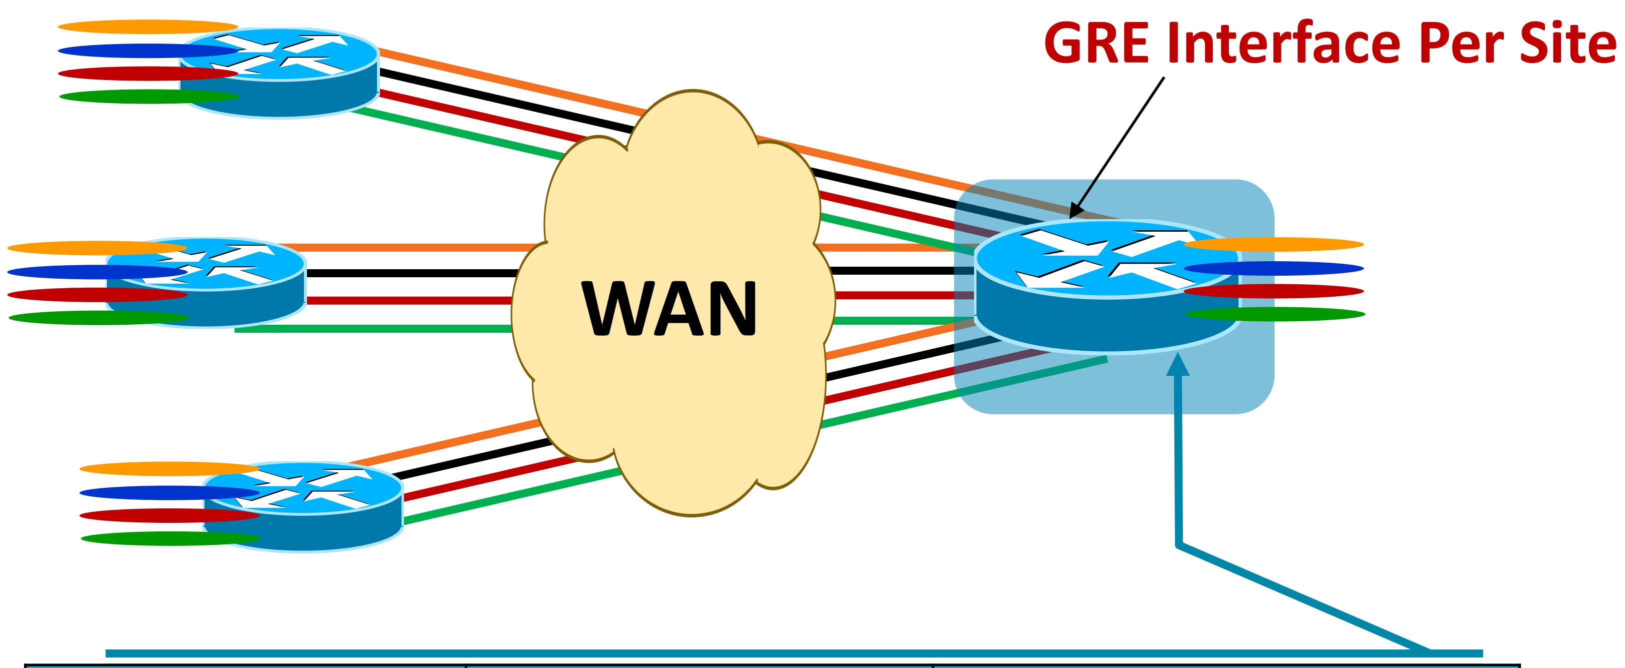

VRFs	Neighbours	GRE Tunnel Interface (N – 1)
50	50	49
100	100	99
250	200	199
500+	1000	999

MPLS VPN Deployment Considerations for WAN Designs (over IP)

EXAMPLE: MPLS VPN over DMVPN (mGRE), MPLS VPN over mGRE (BGP)

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?

Example: 50 – 1000 Sites

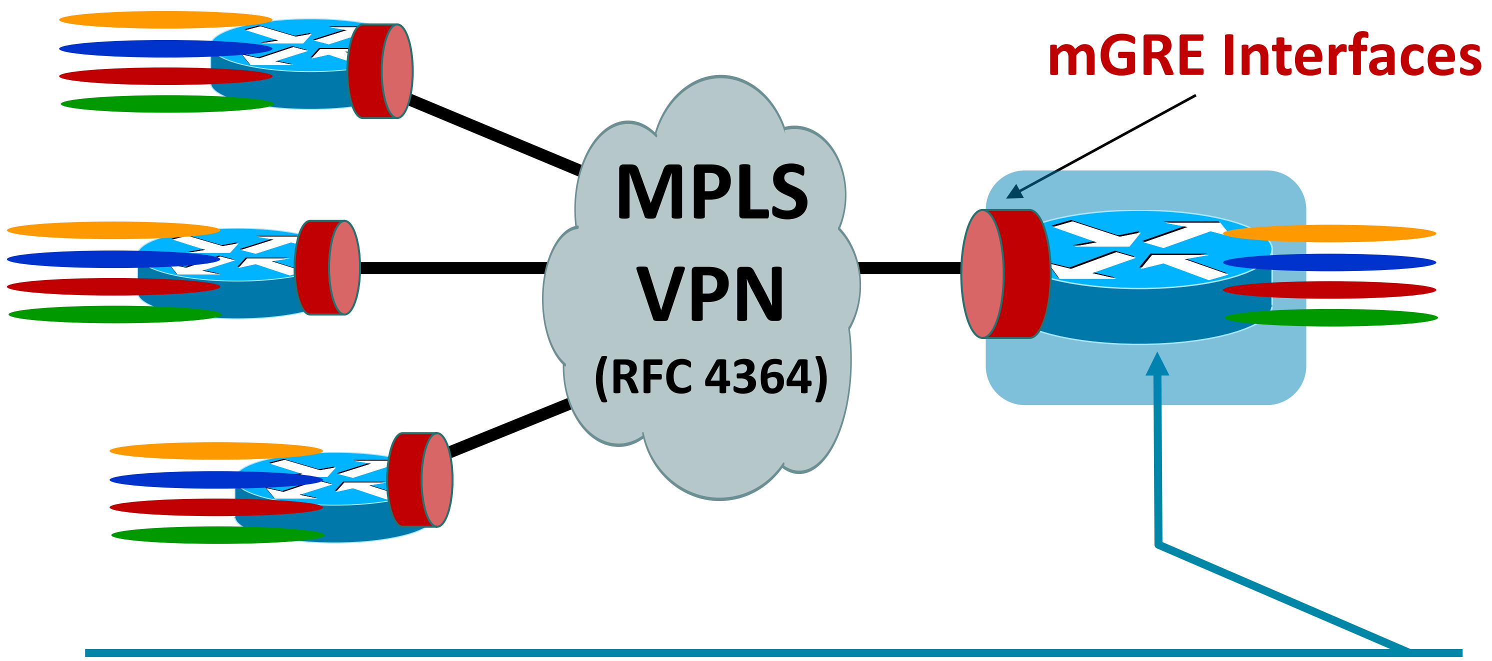

VRFs	Neighbours	GRE Tunnel Interface
50	50	1
100	100	1
250	200	1
500+	1000	1

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) OSPF – 600 (7200) BGP – 1800 (ASR 1K)	1000+ (Manually Intensive)
Supports IPSec Encryption	Yes (GET, SA)	Yes	Yes
Supports MVPN Multicast *	Yes	* 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**

Cisco L3 Virtualisation

Platforms and Feature Support for WAN and Branch

Feature \ Platform	Cisco ISR-G2	Cisco 7200	ASR 1000	Catalyst 6500	Cisco 7600
VRF Lite	S	S	S	S	S
VRF Lite over GRE	S	S	S	S	S
VRF Lite over DMVPN	S	S	S	S	S
MPLS-VPN	S	S	S	S	S
MPLS VPN over GRE (P2P)	S	S	S	S (SIP-400), SUP-2T	S (SIP-400, ES+)
MPLS VPN over DMVPN (mGRE)	S	S	S	S (SIP-400), SUP-2T	S (SIP-400, ES+)
MPLS VPN over mGRE (BGP)	S	S	S	S (integrated into EARL8)	S (SIP-400, ES+)

S = Supported Today   R = Roadmap

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

QoS with GRE, MPLS over GRE

ToS Reflection Behaviour for "transit traffic" Through the Router

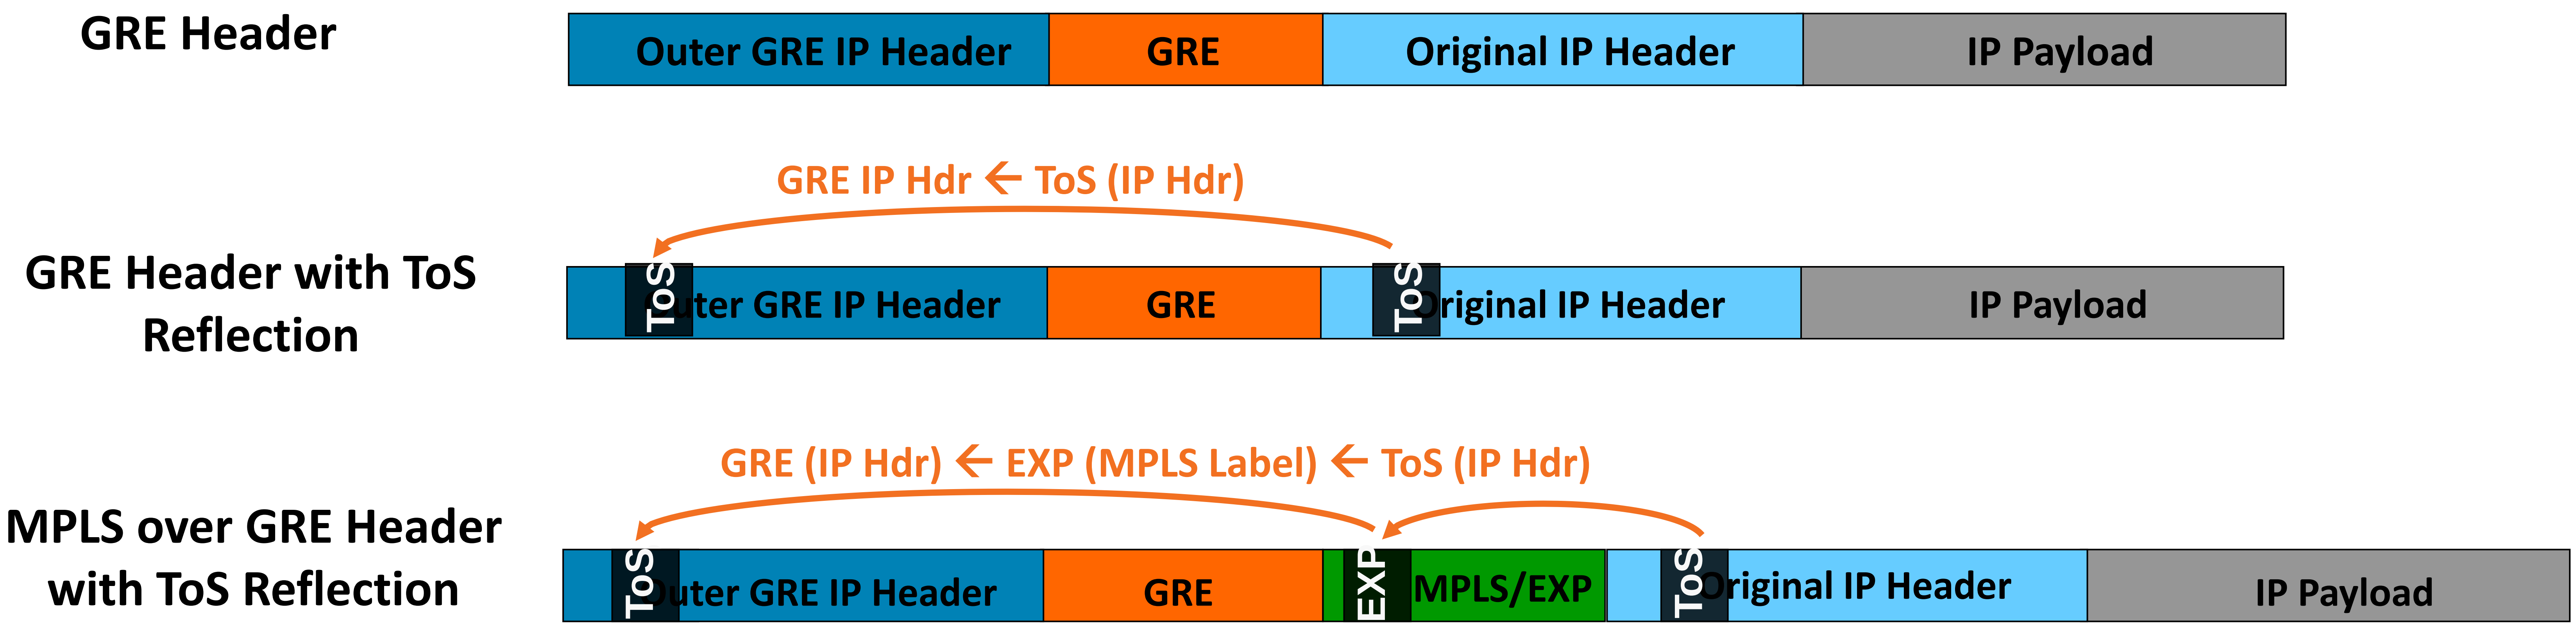

- Router will copy original ToS marking to outer GRE header
- For MPLS over GRE, the EXP marking is copied to the outer header of the GRE tunnel
- This allows the IPv4 "transport" to perform QoS on the multi-encapsulated packet

- **Caveats:**
- **Traffic originating on the router (SNMP, pak_priority for routing, etc…), could have different behaviour**
- **See following link for impact**

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. Prioritised 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

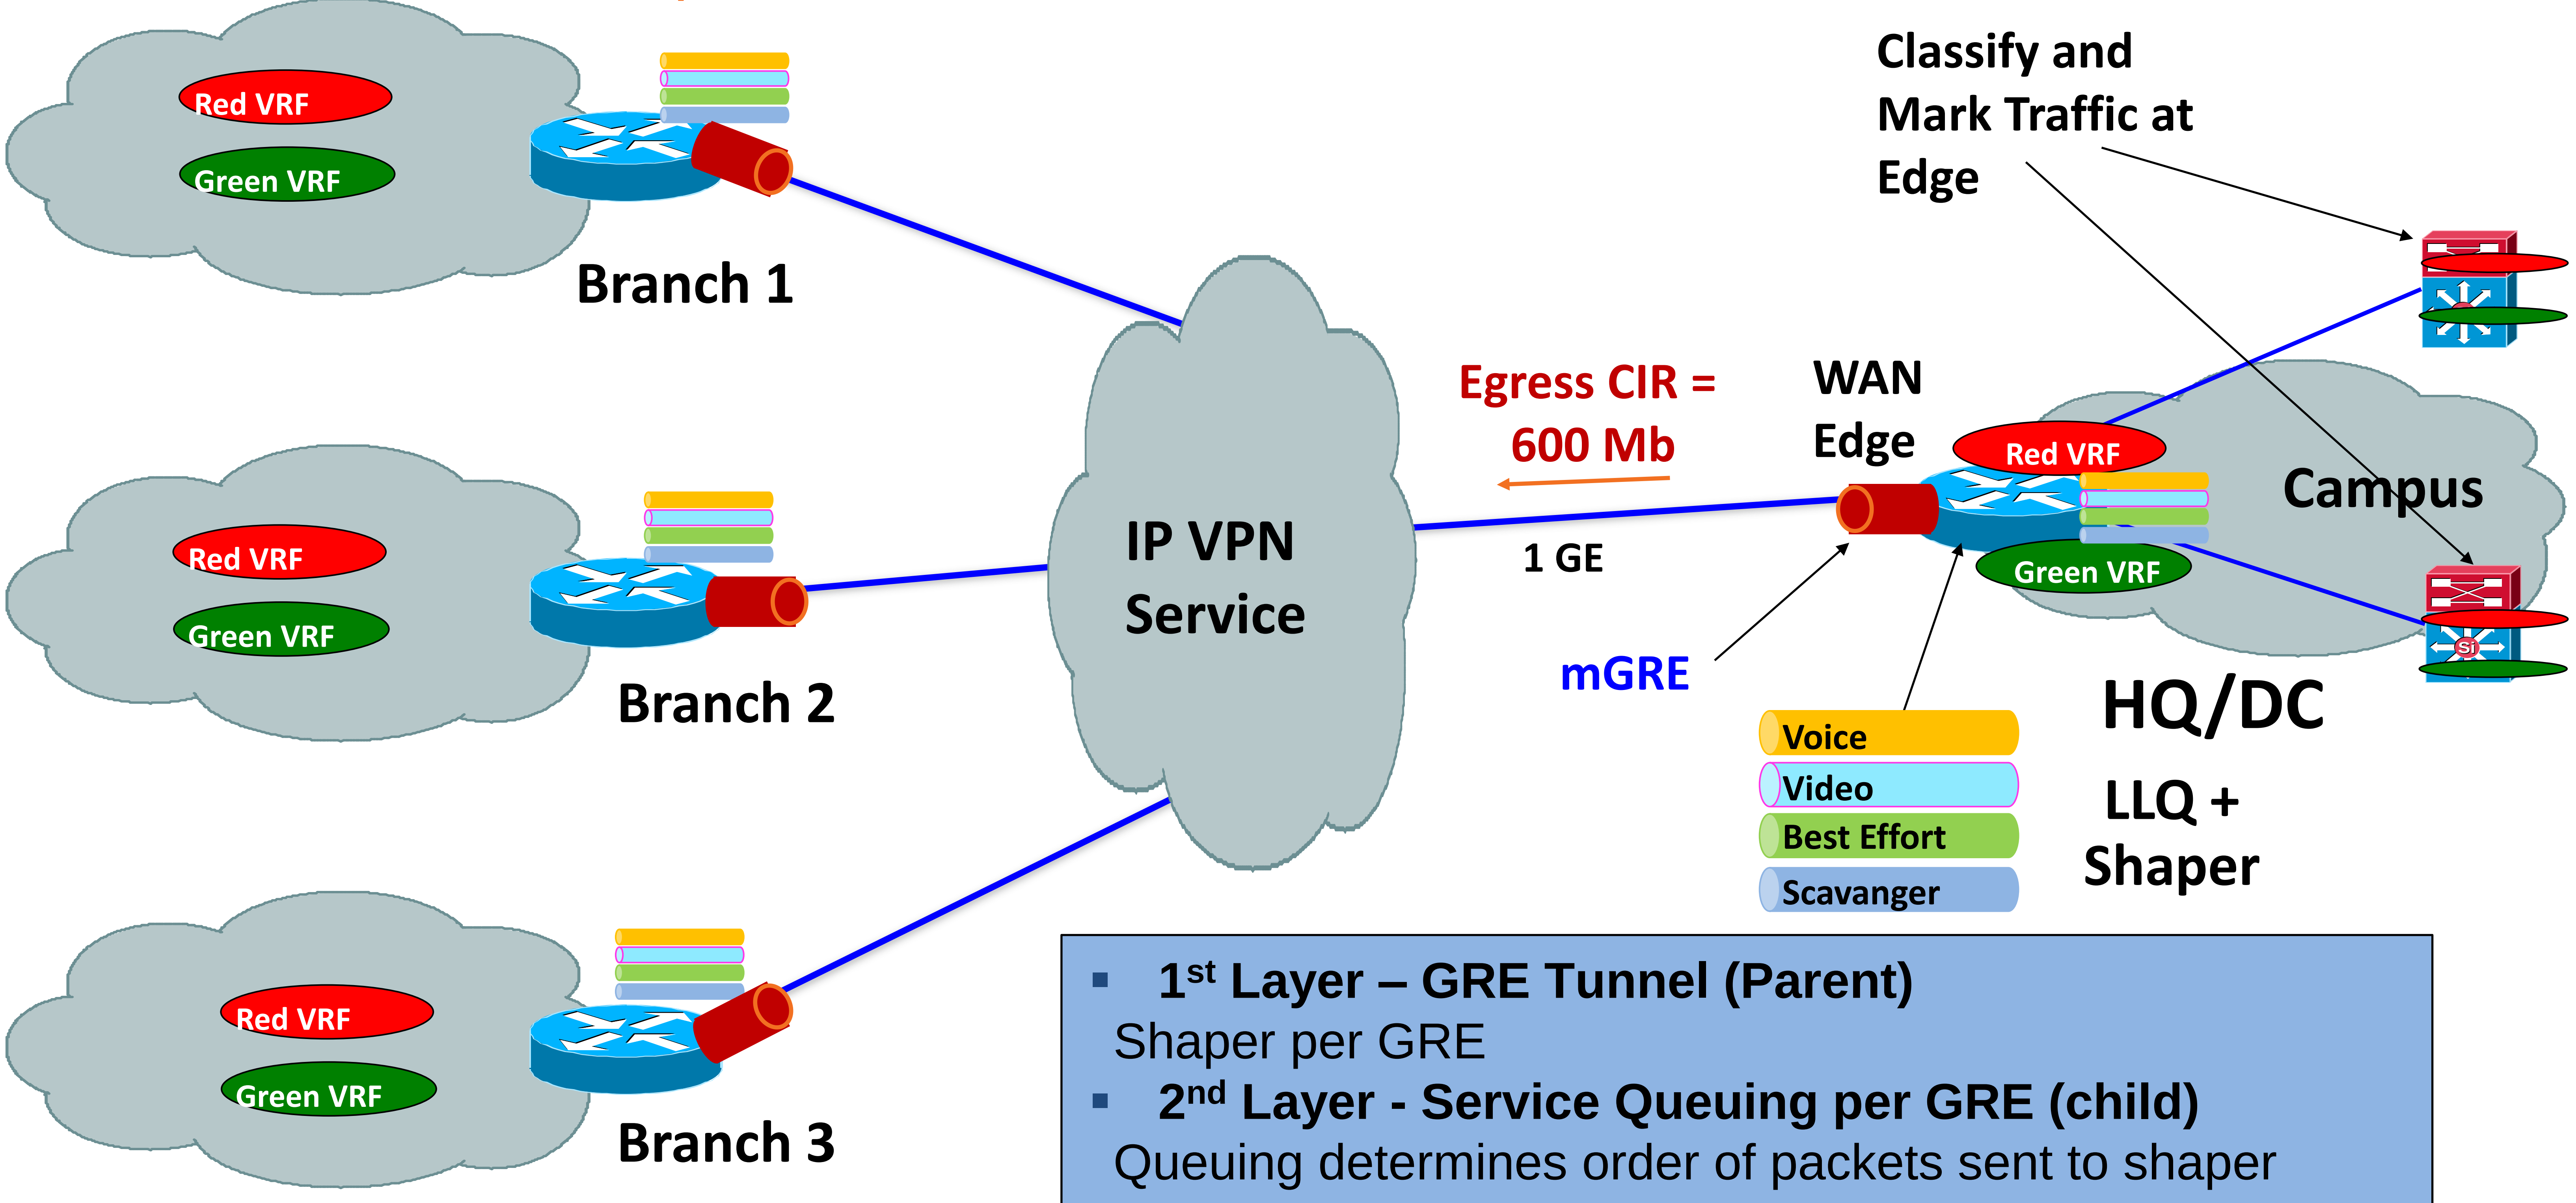

- **1st Layer – GRE Tunnel (Parent)**
  Shaper per GRE
- **2nd Layer - Service Queuing per GRE (child)**
  Queuing determines order of packets sent to shaper
- **H-QoS policy applies to main interface (not mGRE)**

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
      priority level 1
    class Video
      police cir percent 20
      priority level 2
    class Scav
       bandwidth remaining ratio 1
    class class-default
     bandwidth remaining ratio 9

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

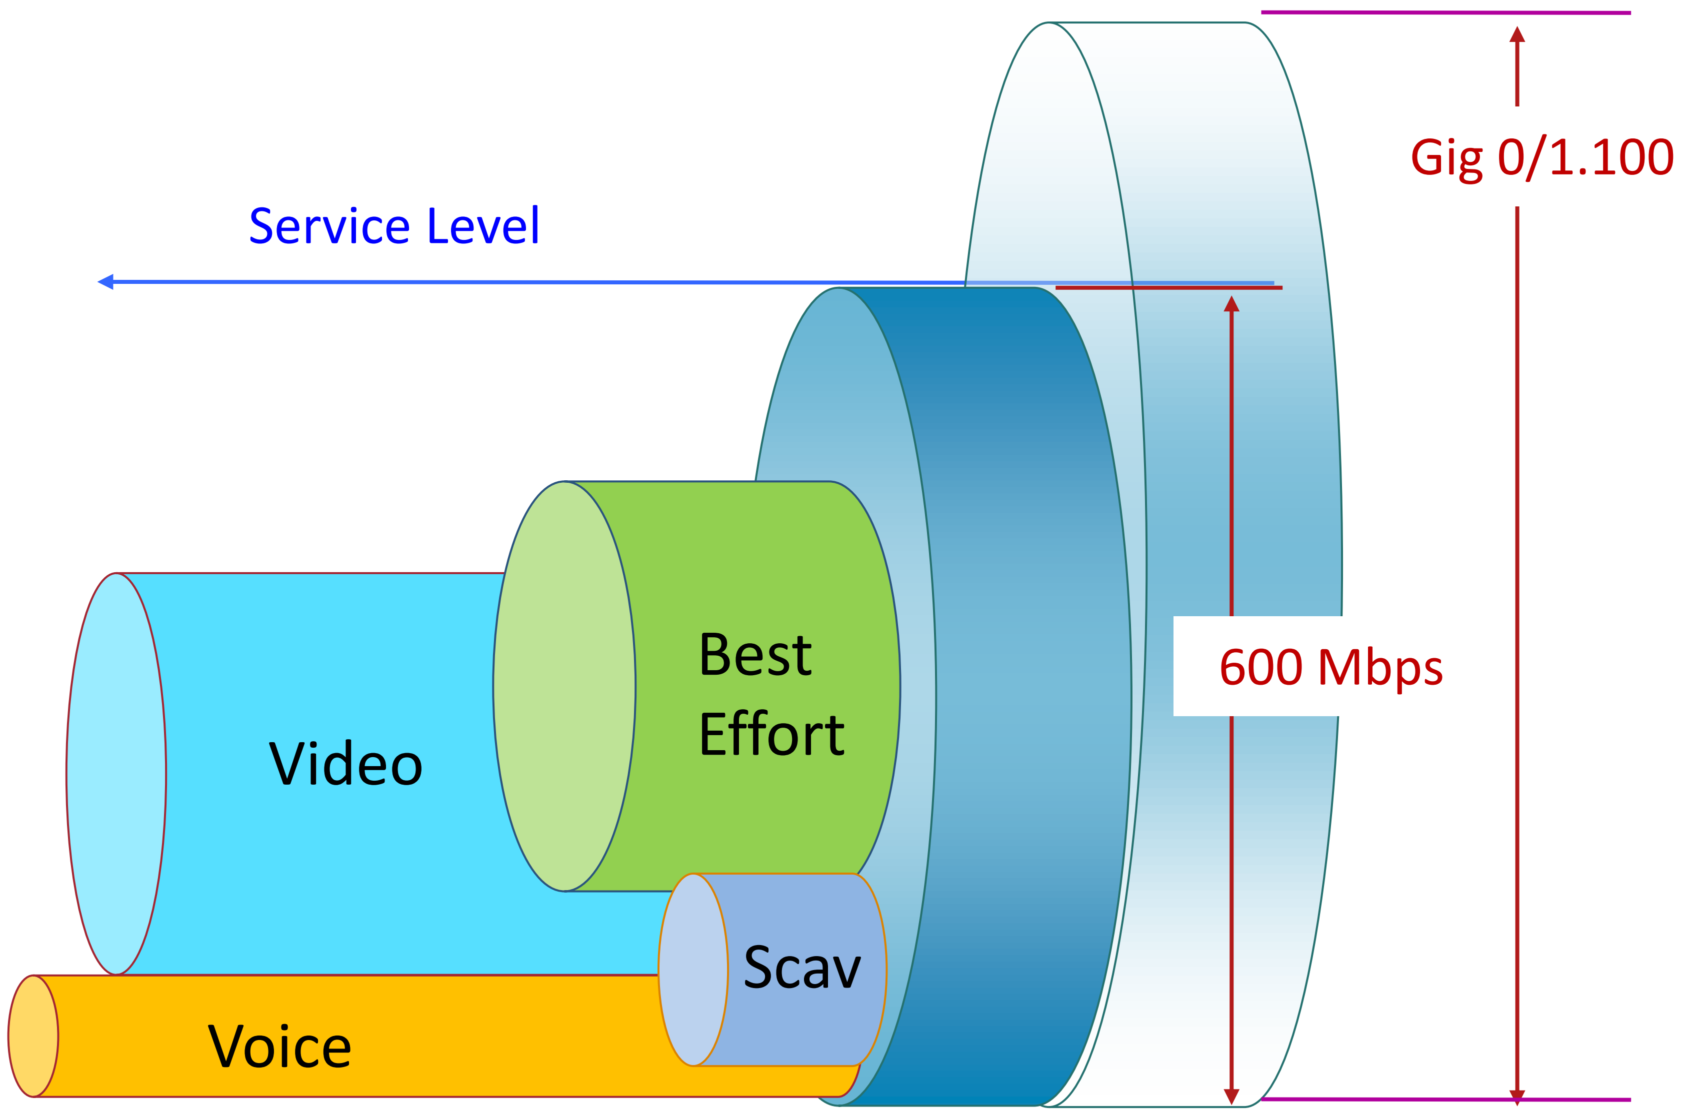

QoS for Virtualisation – Summary

- **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, when deploying QoS in virtualised enterprise network designs**

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
  - Easy Virtual Networking (EVN)
  - Easy Virtual Networking (EVN) + WAN Virtualisation
  - Locator ID Separation Protocol (LISP) – Network Virtualisation for Multi-tenancy
  - Using GET VPN Encryption for Multipoint GRE (mGRE) Solutions
  - MTU Caveats and Solutions for IP Tunneled Environments
- Summary

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
  - **Easy Virtual Networking (EVN)**
  - Easy Virtual Networking (EVN) + WAN Virtualisation
  - Locator ID Separation Protocol (LISP) – Network Virtualisation for Multi-tenancy
  - Using GET VPN Encryption for Multipoint GRE (mGRE) Solutions
  - MTU Caveats and Solutions for IP Tunneled Environments
- Summary

What Is EVN?

Easy Virtual Network (EVN)

http://www.cisco.com/go/evn

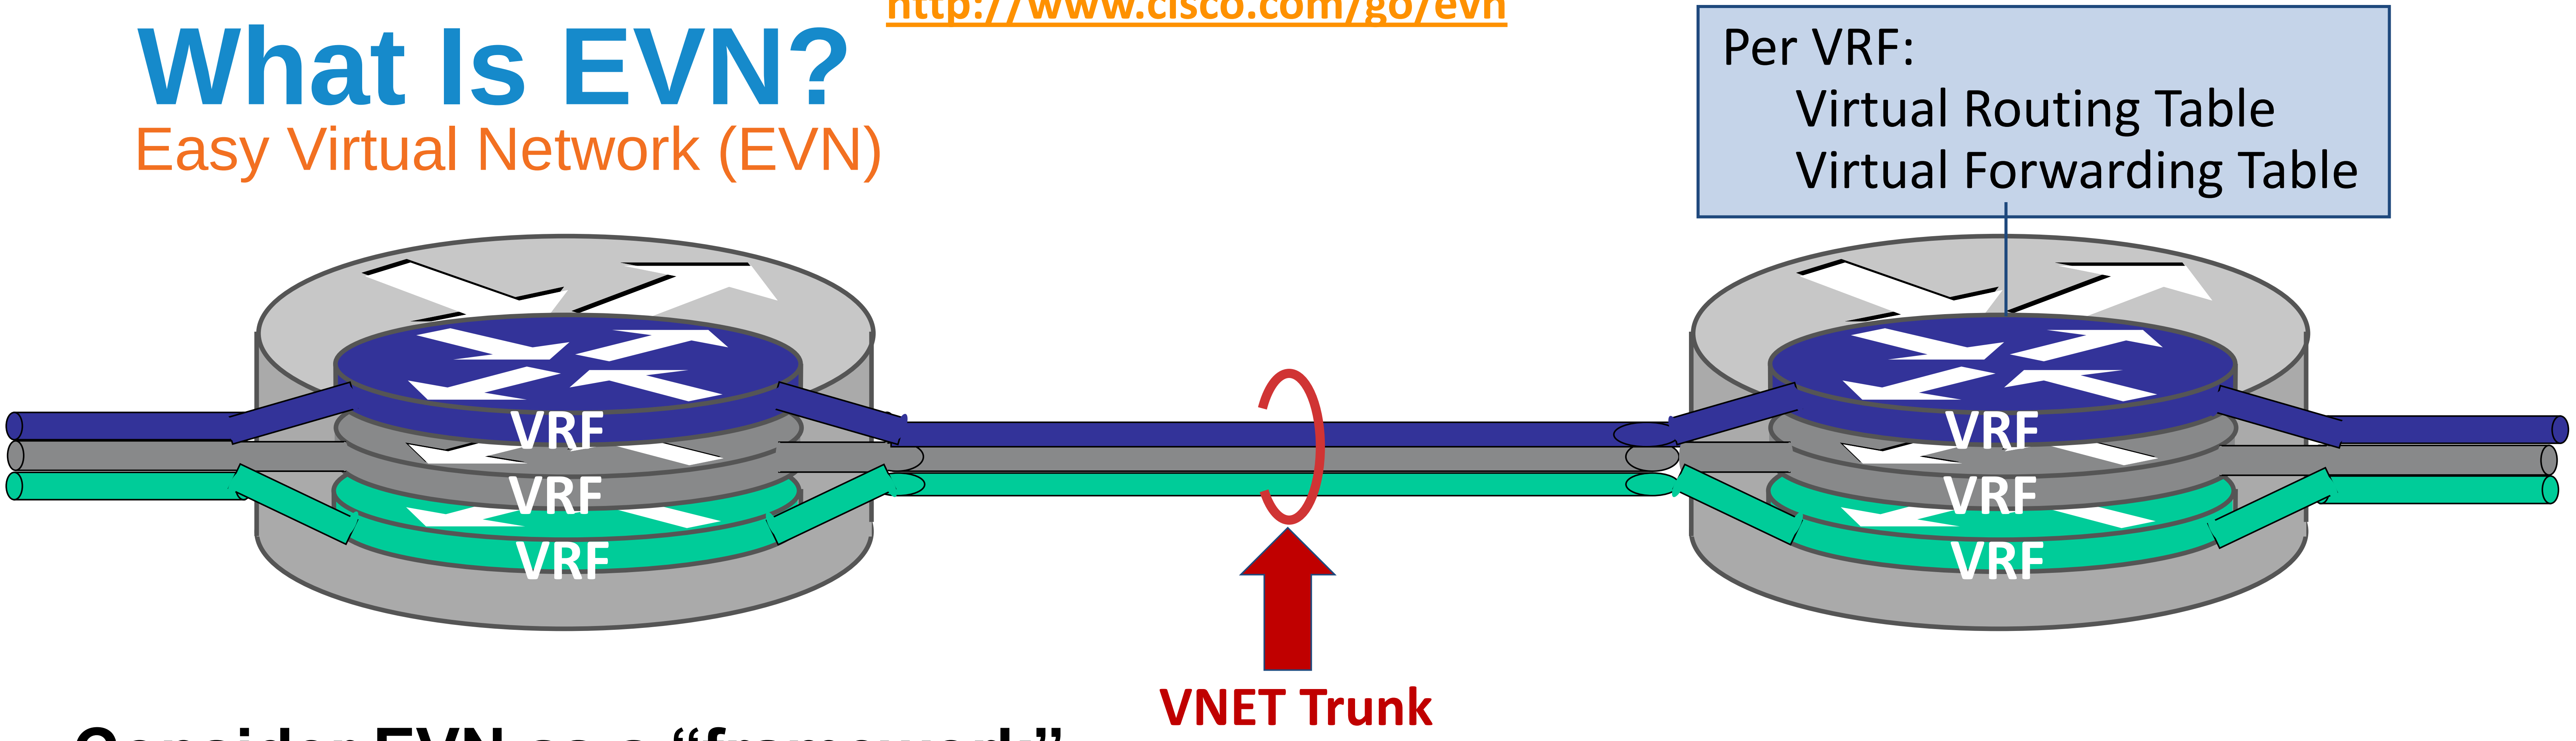

Consider EVN as a "framework"

1. Offers a dynamic way to configure the **"trunk"** between two devices for carrying multiple VRF's
2. Makes the **IOS CLI VRF "context aware"** for configuration ("show", "debug", "traceroute", etc…)
3. Simplifies **route replication** configuration where a "shared" VRF is required (vs. complex BGP import/export commands)

- EVN (like VRF-Lite) still leverages:
  - VRF aware **routing (RIB) and forwarding (FIB)**
  - VRF aware routing protocol processes (EIGRP, OSPF, BGP, RIPv2, static)

VRF-Lite and VNET Trunk Compatibility in EVN

```
ip vrf red
 rd 101:101

ip vrf green
 rd 102:102
```

Switch running VRF-Lite

Switch running EVN

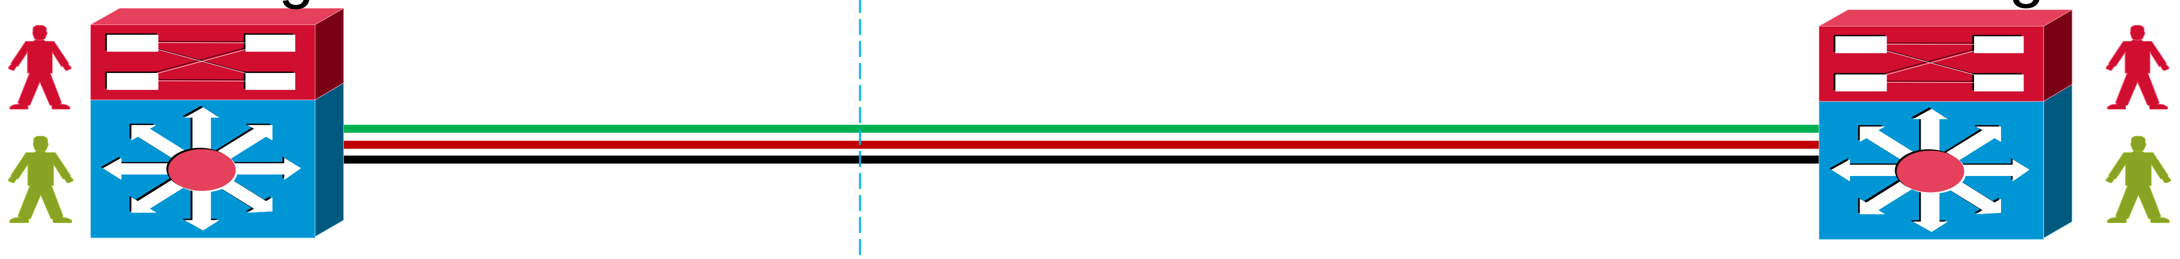

VRF-Lite Sub-interface Config

```
interface TenGigabitEthernet1/1
 ip address 10.122.5.29 255.255.255.252
 ip pim query-interval 333 msec
 ip pim sparse-mode
 logging event link-status

interface TenGigabitEthernet1/1.101
 description Subinterface for Red VRF
 encapsulation dot1Q 101
 ip vrf forwarding Red
 ip address 10.122.5.29 255.255.255.252
 ip pim query-interval 333 msec
 ip pim sparse-mode
 logging event subif-link-status

interface TenGigabitEthernet1/1.102
 description Subinterface for Green VRF
 encapsulation dot1Q 102
 ip vrf forwarding Green
 ip address 10.122.5.29 255.255.255.252
 ip pim query-interval 333 msec
 ip pim sparse-mode
 logging event subif-link-status
```

VNET Trunk Config

```
interface TenGigabitEthernet1/1
 vnet trunk
 ip address 10.122.5.30 255.255.255.252
 ip pim query-interval 333 msec
 ip pim sparse-mode
 logging event link-status
```

Both Routers Have VRFs Defined

VNET Router Has Tags

```
Global Config:
vrf definition red
 vnet tag 101

vrf definition green
 vnet tag 102
```

* RouterEVN# Show derived-config (will display the config beyond what EVN displays from a simplification perspective)

VRF Simplification - Trunk Advantage

VRF-Lite Subinterfaces

```
interface TenGigabitEthernet1/1.101
 description 10GE to core 3
 encapsulation dot1Q 101
 ip vrf forwarding Red
 ip address 10.122.5.31 255.255.255.254
 ip pim query-interval 333 msec
 ip pim sparse-mode
 logging event subif-link-status

interface TenGigabitEthernet1/1.102
 description 10GE to core 3
 encapsulation dot1Q 102
 ip vrf forwarding Green
 ip address 10.122.5.31 255.255.255.254
 ip pim query-interval 333 msec
 ip pim sparse-mode
 logging event subif-link-status
```

VNET Trunks

```
interface TenGigabitEthernet1/1
 description 10GE to core 3
 vnet trunk
 ip address 10.122.5.31 255.255.255.254
 ip pim query-interval 333 msec
 ip pim sparse-mode
 logging event link-status
```

1 Point-to-Point Subinterface Configuration, **per VRF per Physical Interfaces**

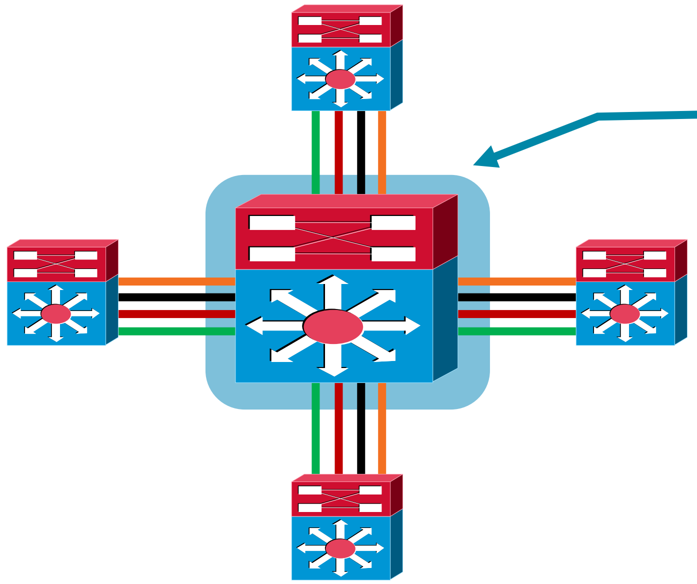

Virtual Networks	Neighbours	VRF Subinterfaces
4	4	16
10	4	40
20	4	80
30	4	120

VRF Simplification - Trunk Advantage

VRF-Lite Subinterfaces

```
interface TenGigabitEthernet1/1.101
 description 10GE to core 3
 encapsulation dot1Q 101
 ip vrf forwarding Red
 ip address 10.122.5.31 255.255.255.254
 ip pim query-interval 333 msec
 ip pim sparse-mode
 logging event subif-link-status

interface TenGigabitEthernet1/1.102
 description 10GE to core 3
 encapsulation dot1Q 102
 ip vrf forwarding Green
 ip address 10.122.5.31 255.255.255.254
 ip pim query-interval 333 msec
 ip pim sparse-mode
 logging event subif-link-status
```

VNET Trunks

```
interface TenGigabitEthernet1/1
 description 10GE to core 3
 vnet trunk
 ip address 10.122.5.31 255.255.255.254
 ip pim query-interval 333 msec
 ip pim sparse-mode
 logging event link-status
```

1 Point-to-Point Subinterface Configuration, **per VRF per Physical Interfaces**

1 **Point-to-Point Trunk Configuration per Physical Interface**

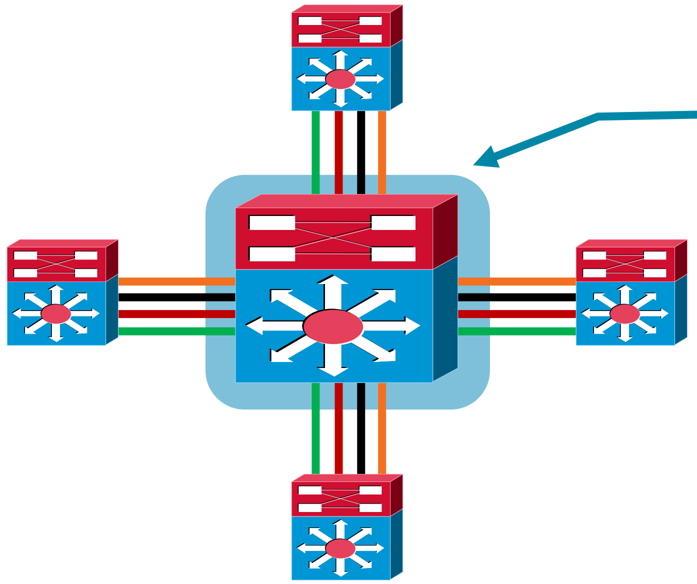

Virtual Networks	Neighbours	VRF Subinterfaces	VRF Trunks
4	4	16	4
10	4	40	4
20	4	80	4
30	4	120	4

EVN - Routing Context – Simplified CLI

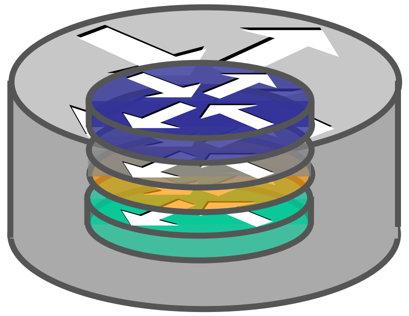

Routing Context

Router# routing-context vrf red

Router%red#

IOS CLI

IOS CLI		Routing Context
Router# show ip route vrf red *Routing table output for red*	→	**Router%red# show ip route** *Routing table output for red*
Router# ping vrf red 10.1.1.1 *Ping result using VRF red*	→	**Router%red# ping 10.1.1.1** *Ping result using VRF red*
Router# telnet 10.1.1.1 /vrf red *Telnet to 10.1.1.1 in VRF red*	→	**Router%red# telnet 10.1.1.1** *Telnet to 10.1.1.1 in VRF red*
Router# traceroute vrf red 10.1.1.1 *Traceroute output in VRF red*	→	**Router%red# traceroute 10.1.1.1** *Traceroute output in VRF red*

Shared Services in Virtualised Networks

Services that you don’t want to duplicate:

- Internet Gateway
- Firewall and NAT - DMZ
- DNS
- DHCP
- Corporate Communications - Hosted Content

Requires IP Connectivity between VRFs

This is usually accomplished through some type of Extranet Capability or Fusion Router/FW

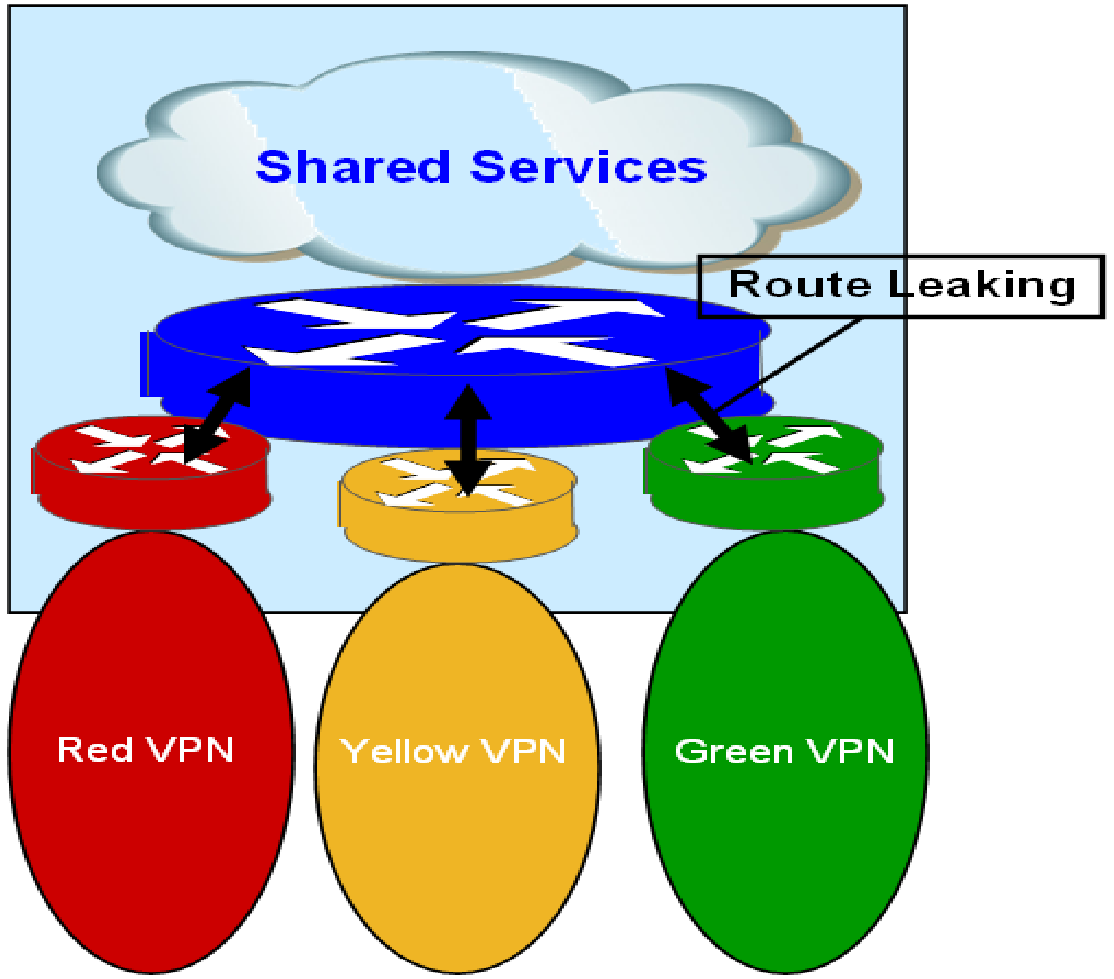

<u>Best Methods for Shared Services</u>

Fusion Router/FW – Internet Gateway, NAT/DMZ

Extranet – DNS, DHCP, Corp Communications

VRF Simplification - Shared Services

Before: Sharing Servers in Existing Technologies

```
ip vrf SHARED
 rd 3:3
 route-target export 3:3
 route-target import 1:1
 route-target import 2:2
!
ip vrf RED
 rd 1:1
 route-target export 1:1
 route-target import 3:3
!
ip vrf GREEN
 rd 2:2
 route-target export 2:2
 route-target import 3:3
!
router bgp 65001
 bgp log-neighbor-changes
 !
 address-family ipv4 vrf SHARED
 redistribute ospf 3
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf RED
 redistribute ospf 1
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf GREEN
 redistribute ospf 2
 no auto-summary
 no synchronization
 exit-address-family
!
```

After: Simple Shared Service Definition

```
vrf definition SHARED
  address-family ipv4
   route-replicate from vrf RED unicast all route-map red-map
   route-replicate from vrf GREEN unicast all route-map grn-map
```

```
vrf definition RED
  address-family ipv4
   route-replicate from vrf SHARED unicast all
```

```
vrf definition GREEN
  address-family ipv4
   route-replicate from vrf SHARED unicast all
```

Route-Replication Advantage:

- No BGP required
- No Route Distinguisher required
- No Route Targets required
- No Import/Export required
- Simple Deployment
- Supports both Unicast/Mcast

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
  - Easy Virtual Networking (EVN)
  - **Easy Virtual Networking (EVN) + WAN Virtualisation**
  - Locator ID Separation Protocol (LISP) – Network Virtualisation for Multi-tenancy
  - Using GET VPN Encryption for Multipoint GRE (mGRE) Solutions
  - MTU Caveats and Solutions for IP Tunneled Environments
- Summary

Extending EVN Over the WAN

Leverage MPLS VPN over mGRE for EVN Extension

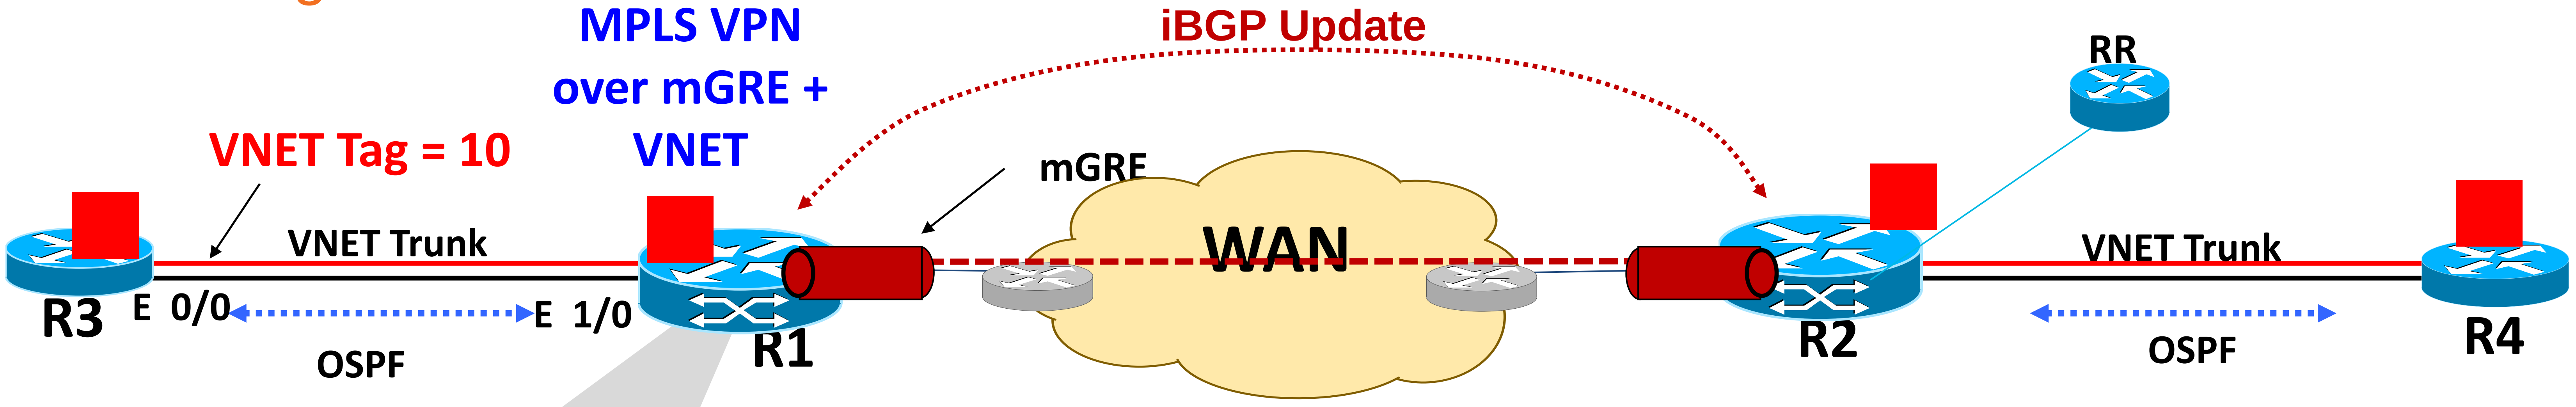

- On MPLS "PE", apply the 'vnet tag' under the "vrf definition"
- This injects the VNET into the VRF and is handled as normal VRF forwarding (over MPLS VPN over GRE in this example)

```
!
vrf definition red
 vnet tag 10

 rd 1:1
 route-target export 1:1
 route-target import 1:1
 !
!
 address-family ipv4
 exit-address-family
!
```

VNET Tag Applied under the "vrf Definition"

Normal 'rd' and 'route-target' Applied in MPLS VPN Case

Injects Routes from VNET Trunk into VRF, Allowing Any VRF over WAN Solution to Be Applied Using VNET

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
  - Easy Virtual Networking (EVN)
  - Easy Virtual Networking (EVN) + WAN Virtualisation
  - **Locator ID Separation Protocol (LISP) – Network Virtualisation for Multi-tenancy**
  - Using GET VPN Encryption for Multipoint GRE (mGRE) Solutions
  - MTU Caveats and Solutions for IP Tunneled Environments
- Summary

What is LISP? (Locator-ID Separation Protocol)

A Next Generation Routing Architecture

IETF Draft: http://tools.ietf.org/html/draft-farinacci-lisp-12

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**

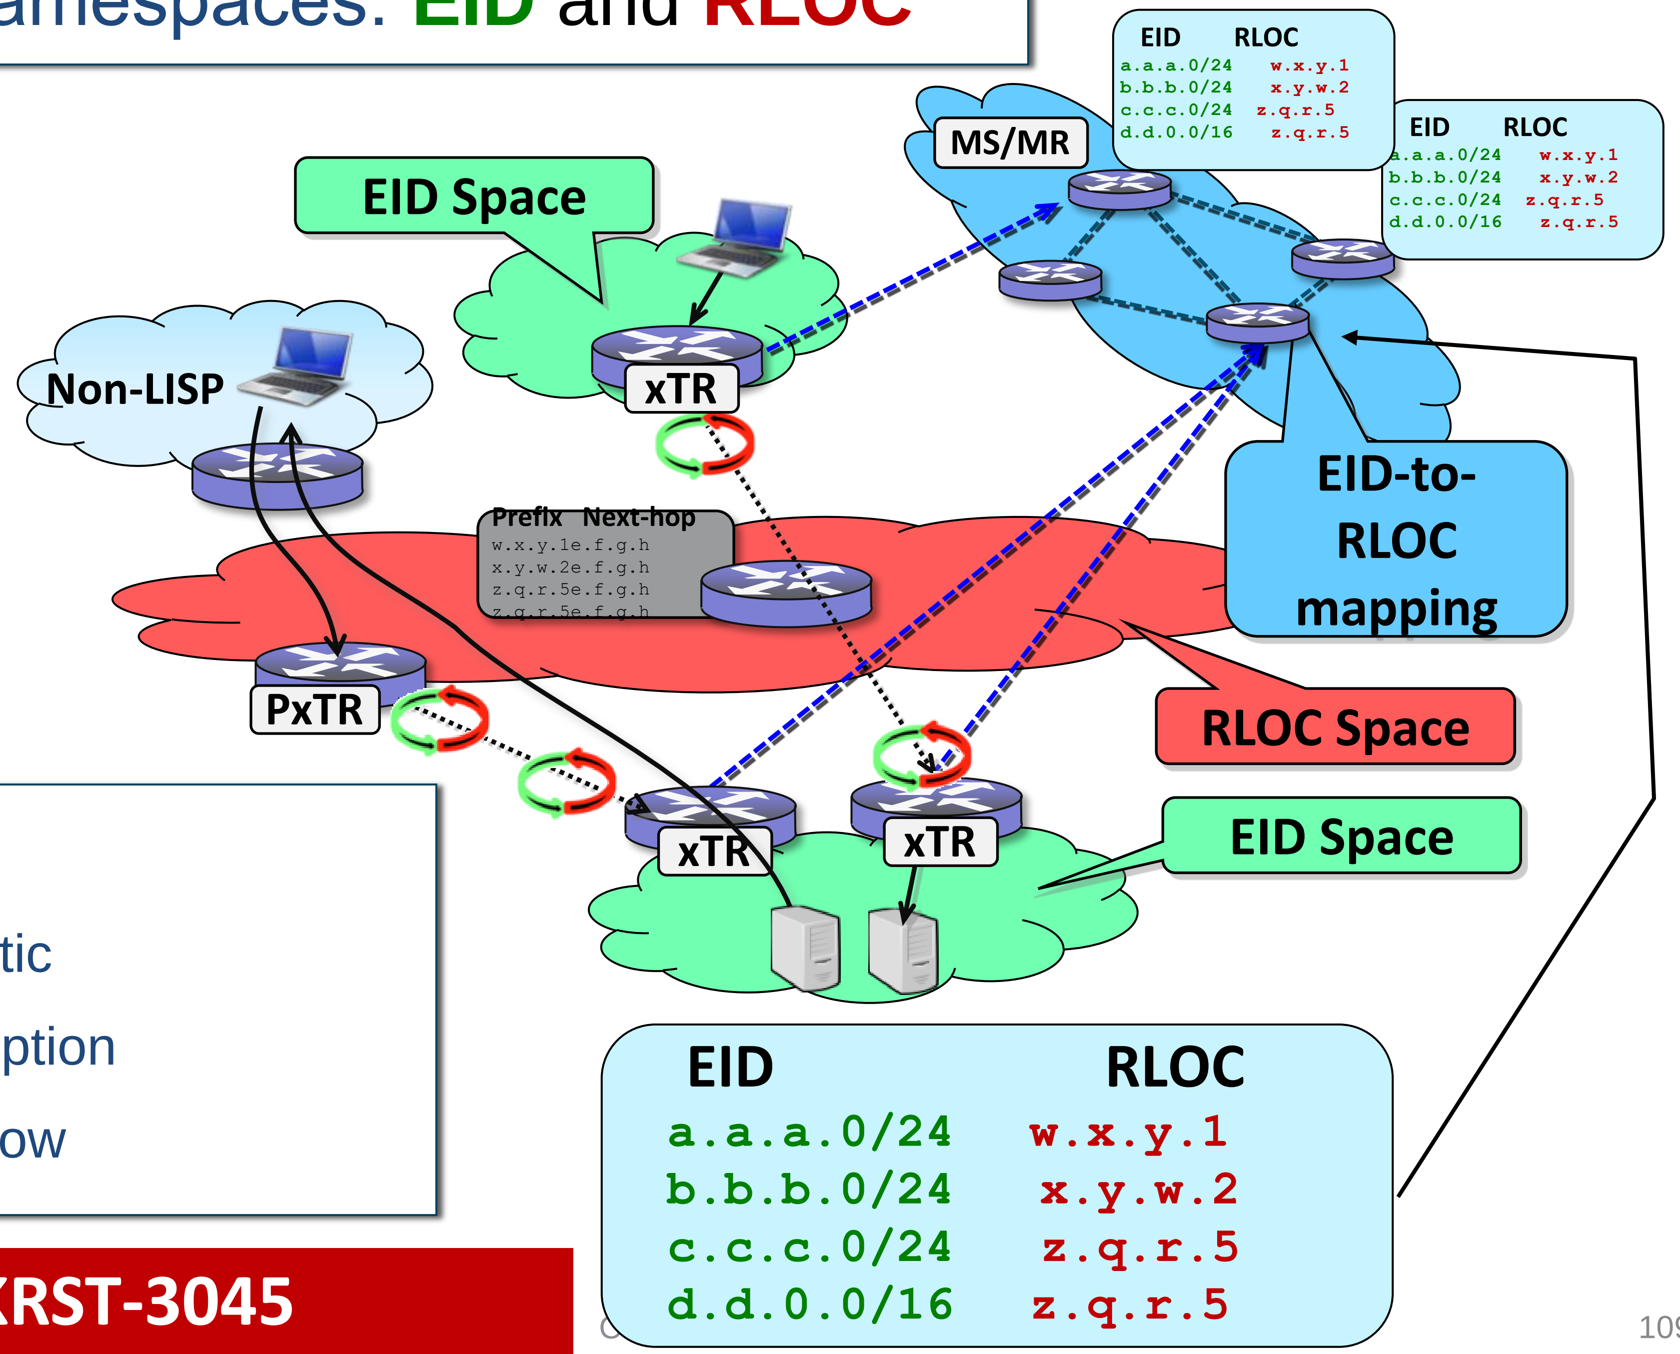

- Network-based solution
- No host changes
- Minimal configuration
- Incrementally deployable
- Support for mobility
- Address Family agnostic
- IPv4 to v6 Transition option
- In Cisco IOS/NX-OS now

More Details on LISP Covered in Session BRKRST-3045

LISP Overview

What do we mean by “location” and “identity”?

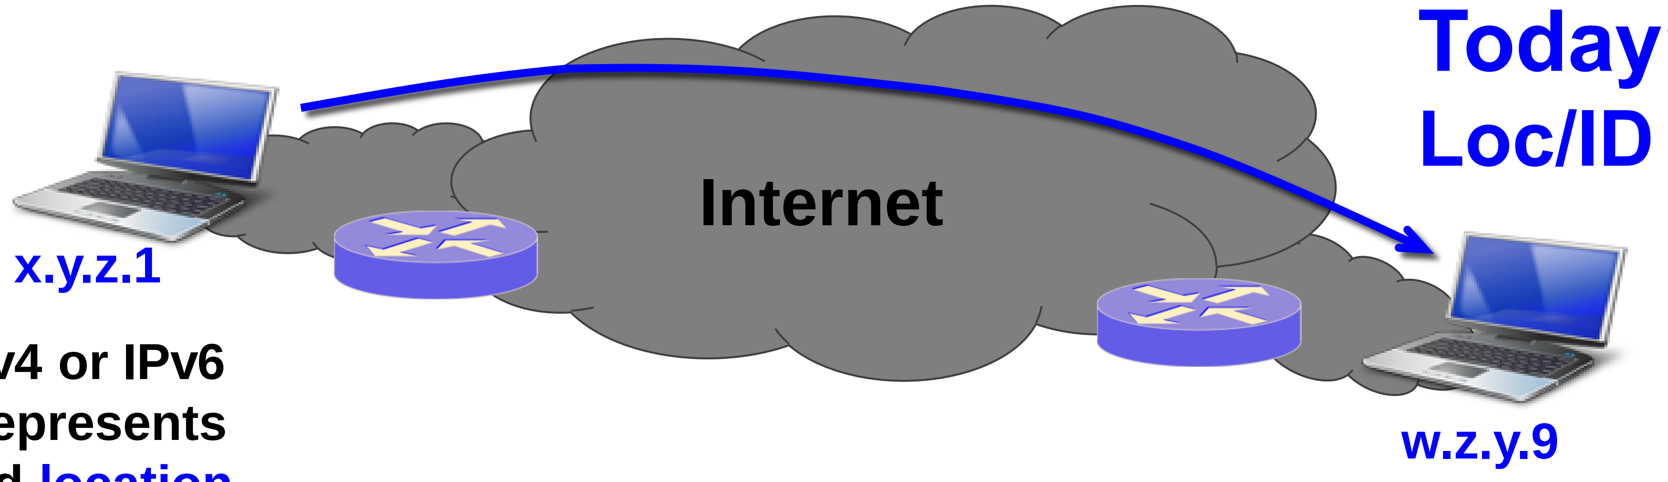

Today’s Internet Behaviour
Loc/ID “overloaded” semantic

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

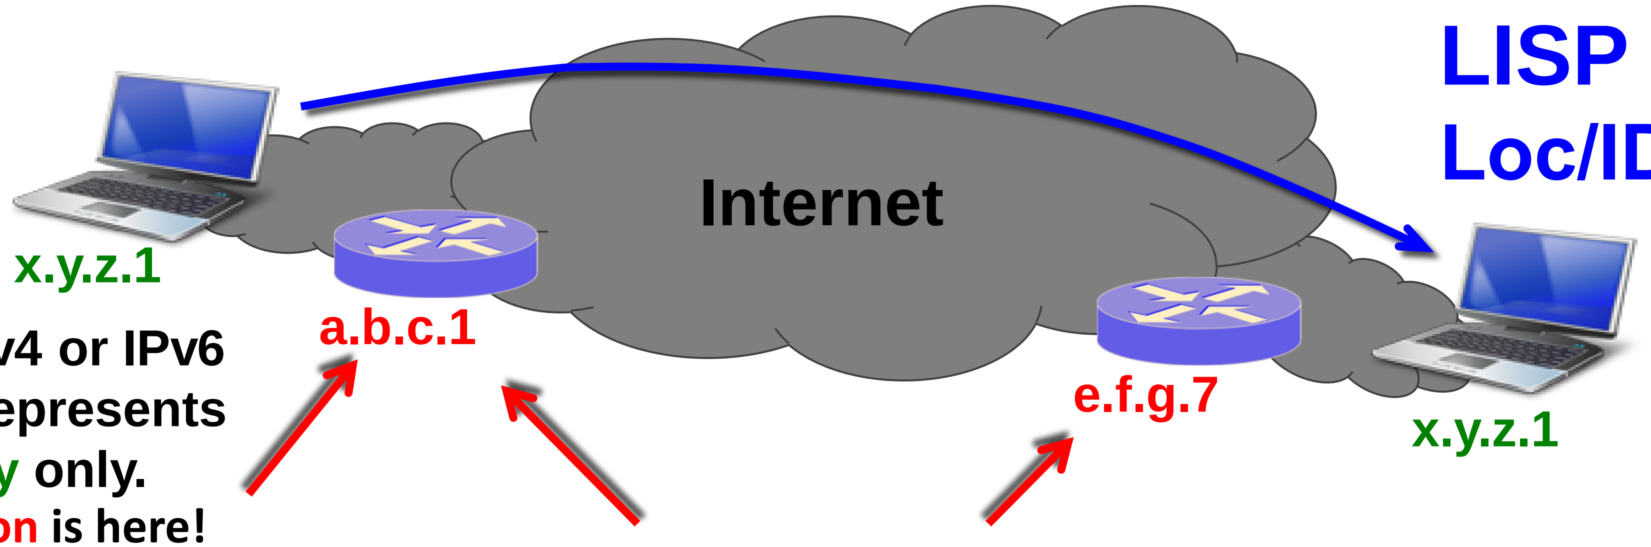

LISP Behaviour
Loc/ID “split”

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 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**

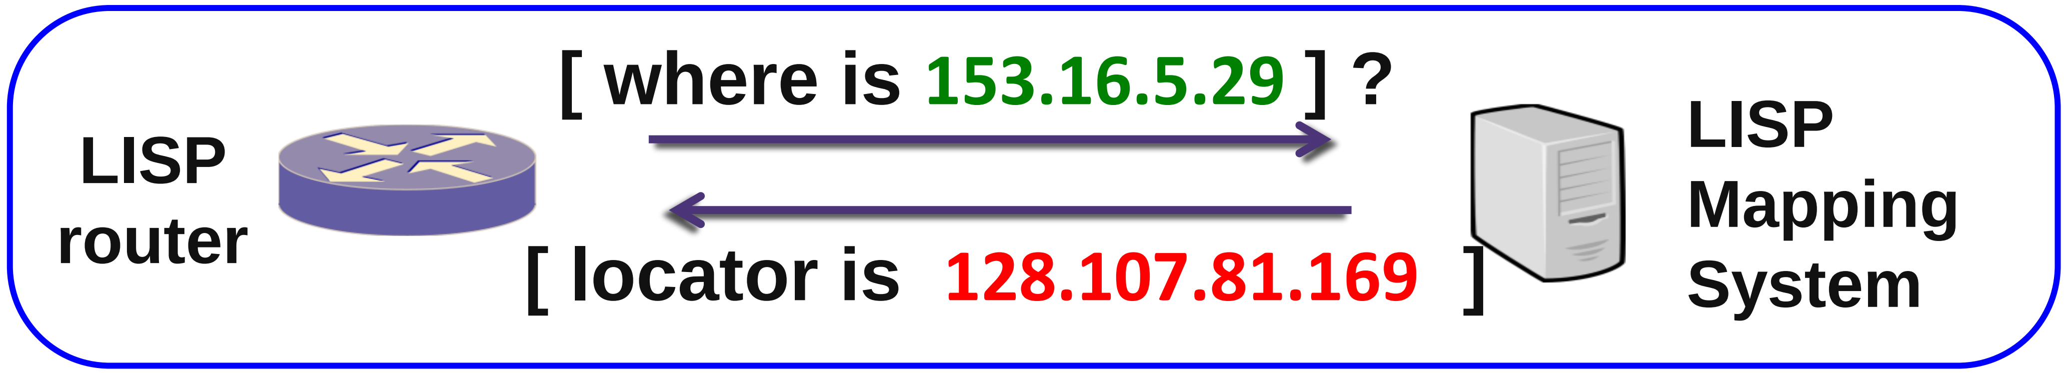

LISP
Identity-to-locator
Mapping Resolution

LISP Overview

What do we mean by “location” and “identity”?

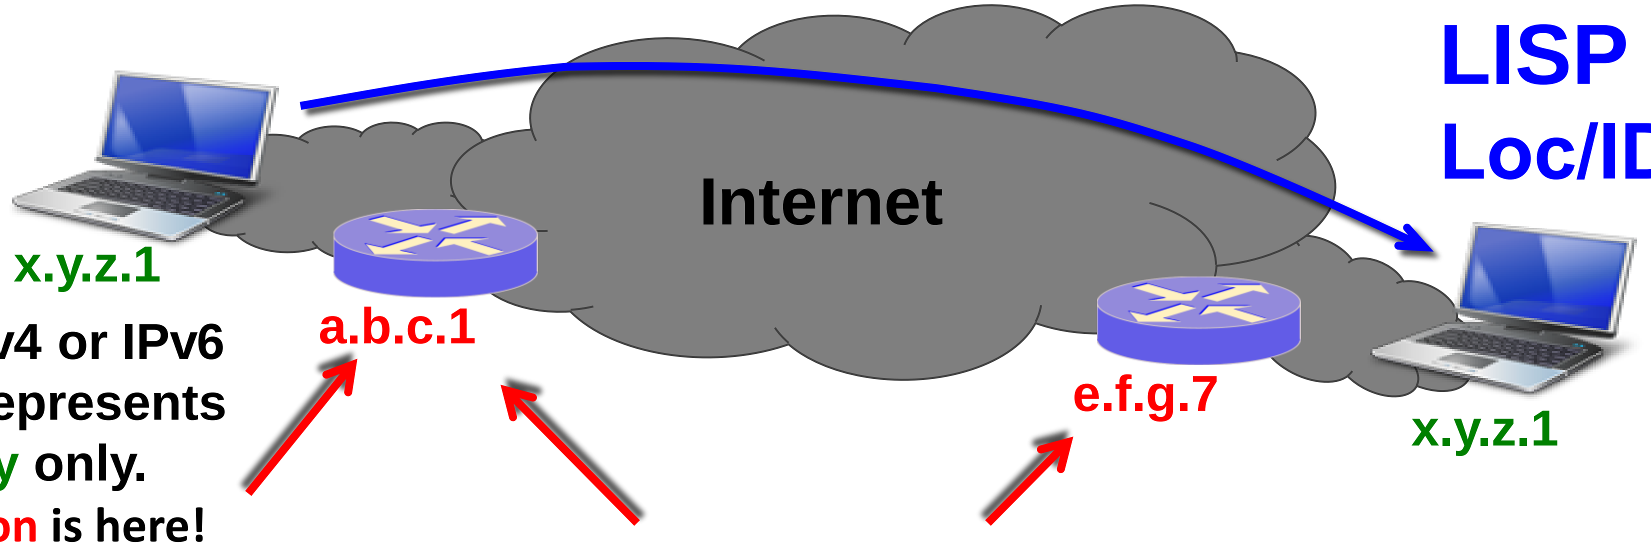

LISP Behaviour
Loc/ID “split”

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 Overview

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

LISP Behaviour
Loc/ID "split"

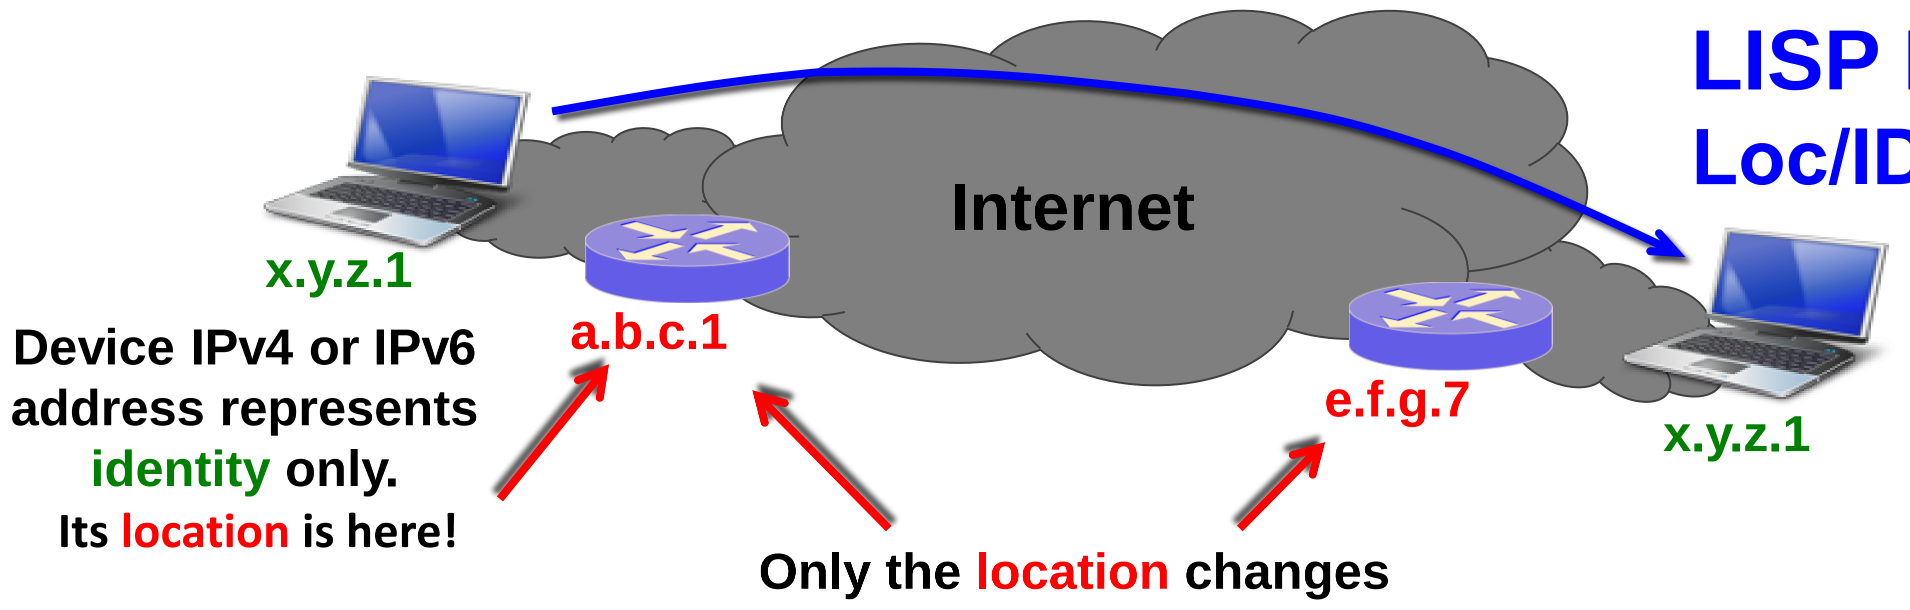

When the device moves, keeps its IPv4 or IPv6 address.
It has the same identity

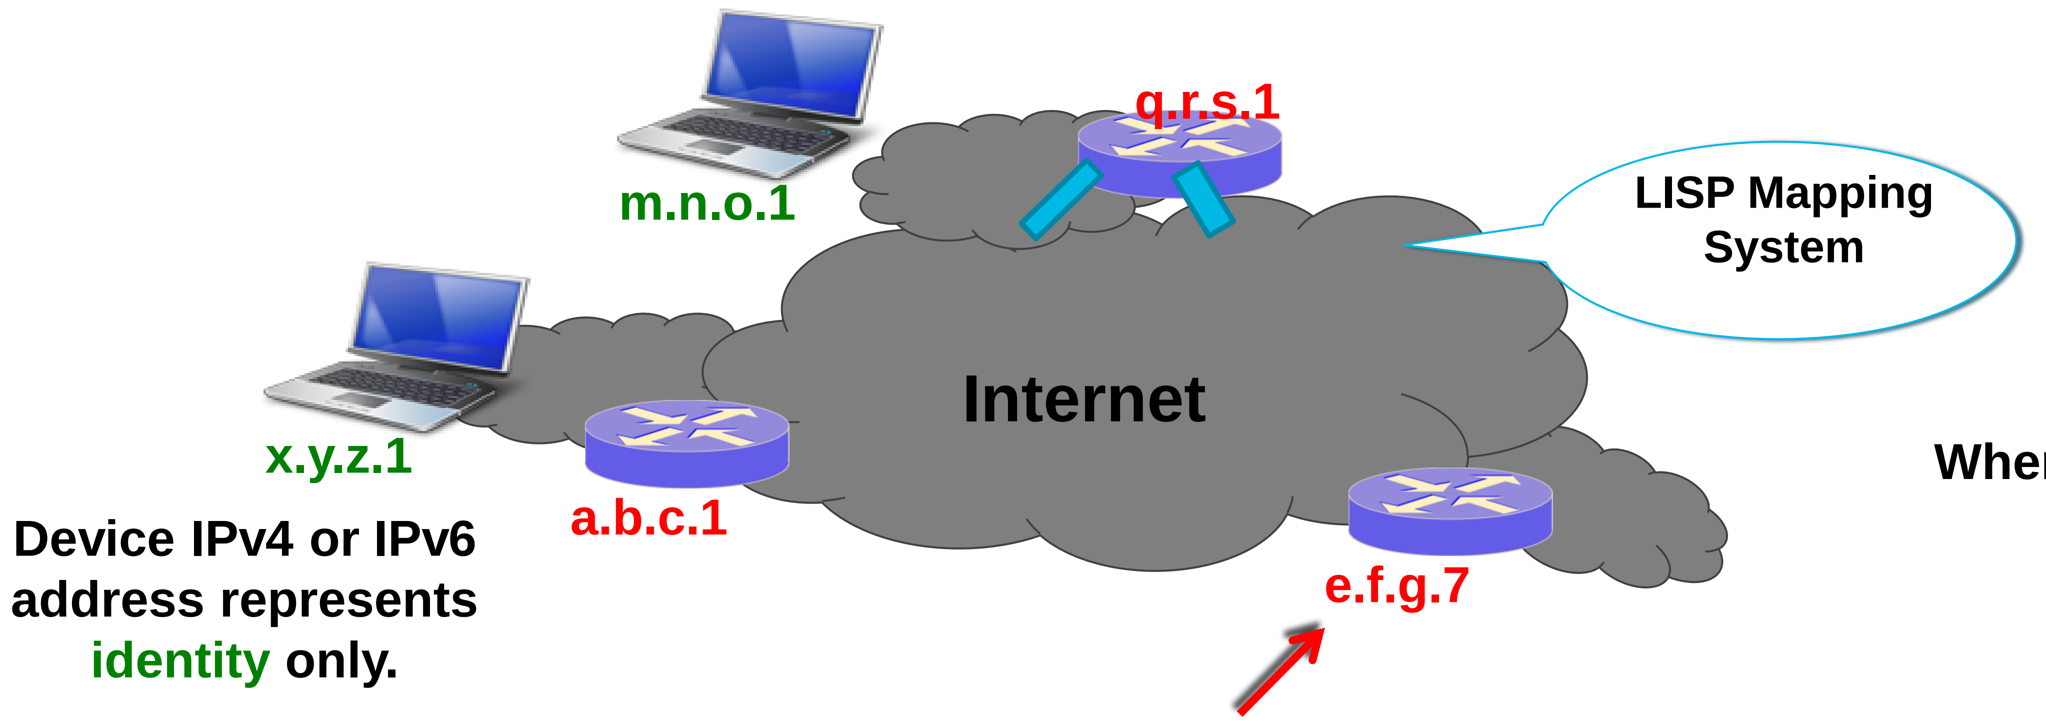

When the device moves, it keeps is IPv4 or IPv6 identity

- Upon move e.f.g.7 notifies Mapping DB of changes
- After notification, q.r.s.1 efficiently routes packets to x.y.z.1 after move

LISP Operations

LISP IPv4 EID/IPv4 RLOC Header Example

draft-ietf-lisp-19

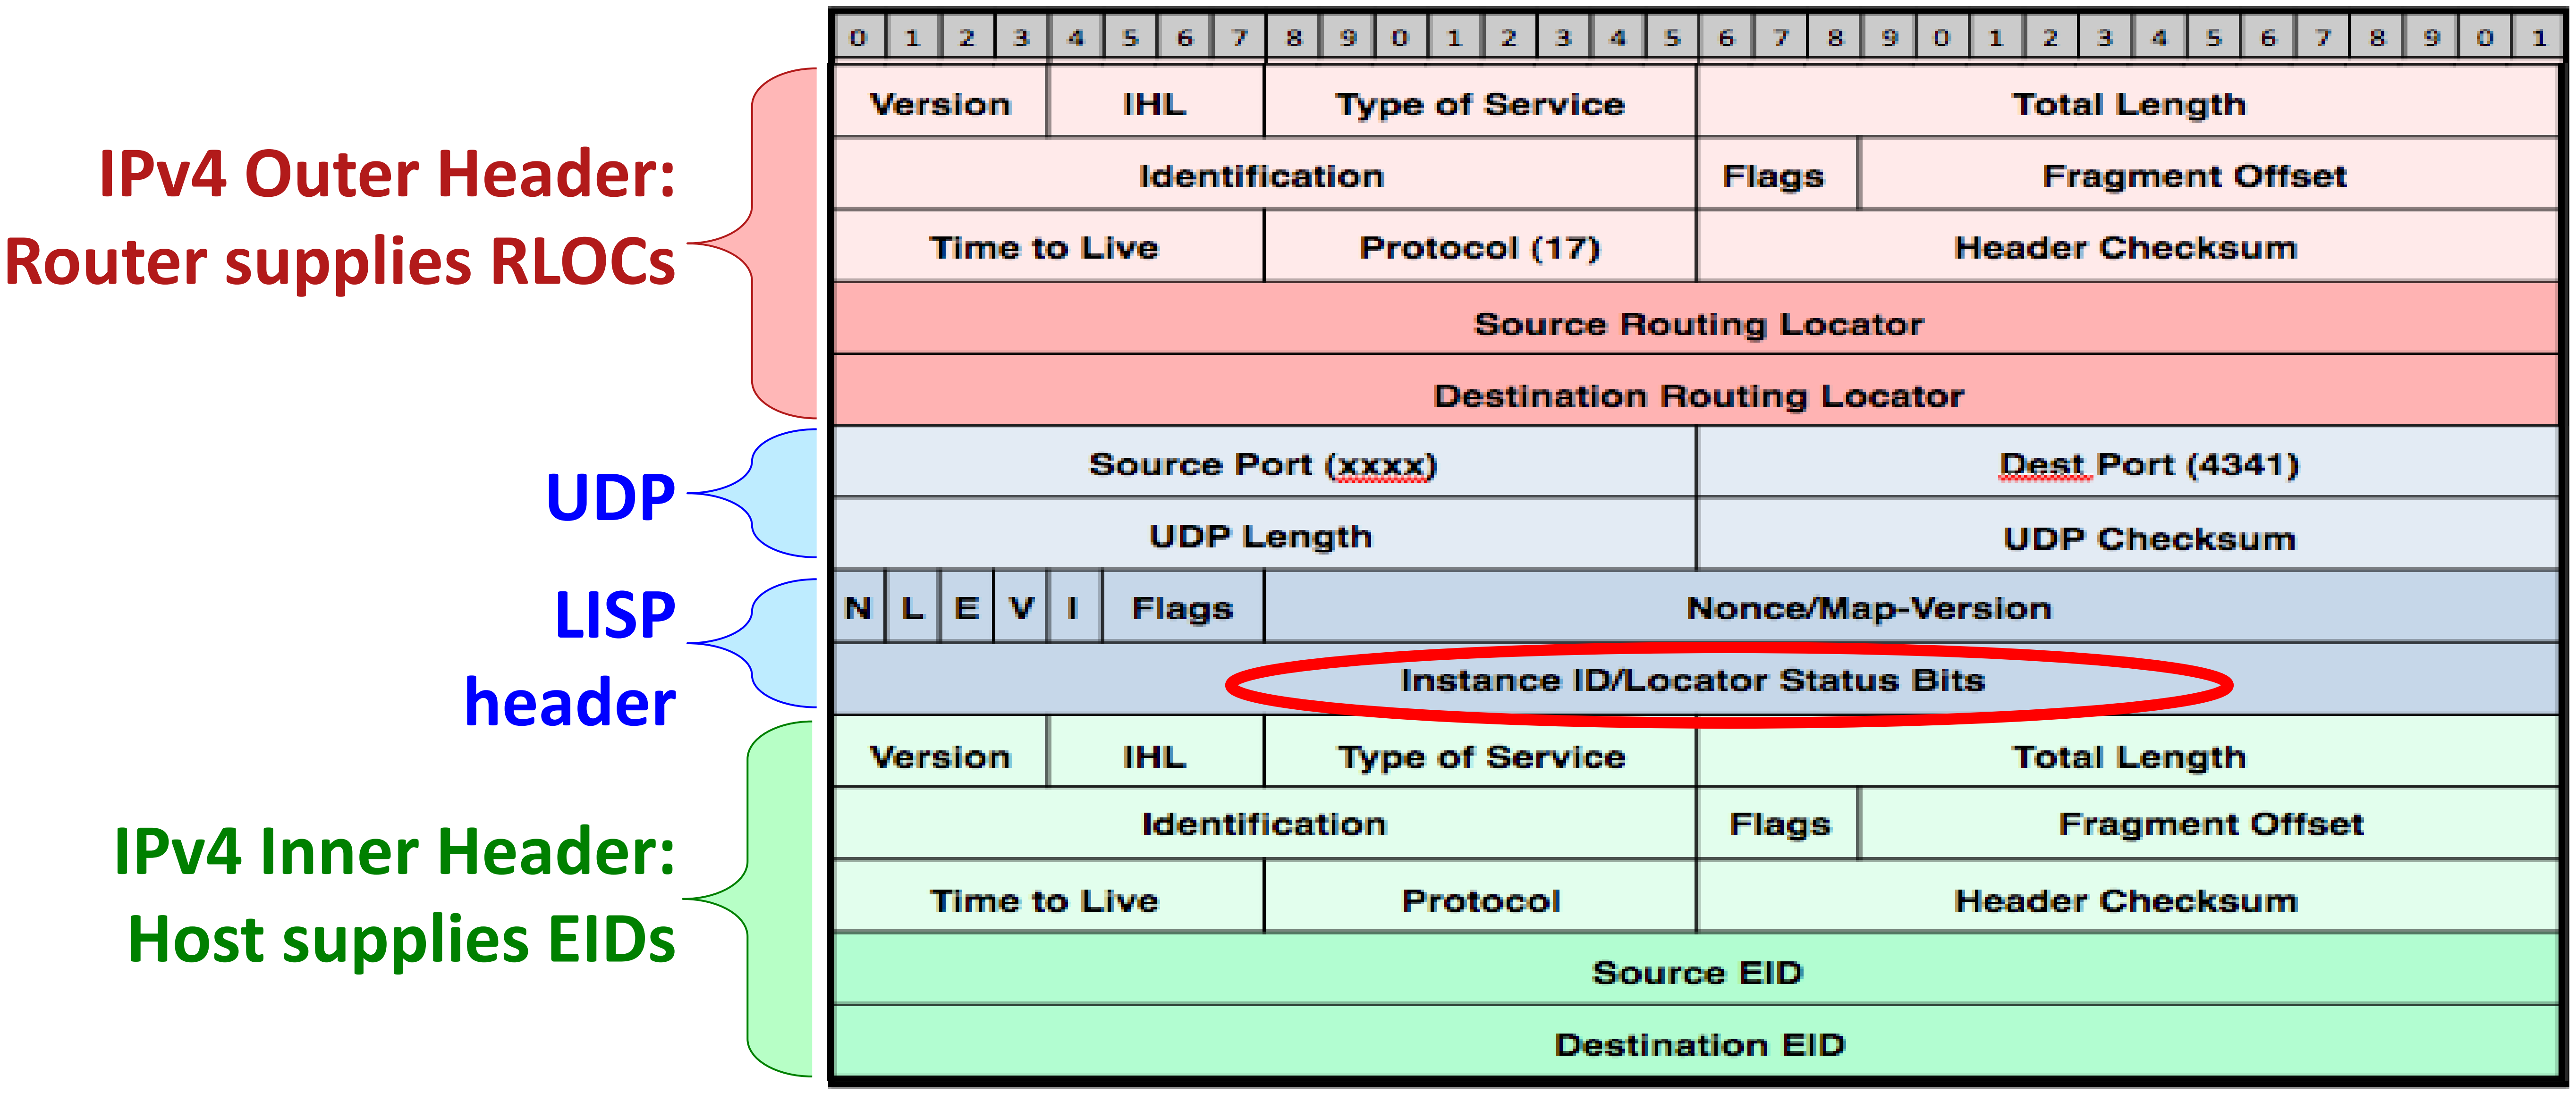

LISP Use Case – Multi-Tenancy

Network Virtualisation “Over the Top”

MDB = Mapping Database System

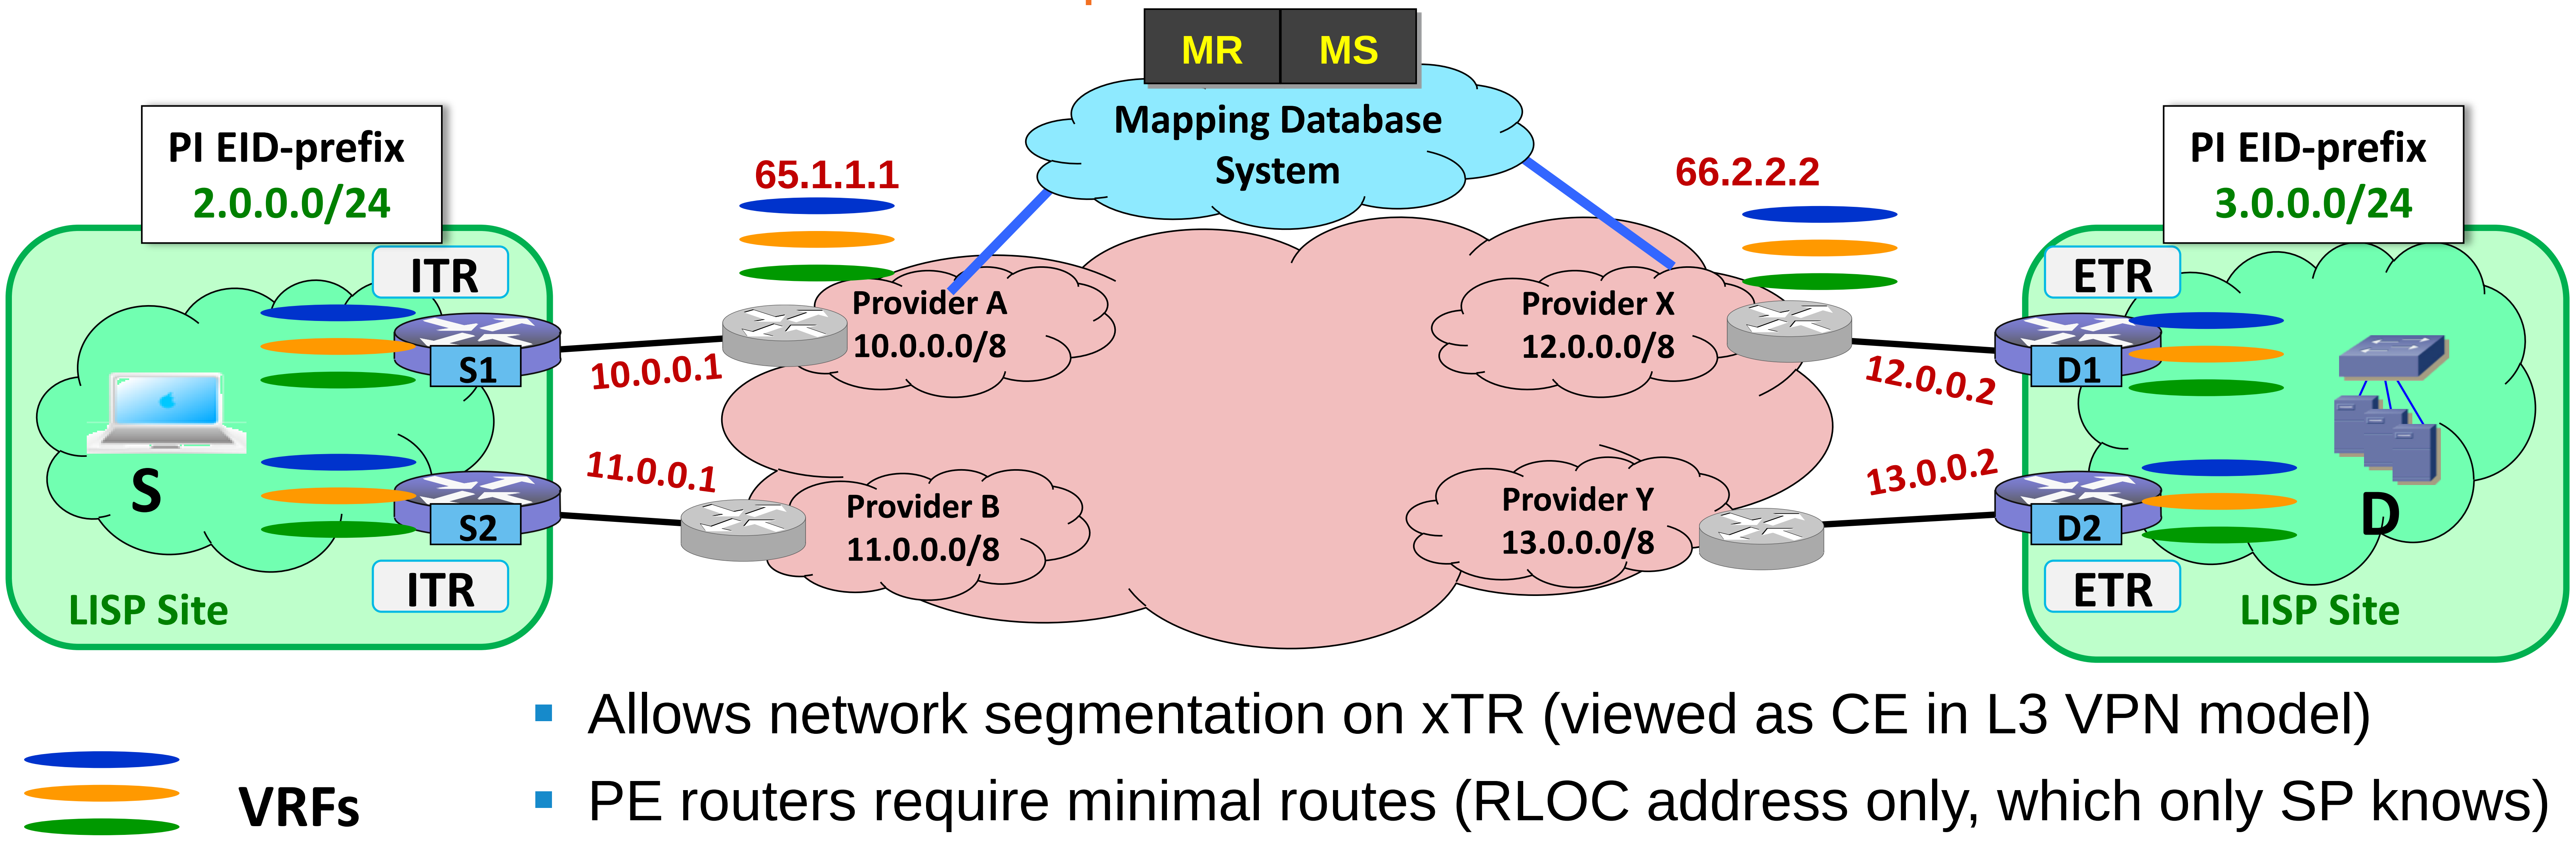

- Allows network segmentation on xTR (viewed as CE in L3 VPN model)
- PE routers require minimal routes (RLOC address only, which only SP knows)
- VRF Segmentation is applied to CE/xTR
- Offers another “over the top” Virtualisation solution (VRF capabilities, and routes are hidden from SP network network)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

Legend:
EIDs -> Green
Locators -> Red
BGP-over-GRE
Physical link

Agenda

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
  - Easy Virtual Networking (EVN)
  - Easy Virtual Networking (EVN) + WAN Virtualisation
  - Locator ID Separation Protocol (LISP) – Network Virtualisation for Multi-tenancy
  - **Using GET VPN Encryption for Multipoint GRE (mGRE) Solutions**
  - MTU Caveats and Solutions for IP Tunneled Environments
- Summary

Group Encrypted Transport (GET) VPN

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

Public/Private WAN

Before: IPSec P2P Tunnels

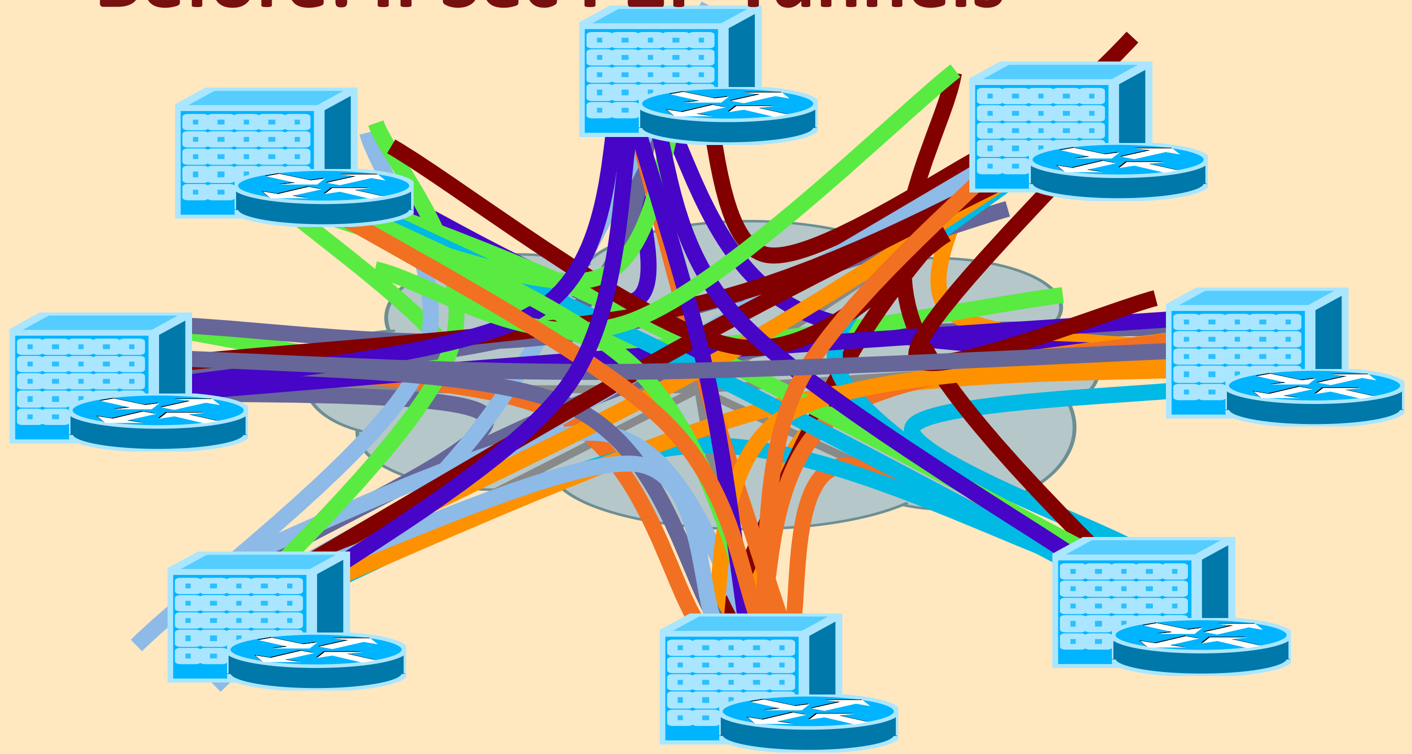

- Scalability—an issue (N^2 problem)
- Overlay routing
- Any-to-any instant connectivity can’t be done to scale
- Limited QoS
- Inefficient Multicast replication

Private WAN

After: Tunnel-Less VPN

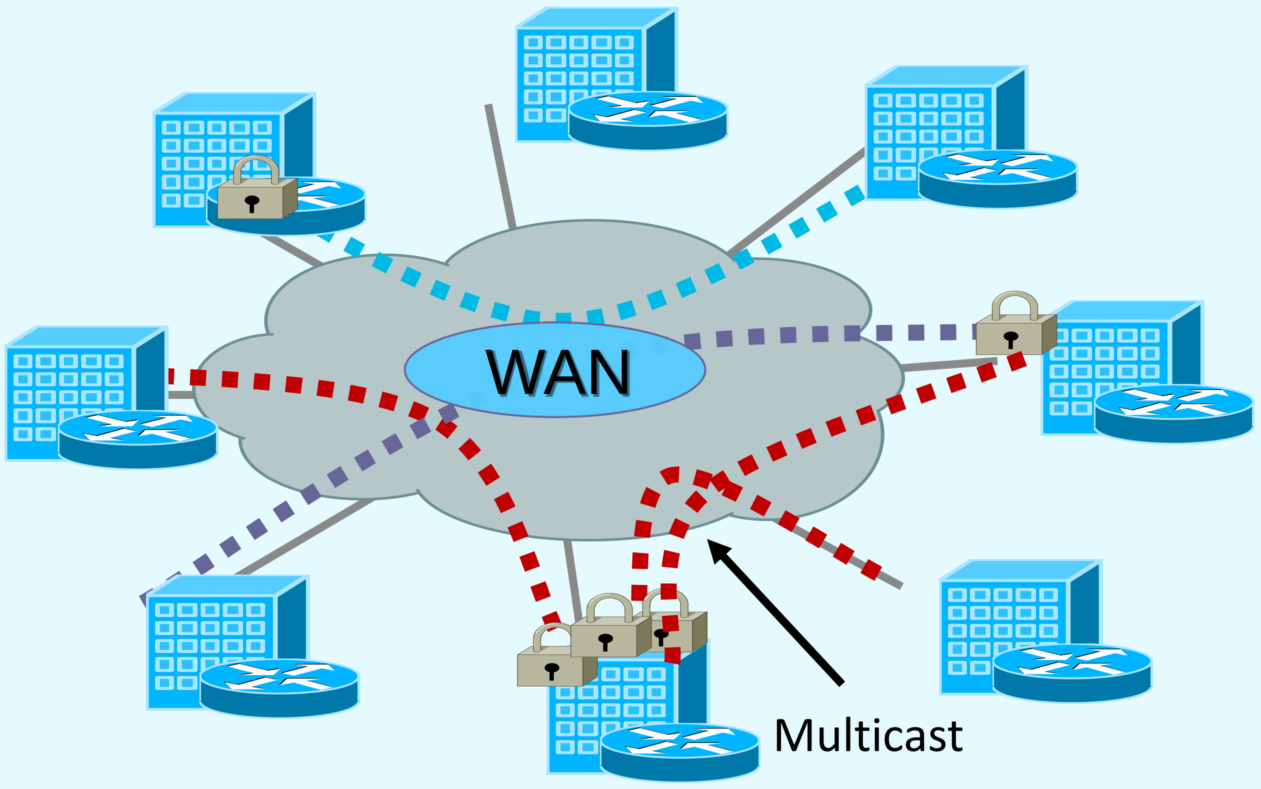

- Scalable architecture for any-to-any connectivity and encryption
- No overlays—native routing
- Any-to-any instant connectivity
- Enhanced QoS
- Efficient Multicast replication

GETVPN Security Devices

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

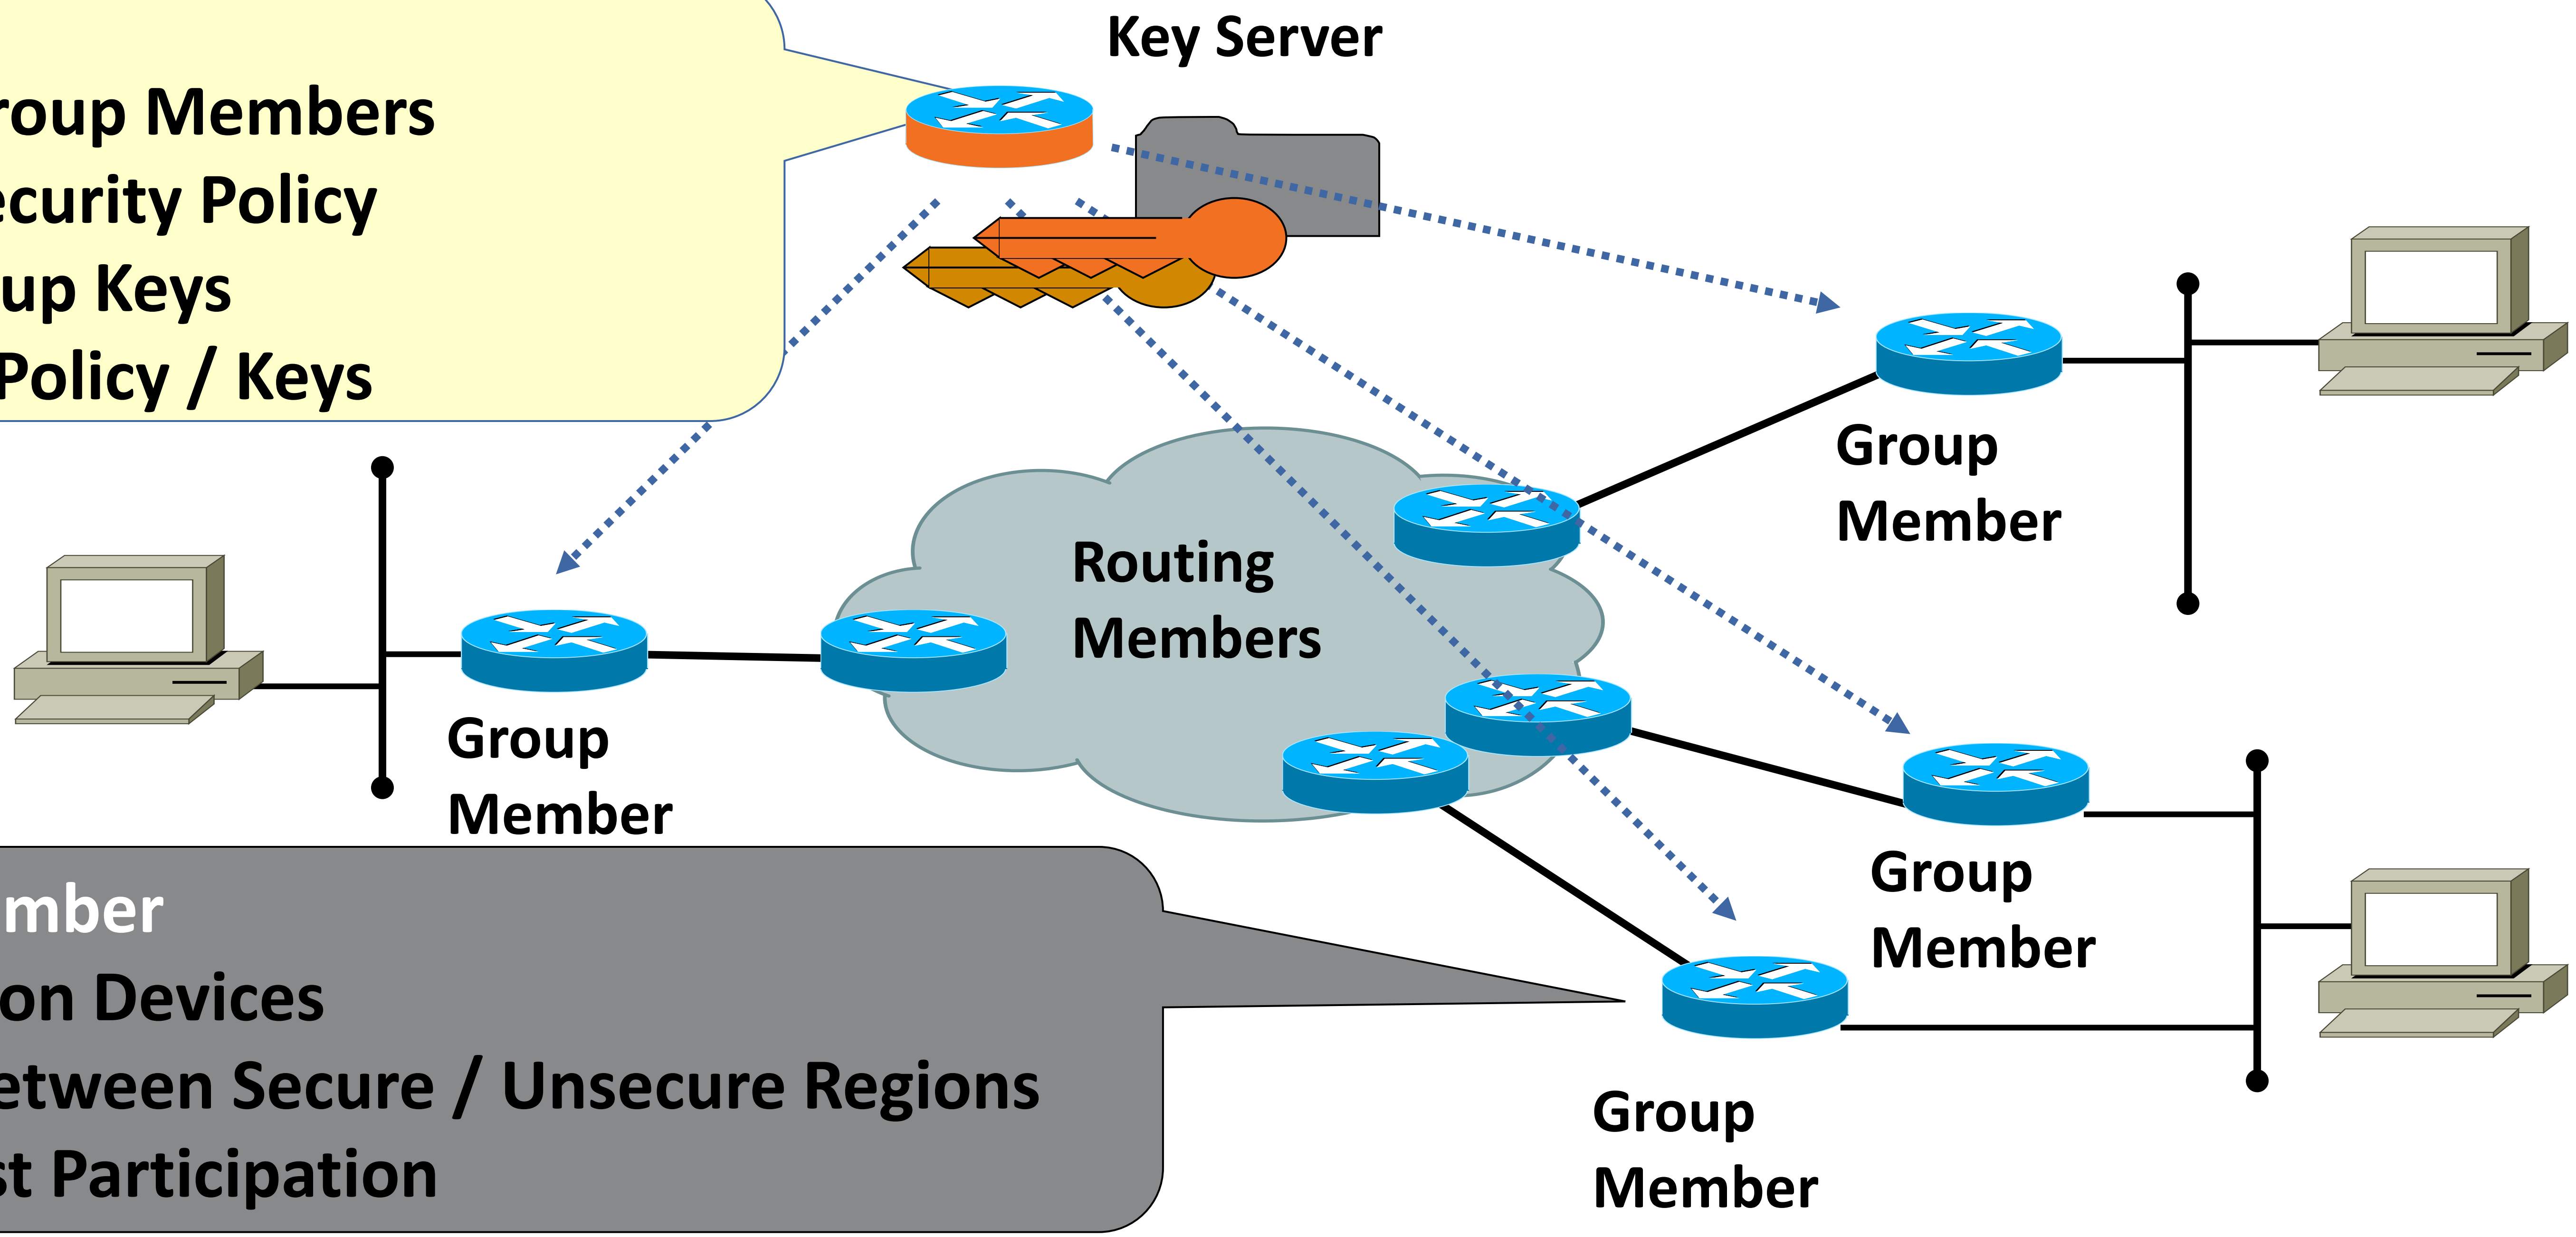

Group Member
- Encryption Devices
- Route Between Secure / Unsecure Regions
- Multicast Participation

GETVPN Security Elements

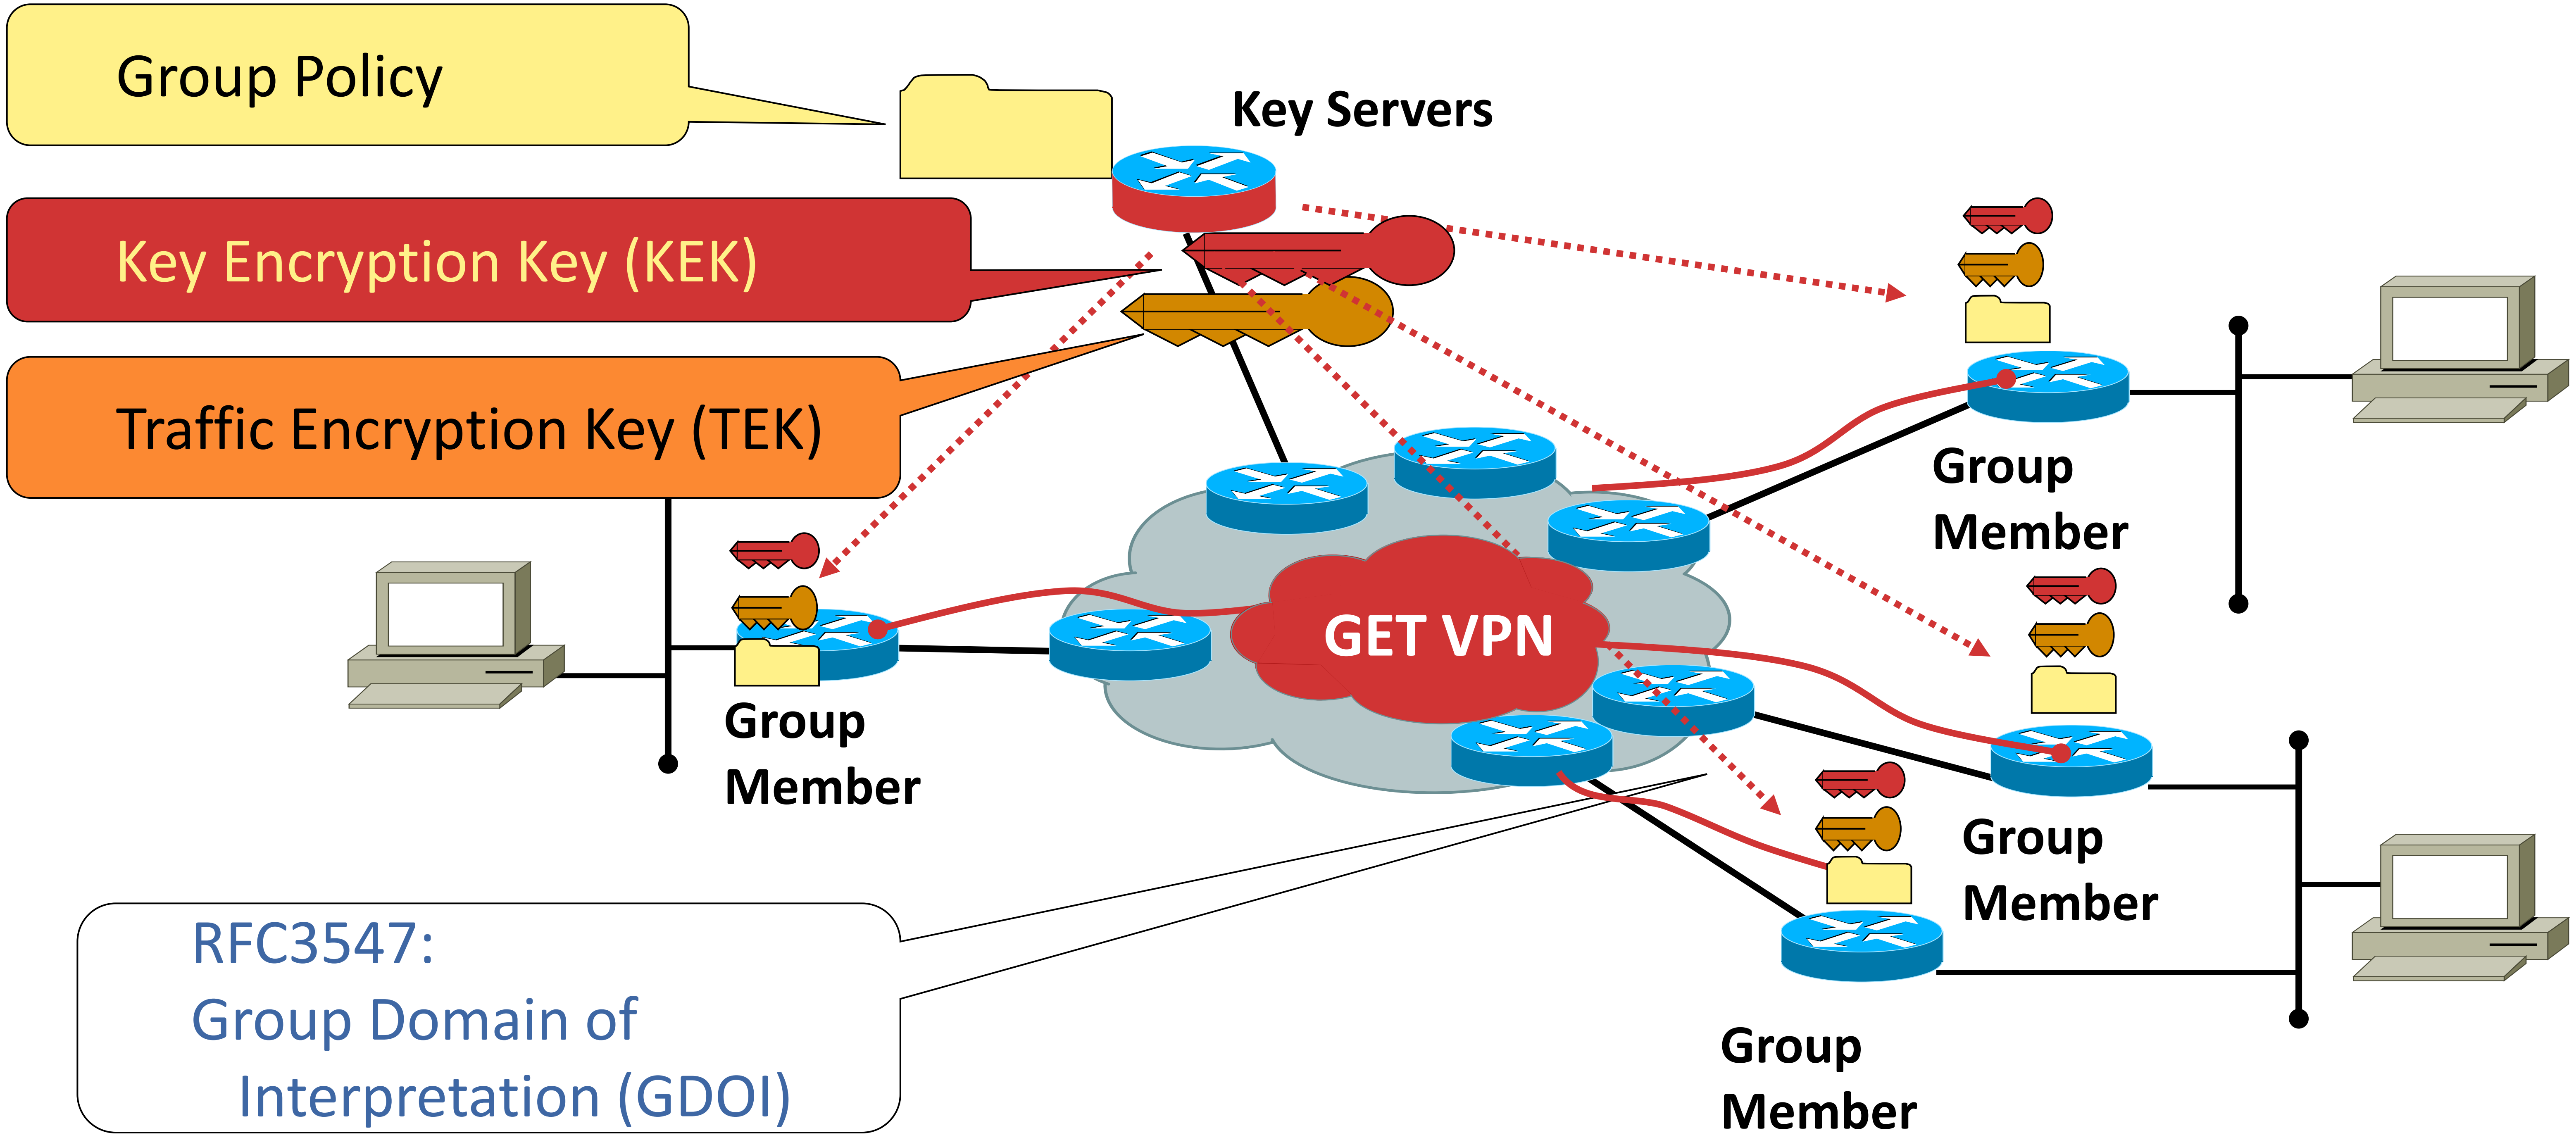

TEK used for Unicast and/or Multicast Encryption

MPLS VPN Over mGRE + GET VPN

VPN Label + Payload Encryption with mGRE Header Preservation

- GET VPN relies on the tunnel header preservation + IPSec Tunnel Mode
- Given all GM's contain the same Group SA (Security Association), the incoming packet (prior to encryption) is shown below:

MPLS/IP Datagram over mGRE *(BEFORE encryption)*

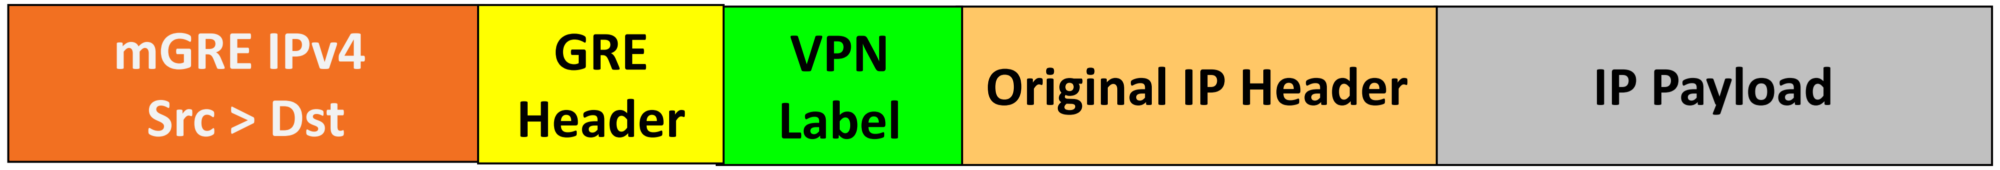

- …. Which, is then converted into the following packet (below) while preserving the original IP source, destination (through address preservation), and MPLS VPN info.

MPLS/IP Datagram over mGRE *(AFTER encryption)*

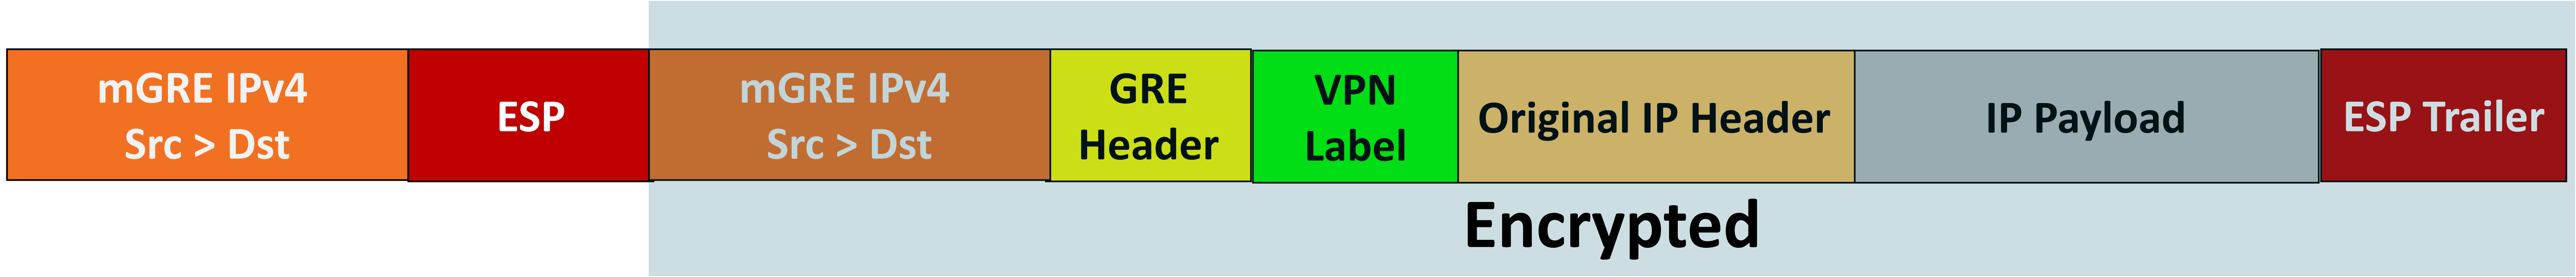

Cisco L3 Virtualisation Solutions + Encryption

Platforms and Feature Support for WAN and Branch

	IPSec SA	GET VPN
VRF Lite	NA	
VRF Lite over GRE	Y	
VRF Lite over DMVPN	Y	
MPLS-VPN	**NA**	
MPLS VPN over GRE (P2P)	Y	
MPLS VPN over DMVPN (mGRE)	Y	
MPLS VPN over mGRE (BGP)	Y	Y

MTU Considerations with GRE Tunnels

Challenges

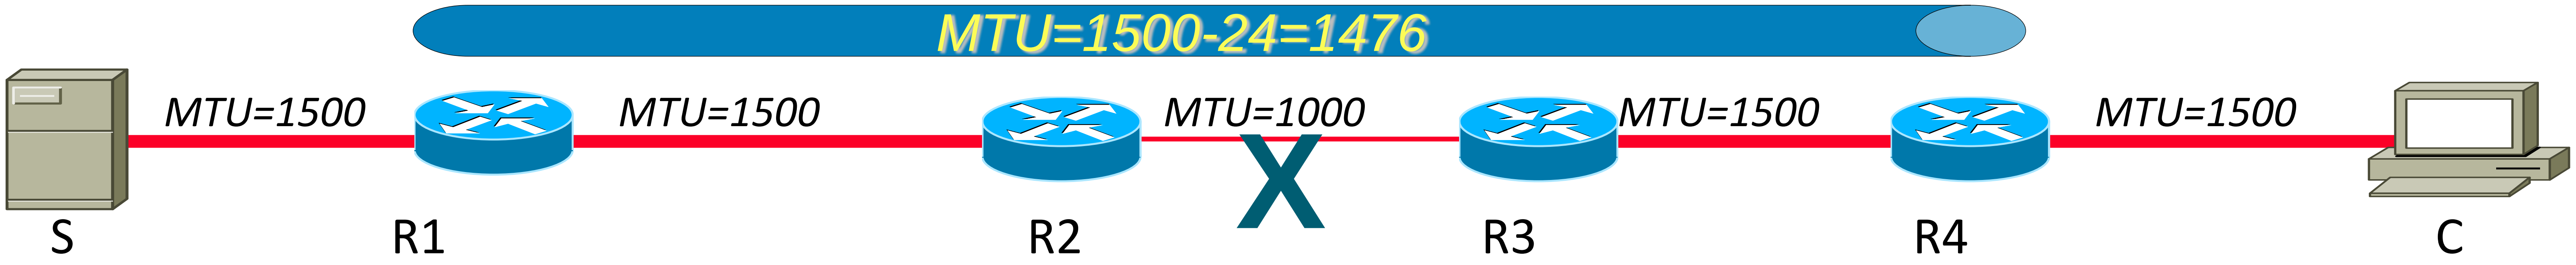

- Fragmentation is unavoidable in some cases
- The use of GRE tunnels increase the chances of MTU issues (i.e. fragmentation) due to the increase in IP packet size GRE adds
- Main Issue: The performance impact to the router when the GRE tunnel destination router must re-assemble fragmented GRE packets
- 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 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

- ✓ 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

- Network Virtualisation Drivers and Building Blocks
- Network Virtualisation Considerations over Common Service Provider Offerings
- Technology and Deployment Solutions Overview for a Virtualised WAN
- Deployment Considerations for QoS over a Virtualised WAN
- Innovations at Cisco in Network Virtualisation Overview
- Summary

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

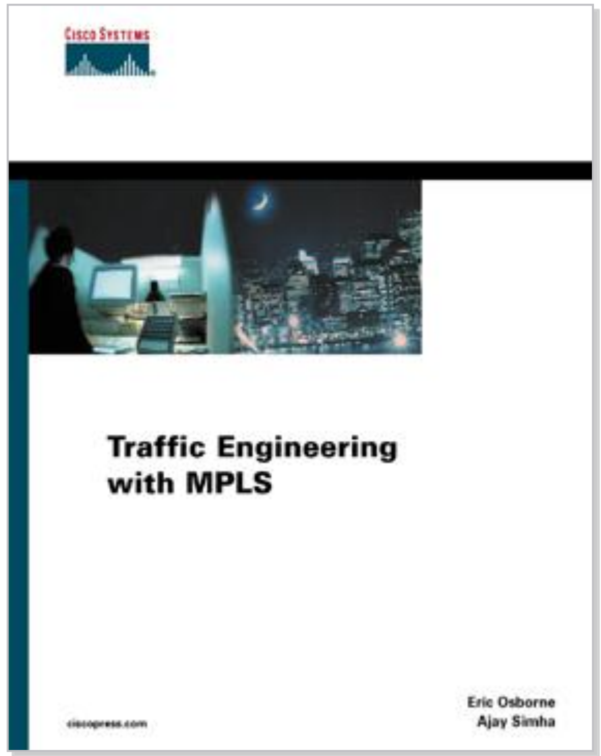

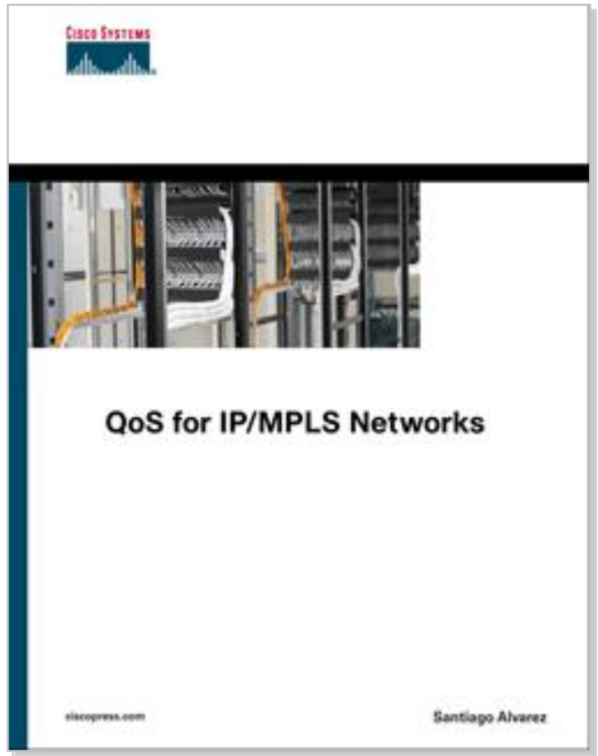

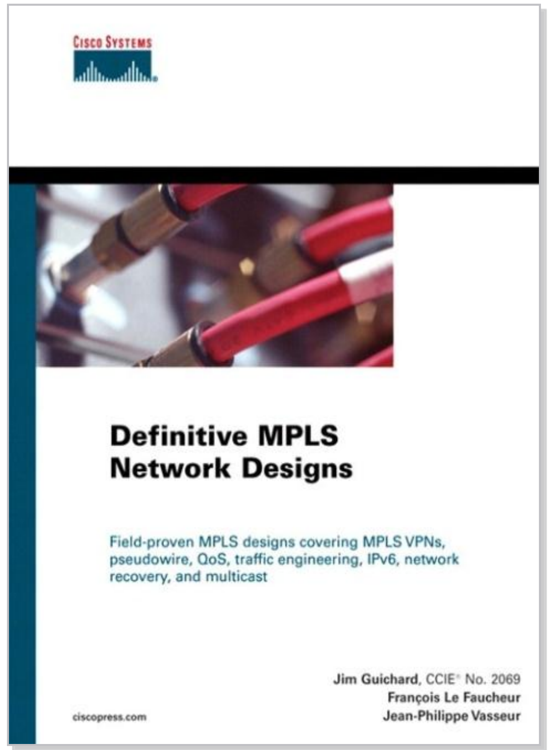

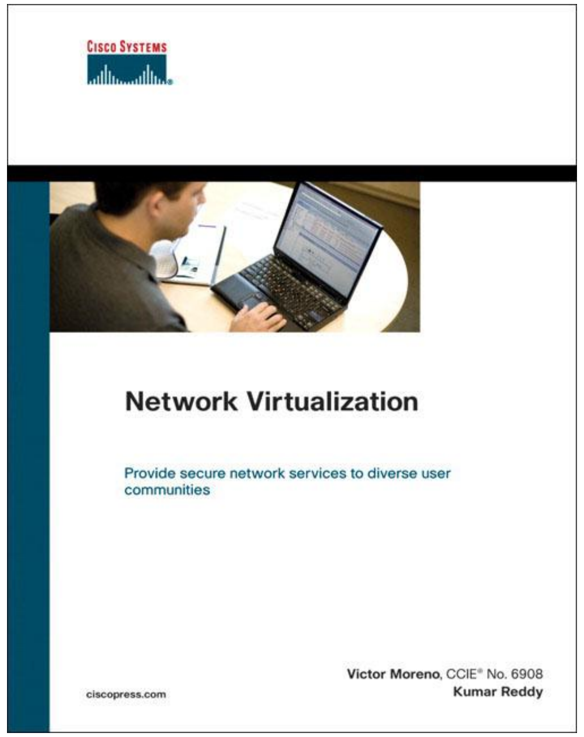

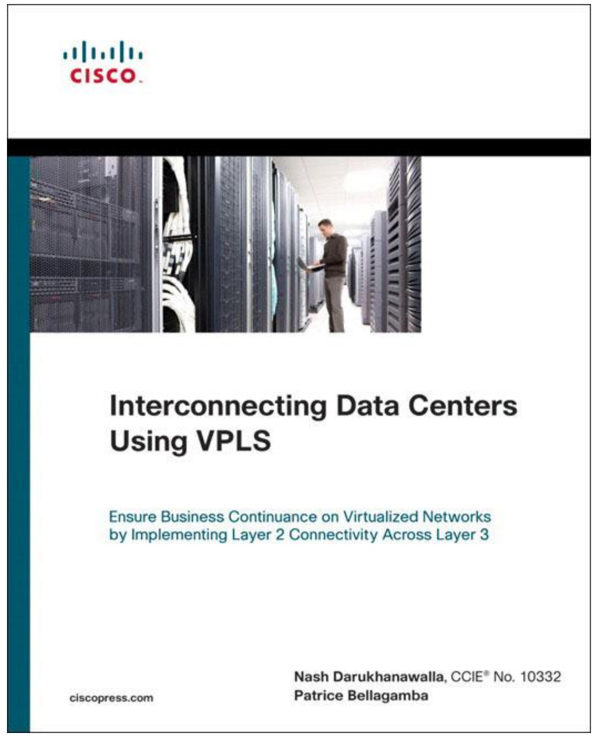

Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2013 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 8 March 12:00pm-2:00pm

Cisco live! 365

Don't forget to activate your Cisco Live 365 account for access to all session material, communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.
www.ciscoliveaustralia.com/portal/login.ww

Backup

Campus-to-WAN Virtualisation Interconnect

Campus-to-WAN Interconnection

Interconnect Virtualisation Policy WAN ←→ Campus

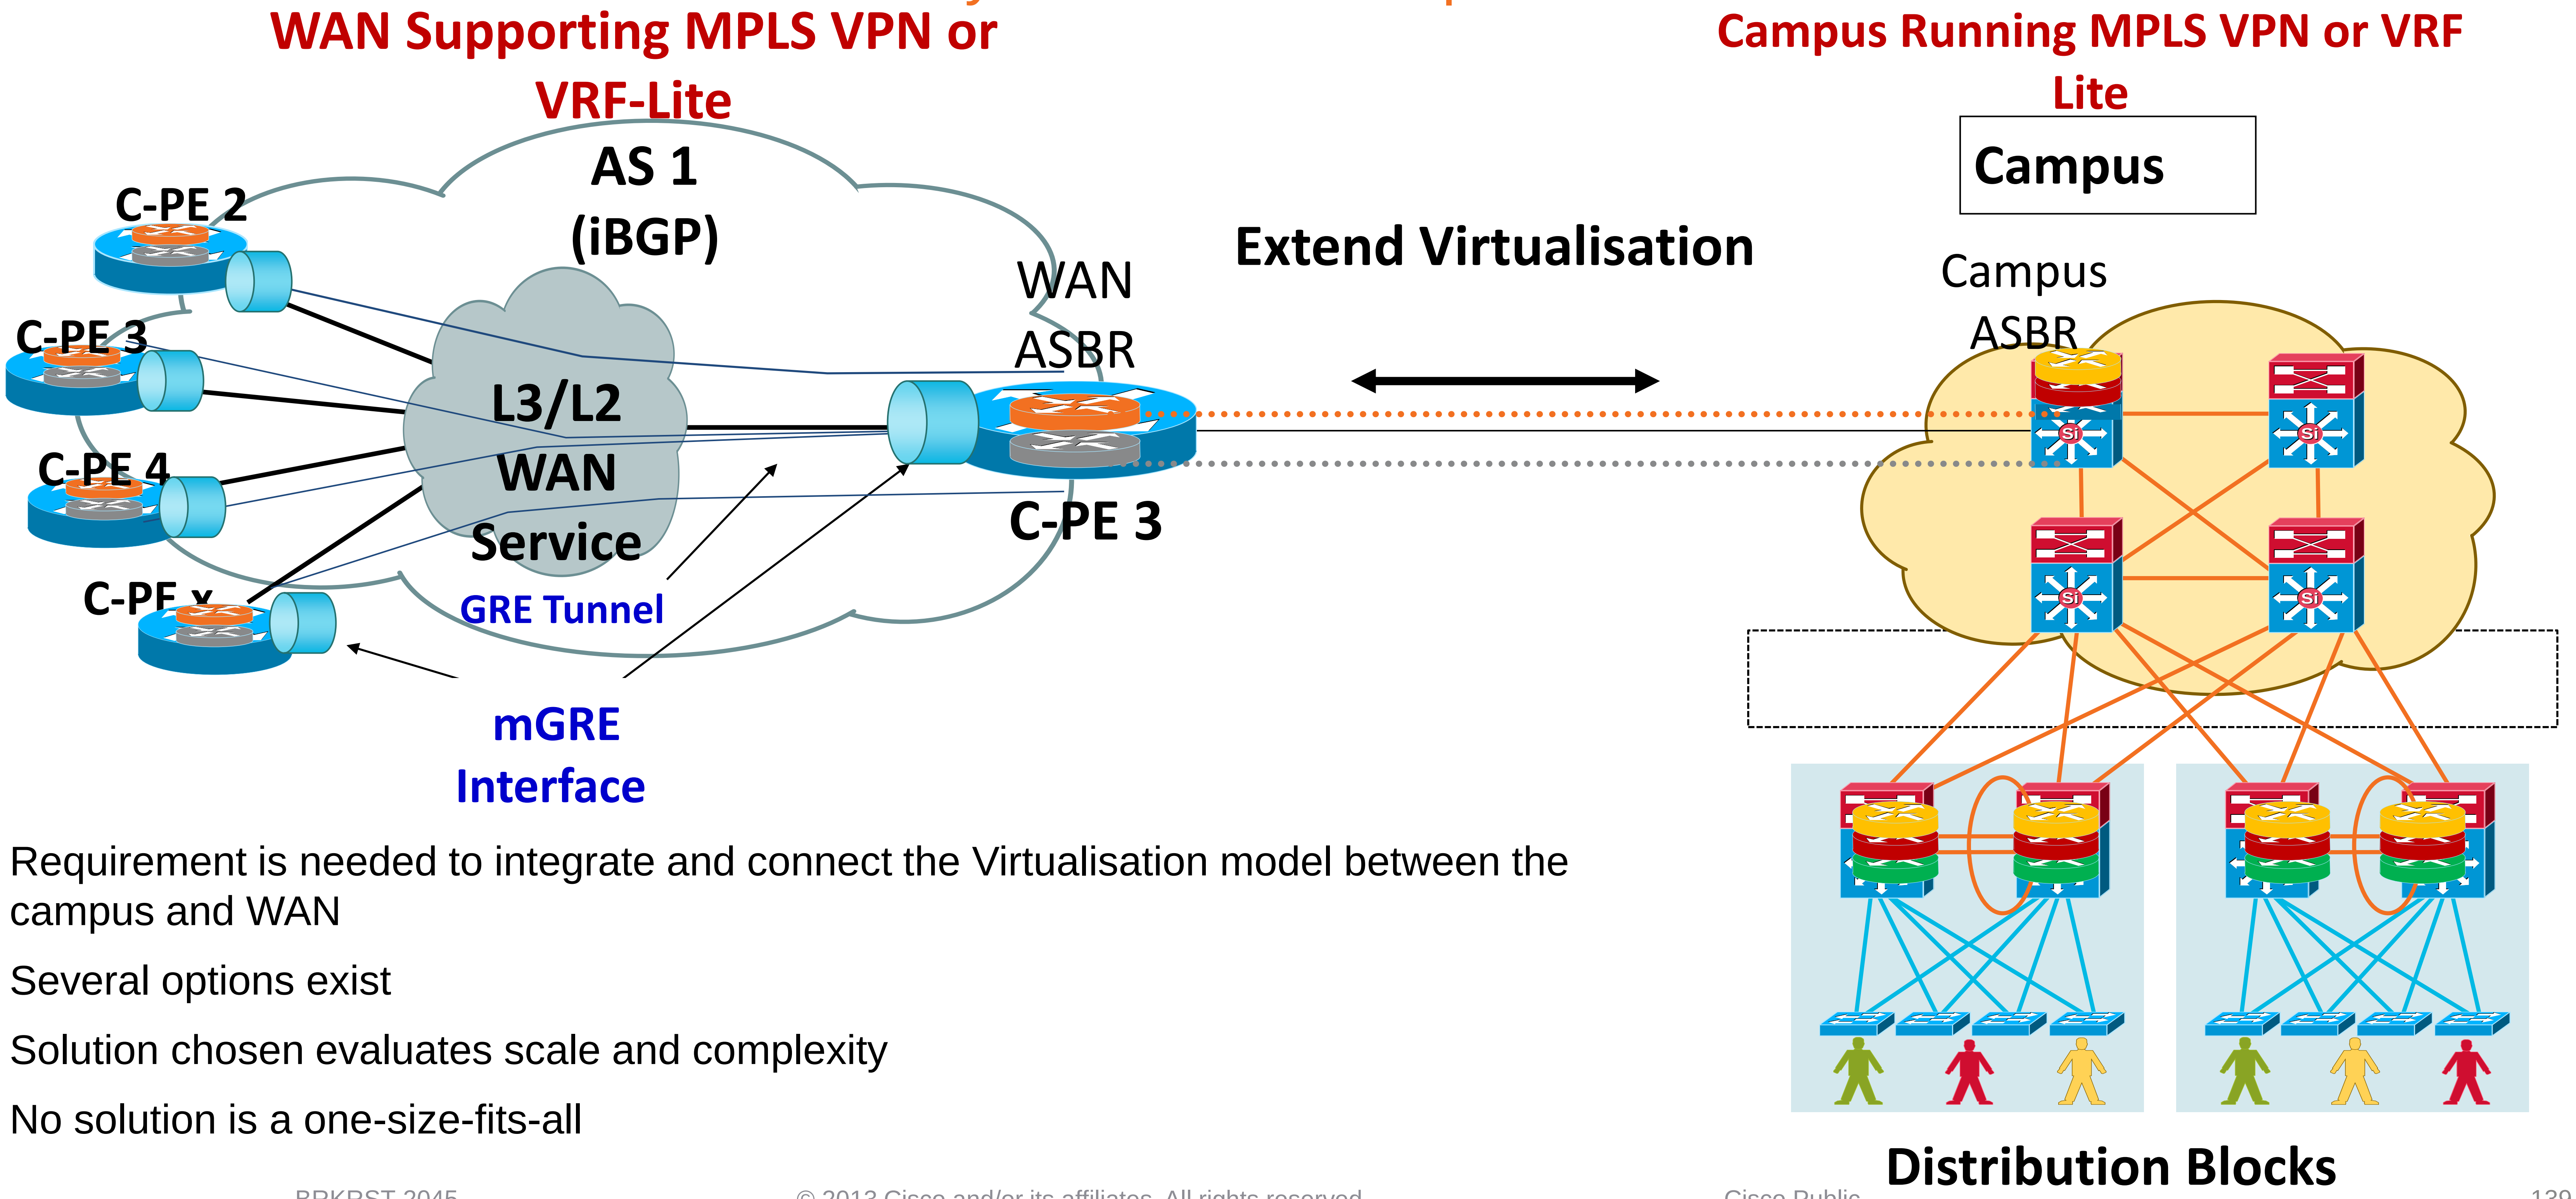

- Requirement is needed to integrate and connect the Virtualisation model between the campus and WAN
- Several options exist
- Solution chosen evaluates scale and complexity
- No solution is a one-size-fits-all

Campus-to-WAN Interconnection

Inter AS Option A (Back to Back VRFs)

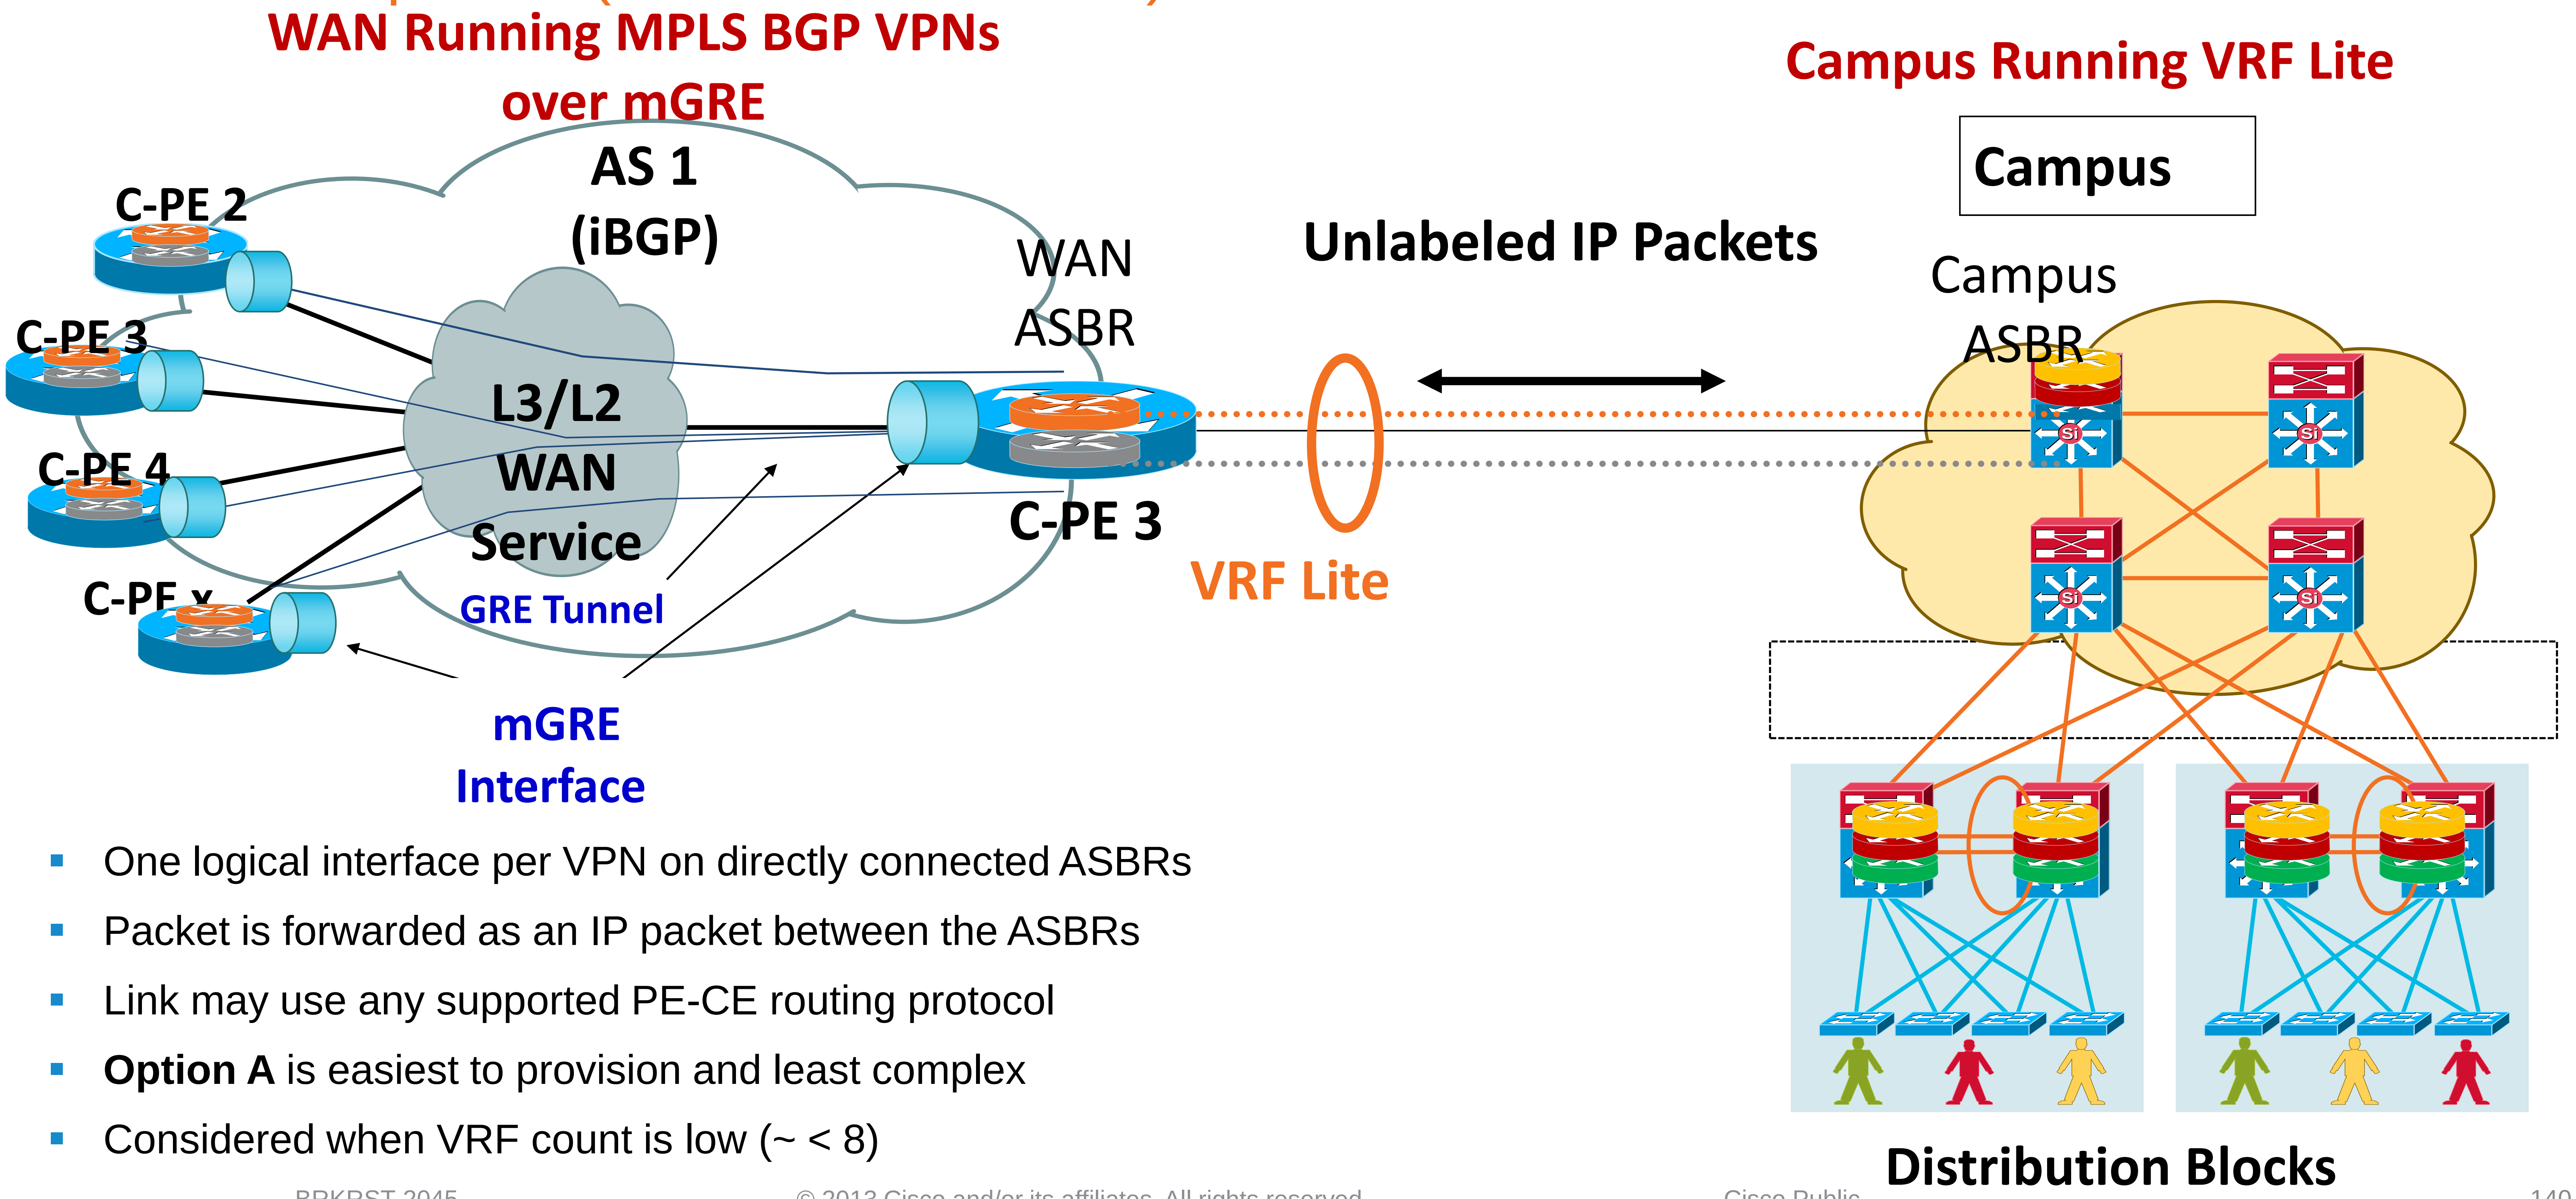

- One logical interface per VPN on directly connected ASBRs
- Packet is forwarded as an IP packet between the ASBRs
- Link may use any supported PE-CE routing protocol
- **Option A** is easiest to provision and least complex
- Considered when VRF count is low (~ < 8)

Campus-to-WAN Interconnection

Inter AS Option B (Medium/Large VRF Deployments)

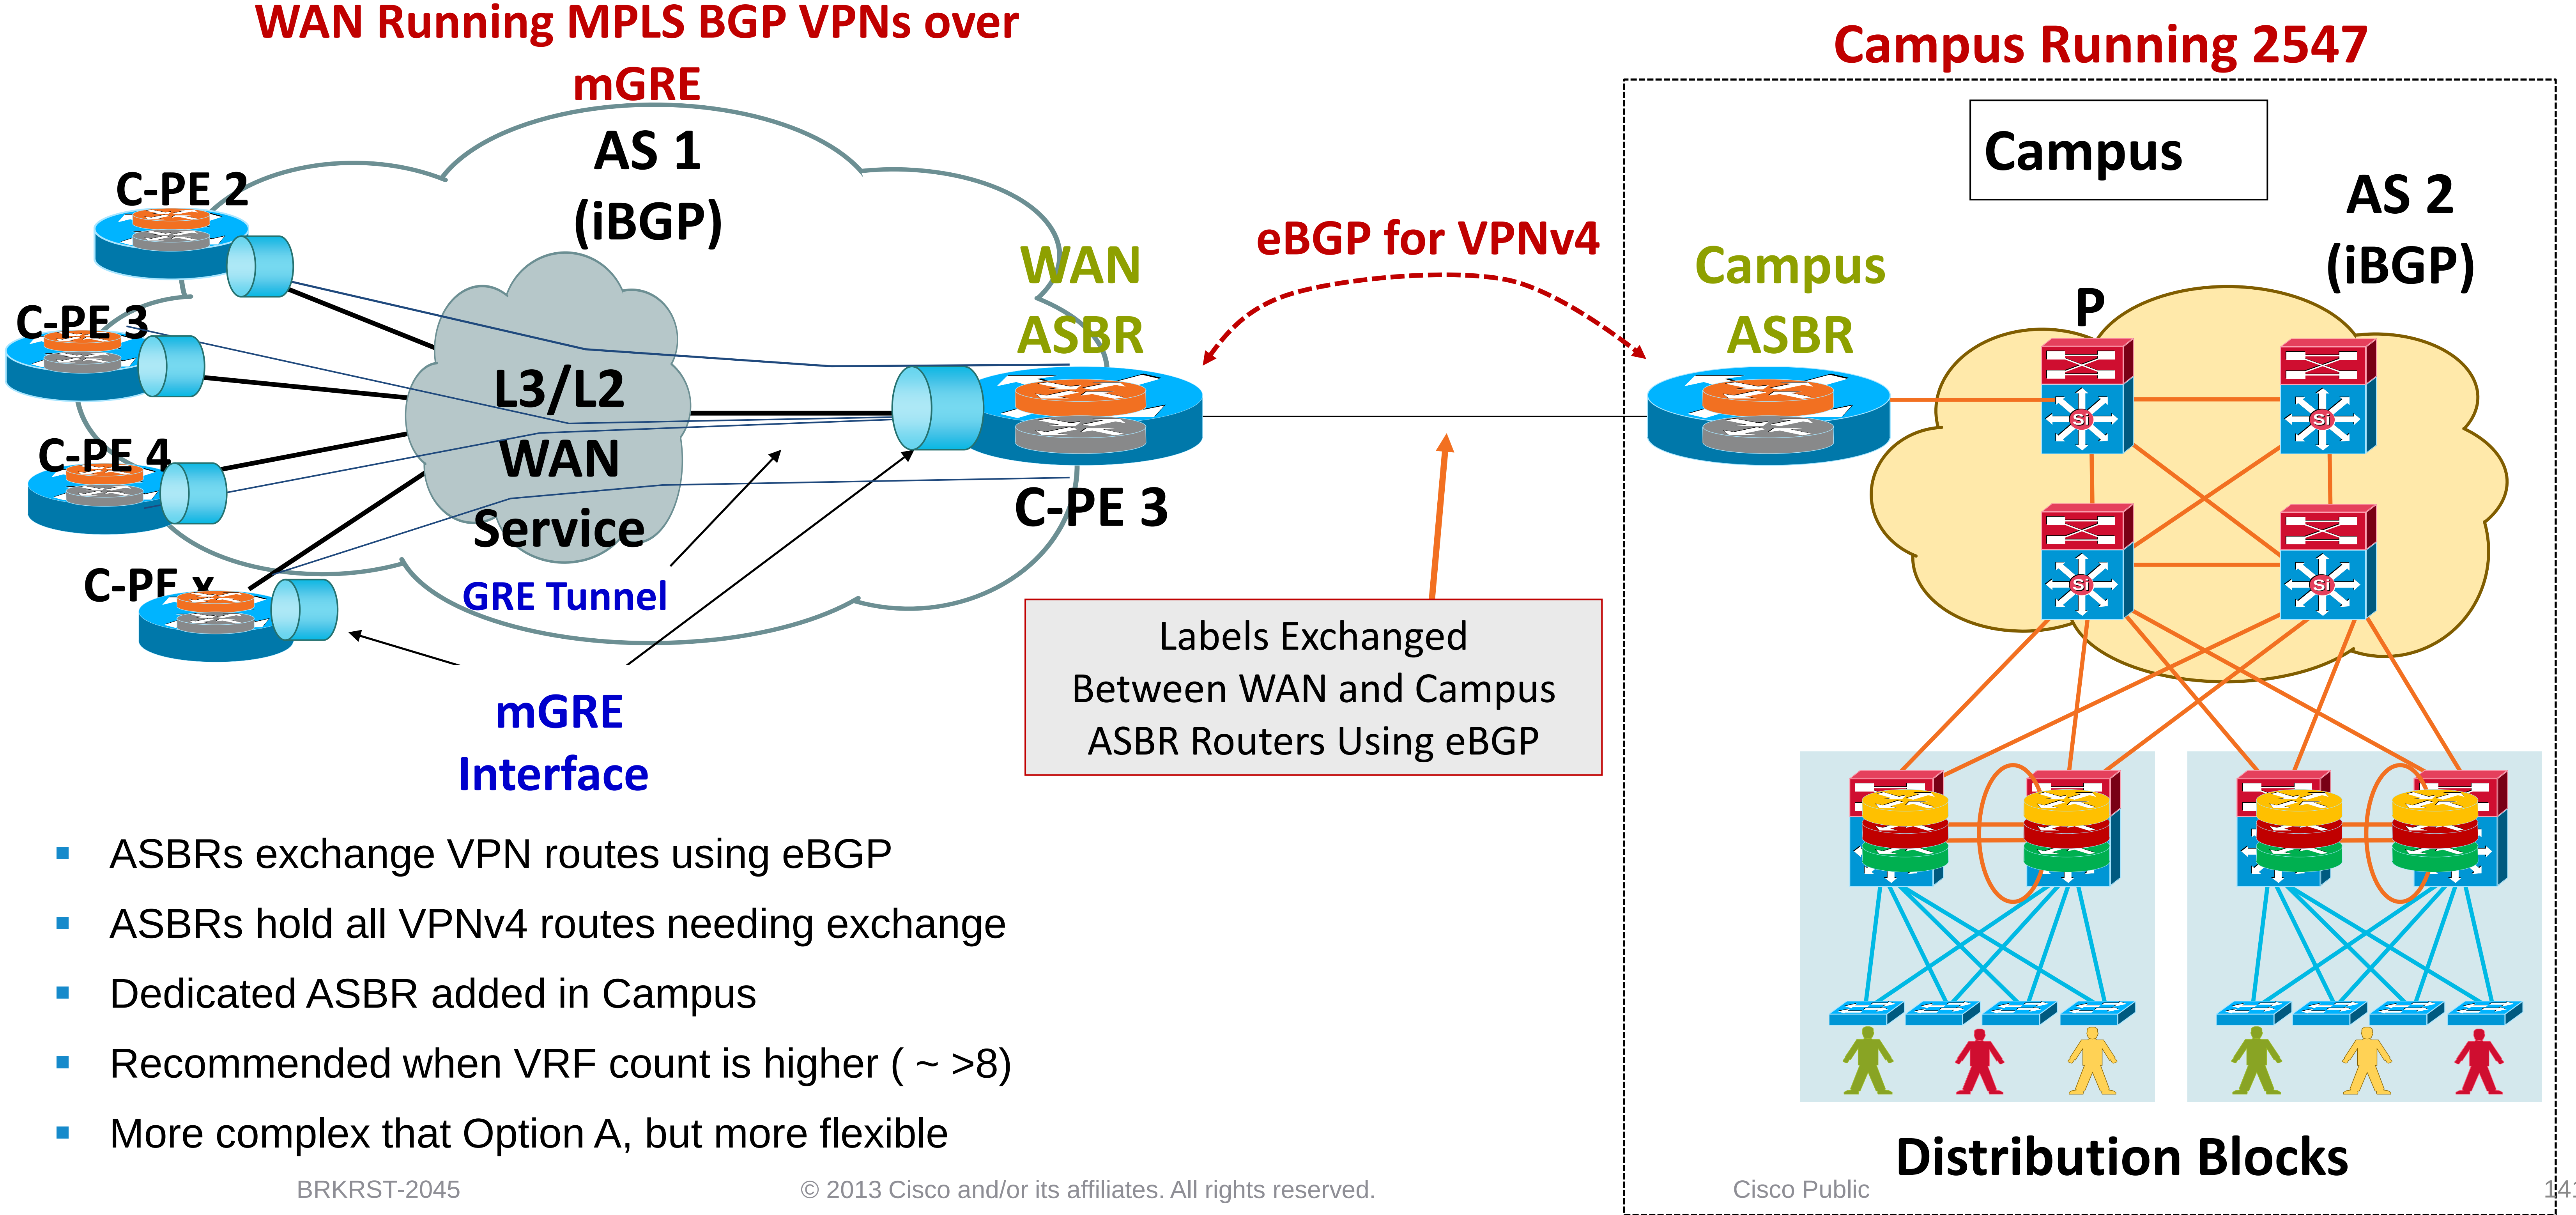

- ASBRs exchange VPN routes using eBGP
- ASBRs hold all VPNv4 routes needing exchange
- Dedicated ASBR added in Campus
- Recommended when VRF count is higher (~ >8)
- More complex that Option A, but more flexible

Campus/WAN Interconnect

Recommendations

- **< 8 VRFs** **Back to Back VRFs (Option A)**
  - VRF lite in the campus
  - Back to Back VRFs with a single AS between Campus and WAN
  - Low VRF count network-wide
- **~8 – 15 VRFs** **Back to Back or Inter AS (Option B)**
  - VRF-Lite or RFC 4364 running in the Campus
  - Dedicate ASBR router in the campus (Core router/switch) to peer to WAN
  - Solution choice dictated by customers operational expertise, change frequency
- **~ > 15 VRFs** **Inter-AS (Option B)**
  - RFC 4364 running in the Campus
  - Dedicate ASBR router in the campus (Core router/switch) to peer to WAN
  - Inter-AS option “B” recommended

WAN Extension Solution (i.e. Options Discussed in This Presentation) Could Also Dictate Choice of Inter-AS Solution

