The bridge to possible

Сетевой марафон Cisco: Классика WAN

День 1. Основные вопросы, возникающие при построении межфилиальной WAN сети. И правильные ответы на них.

Денис Коденцев
Старший Архитектор, CCIE
19 апреля 2021

О чем пойдет речь?

- Общая архитектура и дизайн WAN
- Варианты топологий и схем подключения WAN
- Вопросы шифрования в WAN
- Обеспечение высокой в WAN
- Сегментация в WAN
- QoS и Application Quality of Experience (AppQoE)

Why do we build Wide Area Networks?

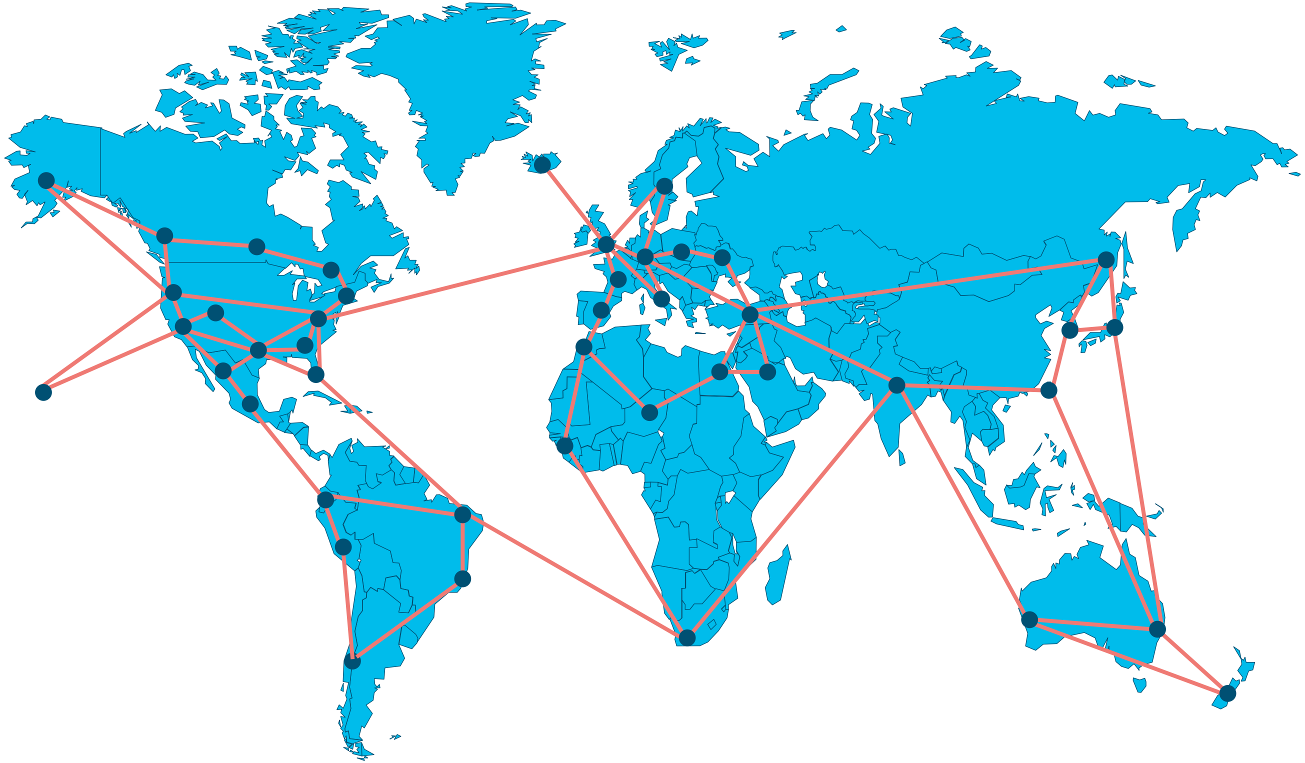

The WAN Technology Continuum

Early Networking

Flat/Bridged

Experimental Networks

Architectural Lessons

Protocols required for Scale & Restoration

Early-Mid 1990s

Multiprotocol

Business Enabling

Architectural Lessons

Route first, Bridge only if must

Mid 1990s-Late 2000s

Large Scale

Mission Critical

Architectural Lessons

Redundancy

Build to Scale

Today

Global Scale

IP Ubiquity

Cloud Connected

Planning

?

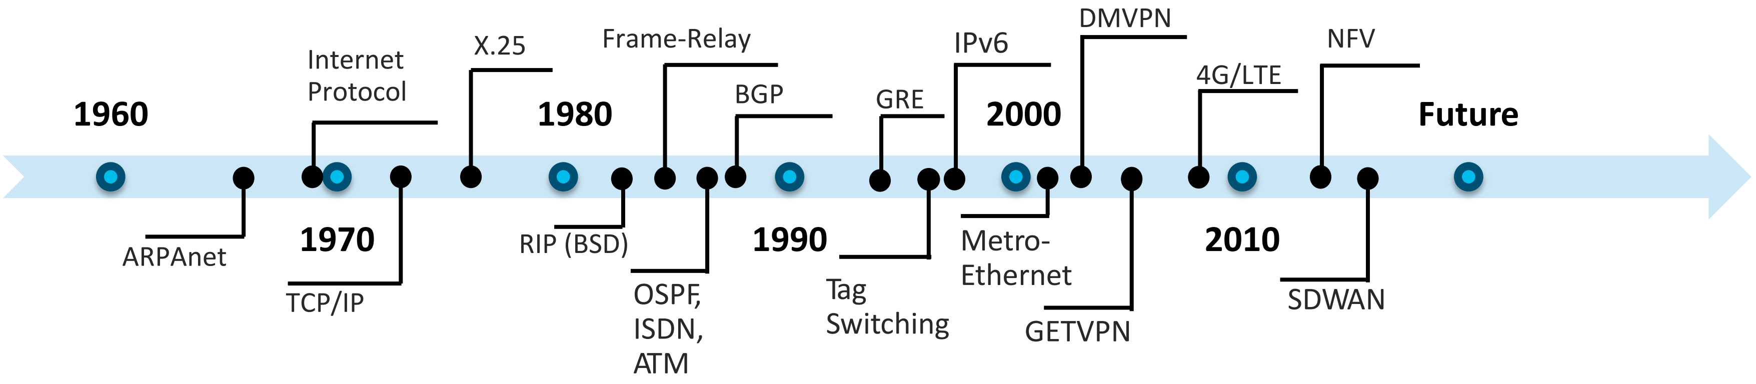

Delivering a positive application experience

Aspects relating to the Enterprise WAN

- Applications: Code usability and quality
- Security: must be foundational
- Operations: large scale, local and wide area networks, connected to Clouds
- Cloud: private or public application hosting
- Services: based on application requirements
- Metrics: must be able to measure the quality of the application experience

Physical Requirements and Constraints

- **Company Locations**
  - 10’s, 100’s, or 1000’s of sites
  - Where in the world
  - Site diversity
    - retail store, campus, large manufacturing plant, etc.
- **Topology Implications**
  - Single or dual connected
  - Geographical dispersity
    - Local, Regional, Global
  - Network role
    - Data Center, Colo Facility, Branch, Remote access, Public/Guest access
- **Operational requirements**
  - Access to resources
  - Transport options
  - Available power
  - Size and quantity of equipment
- **Risks associated with the Business and Technical requirements**

Technical Requirements and Constraints

- **Application requirements**
  - Bandwidth, Latency, Jitter
  - Connectivity and Protocols
    - L2 or L3, IPv4 or IPv6, Multicast,
- **Policy and Compliance**
  - Security
  - Segmentation
  - Encryption
- **Performance and Resiliency**
  - Quality-of-Experience
  - High Availability
    - Convergence and Recovery
    - Device quantities and capabilities
- **Existing Network Infrastructure**
  - Greenfield or Brownfield
  - Available documentation
  - Current designs and technologies

Business Requirements and Constraints

- **Business Environment**
  - Market transitions
  - Competitive pressures
  - Project goals
  - Mergers and acquisitions

- **Workforce Productivity**
  - User experience
  - Access to resources
  - Employee satisfaction

- **Costs**
  - OPEX and CAPEX
  - Lifecycle and ROI
  - IT Capabilities
  - Opportunity costs

- **Compliance and Policy**
  - Government and Industry Regulations
  - Security mandates
  - Reputation and perception

When designing a Wide Area Network…

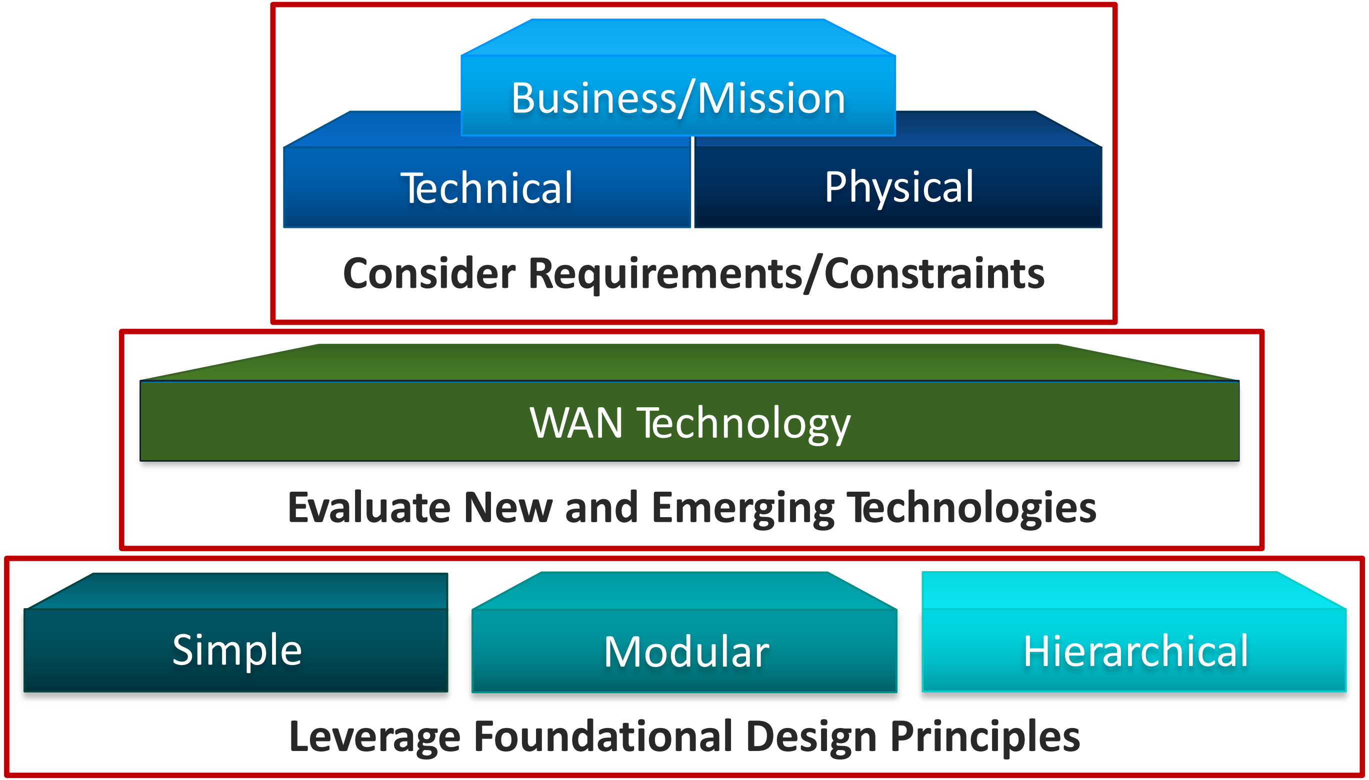

WAN Locations and Devices

- Organization sites
  - Headquarters Campus
  - Branch Office
  - Retail store
  - Factory, etc.
- Remote Access
  - Mobile workers
  - Home office
- Cloud
  - Private Data Center
  - Public IaaS
  - SaaS
  - Colocation Facility

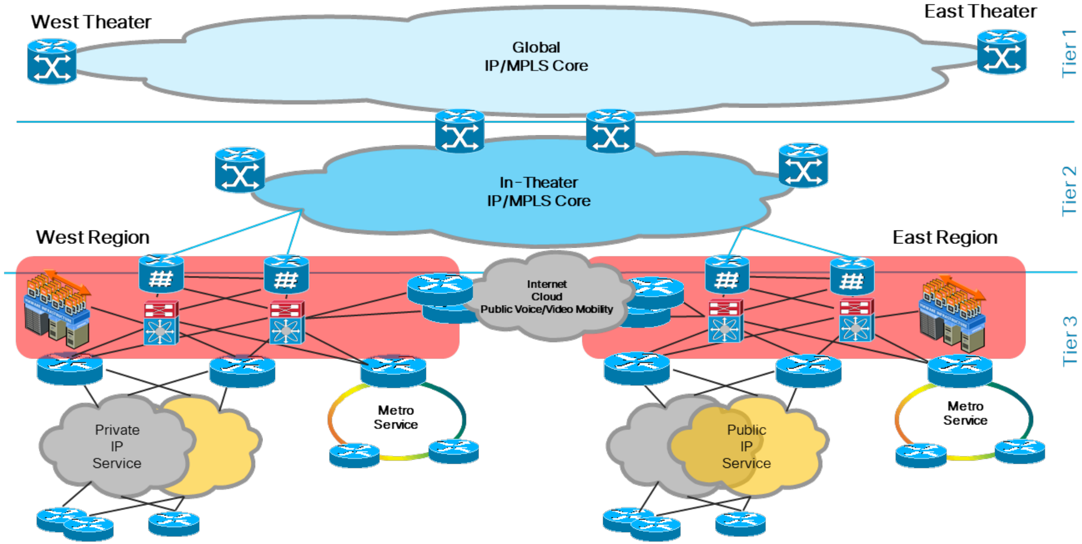

- Physical devices
  - Router/CPE
  - Firewall
  - Multi-purpose compute
  - Client devices
- Virtualized Network Functions
  - Virtual router
  - Virtual Firewall
  - etc…

Cisco offers the broadest portfolio for WAN transformation

Resilient Infrastructure to deliver IT agility

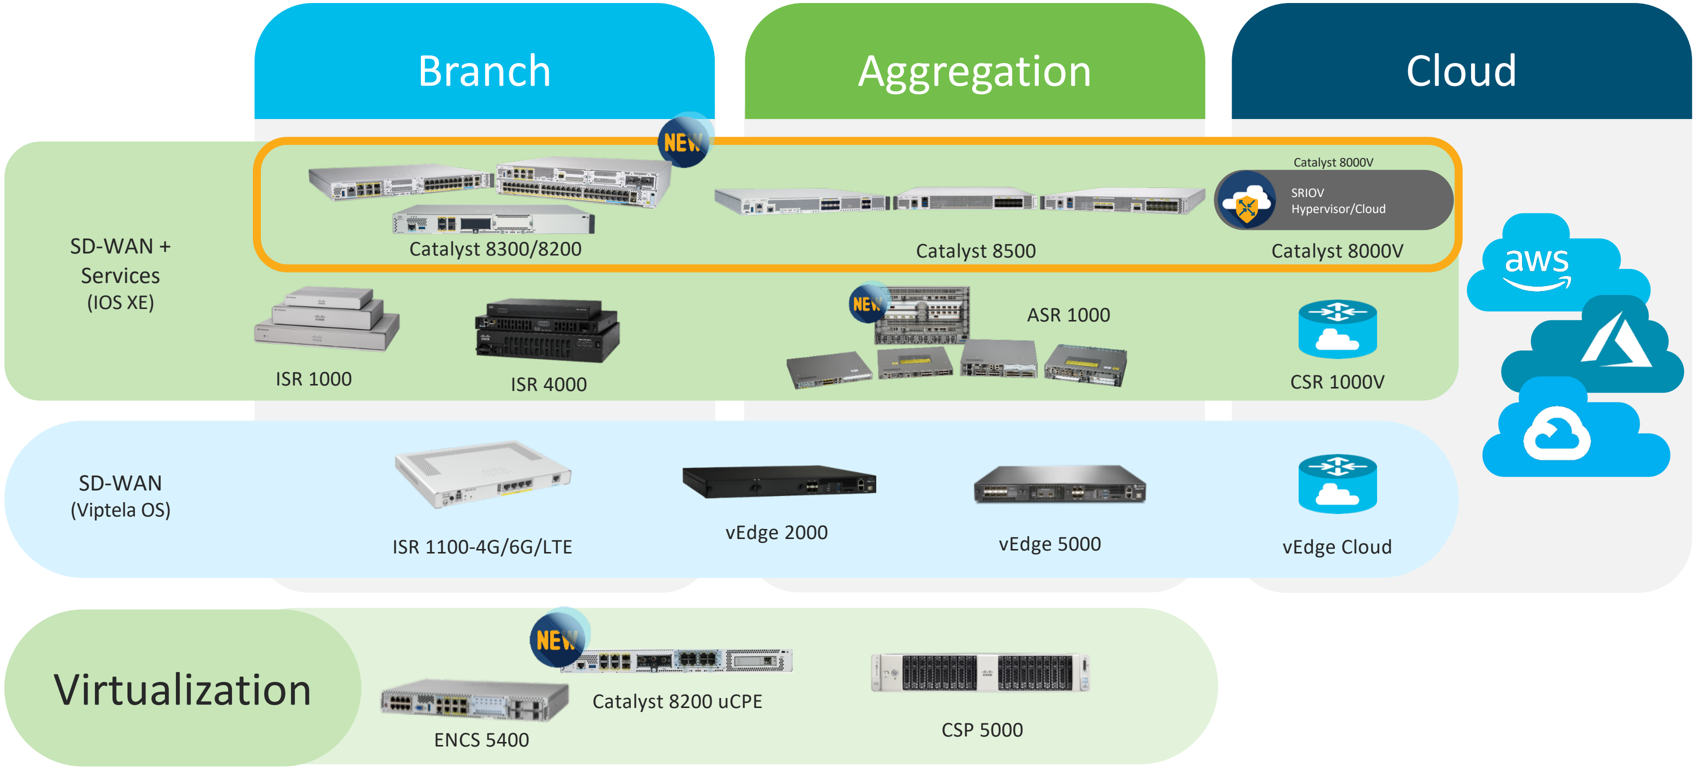

Топологии и схемы подключения

Hierarchical Network Design

Without a Rock Solid Foundation the Rest Doesn't Matter

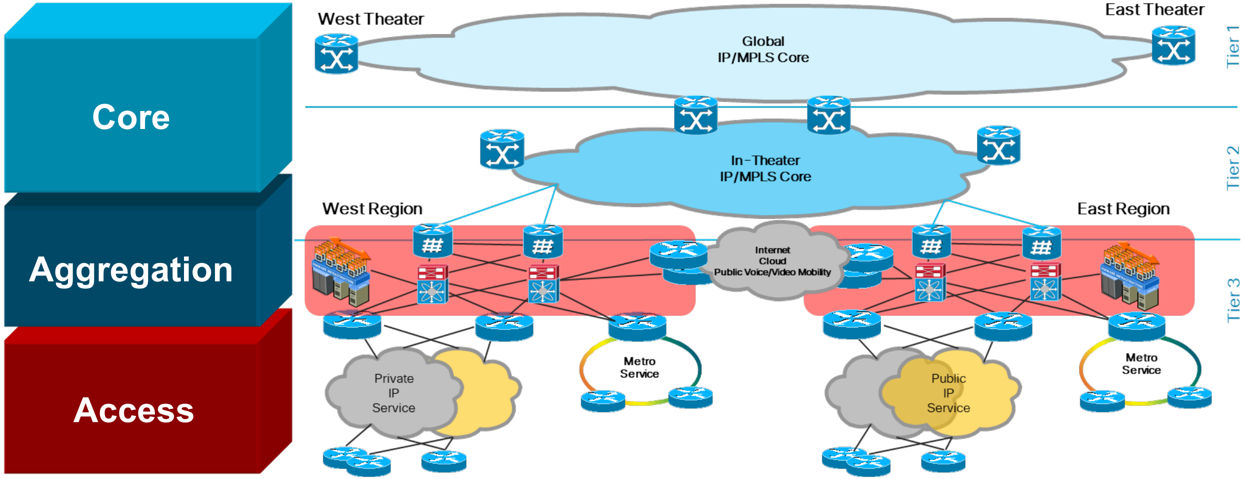

- Hierarchy—each layer has specific role
- Modular topology—building blocks
- Easy to grow, understand, and troubleshoot
- Creates small fault domains— clear demarcations and isolation
- Promotes load balancing and redundancy
- Promotes deterministic traffic patterns
- Incorporates balance of both Layer 2 and Layer 3 technology, leveraging the strength of both
- Utilizes Layer 3 routing for load balancing, fast convergence, scalability, and control

WAN Connection and Transport Technologies

- Dark Fiber
  - Highest flexibility, control, and security but only point-to-point connectivity
  - Most costly unless owned by the organization
- MPLS
  - Widely available service with flexible bandwidth options
  - Provider manages complex WAN routing with QoS SLAs
  - Offers simplicity with global scale if the organization can afford it
- Metro Ethernet
  - Layer 2 Ethernet connectivity service between up to hundreds of locations within a specific geographic region
  - Organization manages its own routing and QoS policies but may offer higher bandwidth at less cost than MPLS

- Broadband
  - Lower cost, high bandwidth Internet connectivity
  - Organization manages a secure overlay VPN between sites but has no control over latency or QoS
  - Available as wired (DSL, Cable) or wireless (3G/4G/5G or satellite)
- Legacy TDM
  - Last resort option
  - Cost comparable to Metro Ethernet but only 1.5Mbps bandwidth
  - Point-to-point layer 2 connectivity and requires non-Ethernet type port on router

MPLS VPN Models

CE = Customer Edge router

PE = Provider Edge router

Technology Options

MPLS Layer-2 VPNs

Point-to-Point Layer-2 VPNs

- CE connected to PE via L2 connection (Eth, FR, ATM, etc.)
- CE-CE L2 p2p connectivity
- CE-CE routing
- No SP involvement

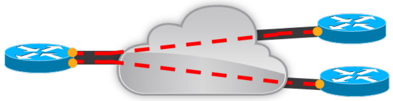

Multi-Point Layer-2 VPNs

- CE connected to PE via Ethernet connection
- CE-CE L2 (Eth) mp connectivity
- CE-CE routing
- No SP involvement

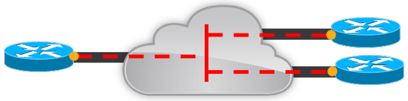

MPLS Layer-3 VPNs

- CE connected to PE via IP-based connection (over any layer-2 type)
  - Static routing
  - PE-CE routing protocol; eBGP, OSPF, IS-IS
- CE has peering relationship with PE
- PEs participate in customer routing
- PEs maintain customer-specific routing tables

Broadband Internet

- Widely available in wired or wireless
- Wired is generally an Ethernet handoff
- High bandwidth to the Internet so creates security vulnerability that must be managed
- Provides access to Public Cloud services such as IaaS and SaaS
- Does not support QoS or Multicast
- IPSec secure connections for private enterprise communication but this restricts some services
- Overlay IP Encapsulation with IPSec creates a secure VPN tunnel between Enterprise locations
- No service guarantee for critical applications but offers a low cost backup or bandwidth augmentation option

Wide Area Network Design Trends

Single Provider Design

- Enterprise will home all sites into a single carrier to provide L3 MPLS VPN connectivity.
- **Pro**: Simpler design with consistent features
- **Con**: Bound to single carrier for feature velocity
- **Con**: Does not protect against MPLS cloud failure with Single Provider

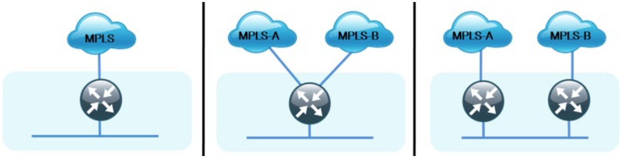

Dual Providers Design

- Enterprise will single or dual home sites into one or both carriers to provide L3 MPLS VPN connectivity.
- **Pro**: Protects against MPLS service failure with Single Provider
- **Pro**: Potential business leverage for better competitive pricing
- **Con**: Increased design complexity due to service implementation differences (e.g. QoS, BGP AS Topology)
- **Con**: Feature differences between providers could force customer to use least common denominator features.

Overlay Network Design

- Overlay tunneling technologies with encryption for provider transport agnostic design
- **Pro**: Can use commodity broadband services for lower cost higher bandwidth service
- **Pro**: Flexible overlay network topology that couples from the physical connectivity
- **Con**: Increased design complexity
- **Con**: Additional technology needed for SLA over commodity transport services

Wide Area Network Design Trends (cont.)

- Hybrid and Overlay Designs
  - Tunneling/encryption enables transport agnostic design
    - + On-demand or permanent backup links
    - + Commodity broadband services offer lower cost, higher bandwidth
    - + Flexible overlay topology independent of physical underlay connectivity
    - – Two “layers” to support
    - – SLA over commodity transport services
    - – Must consider potential for fragmentation

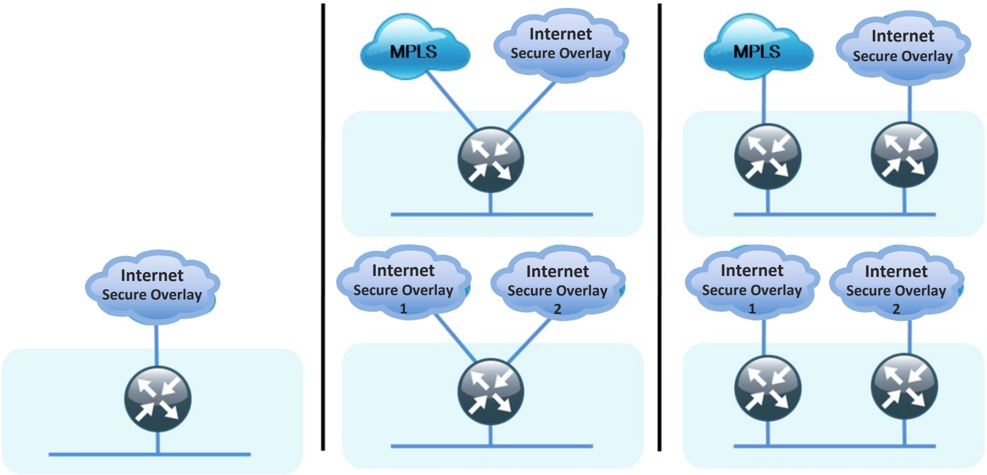

Types of Overlay Service

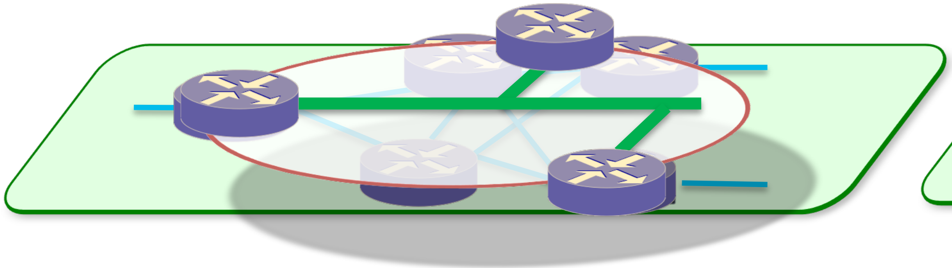

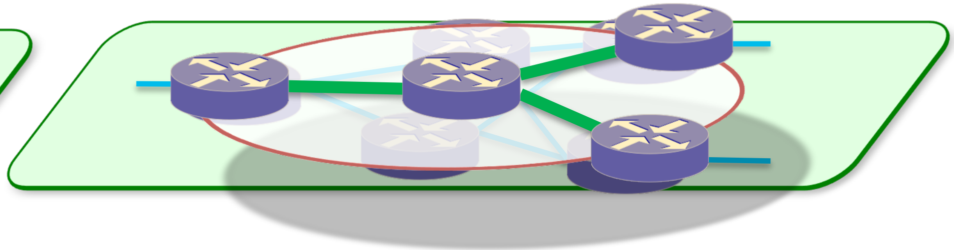

Layer 2 Overlays

- Virtual Extensible LAN (VXLAN)
  - MAC-in-UDP encapsulation
  - 24-bit segment ID for up to 16M logical networks
- Other L2 overlay technologies
  - MPLS-over-GRE/mGRE, L2TPv3, OTV

Layer 3 Overlays

- IPSec—Encapsulating Security Payload (ESP)
  - Strong encryption
  - IP Unicast only
- Generic Routing Encapsulation (GRE)
  - IP Unicast, Multicast, Broadcast
  - Multiprotocol support
- Other L3 overlay technologies
  - MPLS-over-GRE/mGRE, LISP

GRE and IPSec Overlay Encapsulation Example

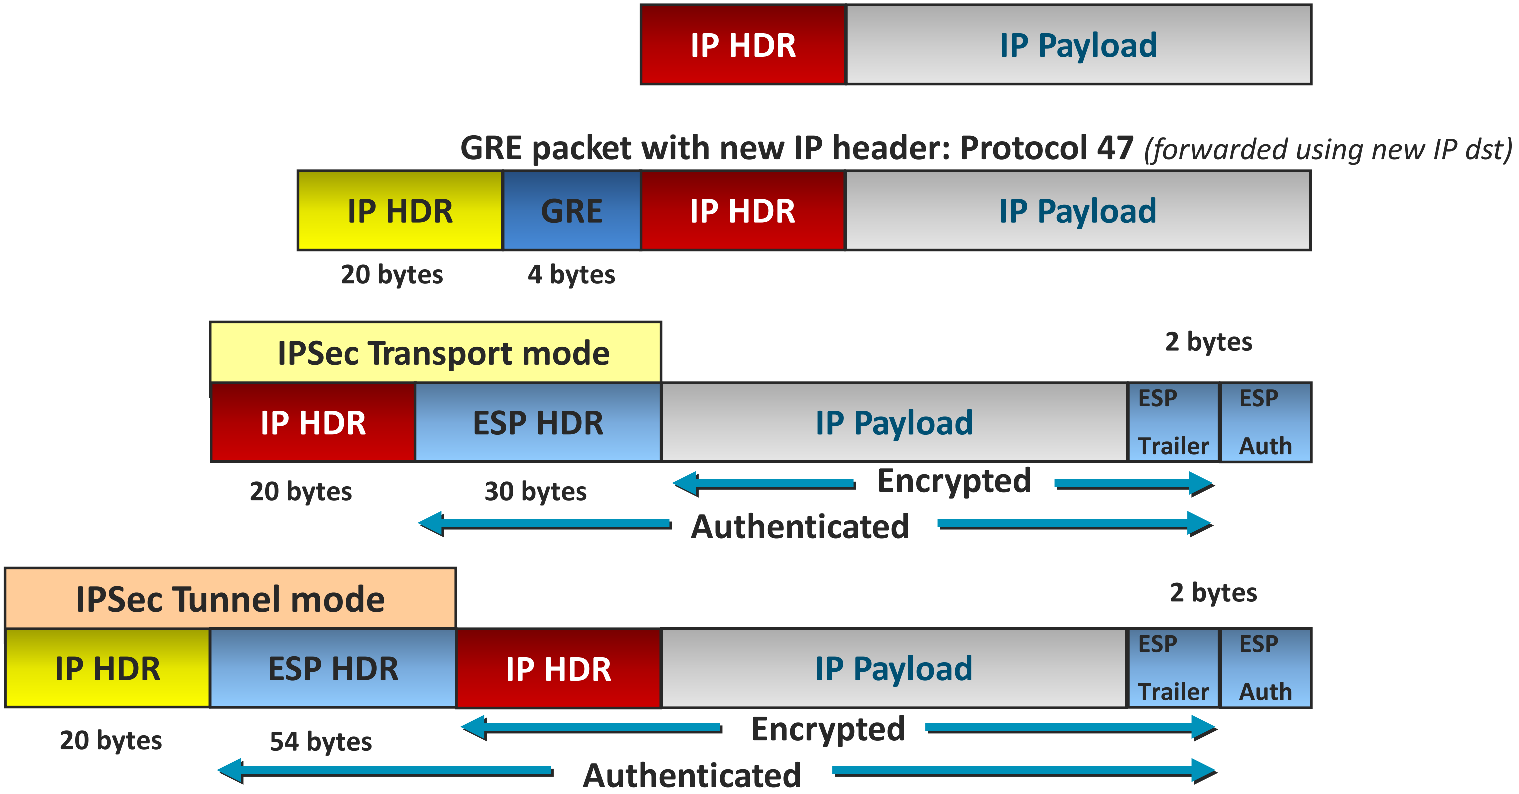

Dynamic Multipoint VPN (DMVPN)

- Branch spoke sites establish an IPsec tunnel to and register with the hub site
- IP routing exchanges prefix information for each site
- BGP or EIGRP are typically used for scalability
- **With WAN interface IP address as the tunnel source address, provider network does not need to route customer internal IP prefixes**
- Data traffic flows over the DMVPN tunnels
- When traffic flows between spoke sites, the hub assists the spokes to establish a site-to-site tunnel
- **Per-tunnel QOS is applied to prevent hub site oversubscription to spoke sites**

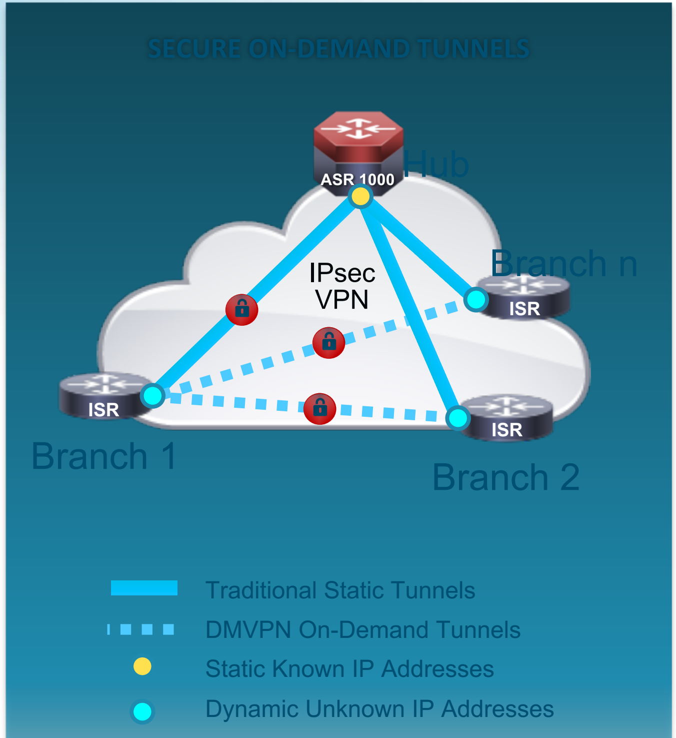

Any-to-Any Encryption

Before and After GETVPN

Public/Private WAN

Before: IPSec P2P Tunnels

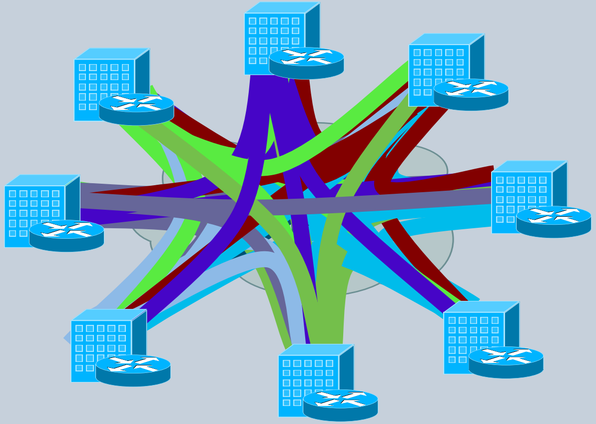

- Scalability—an issue (N^2 problem)
- Overlay routing
- Any-to-any instant connectivity can't be done to scale
- Limited QoS
- Inefficient Multicast replication

Private WAN

After: Tunnel-Less VPN

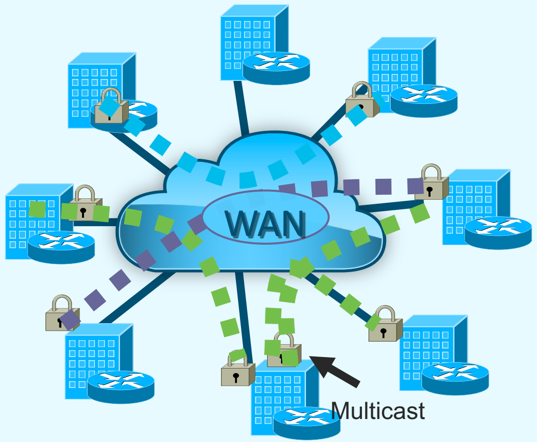

- Scalable architecture for any-to-any connectivity and encryption
- No overlays—native routing
- Any-to-any instant connectivity
- Enhanced QoS
- Efficient Multicast replication

Group Security Functions

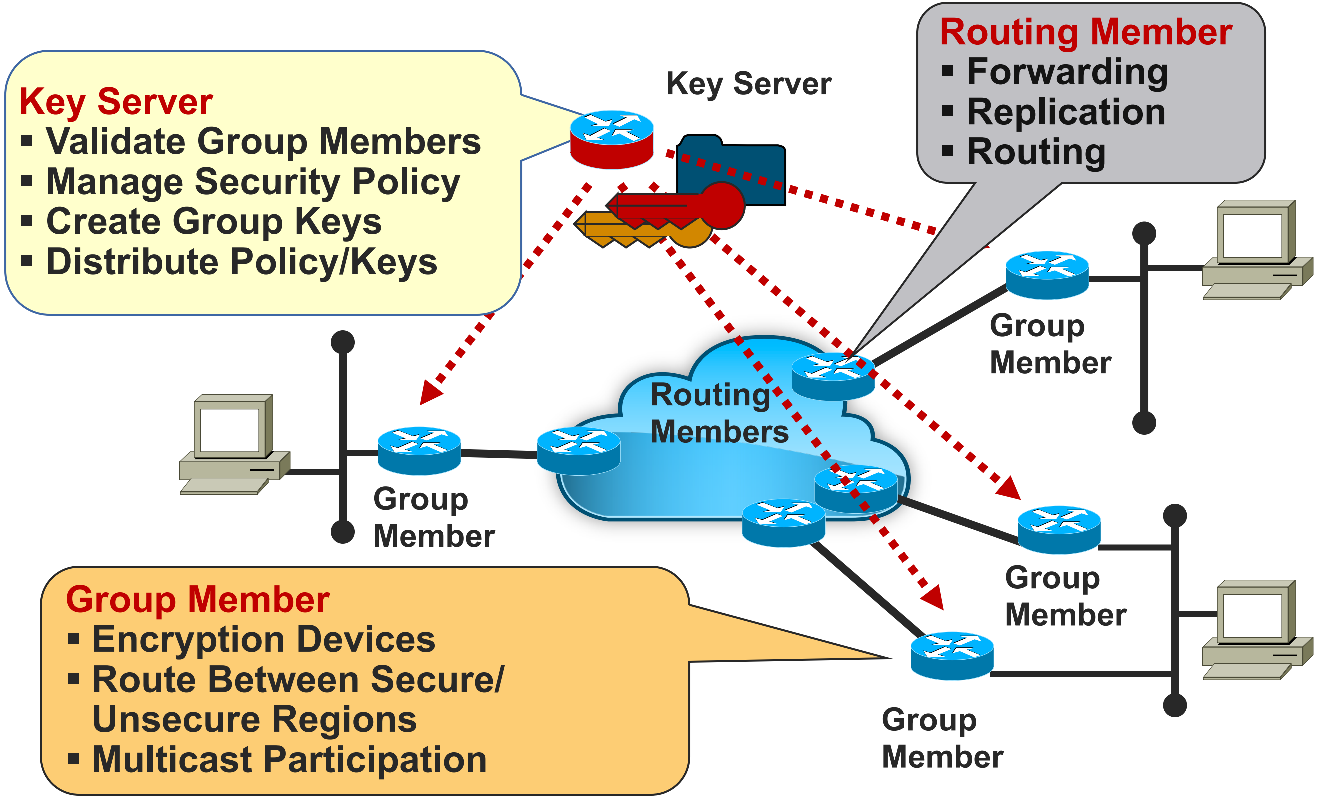

GETVPN - Group Key Technology Operation Example

- **Step 1**: Group Members (GM) “register” via GDOI (IKE) with the Key Server (KS)
  - KS authenticates and authorizes the GM
  - KS returns a set of IPsec SAs for the GM to use

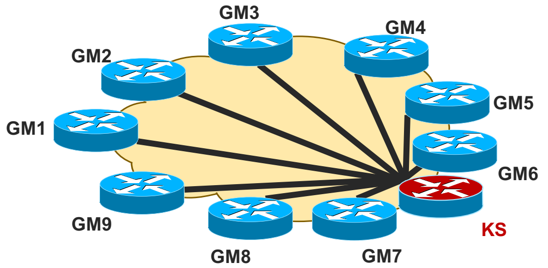

- **Step 2**: Data Plane Encryption
  - GM exchange encrypted traffic using the group keys
  - The traffic uses IPSec Tunnel Mode with “address preservation”

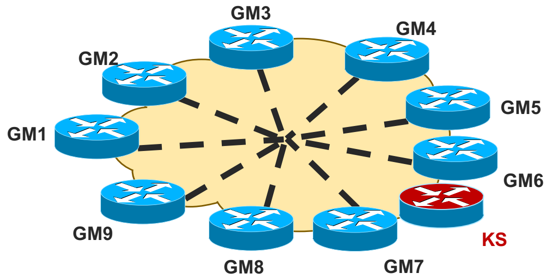

- **Step 3**: Periodic Rekey of Keys
  - KS pushes out replacement IPsec keys before current IPsec keys expire; This is called a “rekey”

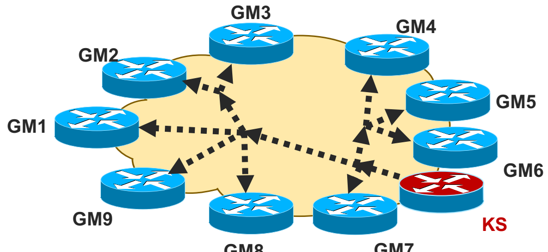

Legacy IPsec VPN Technologies Comparison

Features	DMVPN	FlexVPN	GET VPN
Infrastructure Network	▪ Public or Private Transport ▪ Overlay Routing ▪ IPv4/IPv6 dual Stack	▪ Public or Private Transport ▪ Overlay Routing	▪ Private IP Transport ▪ Flat/Non-Overlay IP Routing
Network Style	▪ Large Scale Hub and Spoke with dynamic Any-to-Any	▪ Converged Site to Site and Remote Access	▪ Any-to-Any; (Site-to-Site)
Failover Redundancy	▪ Active/Active based on Dynamic Routing	▪ Dynamic Routing or IKEv2 Route Distribution ▪ Server Clustering	▪ Transport Routing ▪ COOP Based on GDOI
Scalability	▪ Unlimited ▪ 3000+ Client/Server	▪ Unlimited ▪ 3000+ Client/Server	▪ 8000 GM total ▪ 4000 GM/KS
IP Multicast	▪ Multicast replication at hub	▪ Multicast replication at hub	▪ Multicast replication in IP WAN network
QoS	▪ Per Tunnel QoS, Hub to Spoke	▪ Per SA QoS, Hub to Spoke ▪ Per SA QoS, Spoke to Spoke	▪ Transport QoS
Policy Control	▪ Locally Managed	▪ Centralized Policy Management	▪ Central or Local Management
Technology	▪ Tunneled VPN ▪ Multi-Point GRE Tunnel ▪ IKEv1 & IKEv2	▪ Tunneled VPN ▪ Point to Point Tunnels ▪ IKEv2 Only	▪ Tunnel-less VPN ▪ Group Protection ▪ IKEv1 & IKEv2

Common WAN Topologies

Design and Deployment Considerations

Design Challenges with Growing Needs and New Innovation

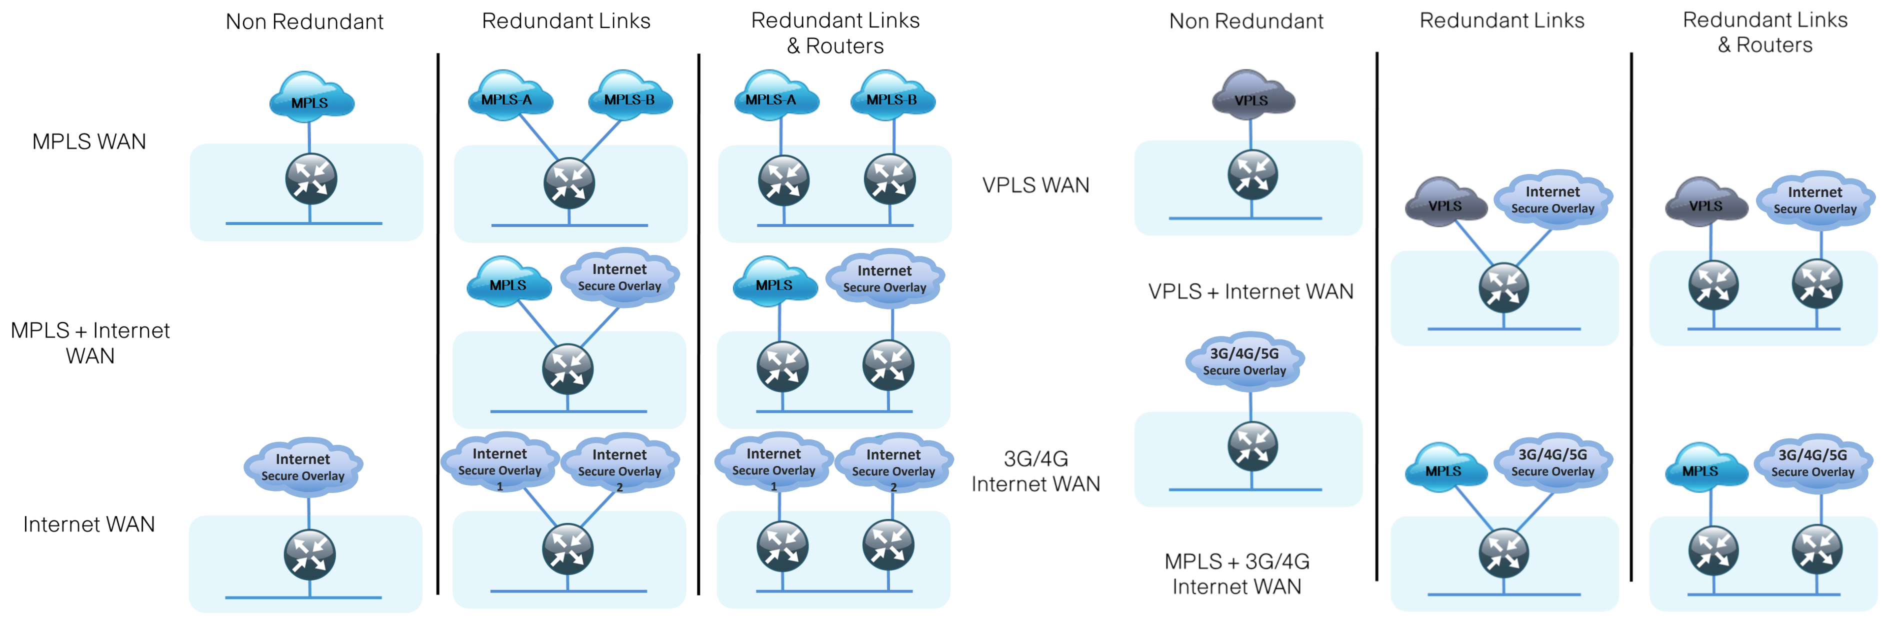

Modern Hierarchical Global WAN Design

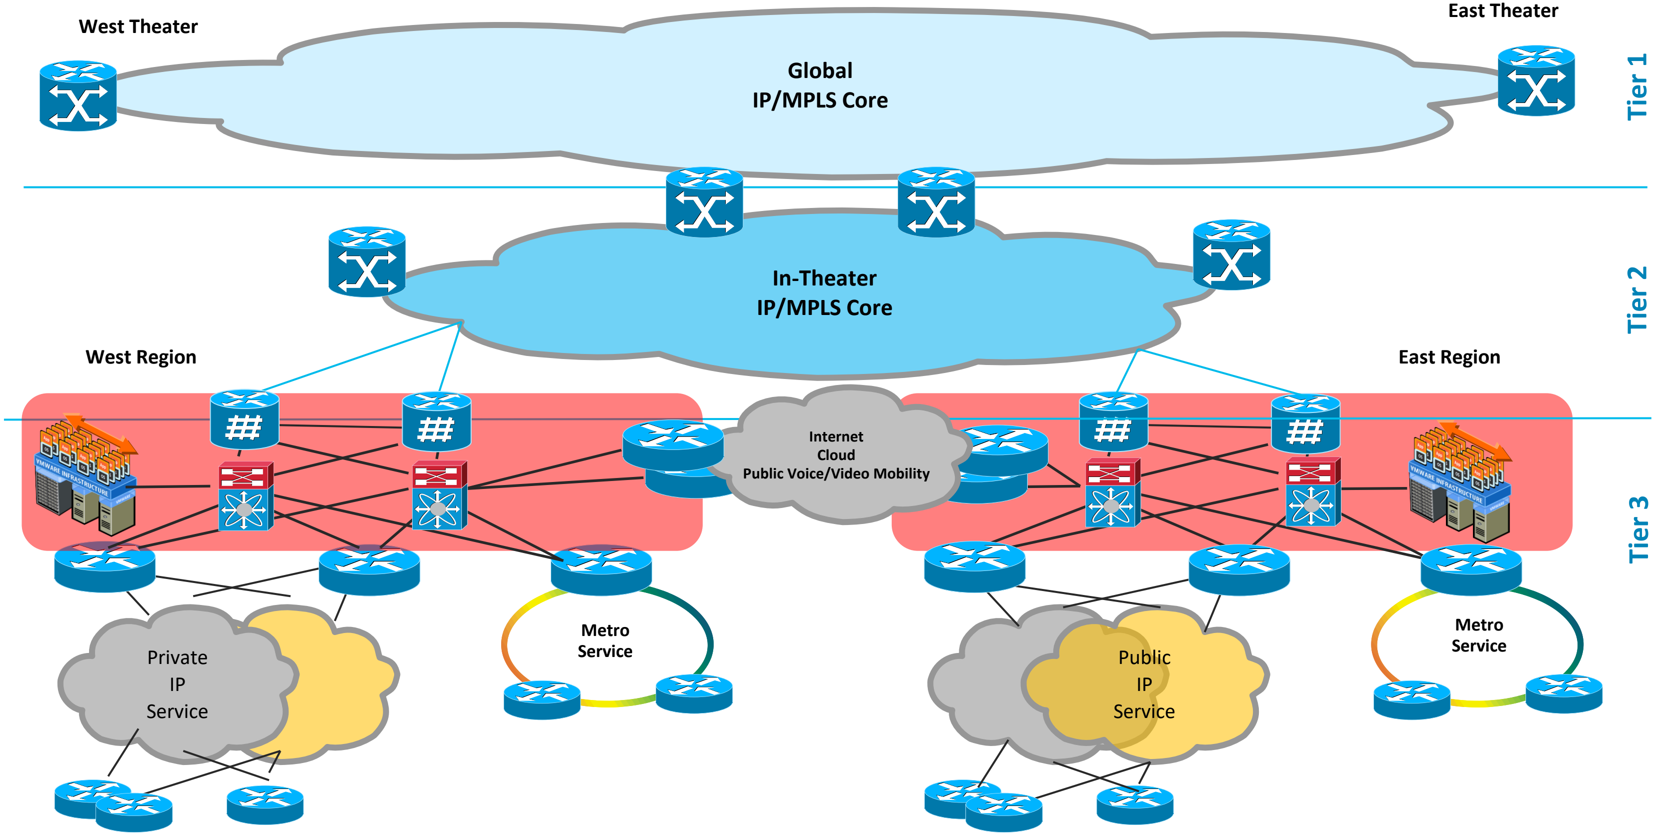

Modern Hierarchical Global WAN Design

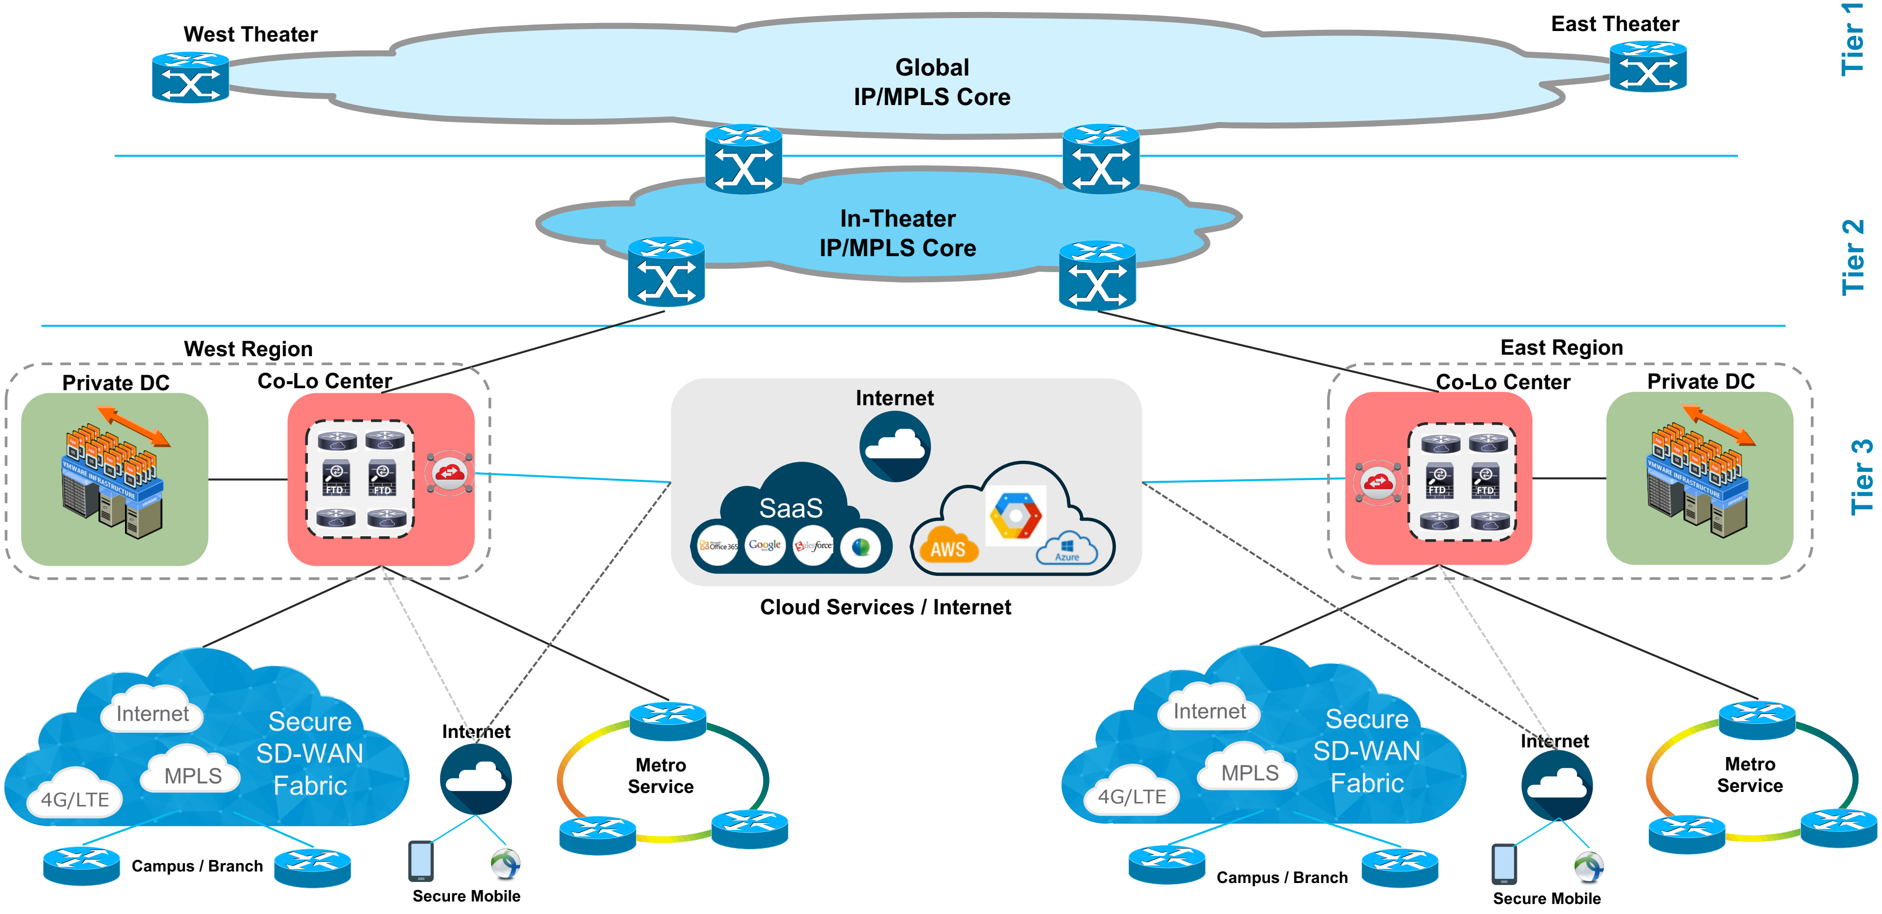

The WAN Technology Continuum

Early Networking	Early-Mid 1990s	Mid 1990s-Late 2000s	Today
Flat/Bridged Experimental Networks	Multiprotocol Business Enabling	Large Scale Mission Critical	Global Scale IP Ubiquity Cloud Connected
Architectural Lessons	**Architectural Lessons**	**Architectural Lessons**	**Architectural Lessons**
Protocols required for Scale & Restoration	Route first, Bridge only if must	Redundancy Build to Scale	**Optimize for application experience** **SDN delivers agility** **Central policy enforcement**

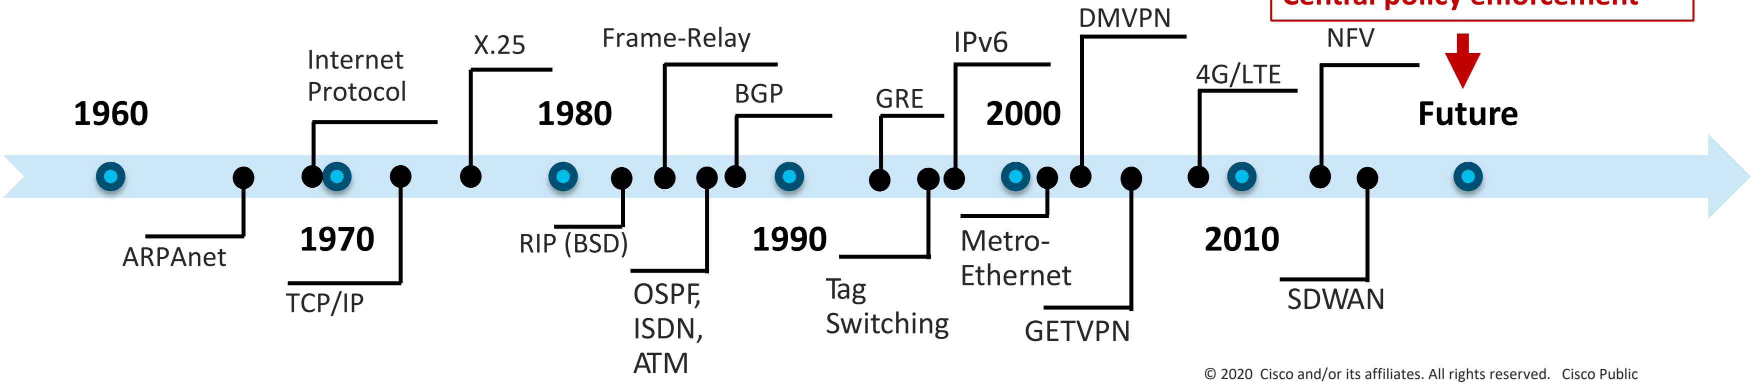

Шифрование

Link Speeds Out-Pacing IP Encryption

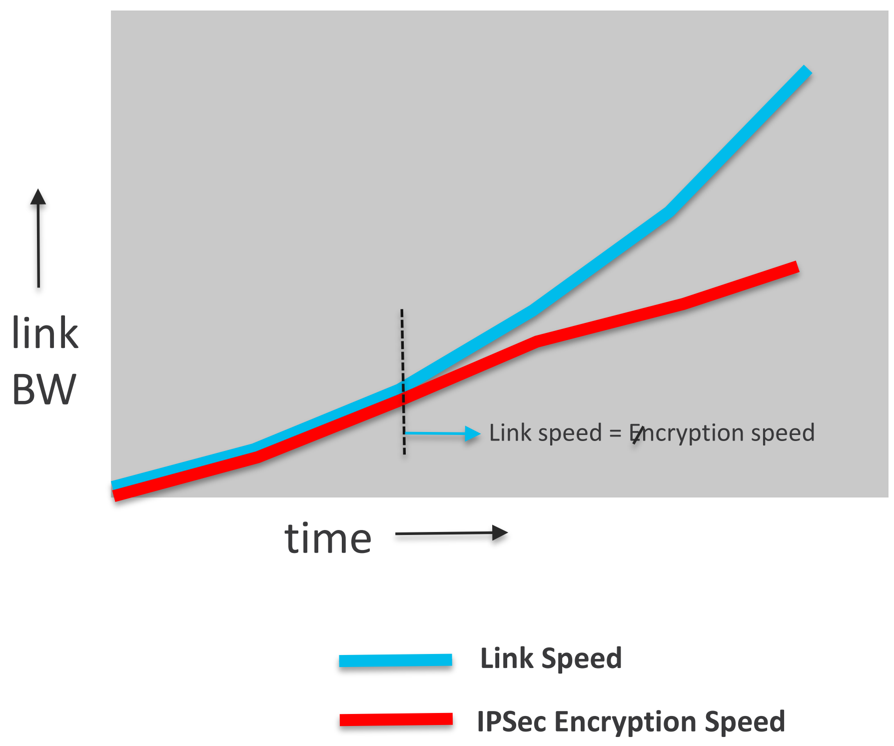

- Bandwidth application requirements out-pacing IP encryption capabilities
- Bi-directional and packet sizes further impact encryption performance
- IPSec engines dictate aggregate performance of the platform (much lower throughput)
- Cost per bit for IPSec much more expensive
- Encryption must align with link speed (100G+) to support next-generation applications

What is MAC Security (MACsec)?

Hop-by-Hop Encryption via IEEE 802.1AE

- Hop-by-Hop Encryption model
  - -Packets are decrypted on ingress port
  - -Packets are in the clear in the device
  - -Packets are encrypted on egress port
- Supports 1/10G, 40G, 100G encryption speeds
- Data plane (IEEE 802.1AE) and control plane (IEEE 802.1x-Rev)
- Transparent to IPv4/v6, MPLS, multicast, routing
- Encryption aligns with Link PHY speed (Ethernet)

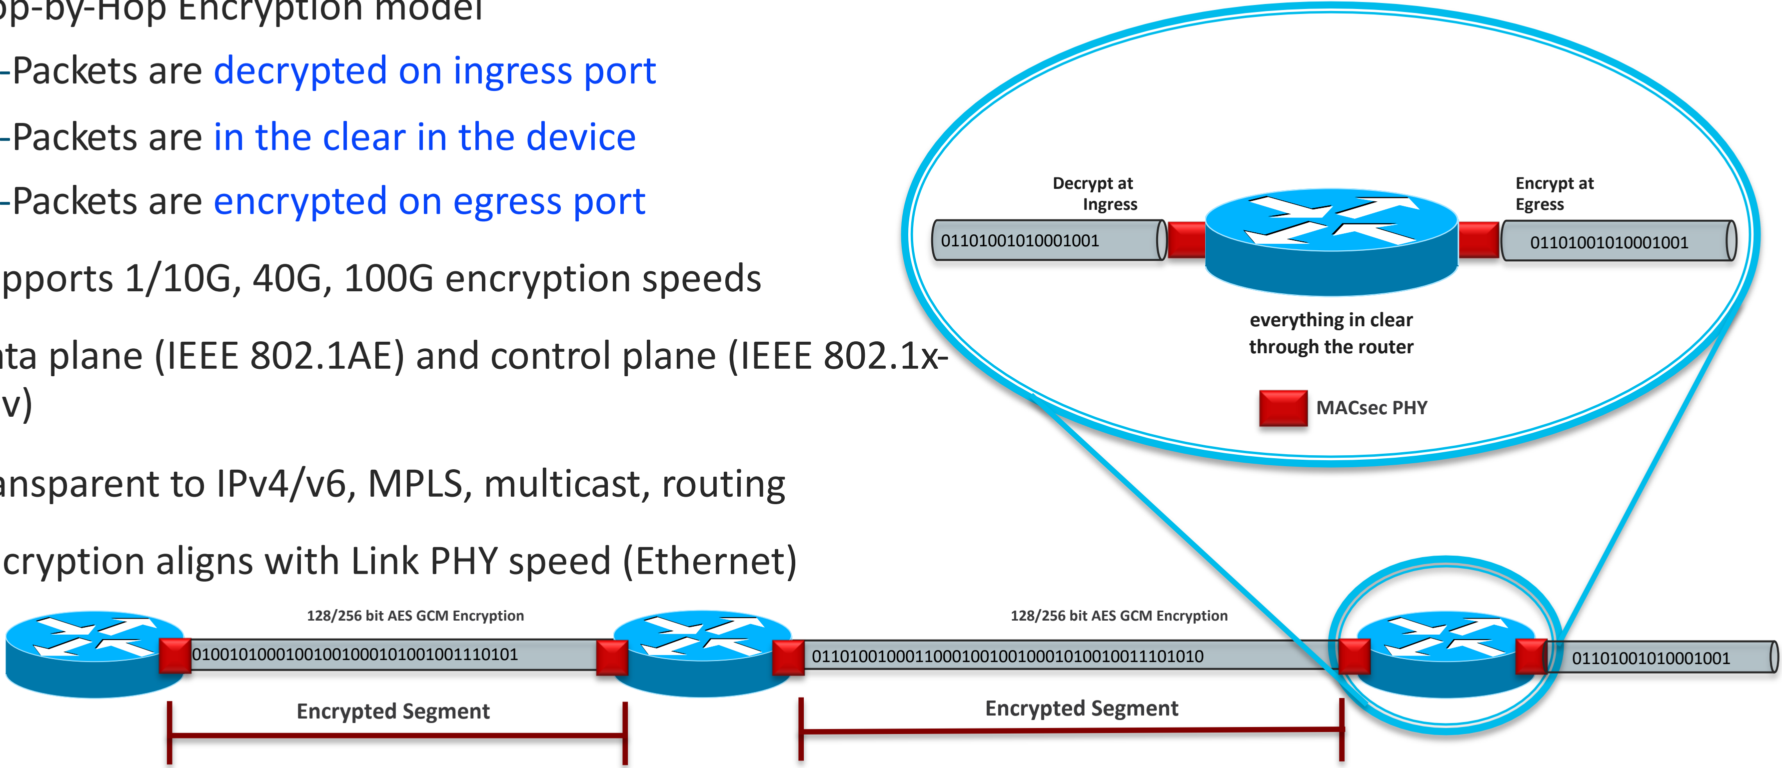

MACSEC vs IPSEC

Category	MACsec	IPsec
Market Positioning	1. Aggregate Deployments such as Regional Hubs 2. Large Branches that require high throughput 3. Data Center Interconnects	1. Small Branches 2. High Scale deployments 3. Low throughput Branches 4. Beyond MetroE (International) Reach
Link Requirement/Topologies	Requires dedicated MetroE EVC circuits for L2 connectivity between sites Point-to-Point, Point-to-MultiPoint	Easily Routable over many commonly available public network Any Topology
Encryption Performance	Per PHY Link Speed (1G, 10G, 40G, 100G)	Constrained by IPsec Crypto engine performance
Services Enablement	No impact to encryption throughput	Impacts encryption throughput
Peers Scale	Limited by hardware resources	Highly Scalable
Throughput	Up to Line Rate on each port (limited only by the forwarding capability)	Aggregate throughput (limited by the encryption throughput)
Configurability	Simple configuration	More complex configuration and policy choices
Layer 3 Visibility for Monitoring	No. Except Layer 2 headers (and optionally VLAN/MPLS Labels) everything else is encrypted	Visible. L3 info can be used for monitoring & policy enforcement purposes

What is “WAN MACsec?

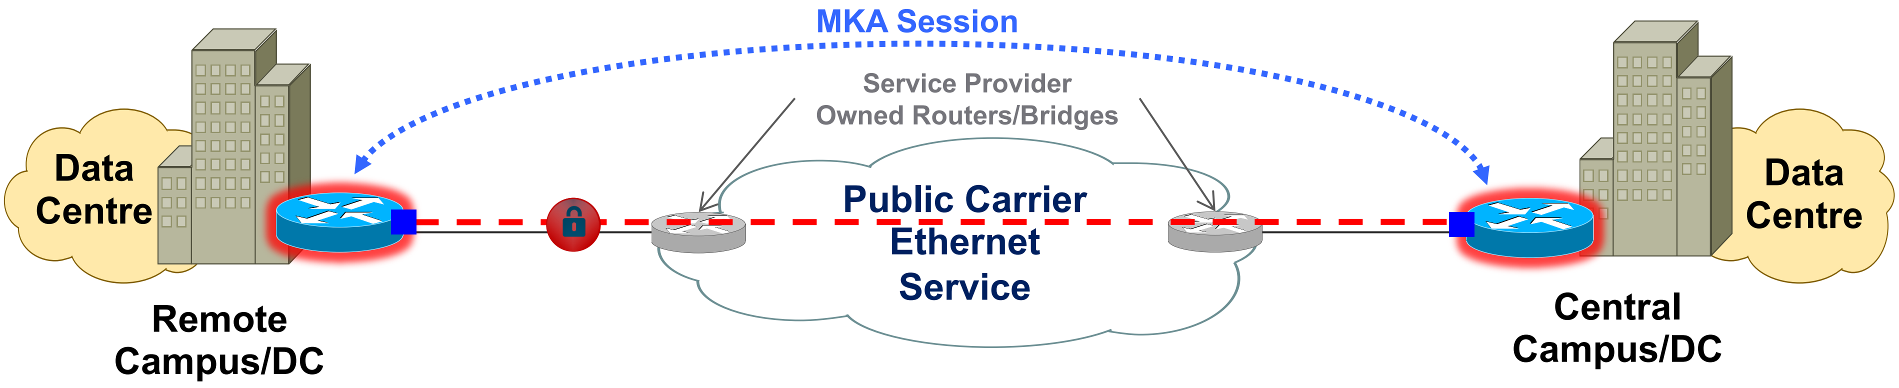

- Leverage MACsec over “public” standard Ethernet transport
- Optimise MACsec + WAN features to accommodate running over public Ethernet transport
- Target “line-rate” encryption for high-speed applications
  - Inter DC, MPLS WAN links, massive data projects
- Targets 100G, but support 1/10/40G as well

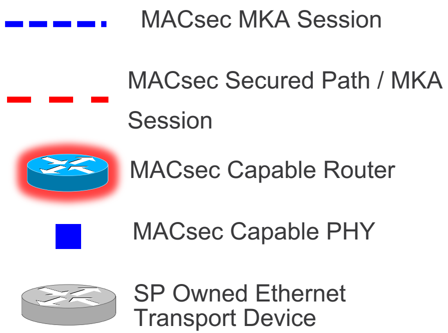

What is "WAN" MACsec?

New Enhancements to 802.1AE for WAN/Metro-E Transport

- **AES-256 (AES/GCM) support – 1/10/40 and 100G rates**
  - Target Next Generation Encryption (NGE) profile that currently leverages public NSA Suite B
- **Standards Based MKA key framework**
  - (defined in 802.1X-2010) within Cisco security development (Cisco "NGE")
- **Ability to support 802.1Q tags in clear**
  - Offset 802.1Q tags in clear before encryption (2 tags is optional)
- **Vital Network Features to Interoperate over Public Carrier Ethernet Providers**
  - 802.1Q tag in the clear
  - Ability to change MKA EAPoL Destination Address type
  - Ability to change MKA Ether-type value
  - Ability to configure Anti-replay window sizes
- **System Interoperability**
  - Create a common MACsec integration among all MACsec platforms in Cisco and Open Standards

802.1AE (MACsec) “Tag in Clear”

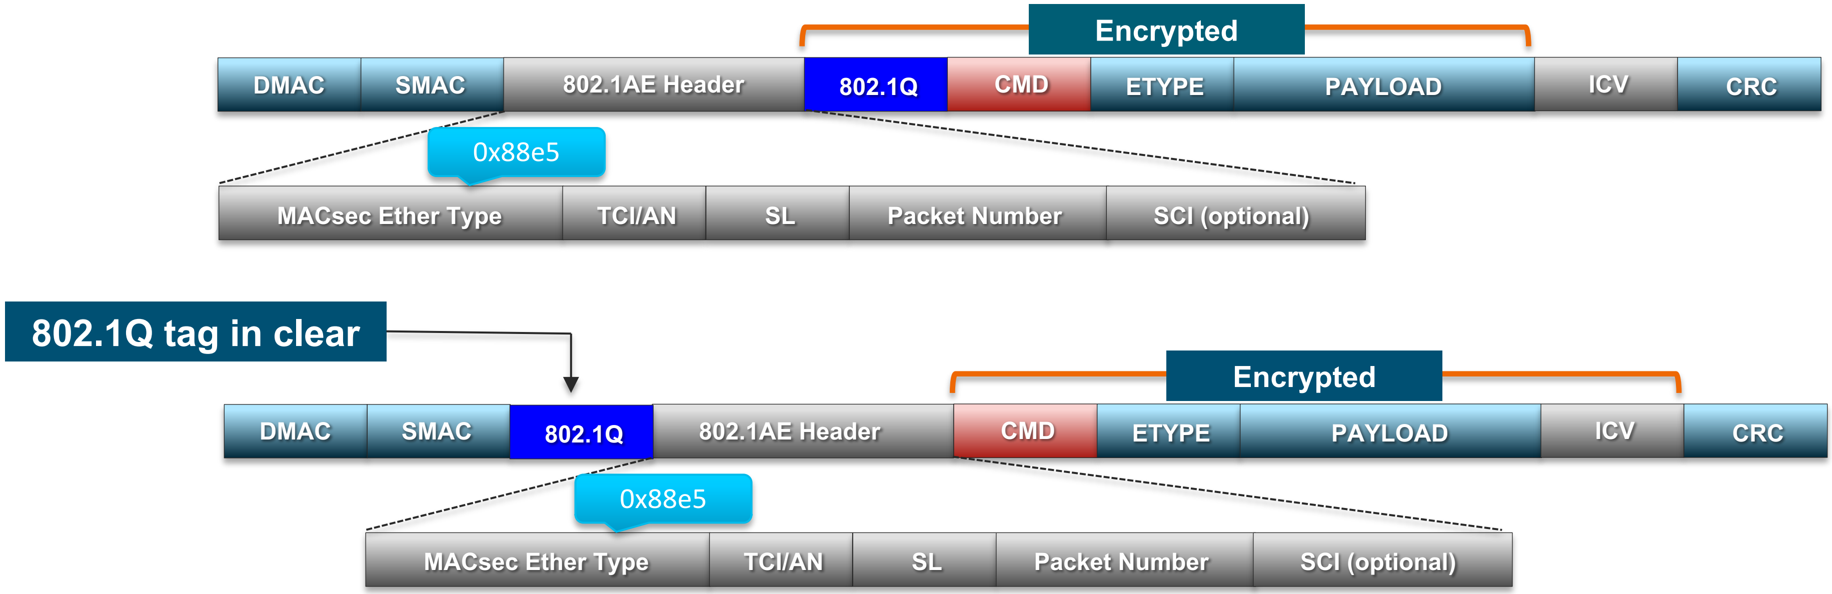

- 802.1Q tag offers significant network design options over the carrier network

WAN MACsec Use Case – 802.1Q Tag in the Clear

- Leverage 802.1Q for logical connectivity to each site
- This is analogous to “channelization” in SONET/SDH
- Router leverages IP sub-interface tag per location

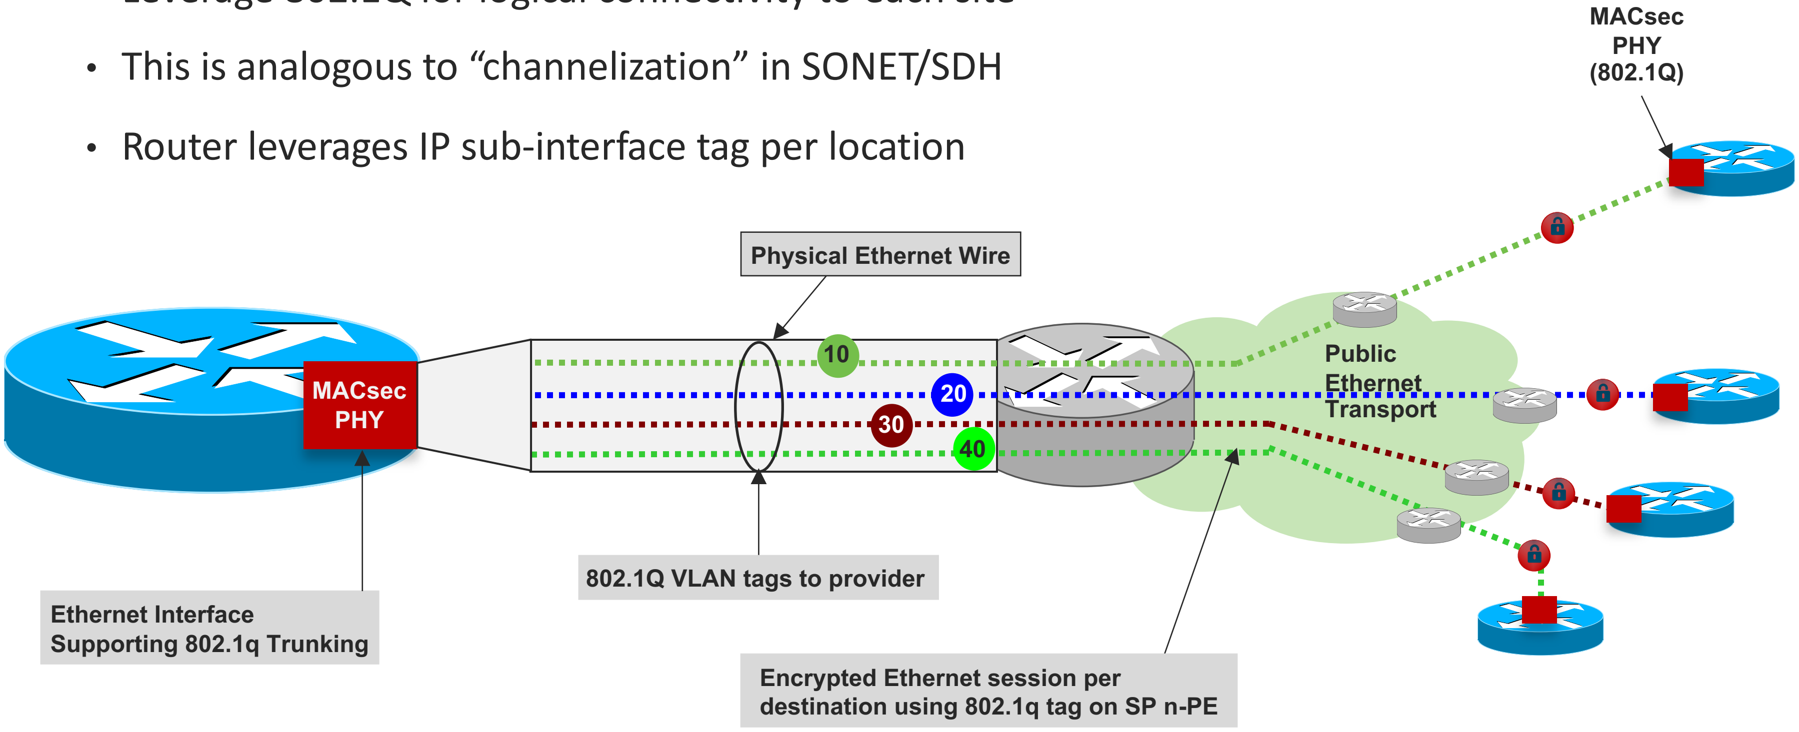

WAN MACsec – 802.1Q Tag in the Clear

Expose the 802.1Q tag “outside” the encrypted payload

- Example:

```
...
interface GigabitEthernet0/0/4
  macsec dot1q-in-clear 1

Interface GigabitEthernet0/0/4.20
  encapsulation dot1Q 20
  ip address 10.3.2.1 255.255.255.0
  mka pre-shared-key key-chain k1
  macsec
!
Interface GigabitEthernet0/0/4.30
  encapsulation dot1Q 30
  ip address 10.3.3.1 255.255.255.0
  mka pre-shared-key key-chain k1
  macsec
```

Allows the ability to leverage MACsec on a per sub-interface basis, exposing the “802.1Q tag” outside the encryption header.

Note: “1” denotes one .1Q tag depth

Hierarchical "Hybrid" MACsec + IPSec Design

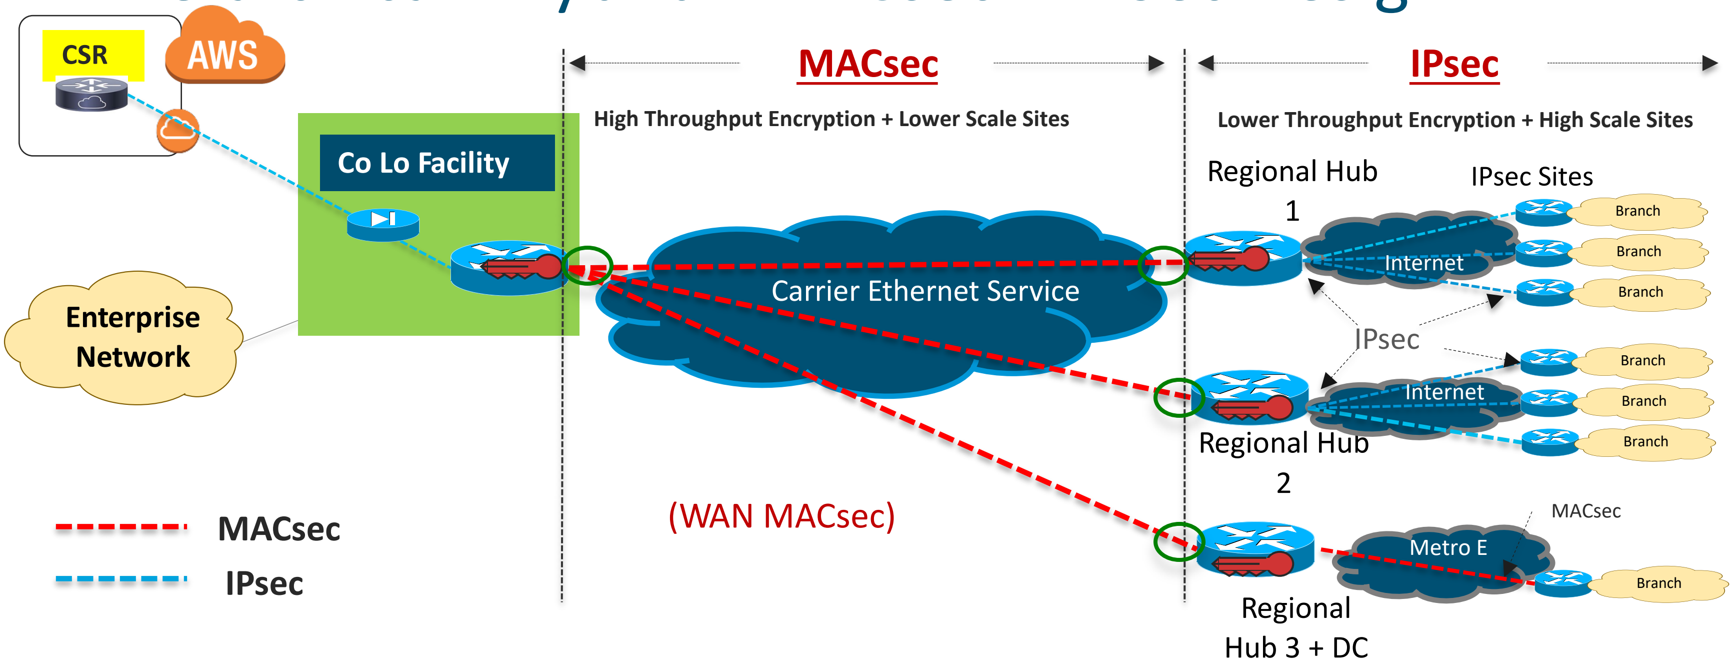

- "Hybrid" design option for mix of scale, performance, leveraging Ethernet services
- **MACsec:** Backbone/Core – Targets Higher BW, Lower Number of Sites
- **IPSec:** Branch/back-haul – Targets Lower BW, high number of sites, cloud (CSR)

White Paper

Innovations in Ethernet Encryption (802.1AE - MACsec) for Securing High Speed (1-100GE) WAN Deployments

Authors

Craig Hill
Distinguished Systems Engineer
U.S. Federal Area

Stephen Orr
Distinguished Systems Engineer
U.S. Public Sector

Introduction

Over the course of the past decade, customer demand for increasing Wide Area Network (WAN) bandwidth has been driving the networking industry to continually innovate in order to increase WAN transport speeds. Thus, we have witnessed the evolution from Asynchronous Transport Mode (ATM) to Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH) and, more recently, innovations in Ethernet and optical. Ethernet and optical have now emerged as the de facto standards and we have seen speeds grow from 10-Gb, 40-Gb, and now to 100-Gb speeds with no end of growth in sight.

Demand for increased bandwidth continues, driven by cloud services, mobile devices, and massive increases in video traffic. With the shift to cloud and mobile services, the need for ever-faster WAN transport speeds continues in order to handle the traffic created by locating applications and data off-premises.

While link speeds and demand for bandwidth continue to increase, the innovation of encryption technologies for securing these high-speed links, specifically for the service providers, cloud providers, large enterprises and governments, has failed to keep up. Furthermore, customers want to simplify their network operations and reduce the amount of protocol layers and complexity they are implementing in these high-speed networks, including the recent interest to hide network layer information in transit (IP addresses and protocol port numbers).

This document provides an in-depth look into:

- How Cisco is addressing this dilemma of link speed bandwidth outpacing the encryption technologies currently available
- Encryption innovations led by Cisco, including a detailed introduction to WAN Media Access Control Security (MACsec)

Previous WAN MACsec Sessions at Cisco Live (CL 365)

BRKRST-2309 – Introduction to WAN MACsec

http://www.cisco.com/c/dam/en/us/td/docs/solutions/Enterprise/Security/MACsec/WP-High-Speed-WAN-Encrypt-MACsec.pdf

Высокая доступность

Goals

- Efficiently utilize available bandwidth
- Dynamically respond to all types of disruptions
- Leverage most effective design techniques that meet the design requirements
- Review today's technology

Redundancy vs. Convergence Time
More Is Not Always Better

- In principle, redundancy is easy
- Any system with more parallel paths through the system will fail less often
- The problem is a network isn't really a single system but a group of interacting systems
- Increasing parallel paths increases routing complexity, therefore increasing convergence times

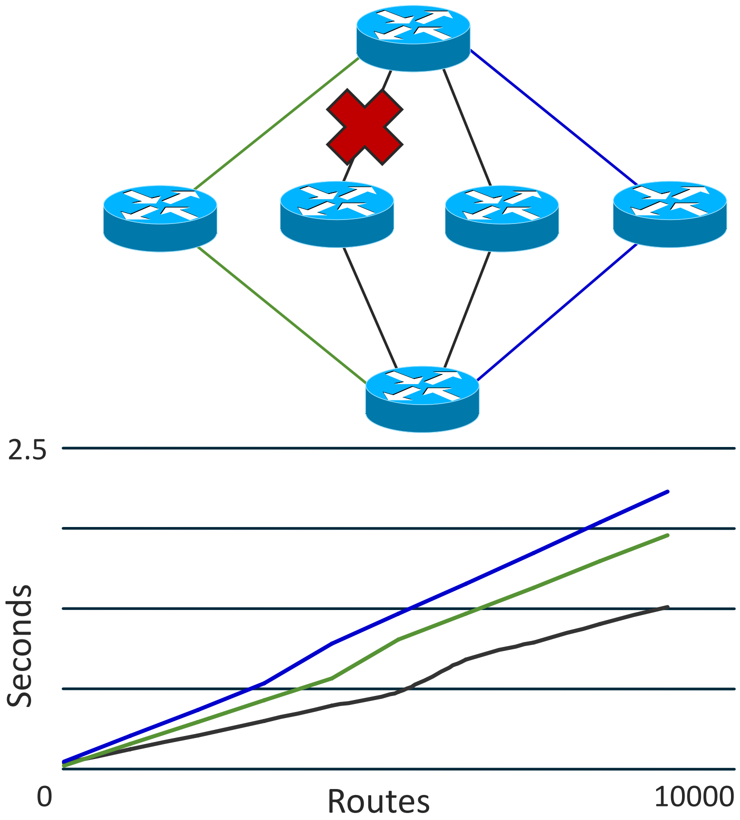

Where Can Outages Occur?

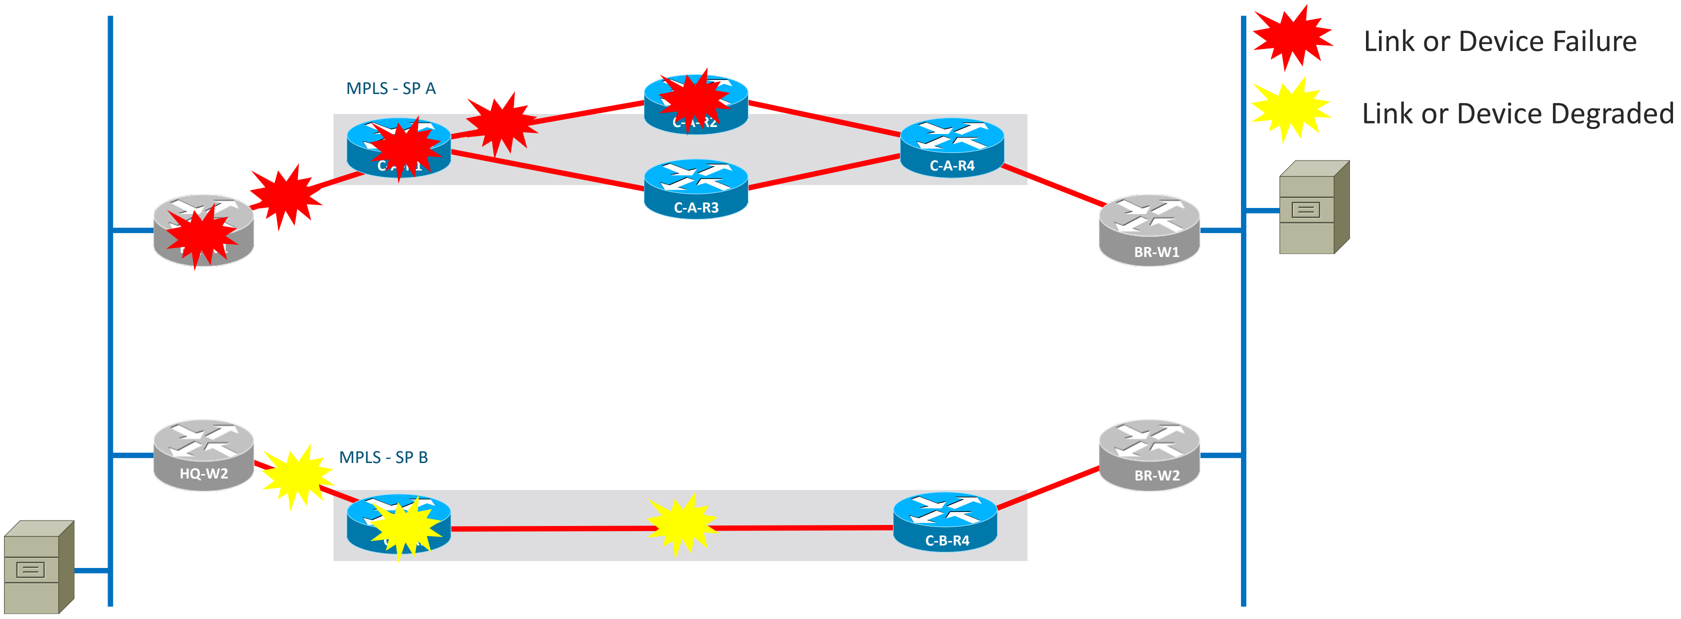

- How does outage manifest?
- How quickly can network detect?
- How long is bidirectional reconvergence?

Defining Availability

- System Availability: a ratio of the expected uptime to the experienced downtime over a period of the same duration
- Branch WAN High Availability: Between 99.99%(4) and 99.999%(5)
- Ultra High Availability: Between 99.9999%(6) and 99.999999%(8)

Availability	Downtime / Year
98.000000%	7.3 Days
99.000000%	3.65 Days
99.500000%	1.825 Days
99.900000%	8.76 Hrs
99.990000%	52.56 Min
99.999000%	5.256 Min
99.999900%	31.536 Sec
99.999990%	3.1536 Sec
99.999999%	.31536 Sec

Branch WAN HA Targets

Ultra HA Targets

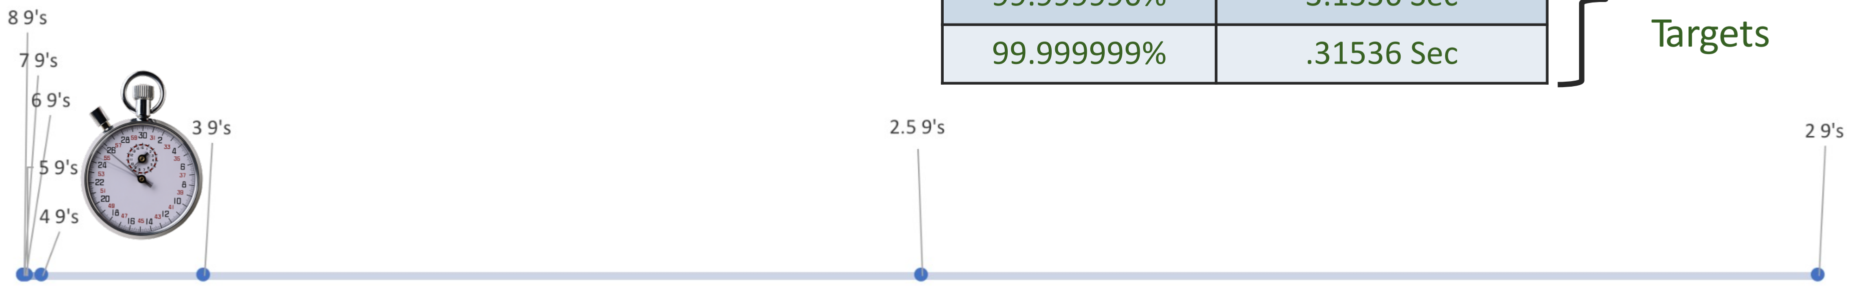

Cisco on Cisco http://cs.co/ithawan

Building Highly Available WANs

Redundancy and Path Diversity

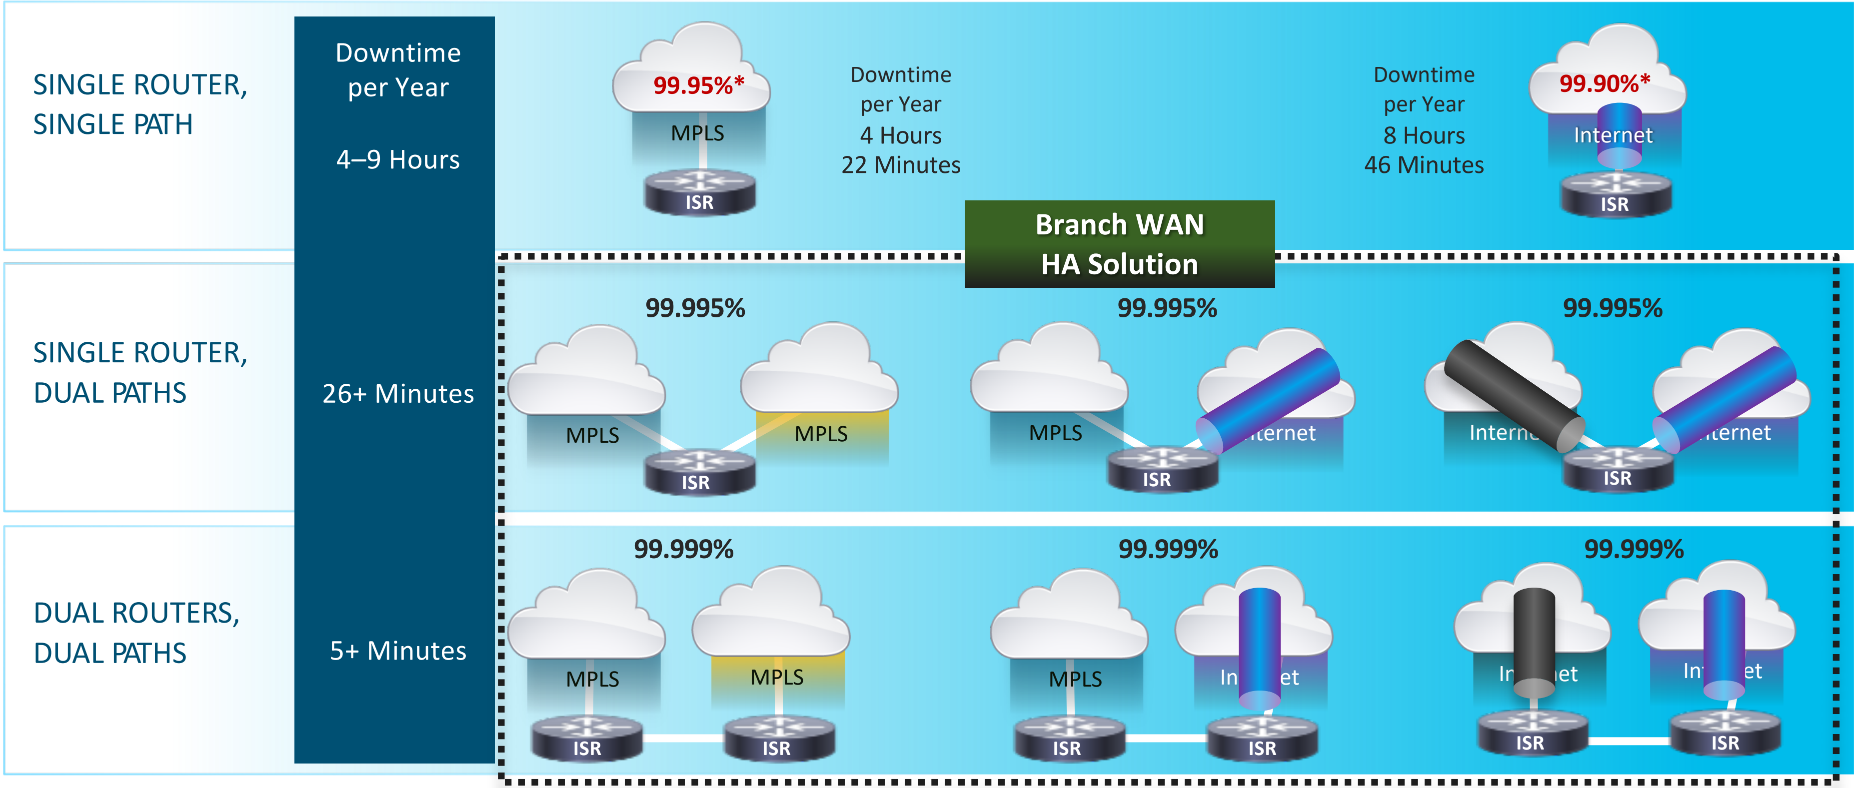

* Typical MPLS and Business Grade Broadband Availability SLAs and Downtime per Year, calculated with Cisco AS DAAP tool.

Load Sharing

- Assume the same routing process attempts to install two routes for the same destination in the RIB
- The routing process may allow the second route to be installed based on its own rules

IGP	OSPF	IS-IS	EIGRP
Route Cost	Must be equal to installed route	Must be equal to installed route	Must be less than the variance times the lowest cost installed route
Maximum Paths	Must be fewer than maximum-paths configured under the routing process (default = 4, maximum = 32)		

Note: BGP default value for maximum-paths = 1

CEF Load Sharing

Per-Destination	Per-Packet[1]
Default behaviour of IOS Universal Algorithm "show cef state"	Requires "ip load-sharing per-packet" interface configuration[1]
Per-flow using destination hash	Per-packet using round-robin method
Packets for a given source/destination session will take the same path	Packets for a given source/destination session may take different paths
More effective as the number of destinations increase	Ensures traffic is more evenly distributed over multiple paths
Ensures that traffic for a given session arrives in order	Potential for packets to arrive out of sequence

[1]Per-Packet Not available in IOS-XE based images

Load Sharing – Equal Cost Multi-Path (ECMP)

```
router#show ip route 192.168.239.0
Routing entry for 192.168.239.0/24
  Known via "eigrp 100", distance 170, metric 3072256, type external
  Redistributing via eigrp 100
  Last update from 192.168.245.11 on Serial0/2/1, 00:18:17 ago
  Routing Descriptor Blocks:
  * 192.168.246.10, from 192.168.246.10, 00:18:17 ago, via Serial2/0
      Route metric is 3072256, traffic share count is 1
      ....
    192.168.245.11, from 192.168.245.11, 00:18:17 ago, via Serial2/1
      Route metric is 3072256, traffic share count is 1
      ....
```

The Traffic Share Count Is Critical to Understanding the Actual Load Sharing of Packets Using These Two Routes

3072256/3072256 = 1

Load Sharing – with EIGRP Variance

```
router#show ip route 192.168.239.0
Routing entry for 192.168.239.0/24
  Known via "eigrp 100", distance 170, metric 3072256, type external
  Redistributing via eigrp 100
  Last update from 192.168.245.11 on Serial0/2/1, 00:18:17 ago
  Routing Descriptor Blocks:
  * 192.168.246.10, from 192.168.246.10, 00:18:17 ago, via Serial2/0
      Route metric is 1536128, traffic share count is 2
      ....
    192.168.245.11, from 192.168.245.11, 00:18:17 ago, via Serial2/1
      Route metric is 3072256, traffic share count is 1
      ....
```

If the Lower Metric Is Less than the Second Metric, the Traffic Share Count Will Be Something Other than 1 (EIGRP with Variance Configured)

3072256/3072256 = 1

3072256/1536128 = 2

2x Faster Link Gets 2 Flows vs. 1 Flow

CEF Hashing and Exact Route

- Now that we have load sharing
- What load-sharing algorithm
- "show cef state"

```
#show cef state
CEF Status:
 RP instance
 common CEF enabled
IPv4 CEF Status:
 CEF enabled/running
 dCEF enabled/running
 CEF switching enabled/running
 universal per-destination load sharing algorithm, id AE3030B1
IPv6 CEF Status:
<snip>
```

- Which exact path are the flows using
- "show ip cef exact-route <src-addr> [src-port] <dest-addr> [dest-port]

```
#show ip cef exact-route 1.1.1.1 2.2.2.2
1.1.1.1 -> 2.2.2.2 =>IP adj out of GigabitEthernet1, addr 10.255.0.1
```

Interface Detection

Carrier-delay

- If a link goes down and comes back up before the carrier delay timer expires, the down state is effectively filtered, and the rest of the software on the router is not aware that a link-down event occurred.
- Imposes a default 2 second pause before processing interface events
- Disabling carrier-delay speeds convergence upon interface events
- Disabling carrier-delay can increase control-plane usage during repetitive interface events (flapping)

Interface Detection

IP Event Dampening

- Imposes a logarithmic delay based on interface events
- Coupled with carrier-delay, dampening protects the control-plane from repetitive events by increasing the delay before processing up events should the interface flap.

```
#conf t
(config-if)#interface GigabitEthernet1
(config-if)#carrier-delay 0
(config-if)#dampening
(config-if)#end
#show dampening interface
1 interface is configured with dampening.
No interface is being suppressed.
Features that are using interface dampening:
 IP Routing
```

Routing Protocol Timers

INFORMATIONAL

	Keepalive (B) Hello (E,I,O) Update (R)	Invalid (R)	Holdtime (B,E,I) Dead (O) Holddown (R)	Flush (R)
BGP	60		180	
EIGRP (< T1)	5 (60)		15 (180)	
IS-IS (DIS)	10 (3.333)		30 (10)	
OSPF (NBMA)	10 (30)		40 (120)	
RIP/RIPv2	30	180	180	240

Note: Cisco Default Values

Routing Protocol Neighbor Behavior

INFORMATIONAL

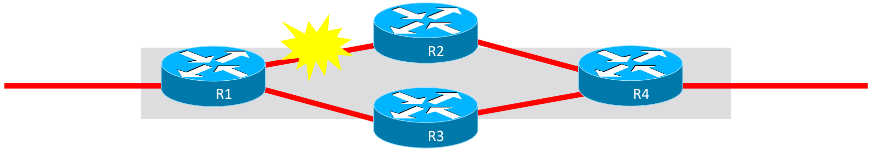

Recovery Times by Protocol

	Link Down Line Protocol Down	Link Up Loss 100%	Link Up Neighbor Down	Link Up Loss ~5%
BGP	~ 1 s	180	180	Never
EIGRP (< T1)	~ 1s	15 (180)	15 (180)	Never
IS-IS (DIS)	~ 1s	30 (10)	30 (10)	Never
OSPF (NBMA)	~ 1s	40 (120)	40 (120)	Never
RIP/RIPv2	~ 1s	240	240	Never

Note: Using Cisco Default Values

Routing Protocol Neighbor Behavior

Adjust Hello Timers

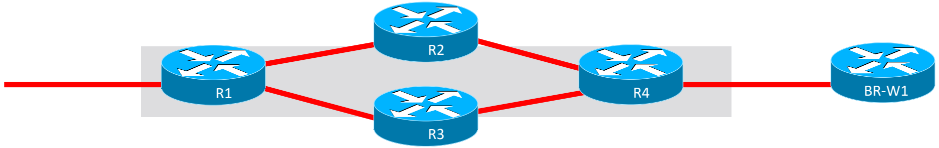


```
R4#show ip bgp vpnv4 vrf cisco neighbor
BGP neighbor is 192.168.101.10, vrf cisco, remote AS 65110, external link
 BGP version 4, remote router ID 192.168.201.10
 BGP state = Established, up for 1d10h
 Last read 00:00:19, hold time is 180, keepalive interval is 60 seconds
```

```
BR-W1#
router bgp 65110
 neighbor 192.168.101.9 timers 7 21
```

```
R4#show ip bgp vpnv4 vrf cisco neighbor
BGP neighbor is 192.168.101.10, vrf cisco, remote AS 65110, external link
 BGP version 4, remote router ID 192.168.201.10
 BGP state = Established, up for 00:01:23
 Last read 00:00:03, hold time is 21, keepalive interval is 7 seconds
```

When Configuring the Holdtime Argument for a Value of Less than Twenty Seconds, the Following Warning Is Displayed:

```
%Warning: A Hold Time of Less than 20 Seconds
Increases the Chances of Peer Flapping
```

Bidirectional Forwarding Detection (BFD)

- Extremely lightweight hello protocol
  - IPv4, IPv6, MPLS, P2MP
- 10s of milliseconds (technically, microsecond resolution) forwarding plane failure detection mechanism.
- Single mechanism, common and standardized
  - Multiple modes: Async (echo/non-echo), Demand
- Independent of Routing Protocols
- Levels of security, to match conditions and needs
- Facilitates close alignment with hardware

Drivers for BFD

- Link-layer detection misses some types of outages
  - e.g. Control Plane failure
- Control Plane failure detection is very conservative
  - 15-180 seconds in default configurations
- Link-layer failure detection is not consistent across media types
  - Less than 50ms on APS- protected SONET
  - A few seconds on Ethernet
  - Several seconds or more on WAN links
- Provides a measure of consistency across routing protocols
- Most current failure detection mechanisms are an order of magnitude too long for time-sensitive applications

Routing Protocol Neighbor Behavior

Bidirectional Forwarding Detection

```
interface GigabitEthernet4
 ip address 10.3.255.9 255.255.255.252
 bfd interval 50 min_rx 50 multiplier 3
router eigrp 1
 network 10.3.0.0 0.0.255.255
 bfd all-interfaces
```

```
R1#show bfd neighbors details
IPv4 Sessions
NeighAddr                              LD/RD         RH/RS     State     Int
10.3.255.10                          4104/1           Up        Up        Gi4
Session state is UP and using echo function with 50 ms interval.
Session Host: Software
OurAddr: 10.3.255.9
Handle: 2
Local Diag: 0, Demand mode: 0, Poll bit: 0
MinTxInt: 1000000, MinRxInt: 1000000, Multiplier: 3
Received MinRxInt: 1000000, Received Multiplier: 3
Holddown (hits): 0(0), Hello (hits): 1000(1371)
Rx Count: 985, Rx Interval (ms) min/max/avg: 34/1978/1226 last: 290 ms ago
Tx Count: 1372, Tx Interval (ms) min/max/avg: 71/1137/879 last: 721 ms ago
Elapsed time watermarks: 0 0 (last: 0)
Registered protocols: EIGRP CEF
Uptime: 00:20:06
Last packet: Version: 1                  - Diagnostic: 0
             State bit: Up              - Demand bit: 0
             Poll bit: 0                - Final bit: 0
             C bit: 0
             Multiplier: 3              - Length: 24
             My Discr.: 1                - Your Discr.: 4104
             Min tx interval: 1000000    - Min rx interval: 1000000
             Min Echo interval: 50000
```

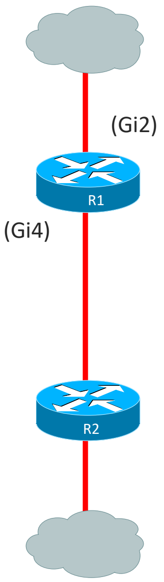

Configured in milliseconds (ms)

Displayed in microseconds (µs)

Routing Protocol Neighbor Behavior

Detecting Unreachable Neighbor (Hello Timers vs. BFD)

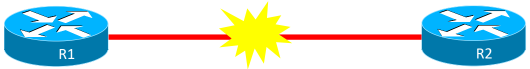

100% Packet Loss (Link Up)

EIGRP Default: Elapsed Time Between 10 – 15 Sec

```
R1#show clock
*09:58:27.716 UTC Sat Jan 27 2018
R1#
*Jan 27 09:58:40.612: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor
10.3.255.10 (GigabitEthernet4) is down: holding time expired
```

12.896 seconds

BFD: Elapsed Time Between 100 - 150 ms with 50ms interval

```
R1#show clock
*09:35:44.408 UTC Sat Jan 27 2018
R1#
*Jan 27 09:35:45.571: %BFDFSM-6-BFD_SESS_DOWN: BFD-SYSLOG: BFD
session ld:4101 handle:2,is going Down Reason: ECHO FAILURE
*Jan 27 09:35:45.575: %BFD-6-BFD_SESS_DESTROYED: BFD-SYSLOG:
bfd_session_destroyed, ld:4101 neigh proc:EIGRP, handle:2 act
*Jan 27 09:35:45.580: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor
10.3.255.10 (GigabitEthernet4) is down: BFD peer down notified
```

1.172 seconds[1]

[1]injecting 100% loss after hitting show clock in the lab

EOT, Static Routing, and DDR

- Enhanced Object Tracking (EOT)
- Static Routing Options
  - Floating Static Routes
  - Reliable Static Routing (RSR) using EOT
- Dial on Demand Routing (DDR)
  - EEM Script
  - DMVPN State Tracking
- More information:
  - http://cs.co/ddrbackup
  - Expands to https://www.cisco.com/c/en/us/support/docs/dial-access/dial-on-demand-routing-ddr/10213-backup-main.html

Enhanced Object Tracking (EOT)

Local Significance

Track Options	Syntax
Line-Protocol State of Interface	**track** *object-number* **interface** *type number* **line-protocol** track 1 interface serial 2/0 line-protocol
IP-Routing State of Interface	**track** *object-number* **interface** *type number* **ip routing** track 2 interface ethernet 1/0 ip routing
IP-Route Reachability	**track** *object-number* **ip route** *IP-Addr/Prefix-len* **reachability** track 3 ip route 10.16.0.0/16 reachability
Threshold* of IP-Route Metrics	**track** *object-number* **ip route** *IP-Addr/Prefix-len* **metric threshold** track 4 ip route 10.16.0.0/16 metric threshold

```
Router#show track 100
Track 100
Interface Serial2/0 line-protocol
Line protocol is Up
1 change, last change 00:00:05
Tracked by:
GLBP FastEthernet0/1 1
```

```
Router#show track 103
Track 103
IP route 10.16.0.0 255.255.0.0
reachability
Reachability is Up (EIGRP)
1 change, last change 00:02:04
First-hop interface is FastEthernet0/0
Tracked by:
GLBP FastEthernet0/1 1
```

IPv6 Support
15.3(3)S
15.4(1)T

* EIGRP, OSPF, BGP, Static Thresholds Are Scaled to Range of (0 – 255)

Enhanced Object Tracking (EOT)

External Significance

Track Options	Syntax
IP SLAs Operation	**track** *object-number* **ip sla** *type number* **state** `track 5 ip sla 4 state`
Reachability of an IP SLAs Host	**track** *object-number* **ip sla** *type number* **reachability** `track 6 ip sla 4 reachability`

Types of IP SLA Probes:

dhcp	http	path-jitter
dns	icmp-echo[1]	tcp-connect[1]
ethernet	icmp-jitter	udp-echo[1]
frame-relay	mpls	udp-jitter[1]
ftp	path-echo	voip

[1]Available for IPv6

Enhanced Object Tracking (EOT)

Compound Operations

Track Options	Syntax
list boolean	**track** *object-number* **list boolean {and|or}** and - both are up for object to be up or - one is up for object to be up track 5 list boolean or object 51 object 52 not ! Negates state of object
list threshold	**track** *object-number* **list threshold {weight|percentage}** track 6 list threshold weight object 61 weight 20 ! Twice as important object 62 ! Default weight 10 object 63 object 64 threshold weight up 30 down 25

Reliable Static Routing

Tracking IP SLA

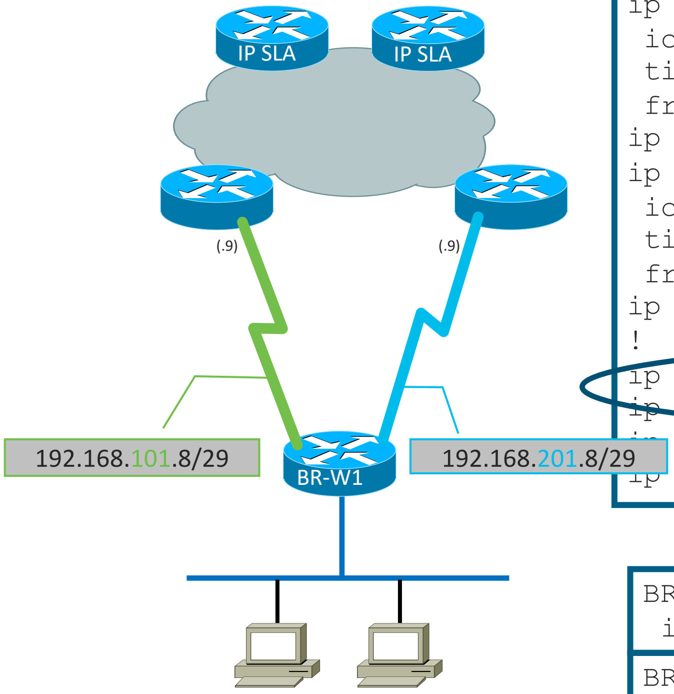

Static Host Route Guarantees probe destination only reachable via desired path

Permanent to guarantee probes only utilize desired path. Stay down when down.

```
track 4 list boolean or
 object 400
 object 401
track 400 ip sla 400 reachability
track 401 ip sla 401 reachability
ip sla 400
 icmp-echo 10.100.100.100 source-ip 10.1.2.120
 timeout 100
 frequency 10
ip sla schedule 400 life forever start-time now
ip sla 401
 icmp-echo 10.100.200.100 source-ip 10.1.2.120
 timeout 100
 frequency 10
ip sla schedule 401 life forever start-time now
!
ip route 10.100.100.100 255.255.255.255 Ethernet 0/1 192.168.101.9 permanent
ip route 10.100.200.100 255.255.255.255 Ethernet 0/1 192.168.101.9 permanent
ip route 10.100.0.0 255.255.0.0 192.168.101.9 track 4
ip route 10.100.0.0 255.255.0.0 192.168.201.9 200
```

```
BR-W1#show ip route track-table
 ip route 10.100.0.0 255.255.0.0 192.168.101.9 track 4 state is [up]
```

```
BR-W1#show ip route 10.100.0.0 255.255.0.0
S    10.100.0.0/16 [1/0] via 192.168.101.9
```

Reliable Static Routing

Tracking IP SLA

```
BR-W1#
*Mar 12 03:57:28.367: %TRACKING-5-STATE: 400 ip sla 400 reachability Up->Down
*Mar 12 03:57:37.374: %TRACKING-5-STATE: 401 ip sla 401 reachability Up->Down
*Mar 12 03:57:38.137: %TRACKING-5-STATE: 4 list boolean or Up->Down
```

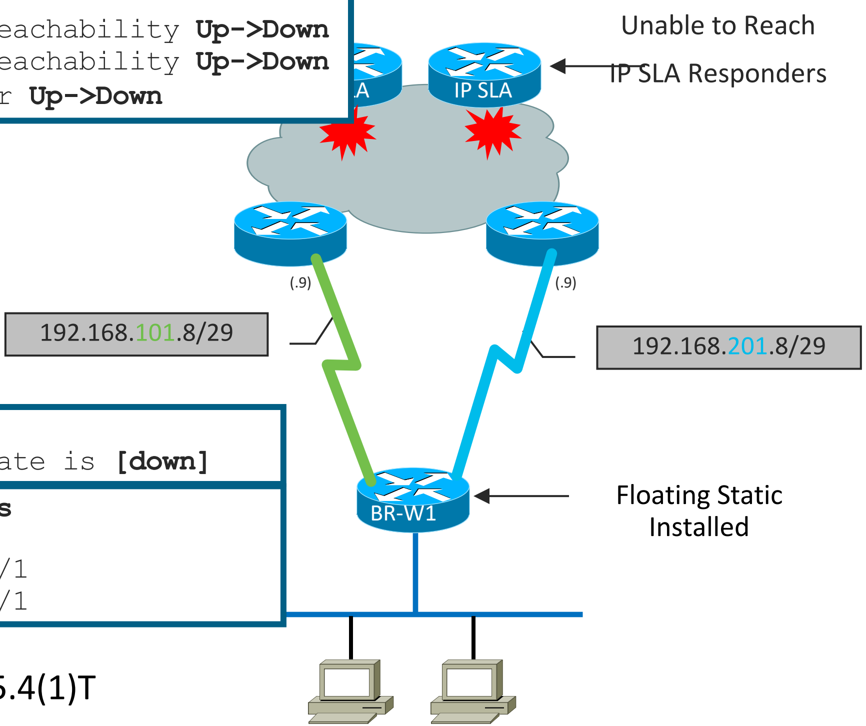


```
BR-W1#show ip route track-table
 ip route 10.100.0.0 255.255.0.0 192.168.101.9 track 4 state is [down]
```

```
BR-W1#show ip route 10.100.0.0 255.255.0.0 longer-prefixes
S    10.100.0.0/16 [200/0] via 192.168.201.9
S    10.100.100.100/32 [1/0] via 192.168.101.9, Ethernet0/1
S    10.100.200.100/32 [1/0] via 192.168.101.9, Ethernet0/1
```

IPv6 Reliable Static Routing added in 15.4(1)T

EEM Script

LTE Backup with Event Tracking

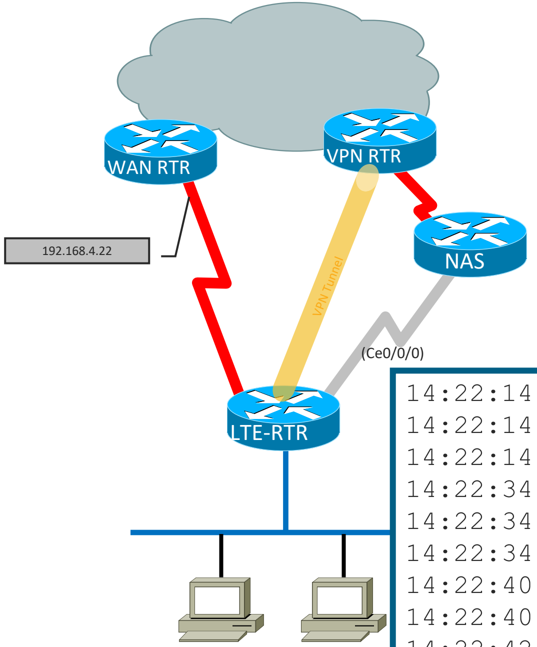


```
ip sla 100
 icmp-echo 192.168.4.22 source-interface GigabitEthernet0/1
 threshold 1000
 frequency 15
ip sla schedule 100 life forever start-time now

track 60 ip sla 100 reachability

event manager applet ACTIVATE-LTE
 event track 60 state down
 action 1  cli command "enable"
 action 2  cli command "configure terminal"
 action 3  cli command "interface cellular0/0/0"
 action 4  cli command "no shutdown"
 action 5  cli command "end"
 action 99 syslog msg "Activating LTE interface"
```

Don't forget to disable

```
14:22:14: %TRACKING-5-STATE: 60 ip sla 100 reachability Up->Down
14:22:14: %SYS-5-CONFIG_I: Configured from console by  on vty0(EEM:ACTIVATE-LTE)
14:22:14: %HA_EM-6-LOG: ACTIVATE-LTE: Activating LTE interface
14:22:34: %LINK-3-UPDOWN: Interface Cellular0/0/0, changed state to up
14:22:34: %DIALER-6-BIND: Interface Ce0/0/0 bound to profile Di1
14:22:34: %LINEPROTO-5-UPDOWN: Line protocol on Interface Cellular0/0/0, changed state to up
14:22:40: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel11, changed state to up
14:22:40: %CRYPTO-6-ISAKMP_ON_OFF: ISAKMP is ON
14:22:42: %DUAL-5-NBRCHANGE: EIGRP-IPv4 201: Neighbor 10.4.36.1 (Tunnel11) is up: new
adjacency
```

http://www.cisco.com/go/cvd/wan VPN Remote Site over 3G/4G/LTE Technology Design Guide

DMVPN Interface State Control

LTE Backup with DMVPN

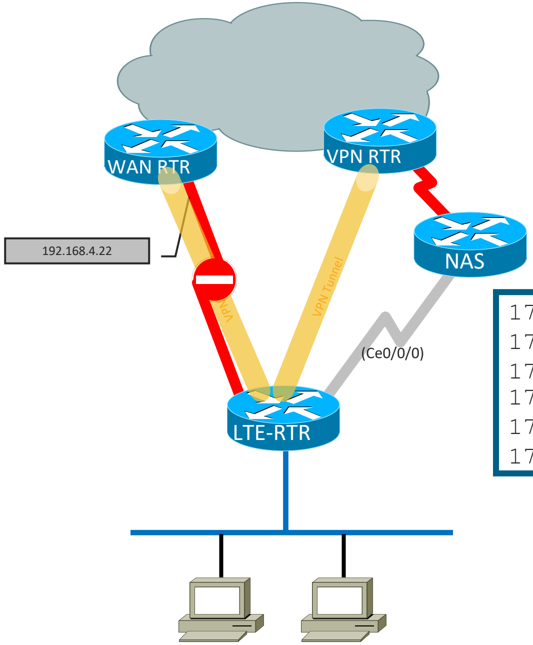


```
track 2 list boolean or
 object 101 not
track 101 interface Tunnel100 line-protocol
interface Tunnel200
 if-state track 2
 tunnel source Cellular0/0/0
end
#show track 2
Track 2
  List boolean or
  Boolean OR is Down
    7 changes, last change 00:07:55
    object 101 not Up
  Tracked by:
    IF-State Control 2
```

```
17:24:18.682: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel100, changed state to down
17:24:18.682: %TRACK-6-STATE: 101 interface Tu100 line-protocol Up -> Down
17:24:18.744: %TRACK-6-STATE: 2 list boolean or Down -> Up
17:24:28.683: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel200, changed state to up
17:24:29.276: %BGP-5-ADJCHANGE: neighbor 192.168.200.12 Up
17:24:37.505: %BGP-5-ADJCHANGE: neighbor 192.168.200.22 Up
```

```
#show track 2
Track 2
  List boolean or
  Boolean OR is Up
    8 changes, last change 00:00:32
    object 101 not Down
  Tracked by:
    IF-State Control 2
```

Hot Standby Routing Protocol (HSRP)

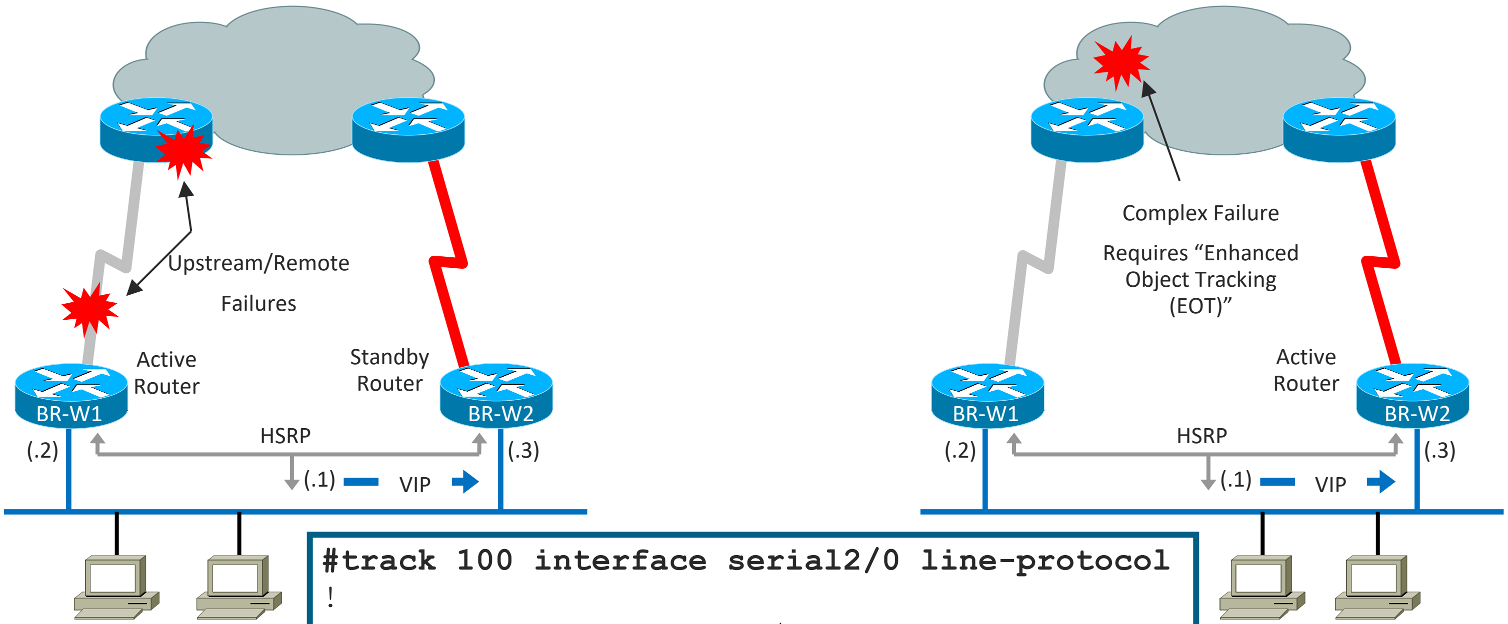


```
#track 100 interface serial2/0 line-protocol
!
interface FastEthernet0/0
 standby version 2
 standby 4 priority 110
 standby 4 track 100 decrement 20
 standby 6 priority 110
 standby 6 track 100 decrement 20
```

Hot Standby Routing Protocol (HSRP) BFD

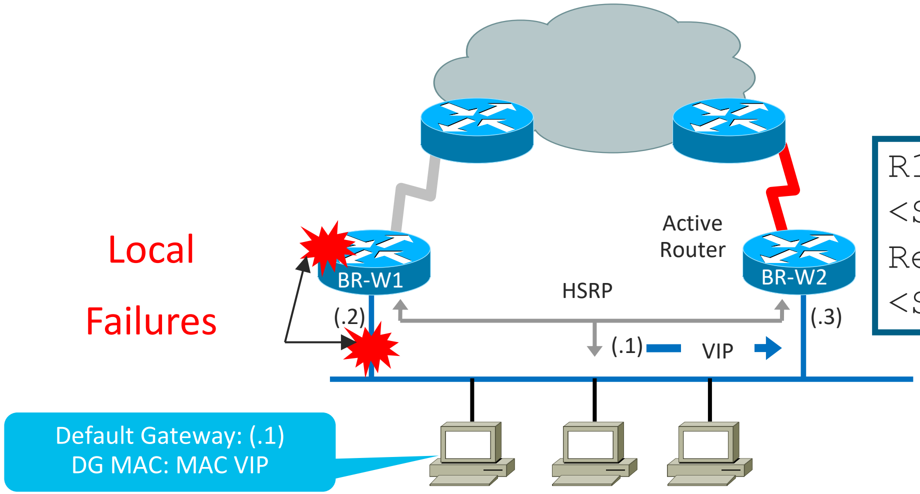


```
interface FastEthernet0/0
 bfd interval 50 min_rx 50 multiplier 3
```

```
R1#show bfd neighbors details
<SNIP>
Registered protocols: HSRP
<SNIP>
```

```
standby bfd all-interfaces ! default
!
interface FastEthernet0/0
 standby bfd ! Required only when all- interfaces disabled
```

Enhanced Object Tracking (EOT)

Tracking IP SLA

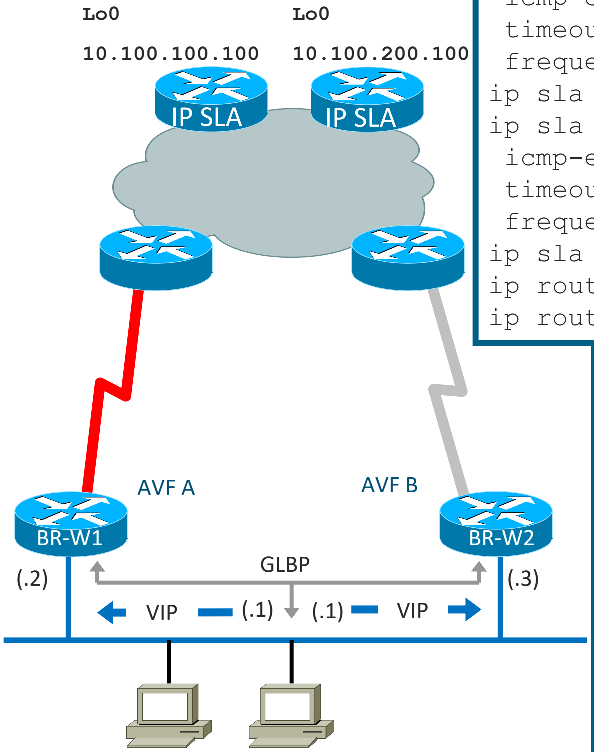


```
ip sla 100
 icmp-echo 10.100.100.100 source-ip 10.1.2.2
 timeout 100
 frequency 10
ip sla schedule 100 life forever start-time now
ip sla 200
 icmp-echo 10.100.200.100 source-ip 10.1.2.2
 timeout 100
 frequency 10
ip sla schedule 200 life forever start-time now
ip route 10.100.100.100 255.255.255.255 FastEthernet0/1 192.168.101.9 permanent
ip route 10.100.200.100 255.255.255.255 FastEthernet0/1 192.168.101.9 permanent
```

```
BR-W1#show ip sla statistics
IPSLA operation id: 100
    Latest RTT: 1 milliseconds
Latest operation start time: *04:42:11.444 UTC Tue Feb 17 2009
Latest operation return code: OK
Number of successes: 46
Number of failures: 0
Operation time to live: Forever
IPSLA operation id: 200
    Latest RTT: 1 milliseconds
Latest operation start time: *04:42:11.356 UTC Tue Feb 17 2009
Latest operation return code: OK
Number of successes: 24
Number of failures: 0
Operation time to live: Forever
```

Enhanced Object Tracking

Tracking IP SLA

```
BR-W1#
track 100 ip sla 100 reachability
track 200 ip sla 200 reachability
track 1 list boolean or
 object 100
 object 200
interface FastEthernet0/0
 ip address 10.1.2.2 255.255.255.0
 glbp 4 ip 10.1.2.1
 glbp 4 priority 110
 glbp 4 preempt
 glbp 4 weighting 110 lower 100
 glbp 4 load-balancing weighted
 glbp 4 weighting track 1 decrement 20
```

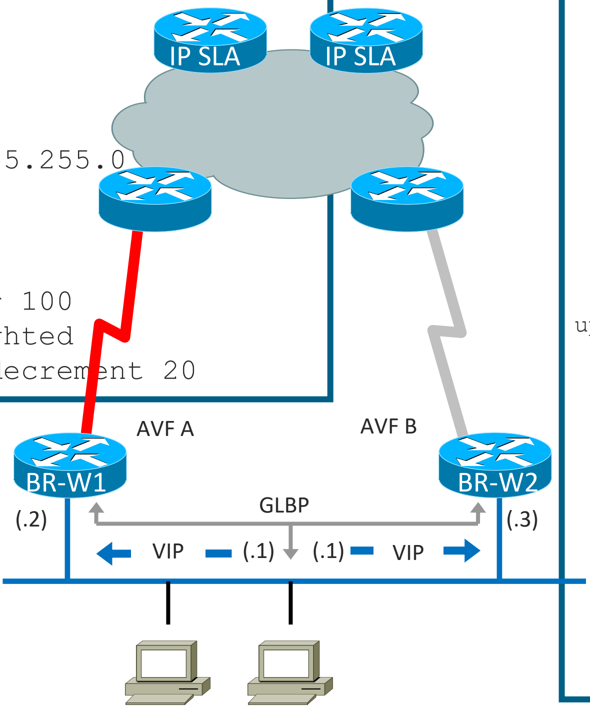


```
BR-W1#show glbp
FastEthernet0/0 - Group 4
  State is Active
    1 state change, last state change 00:09:59
  Virtual IP address is 10.1.2.1
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.336 secs
  Redirect time 600 sec, forwarder timeout 14400 sec
  Preemption enabled, min delay 0 sec
  Active is local
  Standby is 10.1.2.3, priority 105 (expires in 7.808 sec)
  Priority 110 (configured)
  Weighting 110 (configured 110), thresholds: lower 100,
upper 110
    Track object 1 state Up decrement 20
  Load balancing: weighted
  Group members:
    aabb.cc00.0110 (10.1.2.2) local
    aabb.cc00.0410 (10.1.2.3)
  There are 2 forwarders (1 active)
  Forwarder 1
    State is Active
    <SNIP>
  Forwarder 2
    State is Listen
    <SNIP>
```

Dual WAN (MPLS—Dual Carrier)

PE-CE Protocol: BGP

- Default behavior: 1-way load sharing
- Load is shared from HQ to Branch

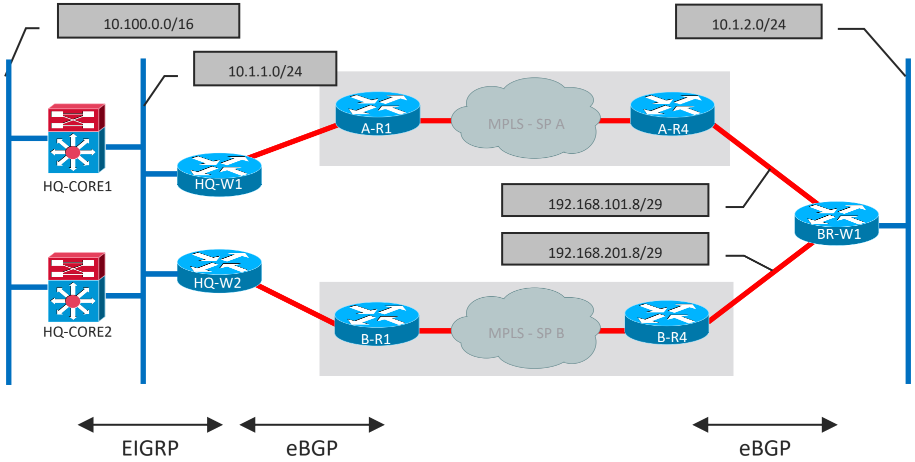


```
HQ-CORE1#show ip route
D EX  10.1.2.0/24 [170/258816] via 10.1.1.110, 02:24:22, Vlan10
                  [170/258816] via 10.1.1.210, 02:24:22, Vlan10
```

- Only one link used Branch to HQ

```
BR-W1#show ip route
B    10.100.0.0/16 [20/0] via 192.168.101.9, 00:34:00
```

Dual WAN (MPLS—Dual Carrier)

PE-CE Protocol: BGP Layer 3 Campus Locations

- IGP (EIGRP examples)
  - Routes redistributed from BGP into IGP (match & tag)
  - BGP routes are treated as IGP external
- BGP
  - No iBGP required between HQ-W1 & HQ-W2 (CE routers)
  - Routes redistributed from IGP into BGP except those tagged as originally sourced from BGP

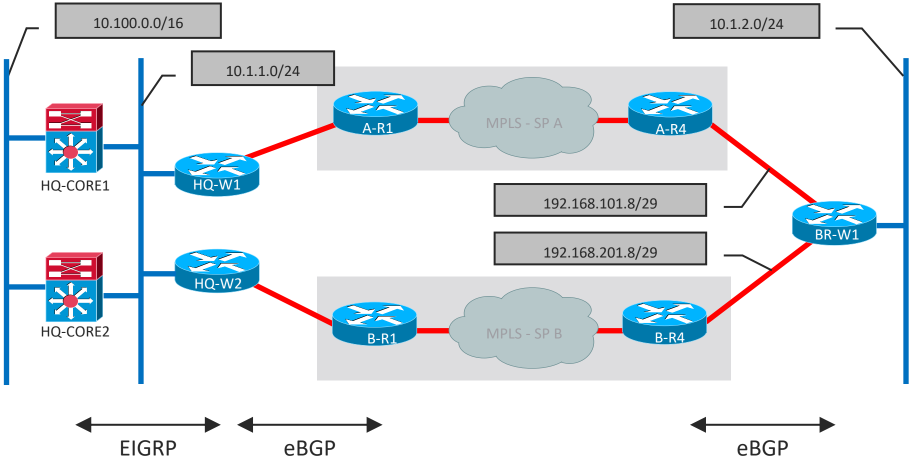

Dual WAN (MPLS—Dual Carrier)

Mutual Route Redistribution Detail

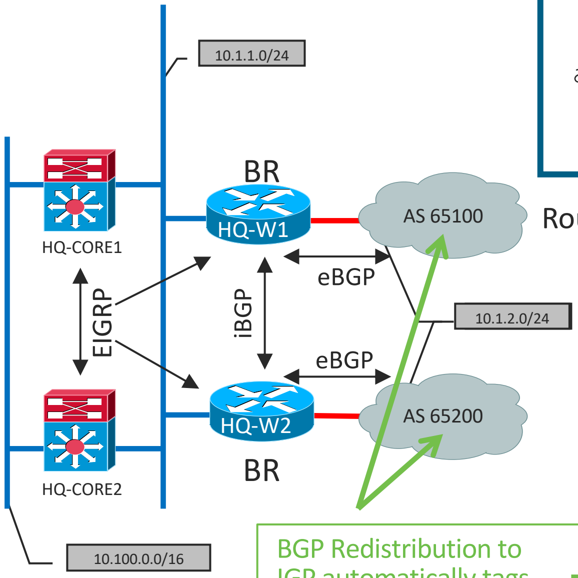

BGP Redistribution to IGP automatically tags routes with neighbor AS Number

Routes into EIGRP

```
HQ-W1#
router eigrp networkers
 address-family ipv4 unicast autonomous-system 65110
  topology base
   redistribute bgp 65110 metric 45000 100 255 1 1500
 address-family ipv6 unicast autonomous-system 65110
  topology base
   redistribute bgp 65110 metric 45000 100 255 1 1500
```

Routes into BGP

```
HQ-W1#
router bgp 65110
 address-family ipv4
  redistribute eigrp 65110 route-map BLOCK-TAGGED-ROUTES
 address-family ipv6
  redistribute eigrp 65110 route-map BLOCK-TAGGED-ROUTES
!
route-map BLOCK-TAGGED-ROUTES deny 10
 match tag 65100 65200
route-map BLOCK-TAGGED-ROUTES permit 20
!
```

Dual WAN (MPLS—Dual Carrier)

PE-CE Protocol: BGP Layer 2 Single Router Branch

- Is it possible to load share from Branch to HQ?
  - maximum-paths 2
- Requires hidden command:
  - bgp bestpath as-path multipath-relax

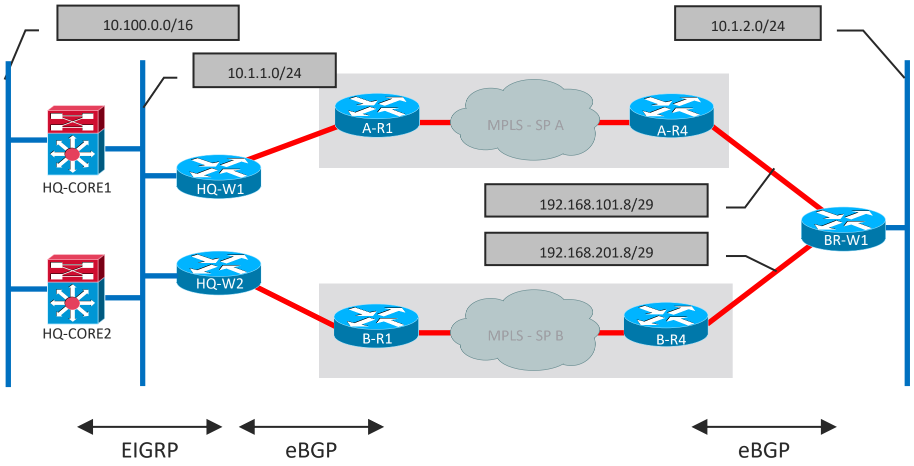


```
router bgp 65110
 bgp bestpath as-path multipath-relax
 address-family ipv4
  maximum-paths 2
 address-family ipv6
  maximum-paths 2
```

```
BR-W1#show ip route
B    10.100.0.0/16 [20/0] via 192.168.201.9, 00:03:44
                   [20/0] via 192.168.101.9, 00:03:44
```

DUAL WAN (MPLS + Internet)

PE-CE Protocol: BGP, Tunnel Protocol: EIGRP

- Headquarters WAN Edge
  - W1 learns Branch route via eBGP
  - W2 learns Branch route via EIGRP
- Headquarters Core
  - W1 redistributes eBGP into EIGRP, results in EIGRP external
  - W2 does not require redistribution, results in EIGRP internal
  - Core1, Core2 install Branch route via W2

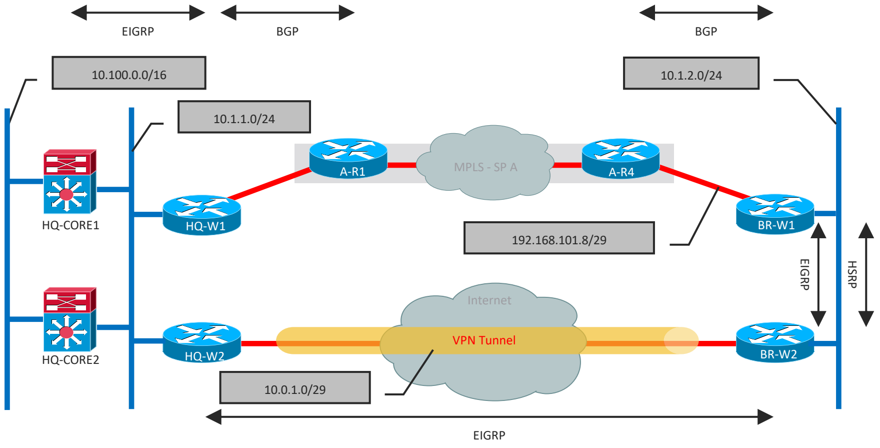

HQ to Branch Traffic Flows Across Tunnel

```
HQ-W1#show ip route
B    10.1.2.0/24 [20/0] via 192.168.101.2, 05:24:01
```

```
HQ-W2#show ip route
D    10.1.2.0/24 [90/26882560] via 10.0.1.2, 00:00:04, Tunnel1
```

```
HQ-CORE1#show ip route
D    10.1.2.0/24 [90/26882816] via 10.1.1.210, 00:02:32, Vlan10
```

DUAL WAN (MPLS + Internet)

PE-CE Protocol: BGP, Tunnel Protocol: EIGRP

- How to force HQ to Branch traffic across MPLS (primary)?
  - Adjust administrative distance
    - For EIGRP routes learned via tunnel
    - Ensure administrative distance is higher than that of EIGRP external (170)

```
HQ-W2#
router eigrp 65110
 network 10.0.1.0 0.0.0.7
 distance 195 10.0.1.0 0.0.0.7
```

Only change is on hub

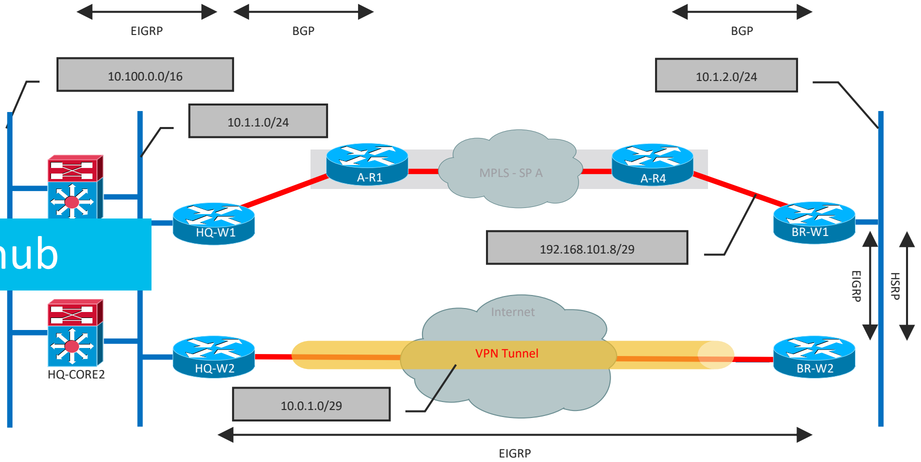

- Redistribute between two EIGRP Processes Forcing External as done between BGP and Campus EIGRP

```
HQ-W2#
Router eigrp 65100
 network 10.0.1.0 0.0.0.7
router eigrp 65110
 redistribute eigrp 65100
```

Requires additional changes

or Proper Pre-Planning

```
HQ-W1#show ip route
B    10.1.2.0/24 [20/0] via 192.168.101.2, 05:24:01
```

```
HQ-W2#show ip route
D EX  10.1.2.0/24 [170/261120] via 10.1.1.110, 00:07:25, GigE0/0
```

```
HQ-CORE1#show ip route
D EX  10.1.2.0/24 [170/258816] via 10.1.1.110, 00:08:44, Vlan10
```

Now:
HQ to Branch Traffic Flows
Across MPLS

DUAL WAN (MPLS + Internet)

MPLS Failure

- Failure within MPLS cloud
  - Dependent on provider
- Worst Case
  - Link up neighbor down
  - Primary dependency BGP timers
  - End to end convergence time as long as BGP Holdtime
- Configuration options
  - BFD for sub-second notification
  - End-to-end Application Restoration as fast as SD-WAN detects

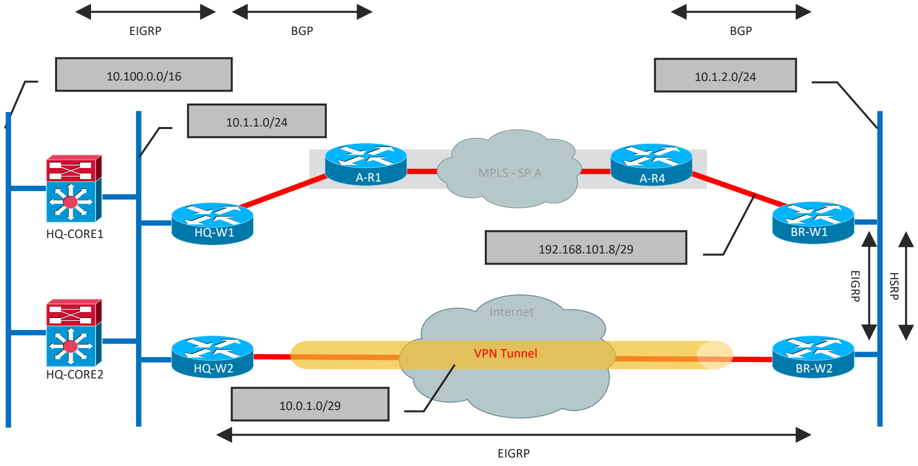

HQ Route Tables

After Failure:
HQ to Branch Traffic
Flows Across Tunnel

```
HQ-W2#show ip route
D    10.1.2.0/24 [195/26882560] via 10.0.1.2, 00:06:46, Tunnel1
HQ-CORE1#show ip route
D    10.1.2.0/24 [90/26882816] via 10.1.1.210, 00:09:18, Vlan10
```

DUAL WAN (MPLS + Internet)

MPLS Failure

- Failure within MPLS cloud
- Suboptimal routing at Branch
  - HSRP primary remains unchanged at BR-W1
  - Use EOT and move HSRP primary to BR-W2

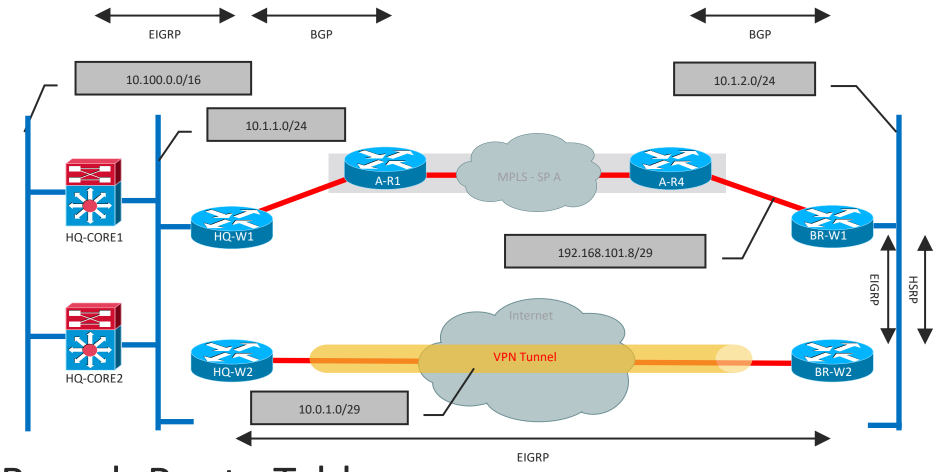

Branch Route Tables

After Failure: Branch to HQ Traffic Flows Across Tunnel

```
BR-W1#show ip route
D     10.100.100.0/24
       [90/26885376] via 10.1.2.220, 00:22:42, FastEthernet0/1
D     10.100.200.0/24
       [90/26885376] via 10.1.2.220, 00:22:42, FastEthernet0/1
```

```
BR-W2#show ip route
D     10.100.100.0/24 [90/26882816] via 10.0.1.1, 01:08:44, Tunnel1
D     10.100.200.0/24 [90/26882816] via 10.0.1.1, 01:08:45, Tunnel1
```

Summary of Convergence Techniques

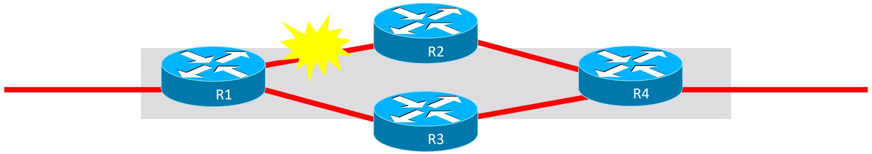

- 🟢 Excellent Option
- 🟡 SubOptimal Option
- 🔴 Bad Option

Effectiveness of Various Techniques for Different Outage Types

	Link Down	Link Up Neighbor Down	Link Up Loss ~5%	Upstream Blackhole	Upstream Brownout
Routing Protocols	🟢	🟡	🔴	🟡	🔴
BFD	🟢	🟢	🟡	N/A[1]	N/A[1]
EOT[2]	🟢	🟡	🔴	🟡	🔴
RSR[3] using EOT (w/IP SLA)	🟢	🟢	🟡	🟢	🟡
SD-WAN	🟢	🟢	🟢	🟢	🟢

[1]BFD Multihop support for Static and BGP routes
[2]Enhanced Object Tracking
[3]Reliable Static Routing

Сегментация

What Is Enterprise L3 “Network” Segmentation?

- Giving One physical network the ability to support multiple L3 virtual networks
- End-user perspective does not change
- Maintains Hierarchy, Virtualizes devices, data paths, and services

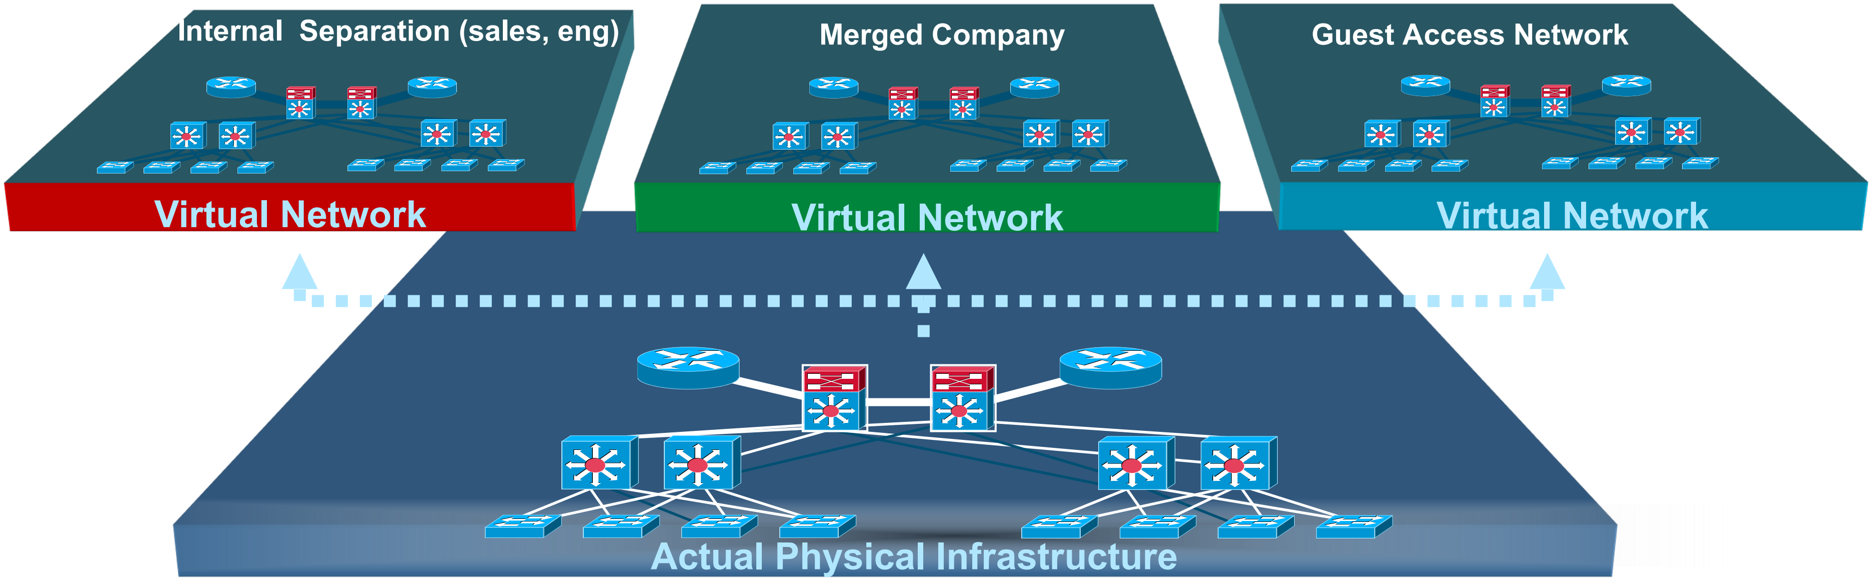

Why L3 Network Segmentation?

Key Drivers and Benefits

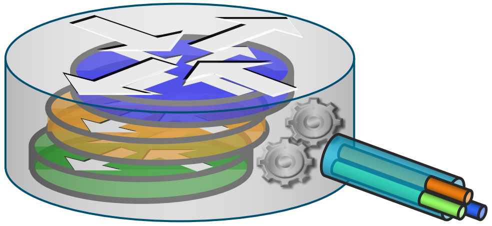

- Cost Reduction
  - Allowing a single physical network the ability to offer multiple virtual networks to tenants
- Simpler OAM
  - Reducing the physical network devices that need to be managed and monitored
- Security
  - Maintaining segmentation of the network for different departments over a single device/Campus/WAN
- Agility
  - Accelerates adding network segments (virtual) over same physical networks
- High Availability
  - Leverage segmentation through clustering devices that appear as one (vastly increased uptime)
- Data Center Applications
  - Offer per/multi-tenant segmentation from the DC into the WAN/campus/Branch and cloud
  - End-to-end Segmentation from-server-to-campus-to-WAN

Why L3 Network Segmentation?

L3 Network Segmentation Use Cases – Current and Evolving

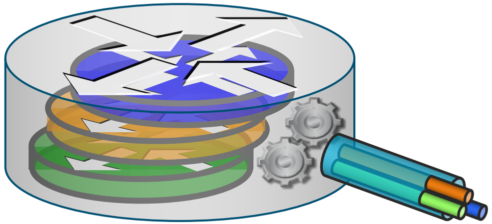

- Multi-Tenant Dwelling requiring Separation
  - Airports – (United, Delta, etc…), Government Facilities – (agencies sharing single building/campus), Intra Organisation segmentation – (sales, engineering, HR, LoB)
  - Company mergers – allowing slow migration for transition, overlapping addressing
  - IoT Device Isolation – segment (IP cameras, badge readers) from the user data
- Security for Isolation
  - Key Fundamental element for Zero Trust Security framework
  - Quarantine Zone – Honey Pot, Steered Traffic as result of DDoS, Anomaly Enforcement
  - Mandates to logically separate varying levels of security (e.g. enclaves)
- Regulation requirements - Health Care – HIPPA | Financial and Transactional – Sarbanes-Oxley, PCI Compliance
- Public Cloud and Key Component of Policy Construct
  - L3 segmentation for "per tenant" - Leveraged in Intent-based network policies

Virtual Routing and Forwarding Instance - VRF

Virtual Routing Table and Forwarding Separate to Customer Traffic

- Logical routing context within the same PE device
- Unique to a VPN
- Allows for customer overlapping IP addresses
- Deployment use cases
  - Business VPN services
  - Network segmentation
  - Data Center access

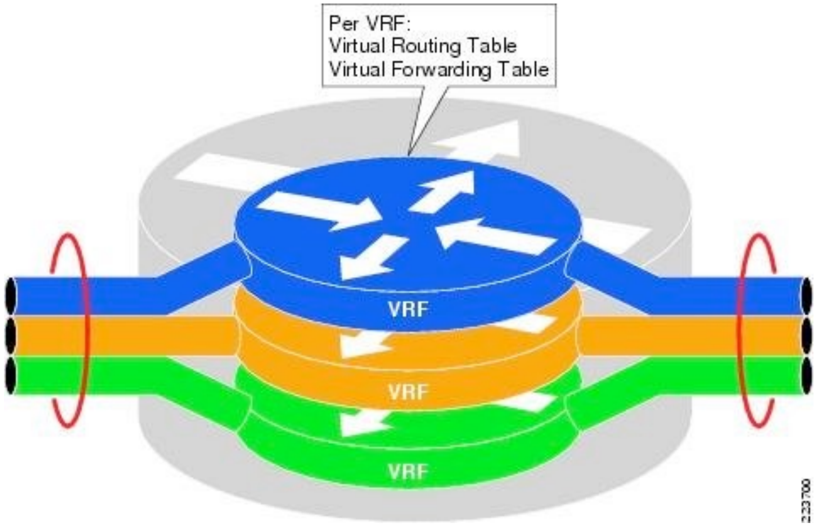

MPLS: The WAN Service Enabler

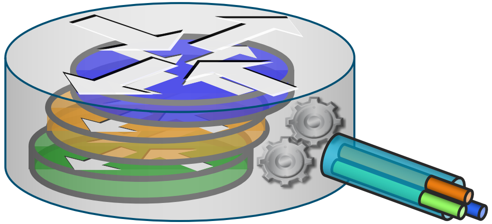

- L3 VPN Services
  - BGP VPN (RFC 4364), VPN over IP, Inter-AS, 6vPE
- L2 VPN Services - PW, VPLS, E-VPN
- Traffic Engineering - Explicit Path Routing
  - Traffic Engineering, disjoint paths, attributes for best path (latency, packet loss)
  - Optimisation of bandwidth, shift to Segment Routing TE (SR-TE)
- Bandwidth Protection Services - LFA, TI-LFA (IP FRR), MPLS TE FRR
- IP Multicast (per VPN/VRF, Rosen, LSM, BIER)
- Interworking with new solutions – VXLAN → L2/L3 VPN
- Leverage Segment Routing for Next-Gen Scale, Central Control, optimised services
  - Offers an “SD-MPLS” solution moving forward

Top Use Cases Today for SR

- Simplicity and complexity reduction in the core
  - Less protocols, reduced state, huge scale, highly programmable
- Protection with integrated TI-LFA FRR
- SR Traffic Engineering made simpler
  - BW optimization and capacity reaction (WAE + collection)
  - Disjointed paths (colored topology, SR Flex Algo)
  - SR-PCE (centralized SR-PCE, end-to-end awareness, multi-domain)
- Low-latency services using Performance Monitoring (PM)
  - Measure real-time per link delay measurement (loss coming in future)
  - Allows path selection based on link delay state, rather just cost
- SR On-Demand Next-Hops (BGP focused, SLA-aware per VPN)
- SR IGP Flexible Algorithms
  - Topology defined by operator, per service

Private IP VPN
“Over the Top”
Solution Options

WAN Segmentation Models

- Self Deployed MPLS Backbone (SD-Core) Supporting MPLS BGP IP VPN Services (RFC 4364)

- Self deployed MPLS BGP IP VPNs “over the top” of an SP Offered IP transport

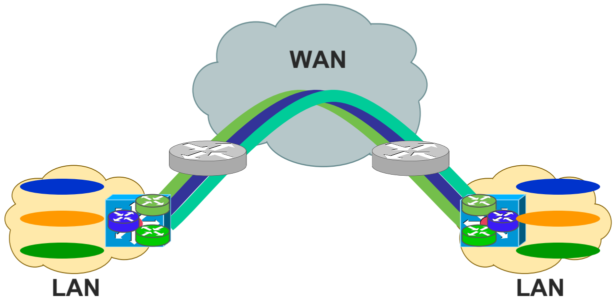

MPLS VPN over IP…

Simplifying MPLS VPN over IP - RFC 4797 + RFC 4364 + RFC 4023

- Customer may not control the WAN transport Between MPLS networks
- Cannot depend on “end to end” label forwarding for transport
- Customer requires encryption for their PE to PE MPLS traffic
  - No native MPLS encryption exists today, must leverage IP
- MPLS over IP allows MPLS VPN solutions to leverage cost effective IP transport

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)

Private MPLS VPNs “Over the Top”

Overview

- Allows enterprises to deploy simpler-to-manage MPLS VPN (v4/v6) solutions over IP
- CE owner (“us” ☺) controls the L3 VPN deployment
- PE (“SP”) provides transport of IP
- Key Benefit?
  1. CE owner can still leverage cost effective L3 transport services, Internet, QoS SLA’s… from the SP
  2. CE owner controls policy, segmentation, topology, encryption… “over the top”
- <u>Target Use cases:</u> simplified “Enterprise controlled” MPLS VPN over IP Transport

Enterprise SD-WAN (Over the Top)

Solutions with/without SDN Controller

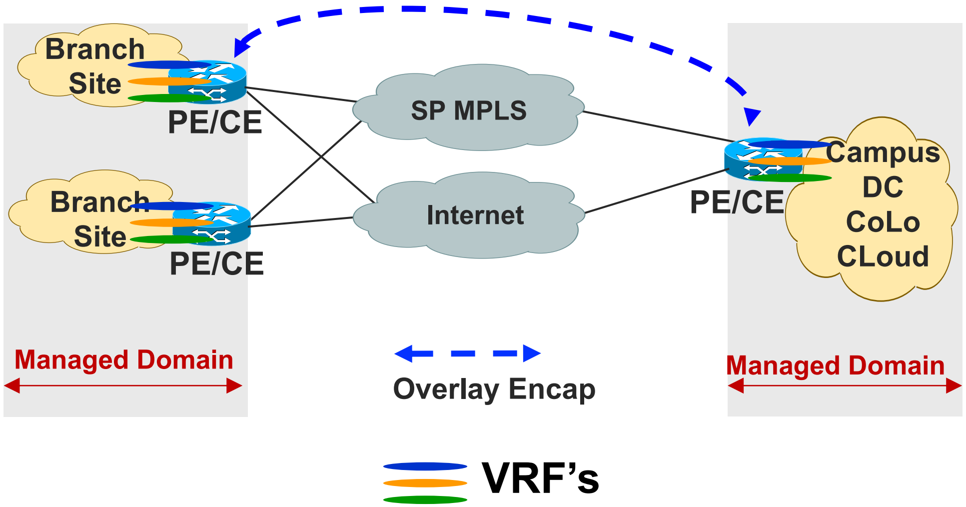

Common Use Cases for MPLS over IP

- State, country or Global based MPLS VPN where transport option is IP only
- The "business requirement" mandates segmentation (refer to L3 segmentation use cases)
  - Wants MPLS VPN but also requires encryption
  - Stitch MPLS VPN's over non MPLS (labeled) transport
  - L3 Manages Services (managed CPE over L3 VPN transport)
- Campus and/or DC networks require "policy" extension over the WAN
  - Cisco SD-Access = VN / Cisco ACI = VRF / Cisco SD-WAN = VPN
- Security focused customer leverages proprietary encryption devices
- Targets customers (enterprise/govt/managed-SP) that desire quick on-demand control of the L3 VPN Segmentation

Primary Components – VPN over IP

- **Segmentation component**
  - Virtual Route Forwarding Instance (VRF)
- **Control Plane component**
  - MP-BGP (RFC 4364)
  - SD-WAN L3 VPN - Overlay Management Protocol (OMP)
  - DMVPN L3 VPN - NHRP
- **Data Plane component**
  - MPLS over GRE/IP-UDP (RFC 4023)

- Service Support of Each Solution: QoS, IPv6 (selective), Encryption, Multicast, etc…

Encapsulation for MPLS in GRE (RFC 4023) (GRE RFC = 2784)

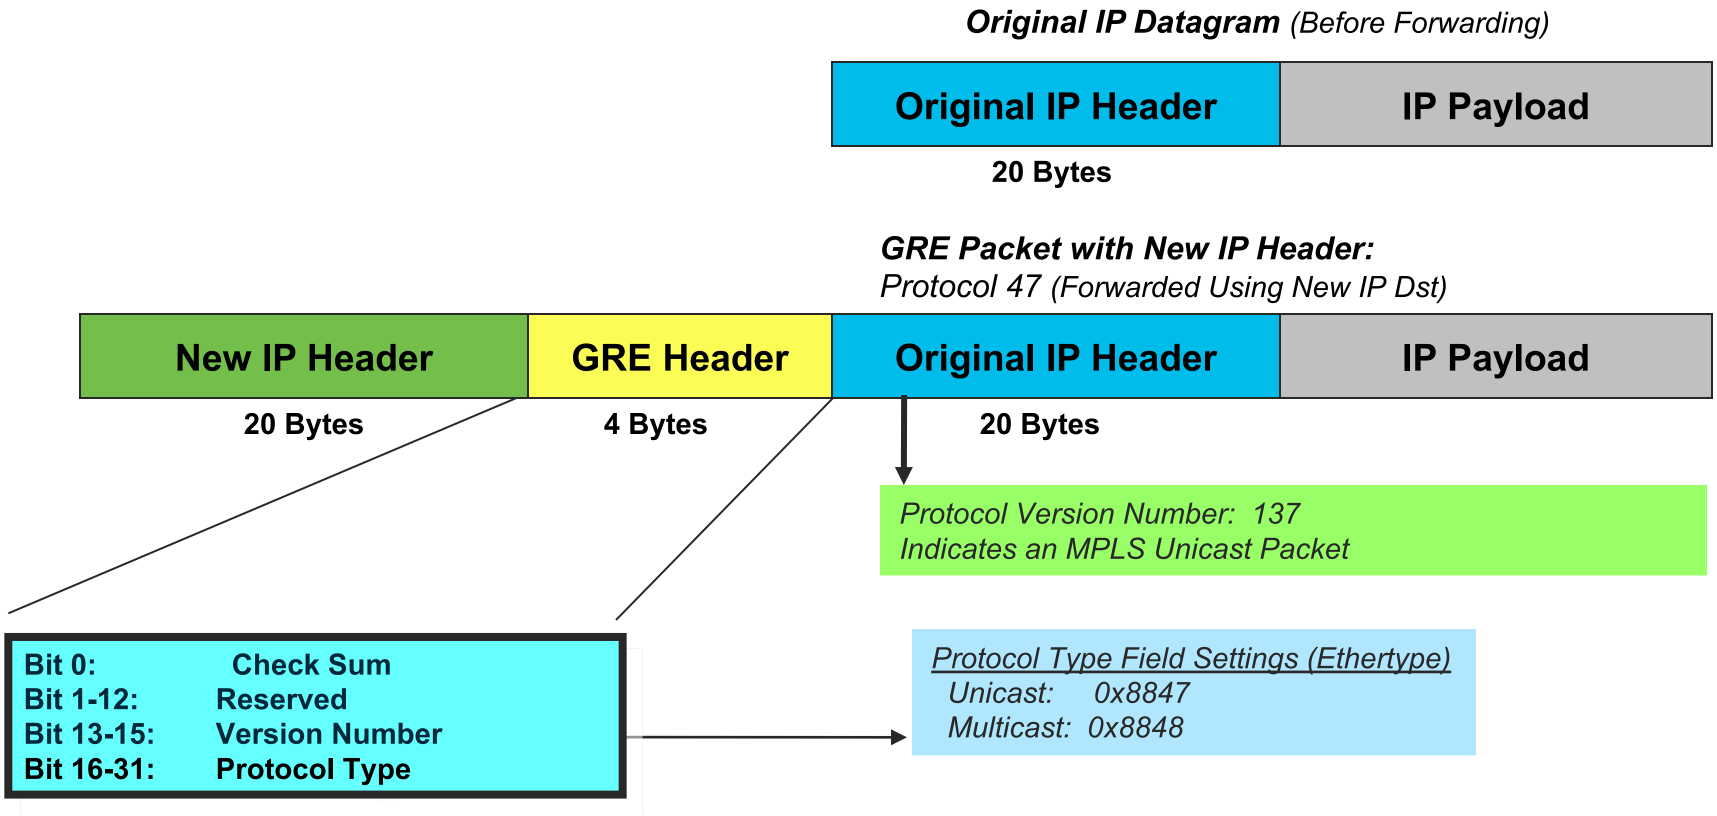

GRE Tunnel Format with MPLS (Reference: RFC 4023)

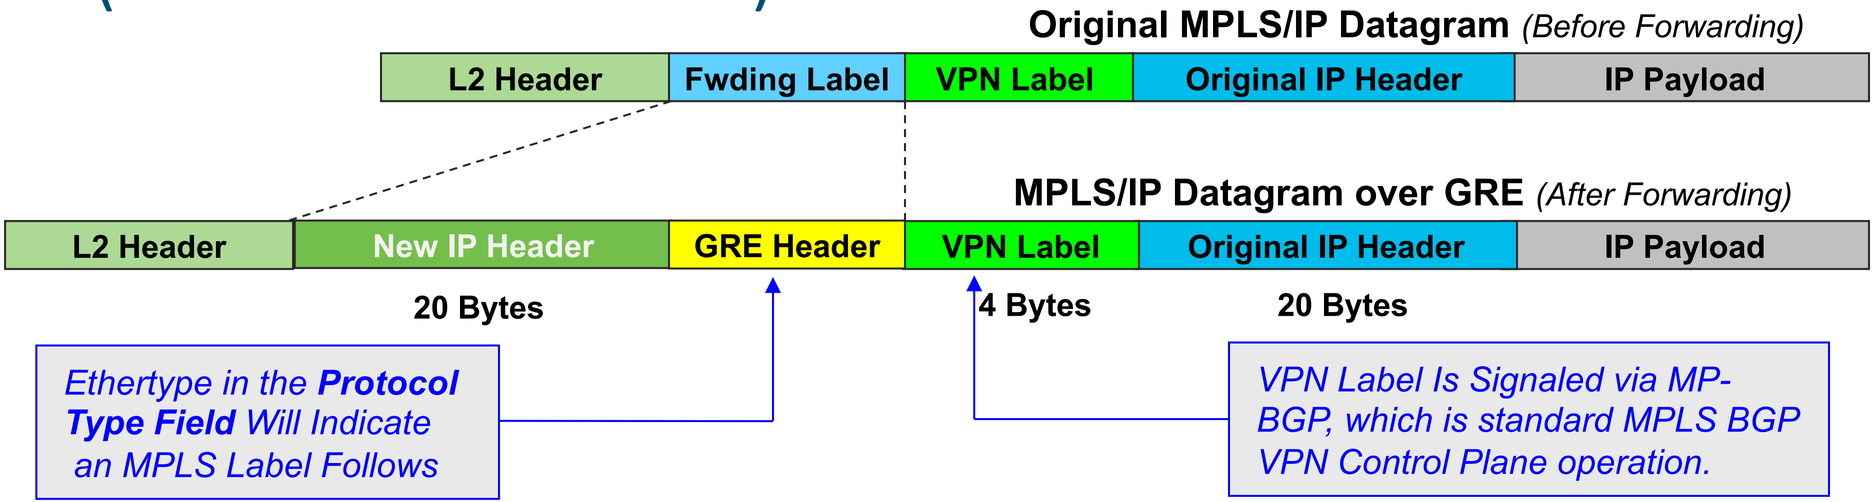

- MPLS Tunnel label (top) is replaced with destination PE’s IP address
- Encapsulation defined in RFC 4023
- Most widely deployed form of MPLS over IP encapsulation

GRE Tunnel Modes

“Stateful” vs. “Stateless” GRE Tunnelling

Point-to-Point GRE

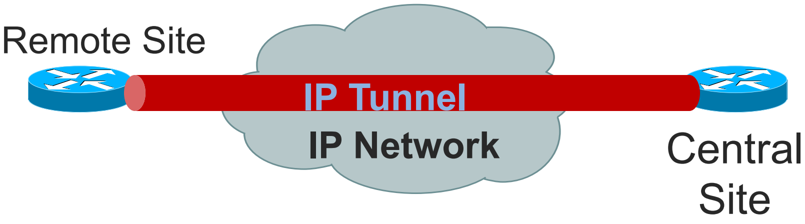

- Source <u>and</u> destination requires manual configuration
- Tunnel end-points are stateful neighbours
- Tunnel destination is explicitly configured
- Creates a logical point-to-point “Tunnel”
- IGP, BGP, and LDP/MPLS run through static tunnel

Multipoint GRE

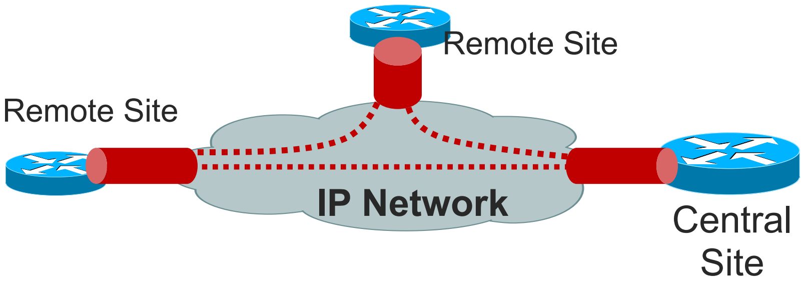

- **<u>Single</u>** multipoint tunnel interface is created per node
- Only the tunnel <u>source</u> is defined
- Tunnel destination is derived dynamically
  - DMVPN – uses NHRP
  - MPLS VPN over mGRE – uses BGP
- Creates an “encapsulation” using IP headers (GRE)

Enhancing the L3 VPN Segmentation Portfolio…

- **VRF Lite Options**
  - Leverage Carrier Ethernet E-LINE/E-LAN services
  - VRF-Lite Over GRE and DMVPN (multipoint)
- **L3 MPLS BGP VPN (RFC 4364)**
  - Over L2 transport (PE-PE, P-P, PE-P)…
  - Over p2p GRE tunnels (operationally not scalable)
- **L3 MPLS BGP VPN (RFC 4364) over DMVPN**
  - MPLS VPN over p2p GRE tunnels (operationally not scalable)
- **MPLS VPN (RFC 4364) over multipoint GRE (mGRE)**
- **Cisco SD-WAN (Viptela) – Secure L3 VPN Segmentation (SD-WAN)**

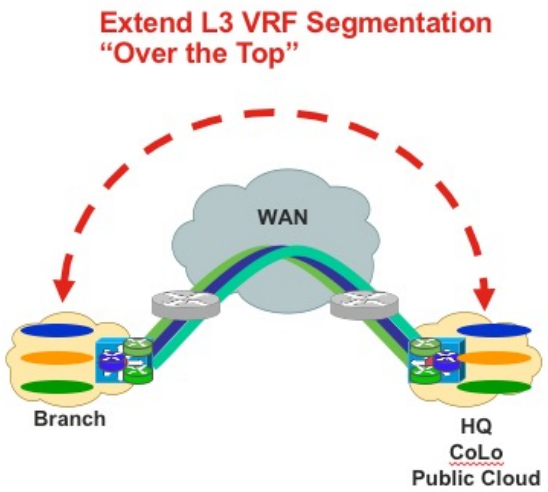

MPLS VPN over Multi Point GRE (mGRE)

Private MPLS VPN "over the top" of SP Offered IP VPN Transport owns CE

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

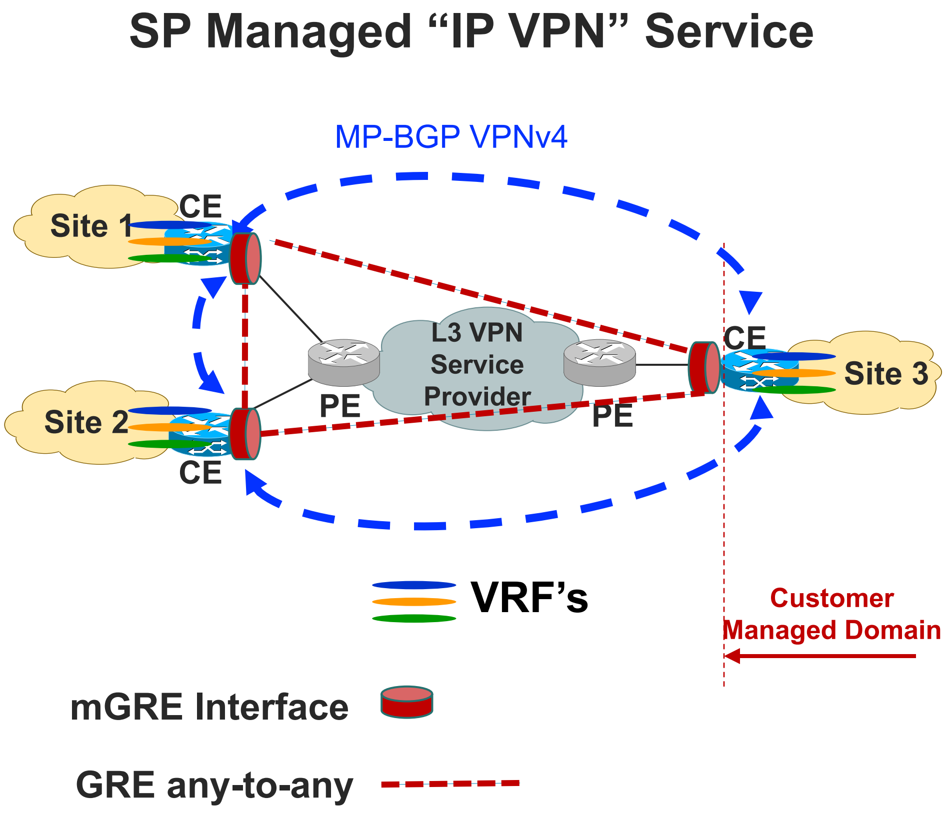

MPLS VPN over Multipoint GRE (mGRE)

Control/Data Plane Example over Service Provider Model

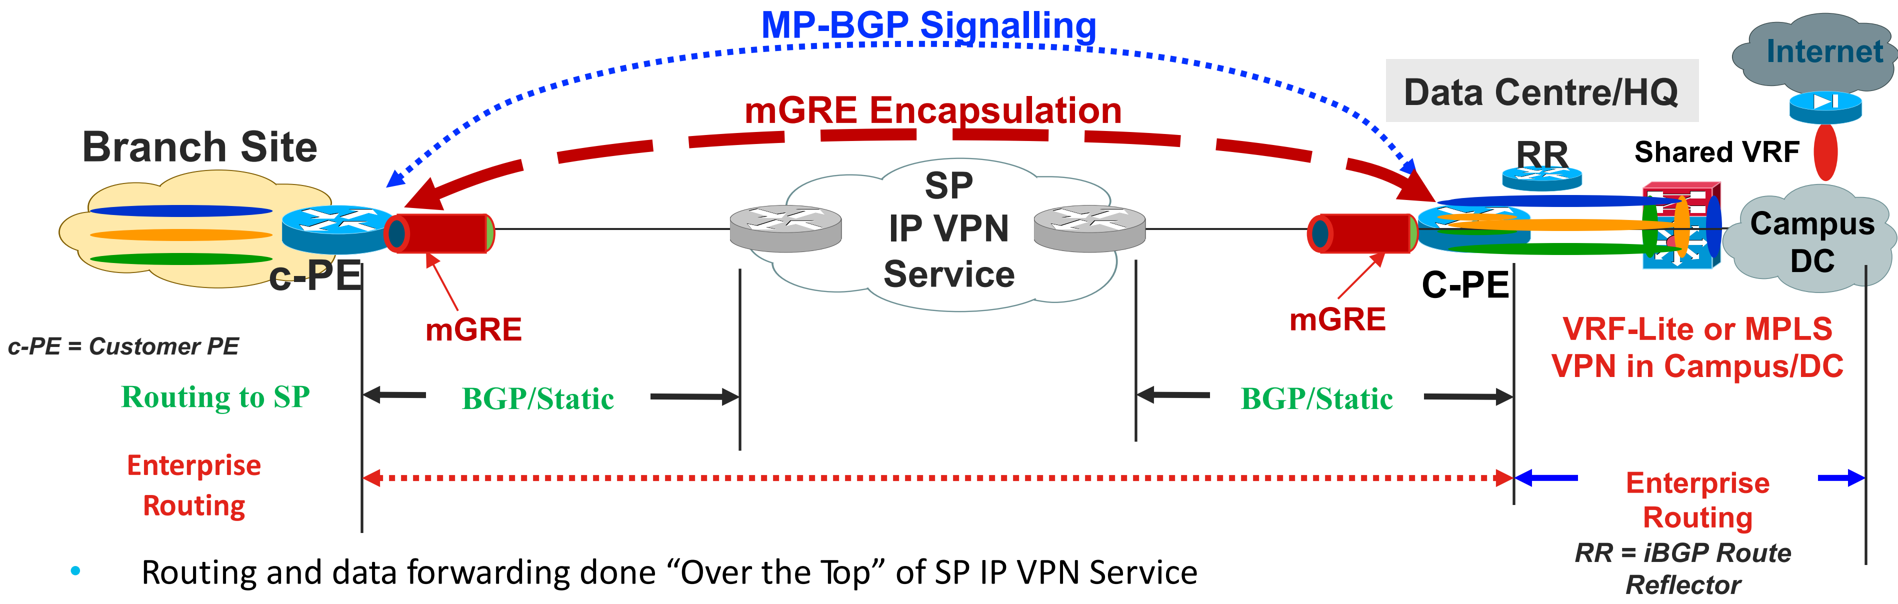

- Routing and data forwarding done “Over the Top” of SP IP VPN Service
- iBGP: (1) Advertise VPNv4 routes, (2) exchange VPN labels
- eBGP: (1) exchange tunnel end point routes with SP (or directly connected)
- Requires advertising a SINGLE IP prefix to SP (e.g. IP tunnel “end points”)

MPLS VPN over mGRE Model

mGRE Interface is Dynamic and De-coupled from Physical Interfaces

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

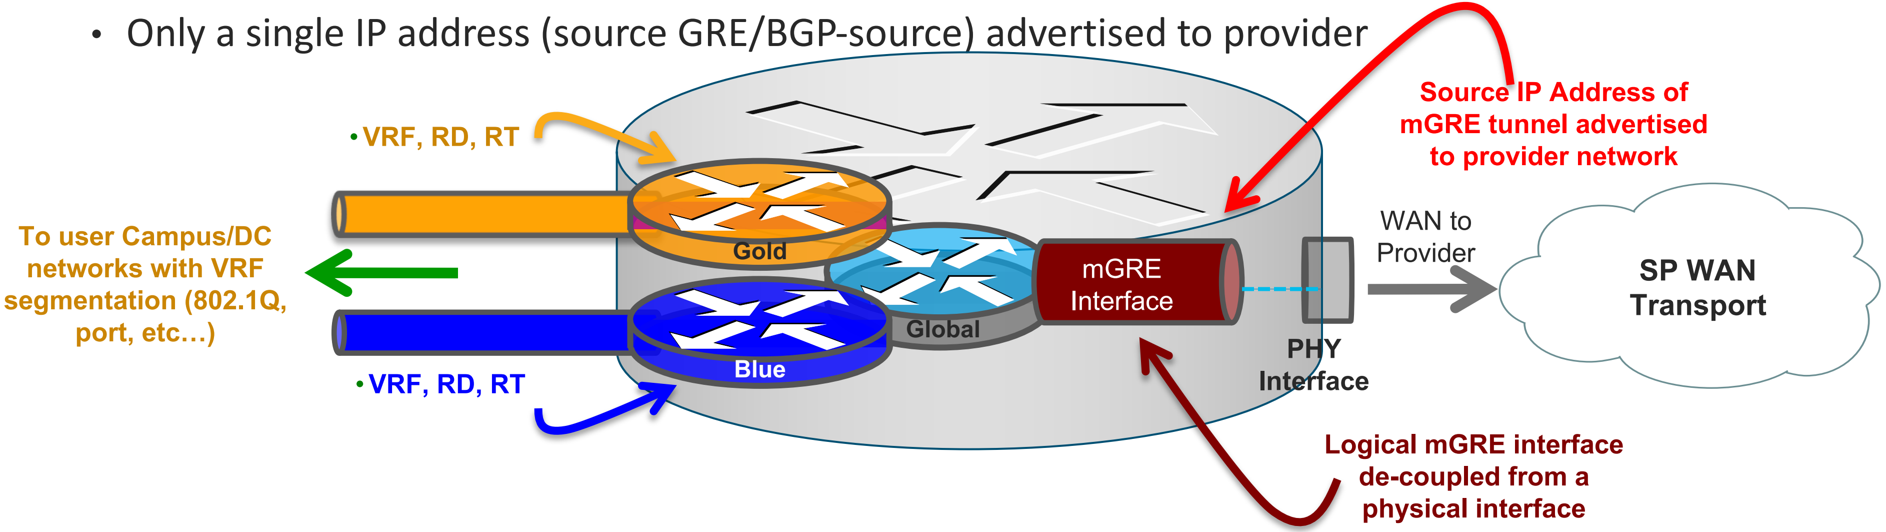

MPLS VPN over Multipoint GRE (mGRE)

VPNv4 Configuration Example

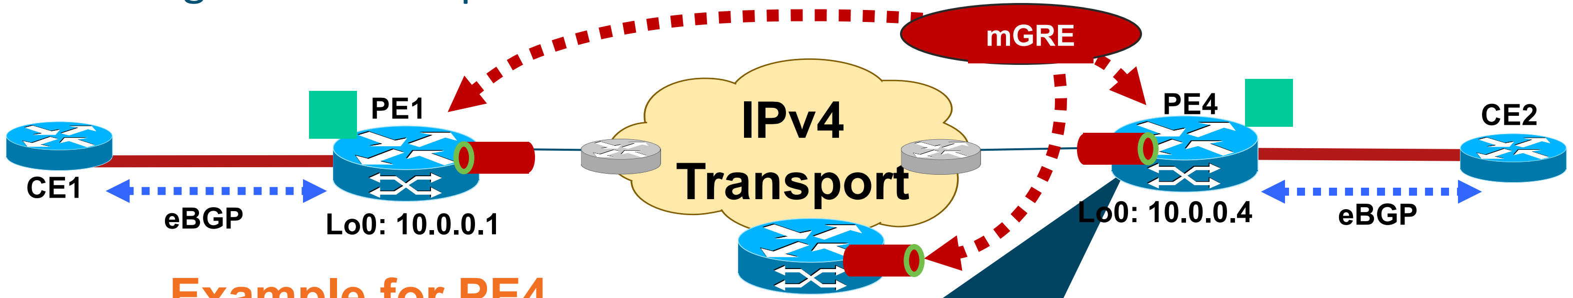

Example for PE4

```
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

Summary

MPLS VPN over Multipoint GRE (mGRE)

GitHub Repo Location
https://github.com/netwrkr95/mpls-mgre-configs

- Simple:
  - Only requires advertising a single IP prefix to SP for mGRE operation
  - **Dynamic** Tunnel endpoint discovery is done via iBGP/route-map (no static GRE tunnel)
  - Solution requires **NO manual configuration of GRE tunnels. LDP NOT required!**
  - E-BGP can/is still be used for route exchange (mGRE end-point) with the SP
- Standards Based - Leverages standard MP-BGP control plane (RFC 4364)
- Flexible - Supports MVPN and IPv6 per MPLS VPN model (MDT and 6vPE respectfully)
- Multi-platform support:
  - ASR 1000 series, ISR/G2, ISR 4xxx, SUP-2T, Cloud Services Router (CSR), Catalyst 8000 ☺
- Supports Inter-AS VPN, Multicast VPN (MVPN), standard QoS/H-QoS
- Supports IPSec for PE-PE encryption (GET VPN or manual SA)
- Scales to 2000 PE's with ASR 1000 series

White Paper

Secure Extension of L3 VPN's over IP-Based 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 innovation can be used to simplify Layer 3 (L3) Virtual Private Network (VPN) deployment and operations within secure government, commercial, and enterprise networks.

The key 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], which is a superset of suite B[1]). 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 segmentation is required.

This paper compares the use of MPLS VPN over the WAN with other network virtualization technologies typically deployed today. It also highlights the advantages of Cisco GET VPN over multipoint IP tunnel-based overlay

Configuration Examples on Github:

https://github.com/netwrkr95/mpls-mgre-configs

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

Сегментация с Cisco SD-WAN

SD WAN Deployment Options

Backbone SD-Core

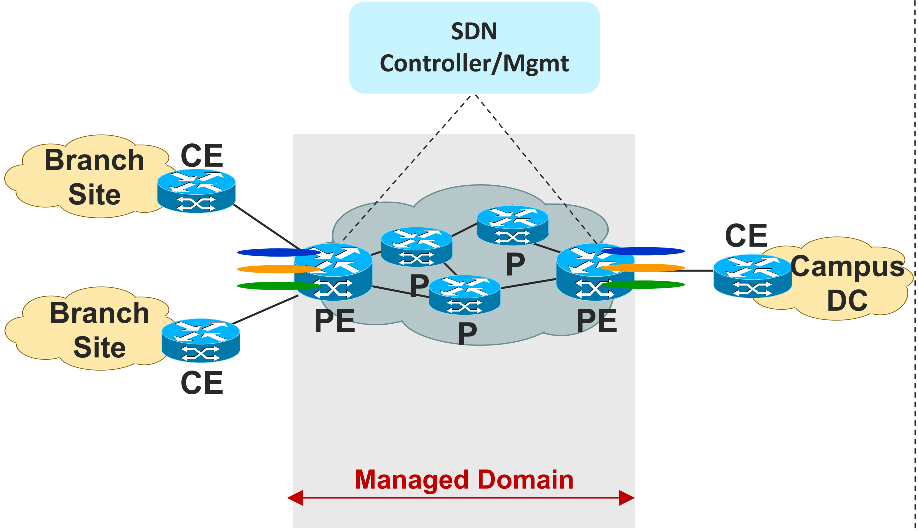

- Targets Service Provider "like" customers who need to control SLA's, rapid service turn up times, tighter granular service options (SR-TE), end-to-end control, provisioning, and visibility
- SR, SR-TE, Centralized WAN controller

Enterprise SD-WAN

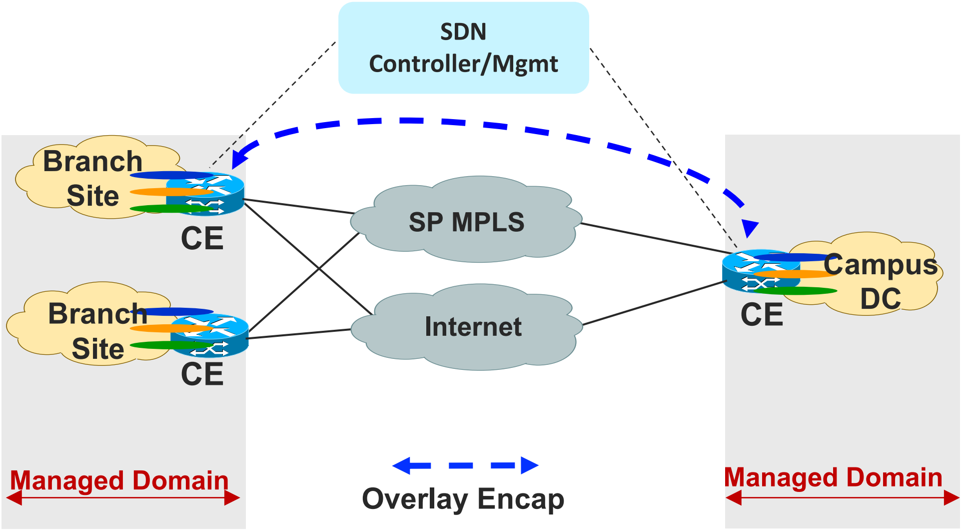

- Targets enterprise customers looking to consume secure WAN transport, with central mgmt., control, and application visibility
- Cisco SD-WAN, MPLS VPN over IP (central controller and/or open tools for automation)

Cisco SD-WAN (Viptela) L3 VPN Segmentation

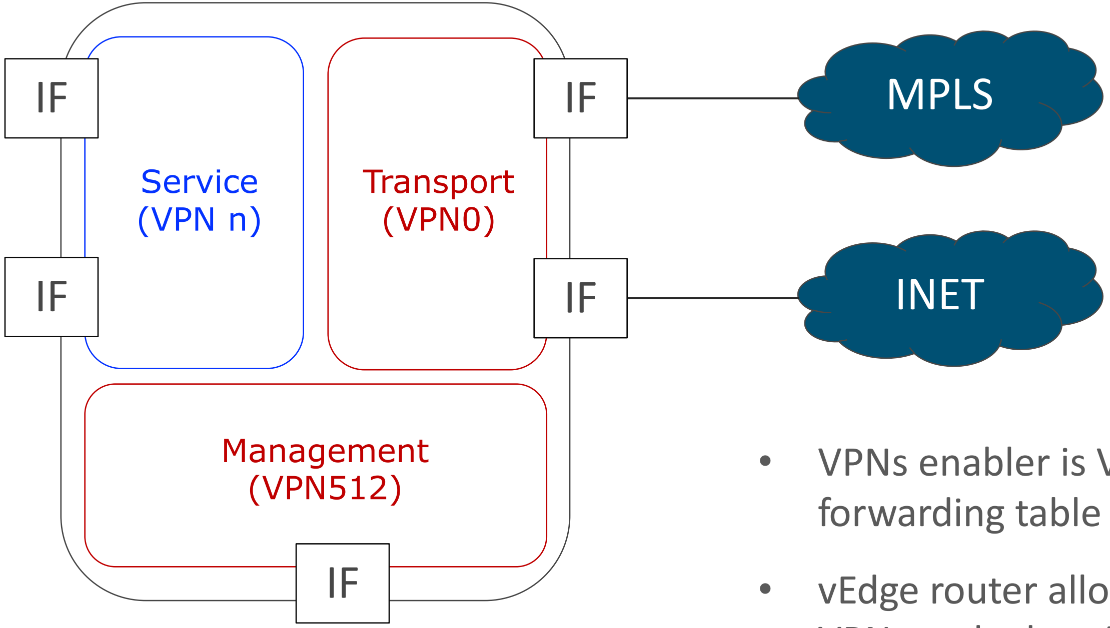

- VPN 0: Transport (locked)
- VPN 512: Mgmt (locked)
- VPN n: open user VPN

- VPNs enabler is VRF's, each VRF having its own forwarding table
- vEdge router allocates label to each of it's service VPNs and advertises it as route attribute in OMP updates
  - VPN Labels used to identify customer VPN in the incoming packets

Secure Segmentation

End-to-End Segmentation

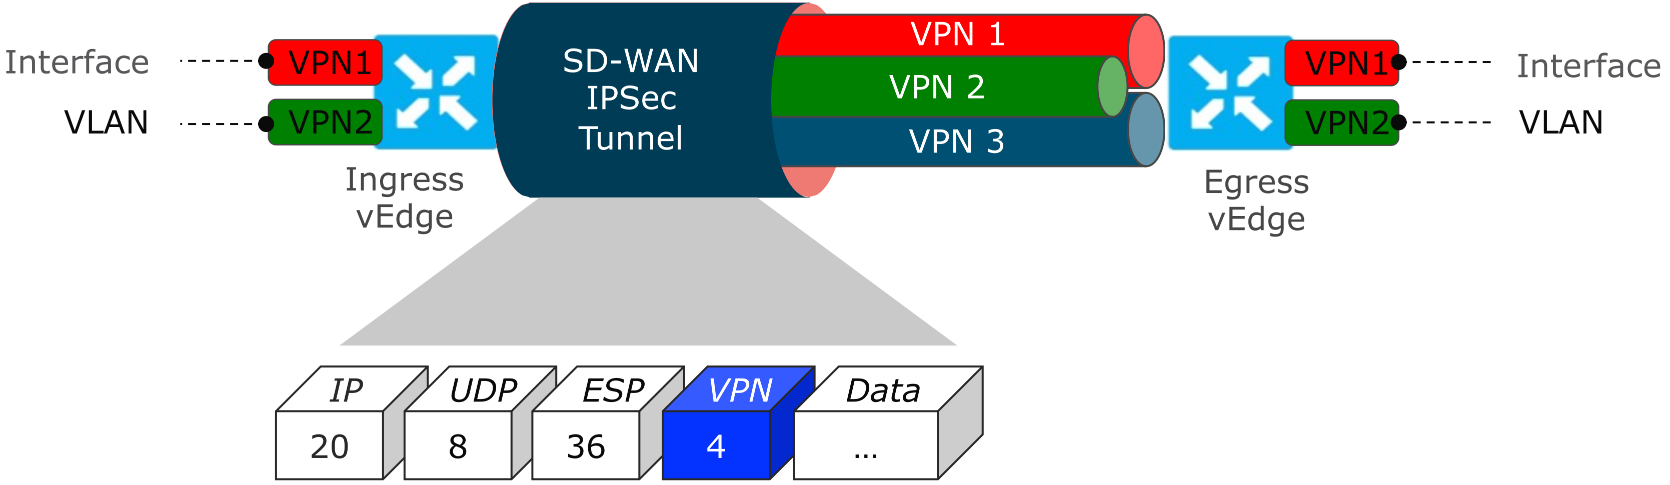

- Segment connectivity across fabric w/o reliance on underlay transport
- vEdge routers maintain per-VPN routing table
- Labels are used to identify VPN for destination route lookup
- Interfaces and sub-interfaces (802.1Q tags) are mapped into VPNs

Cisco SD-WAN (Viptela) L3 Segmentation

Per L3 VPN Topology and Mapping

- Isolated virtual private networks across any transport
- VPN mapping is based on physical vEdge Router interface, 802.1Q VLAN tag or a mix of both
- VPN isolation is carried over all transports
  - https://tools.ietf.org/html/rfc4023

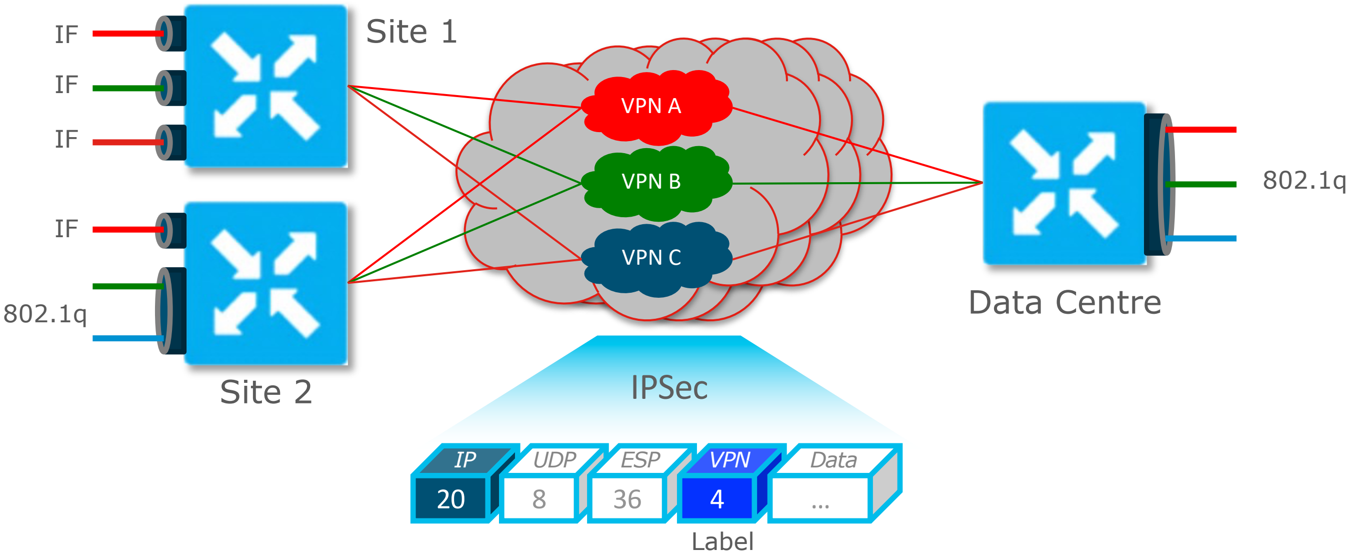

Per L3 VPN Topology (Examples)

Full-Mesh

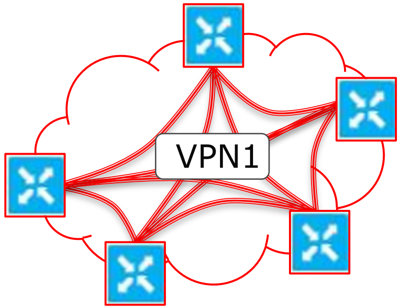

Hub-and-Spoke

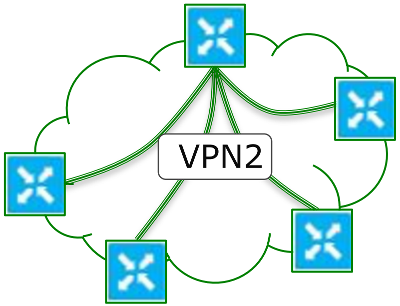

Partial Mesh

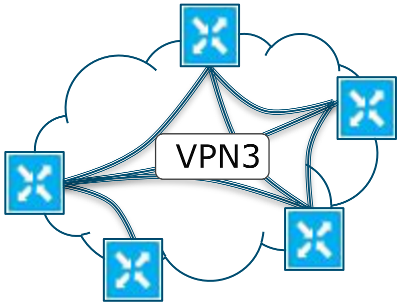

Point-to-Point

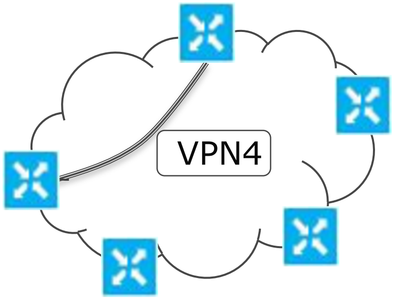

- Each VPN can have it's own topology
  - Full-mesh, hub-and-spoke, partial-mesh, point-to-point, etc…
- VPN topology is influenced by leveraging control policies
  - Filtering TLOCs or modifying next-hop TLOC attribute for routes
- Customer mission, business, and applications can drive a certain topology:
  - Applications in single cloud or on-prem can benefit from hub-spoke
  - voice takes full-mesh topology
  - Security compliance - PCI data takes hub-and-spoke topology

Multi-Topology

Secure and Flexible Traffic Forwarding Options

- Arbitrary per-VPN topology
- Topology reflects desired traffic forwarding patterns, e.g. voice and video full-mesh, business apps hub-and-spoke
- vSmart controls VPN topology through control plane advertisements
- vEdge routers can participate in multiple topologies at the same time

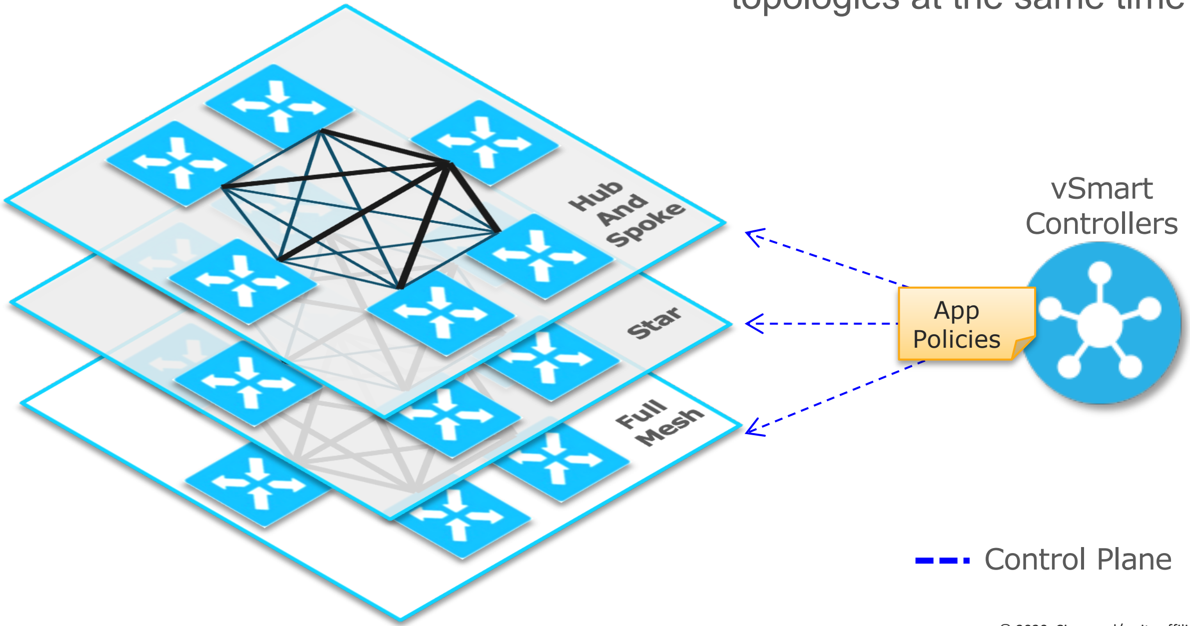

Common Use Cases for L3 VPN over SD-WAN

- State, Country or Global based MPLS VPN where transport option is IP only
- The ”business requirement” mandates segmentation (refer to L3 segmentation use cases)
  - L3 VPN + encryption
  - L3 VPN over (e.g. transparently) non-MPLS (e.g. IP) transport, including Internet
  - L3 VPN Manages Services offering (managed CPE over L3 VPN/IP transport)
  - L3 VPN over proprietary encryption (external) devices (Government, Defense)
  - L3 VPN extension into the public cloud (per application segmentation)
- Extend Campus/DC ”policy” over the WAN
  - Cisco SD-Access = VN / Cisco ACI = VRF / Cisco SD-WAN = VPN
- Targets customers requiring “on-demand, self-deployed” L3 VPN turn-up

Summary of L3 VPN over IP WAN Techniques

Strengths / Weaknesses to help evaluate decision criteria

	MPLS VPN o mGRE	Cisco SD-WAN (Viptela)
Routers only (no controller req)	🟢	🔴
Controller Based (central) routing calculations	🟡	🟢
Native VPN Multicast (MVPN)	🟢	🟡
Application Awareness	🟡	🟢
Transport Agnostic (Internet)	🔴	🟢
Large Scale VRF (>64)	🟢	🟡
"SD-WAN" Requirement (RFP)	🔴	🟢
Per VPN Topology (p2p, mesh)	🔴	🟢

🟢 Excellent Option

🟡 SubOptimal Option

🔴 Bad Option

Enterprise WAN L3 Segmentation Solutions

Let's Recap

- Fully understand the application and network service requirements needed
  - Pace of Service turn-up times, transport available, operational expertise
- **Self Deployed MPLS backbone target**:
  - larger-scale, TE required, L2 VPN, tight control
- **Layer 3 Segmentation over IP:**
  - **MPLS VPN over mGRE**: simple MPLS VPN over IP, customer not ready for full-blown SD-WAN yet
  - **Cisco SD-WAN**: applications scattered across multiple locations (on-prem, public cloud, SaaS), leverage Internet as transport, cloud managed controller interest
- Assure the solution chosen suits the operational skill set of the IT org
- Keep is simple whenever possible

QoS

Enabling QoS

Elements that Affect End-to-End Delay

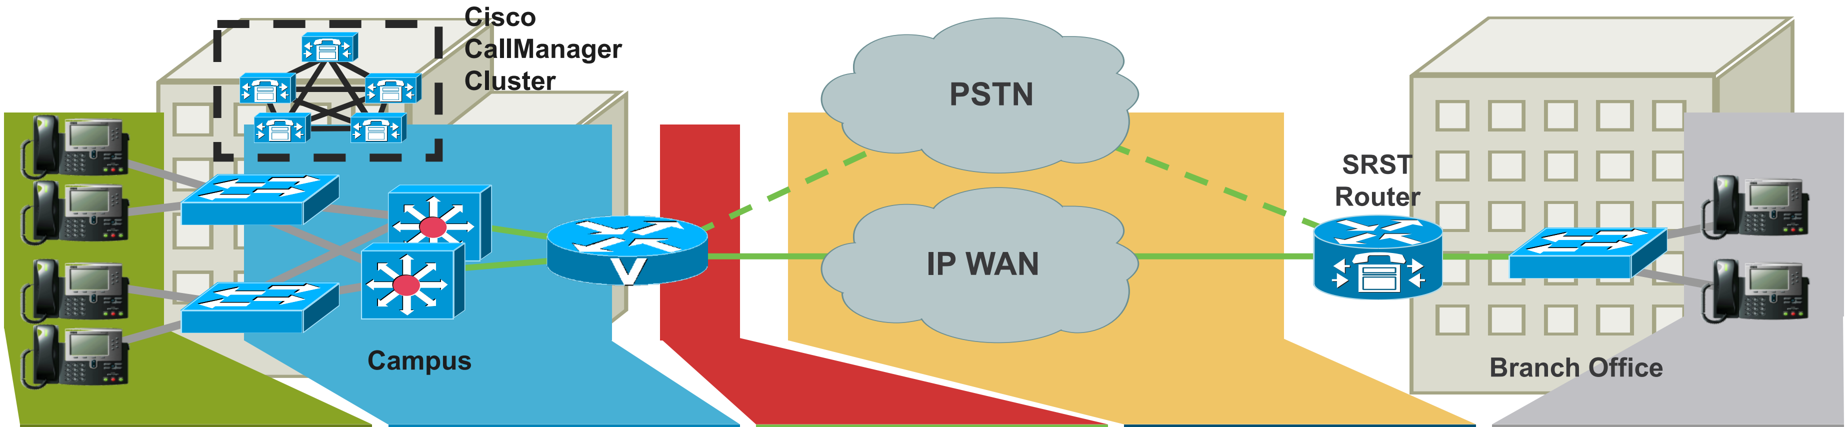

CODEC	Queuing	Serialization	Propagation and Network	Jitter Buffer
G.729A: 25 ms	Variable (Can Be Reduced Using LLQ)	Variable (Can Be Reduced Using LFI)	6.3 µs/Km + Network Delay (Variable)	20–50 ms

End-to-End Delay (Should Be < 150 ms)

The WAN Is the Barrier to Branch
Application Performance

- Applications are designed to work well on LAN's
  - High bandwidth
  - Low latency
  - Reliability
- WANs have opposite characteristics
  - Low bandwidth
  - High latency
  - Packet loss

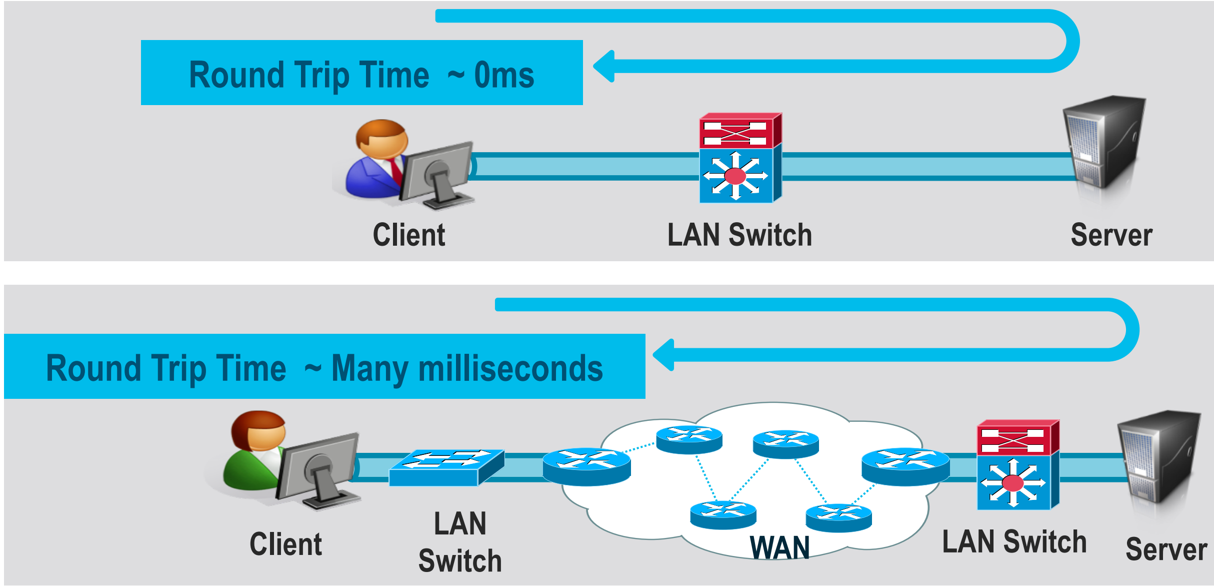

WAN Packet Loss and Latency =
Slow Application Performance =
Keep and manage servers in branch offices ($$$)

Enabling QoS in the WAN

Traffic Profiles and Requirements

Voice

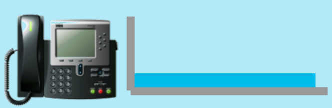

- Smooth
- Benign
- Drop sensitive
- Delay sensitive
- UDP priority

Bandwidth per call depends on codec, Sampling-Rate, and Layer 2 Media

- Latency ≤ 150 ms
- Jitter ≤ 30 ms
- Loss ≤ 1%
- Bandwidth (30-128Kbps)
- One-Way Requirements

SD Video Conf

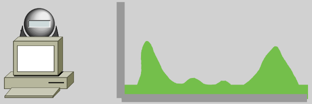

- Bursty
- Greedy
- Drop sensitive
- Delay sensitive
- UDP priority

SD/VC has the same requirements as VoIP, but traffic patterns and BW varies greatly

- Latency ≤ 150 ms
- Jitter ≤ 30 ms
- Loss ≤ 0.05%
- Bandwidth (1Mbps)
- One-Way Requirements

Telepresence

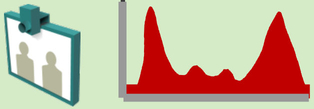

- Bursty
- Drop sensitive
- Delay sensitive
- Jitter sensitive
- UDP priority

HD/VC has tighter req's than VoIP for jitter and BW varies based on the resolutions

- Latency ≤ 200 ms
- Jitter ≤ 20 ms
- Loss ≤ 0.10%
- Bandwidth (5.5-16Mbps)
- One-Way Requirements

Data

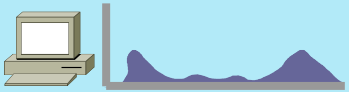

- Smooth/bursty
- Benign/greedy
- Drop insensitive
- Delay insensitive
- TCP retransmits

Traffic patterns for Data vary across applications

Data Classes:

- Mission-Critical Apps
- Transactional/Interactive Apps
- Bulk Data Apps
- Best Effort Apps (Default)

QoS Tools and Techniques

Classifying and Marking

- Network Based Application Recognition (NBAR2)
- Application Visibility and Control (AVC)
- Layer 2 or 3 marking of CoS/EXP or DSCP/IP precedence

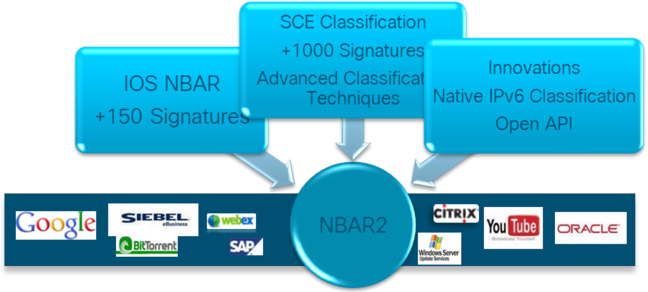

- New DPI engine provides Advanced Application Classification and Field Extraction Capabilities from Service classification engine

Policing and Markdown

- Define traffic metering contracts
- Markdown out-of-contract flows
- Conform, Exceed, Violate actions

Scheduling

- Re-order and selectively drop during congestion
- Class Based Weighted Fair Queuing (CBWFQ)
- Low Latency Queuing (LLQ) and Multi-LLQ

Link-specific tools

- Traffic Shaping and Hierarchical QoS (HQoS)
- Compression
- Fragmentation and Interleaving

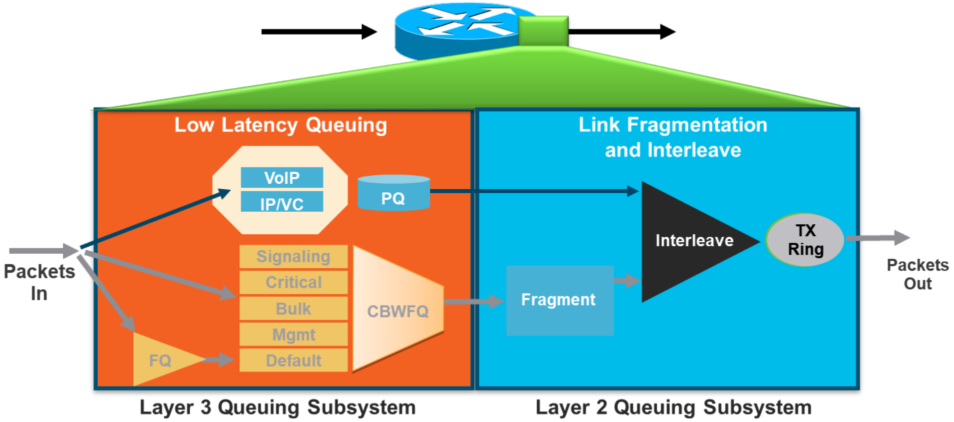

Hierarchical QoS For Subrate Service

H-QoS Policy on WAN Interface, Shaper = CIR

Two Levels MQC

```
Policy-map PARENT
  class class-default
   shape average 150000000
   service-policy output CHILD

Interface gigabitethernet 0/1
 service-policy output PARENT
```

```
Policy-map CHILD
 class VOICE
  priority percent 10
 class VIDEO
  priority percent 23
 class CRITICAL-DATA
  bandwidth percent 15
  random-detect dscp-based
 class DATA
  bandwidth percent 19
  random-detect dscp-based
 class SCAVENGER
  bandwidth percent 5
 class NETWORK-CRITICAL
  bandwidth percent 3
  service-policy MARK-BGP
 class class-default
  bandwidth percent 25
  random-detect
```

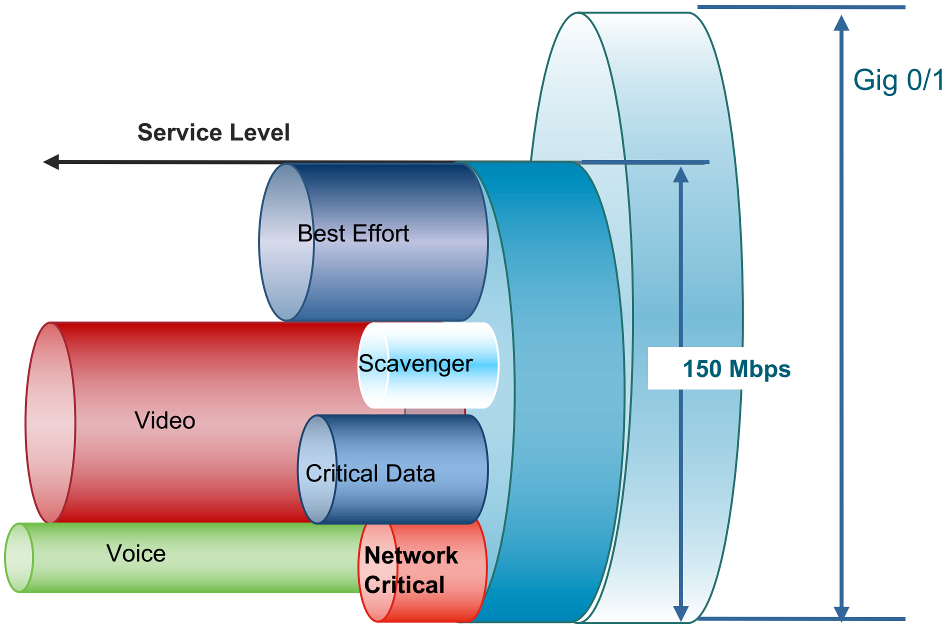

SP-Managed MPLS Services

- Enterprise customers may need to re-mark traffic prior to forwarding to the MPLS provider. This ensures markings conform to the admission criteria defined by the provider, allowing traffic to be serviced by the appropriate queue within the provider network.
- The same concept applies to traffic ingression the enterprise network from the provider cloud.
- Certain applications may need to be re-marked to ensure the enterprise QoS strategy is properly applied.

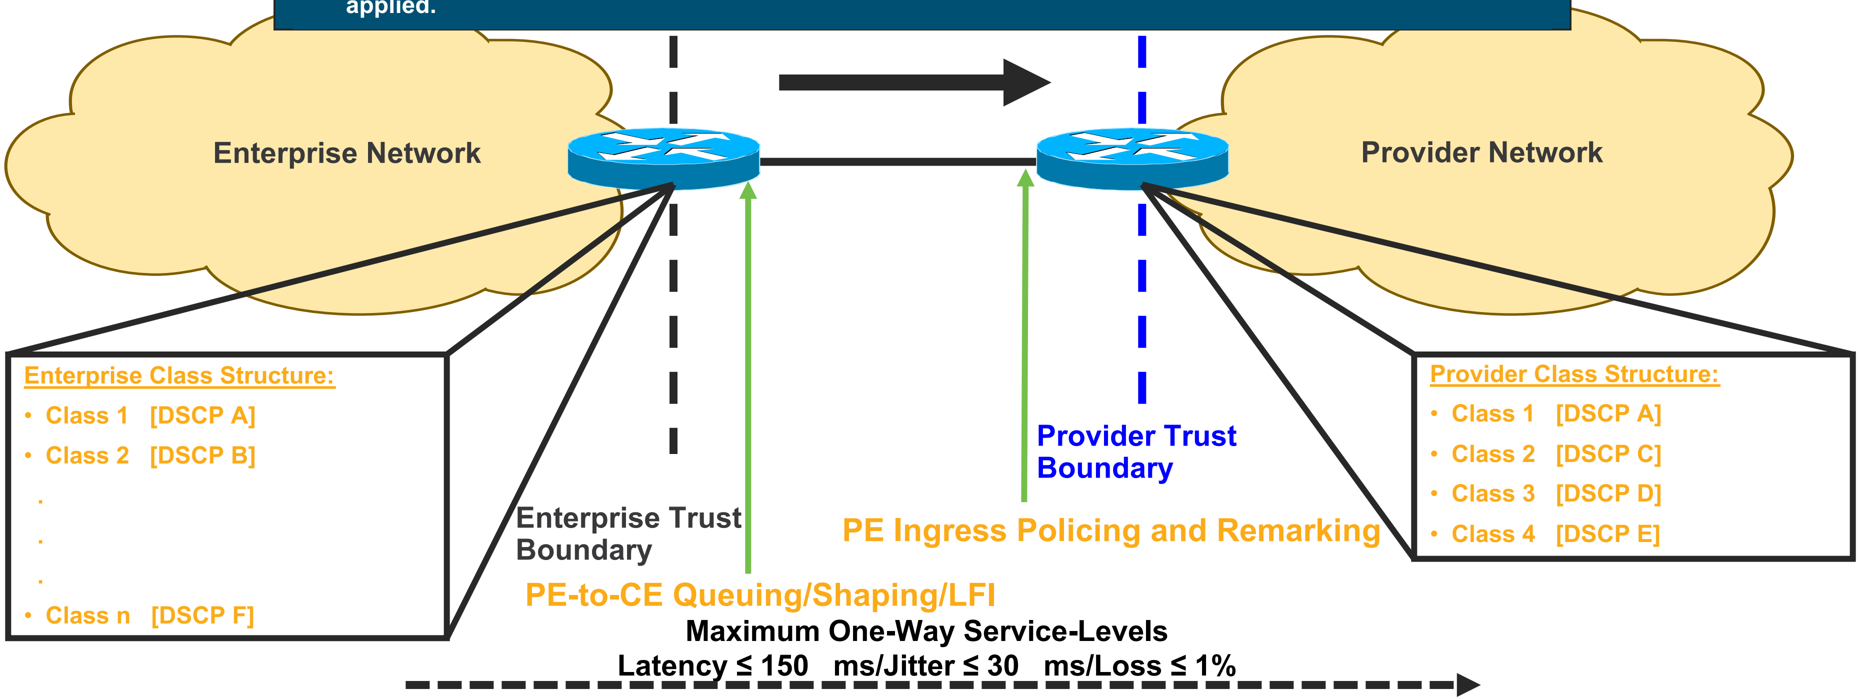

Enterprise-to-Service Provider Mapping

Five-Class Provider-Edge Model Remarking Diagram

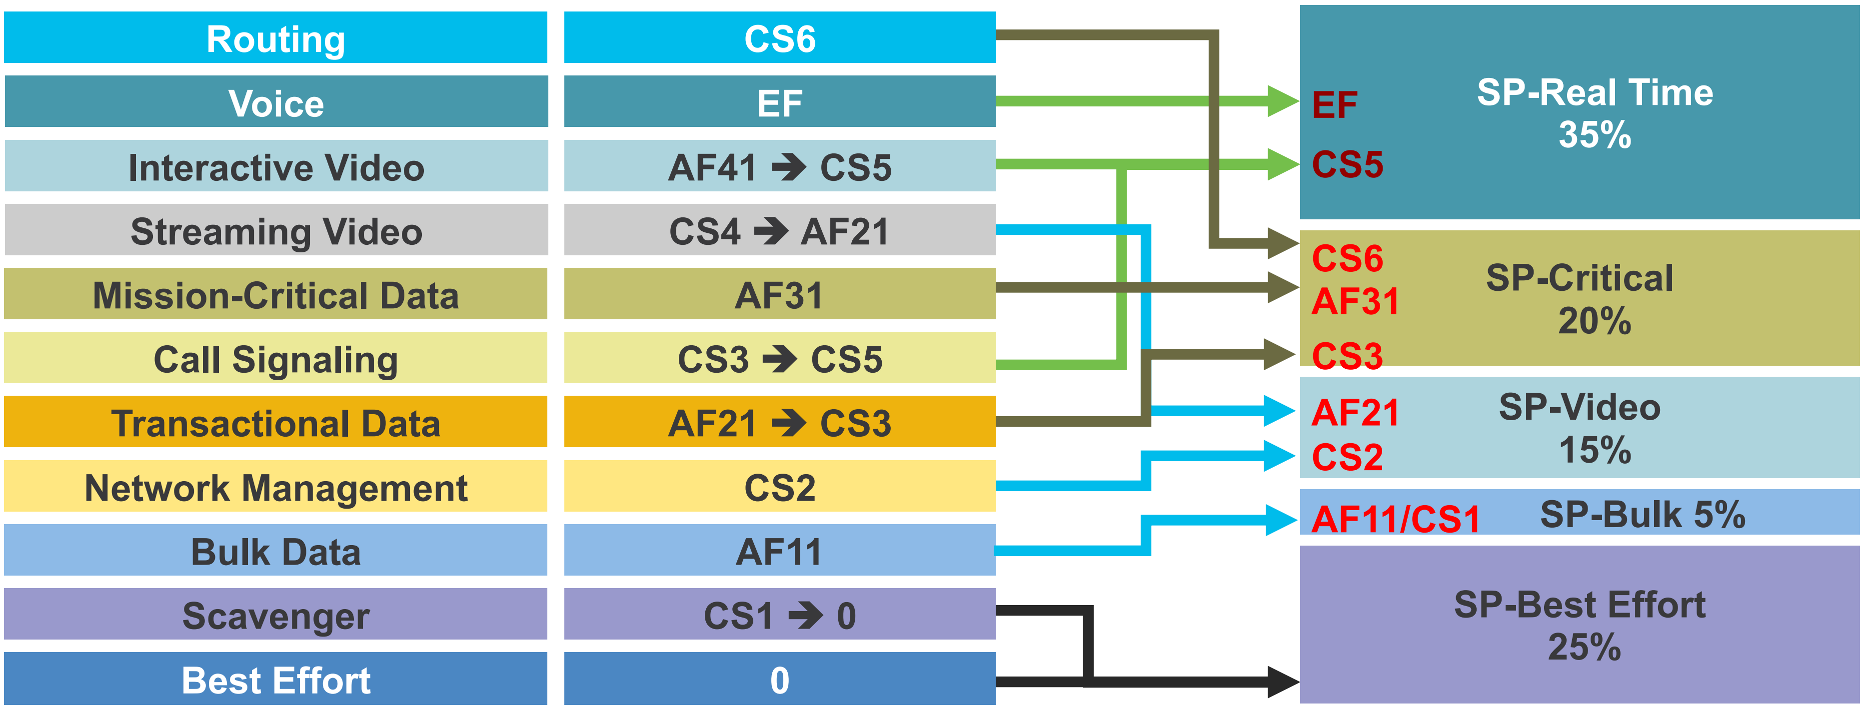

MPLS Short Pipe Mode DiffServ Tunneling

Short Pipe Mode Operation

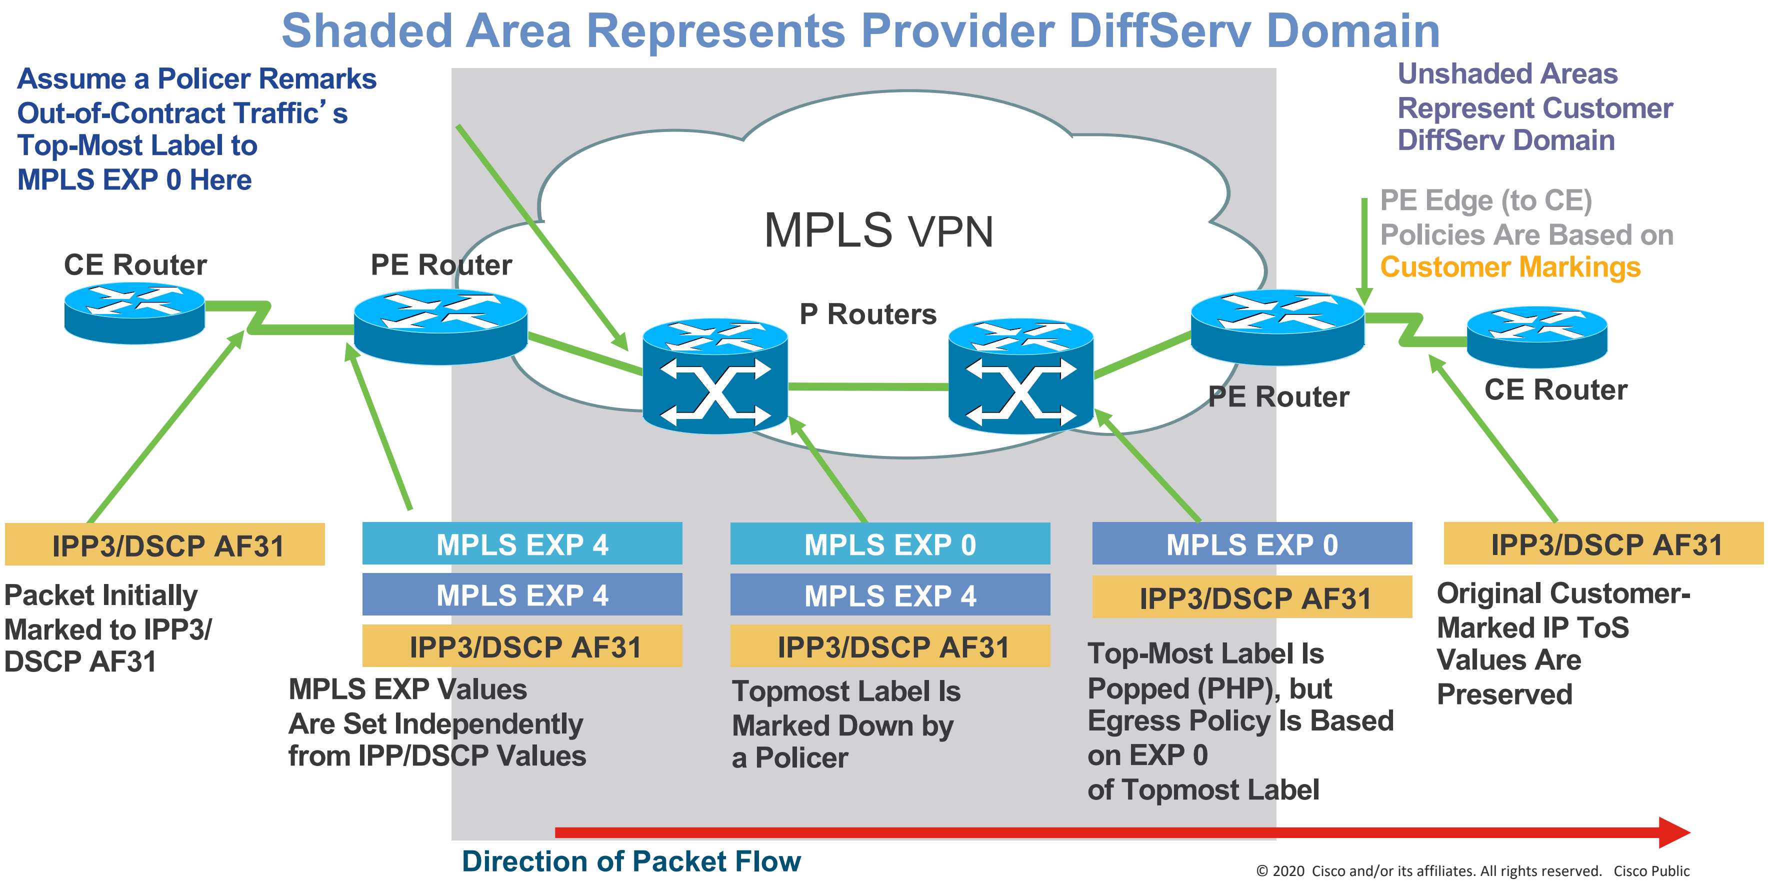

MPLS VPN QoS Considerations

MPLS VPN Port QoS Roles

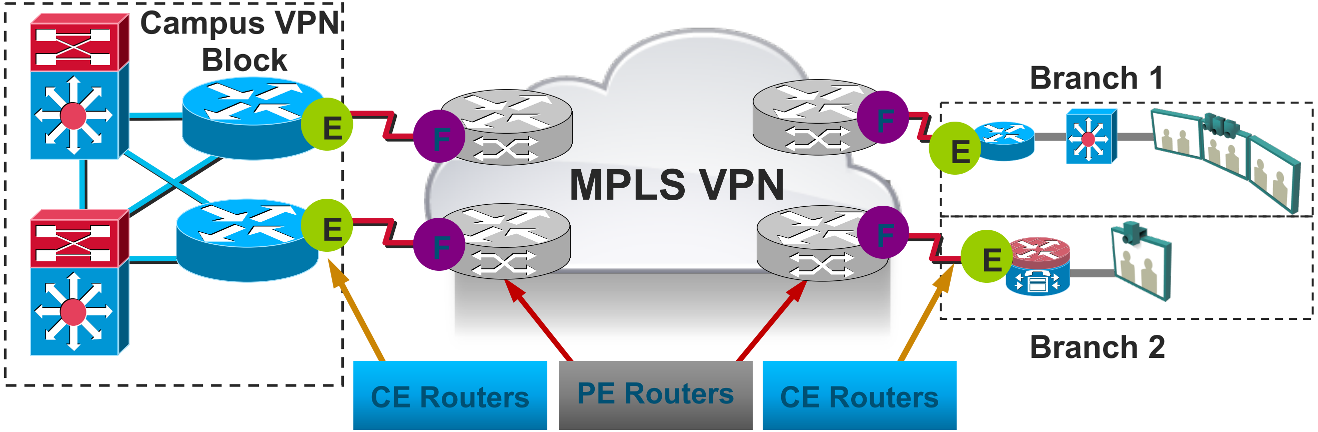

Enterprise Subscriber (Unmanaged CE Routers)

E **Outbound Policies:**

HQoS Shaper (if required)

≤ 33% of BW:
- \+ LLQ for VoIP (EF)
- \+ LLQ or CBWFQ for RT-Interactive (CS4)

\+ Remark RTI (if necessary)

\+ CBWFQ for Signaling (CS3)

\+ Remark Signaling (if necessary)

Inbound Policies:

Trust DSCP

\+ Restore RT-Interactive to CS4 (if necessary)

\+ Restore Signaling to CS3 (if necessary)

Service Provider:

F **Outbound Policies:**

\+ LLQ for Real-Time

\+ CBWFQ for Critical Data

Inbound Policies:

Trust DSCP

Police on a per-Class Basis

WAN QoS – Implementing Per Site Traffic Shaping

500 Mbps in to WAN can easily overrun the lower speed committed rates at remote sites

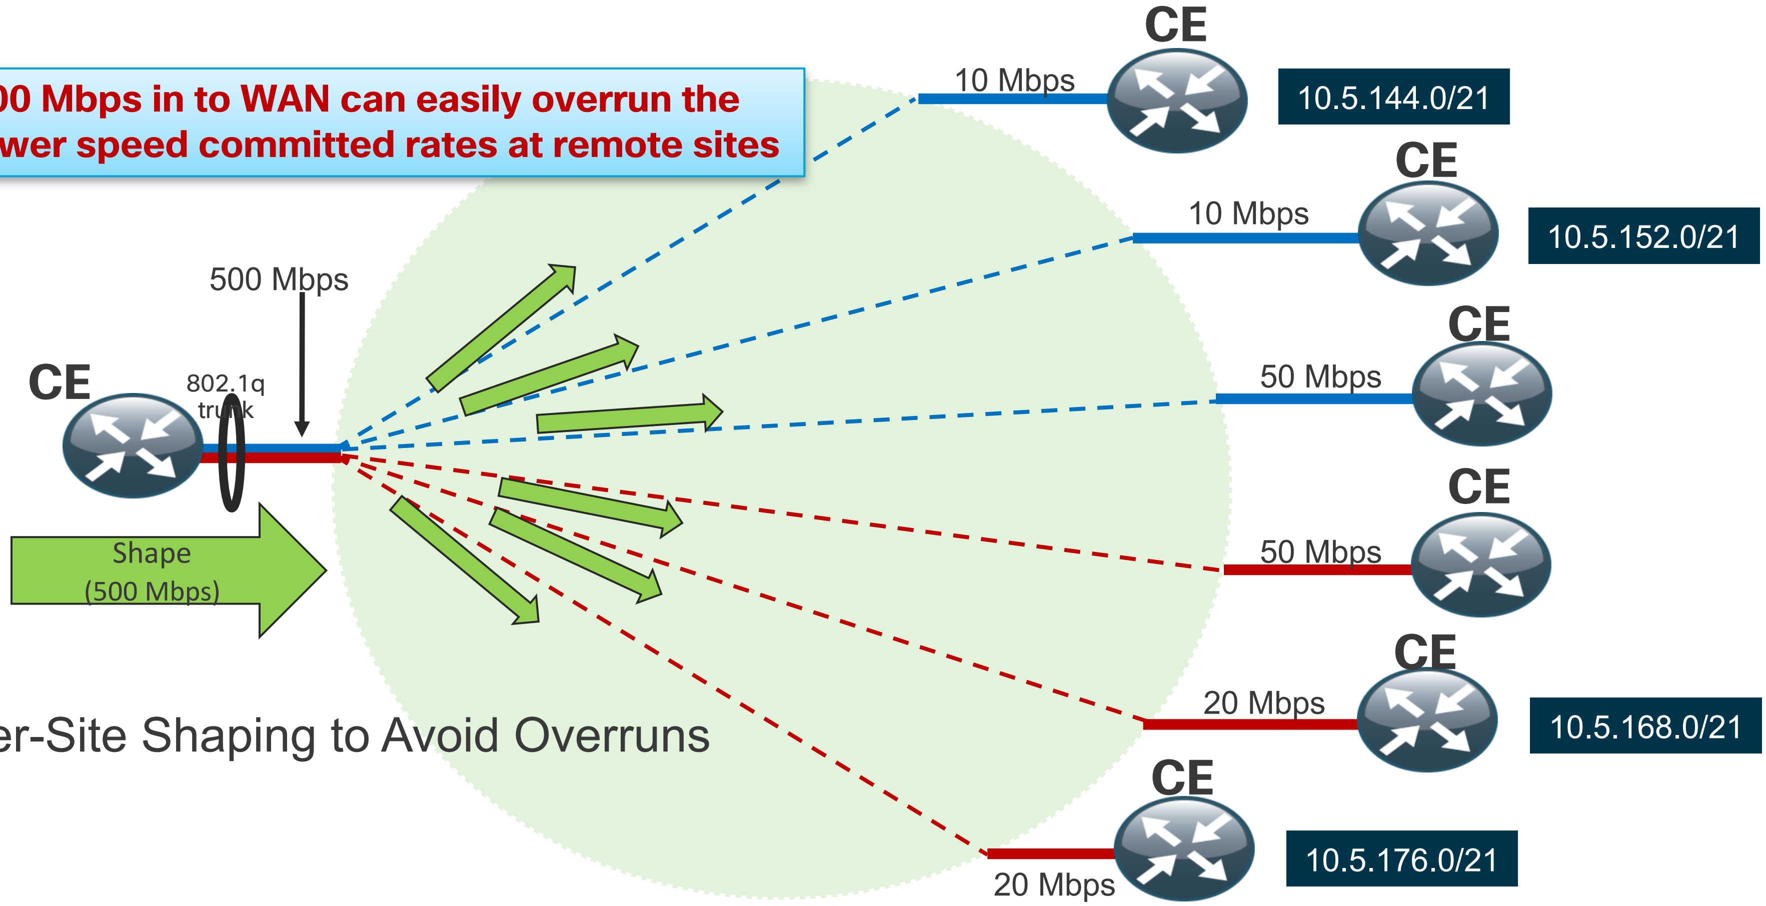

Per-Site Shaping to Avoid Overruns

WAN Quality of Service:

Implementing Per Site Traffic Shaping

```
policy-map POLICY-MAP-Br210
 class VOICE
  priority percent 10
 class INTERACTIVE-VIDEO
  priority percent 23
 class CRITICAL-DATA
  bandwidth percent 15
  random-detect dscp-based
 class DATA
  bandwidth percent 19
  random-detect dscp-based
 class SCAVENGER
  bandwidth percent 5
 class NETWORK-CRITICAL
  bandwidth percent 3
  service-policy MARK-BGP
 class class-default
  bandwidth percent 25
  random-detect
```

Per Destination Service Policies

```
policy-map POLICY-MAP-Br212
 class VOICE
  priority percent 10
 class INTERACTIVE-VIDEO
  priority percent 23
 class CRITICAL-DATA
  bandwidth percent 15
  random-detect dscp-based
 class DATA
  bandwidth percent 19
  random-detect dscp-based
 class SCAVENGER
  bandwidth percent 5
 class NETWORK-CRITICAL
  bandwidth percent 3
  service-policy MARK-BGP
 class class-default
  bandwidth percent 25
  random-detect
```

```
ip access-list extended Br210-10.5.144.0
 permit ip any 10.5.144.0 0.0.7.255
!
class-map match-all CLASS-MAP-Br210
 match access-group name Br210-10.5.144.0
```

Per Destination Class Maps

```
ip access-list extended Br212-10.5.168.0
 permit ip any 10.5.168.0 0.0.7.255
!
class-map match-all CLASS-MAP-Br212
 match access-group name Br212-10.5.168.0
```

```
policy-map POLICY-MAP-L2-WAN-BACKBONE-WITH-PER-SITE-SHAPERS
 class NETWORK-CRITICAL
  bandwidth percent 3
 class CLASS-MAP-Br210
  shape average 10000000
  service-policy POLICY-MAP-Br210
 class CLASS-MAP-Br212
  shape average 20000000
  service-policy POLICY-MAP-Br212
```

Shape to 10 Mbps to BR210

Shape to 20 Mbps to BR212

WAN Quality of Service: Implementing Per Site Traffic Shaping

```
policy-map POLICY-MAP-L2-WAN-BACKBONE-WITH-PER-SITE-SHAPERS
 class NETWORK-CRITICAL
  bandwidth percent 3
 class CLASS-MAP-Br210
  shape average 10000000
  service-policy POLICY-MAP-Br210
 class CLASS-MAP-Br212
  shape average 20000000
  service-policy POLICY-MAP-Br212
```

Shape to 10 Mbps to BR210

Shape to 20 Mbps to BR212

```
policy-map WAN-INTERFACE-G0/0/4
 class class-default
  shape average 500000000
  service-policy POLICY-MAP-L2-WAN-BACKBONE-WITH-PER-SITE-SHAPERS
```

Shape to 500 Mbps aggregate

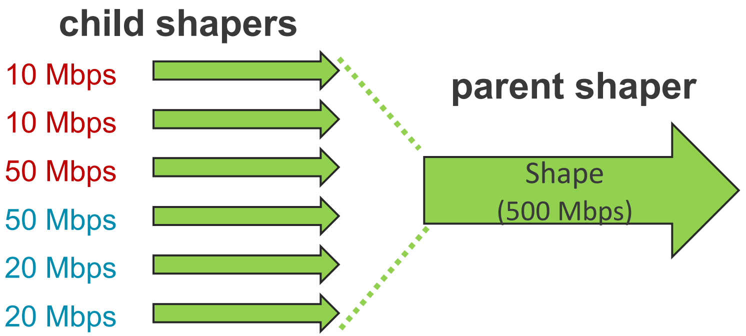

Per Site Traffic Shaping to Avoid Overruns

DMVPN Per-Tunnel QoS

- User NHRP group to dynamically provision HQoS policy on a DMVPN hub per-spoke basis

    Spoke: Configure NHRP group name

    Hub: NHRP group name mapped to QoS template policy

  Multiple spokes with same NHRP group mapped to individual instances of same QoS template policy

- GRE ,IPsec & L2 header are included in calculations for shaping and bandwidth.
- Queuing and shaping is performed at the outbound physical interface
- Can be used with DMVPN with or without IPSec.

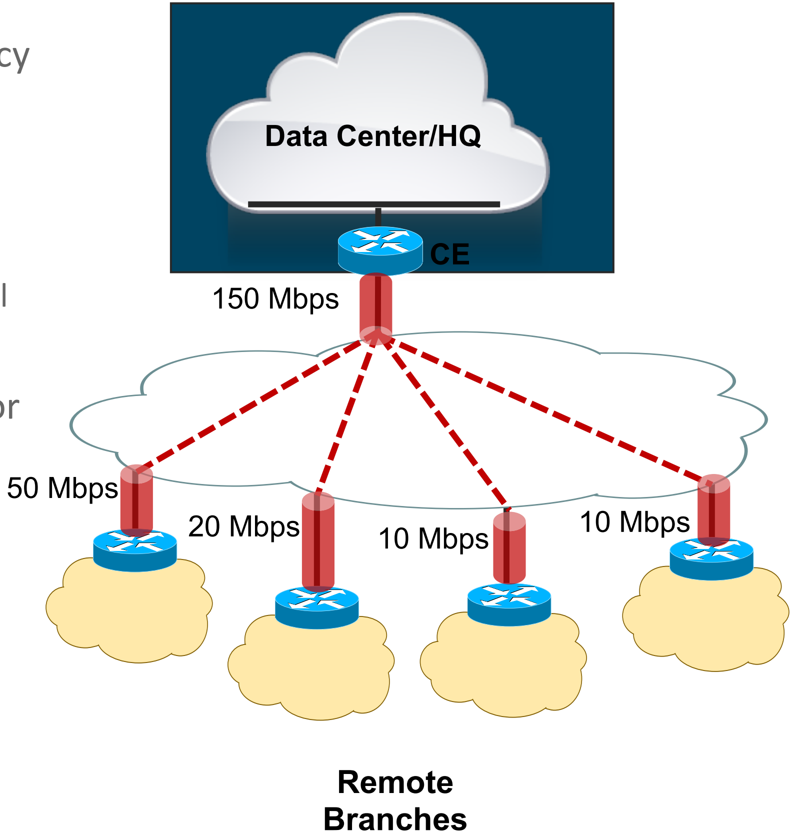

Intelligent Path Control

Performance Routing

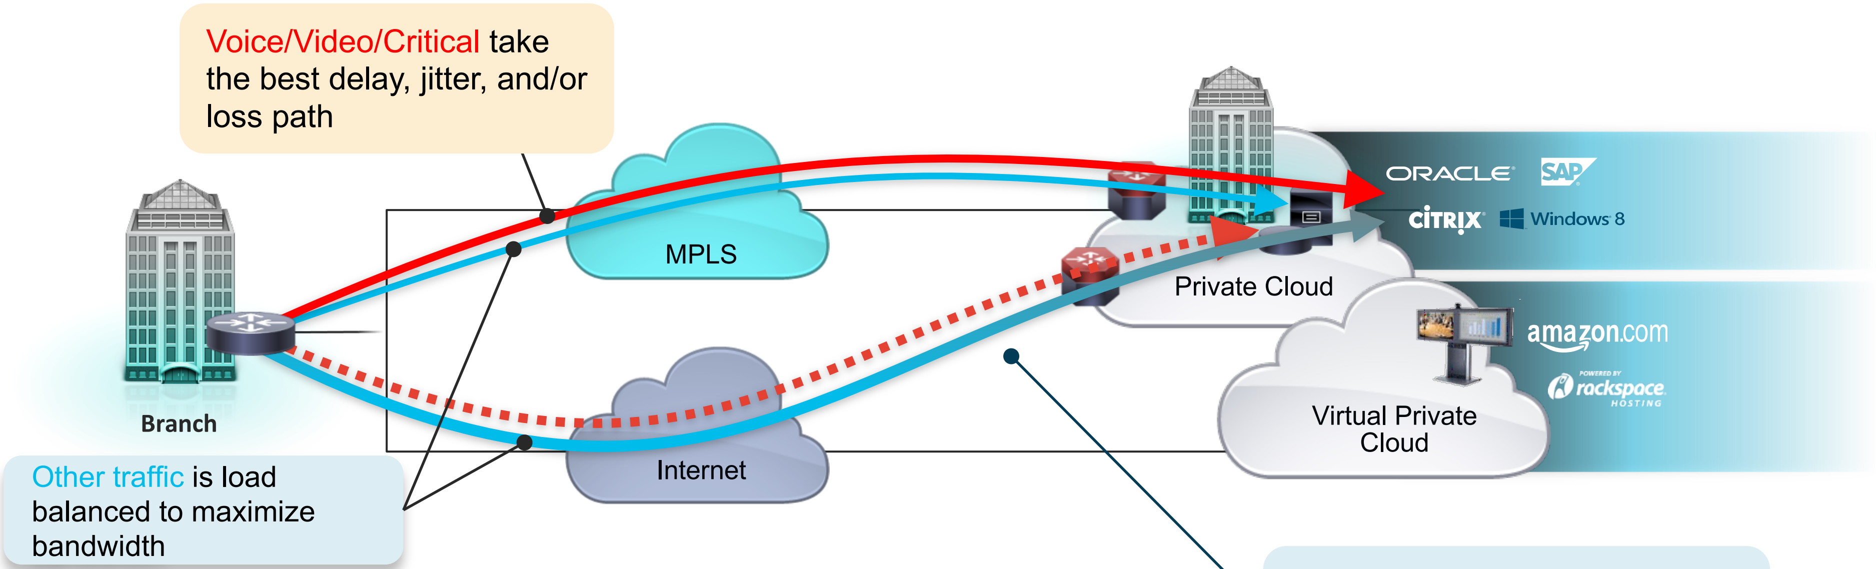

- PfR monitors network performance and routes applications based on application performance policies
- PfR load balances traffic based upon link utilization levels to efficiently utilize all available WAN bandwidth

Подведем итоги.

Part 1: WAN Architectures and Design Principles

Key Takeaways

- The goal is for a simple, modular, hierarchical, structured design
- Business, technical, and physical requirements and constraints must all be considered
- Desired WAN availability and services have design implications
- Evolving technology is driving new WAN designs
- Leveraging Internet, Cloud, and CoLo now fundamental

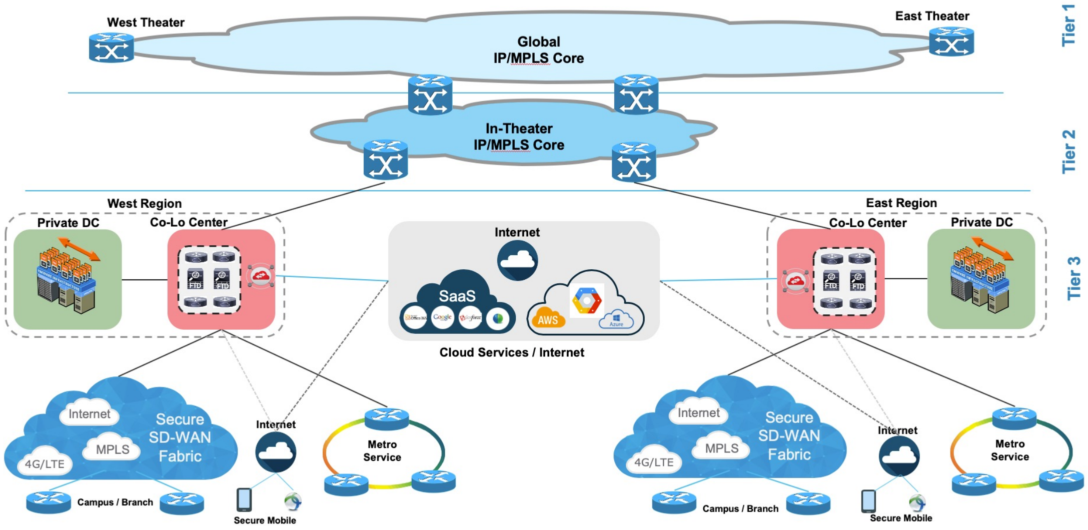

Part 2: Highly Available WAN Design
Key Takeaways

- Network design should target how the applications survive a variation of outages.
- Leverage load sharing capabilities for more resiliency and application performance
- End-to-end convergence time is the goal, and can be affected by localized topology changes
- Consider IP SLA based monitoring and SD-WAN for real-time path selection
- Effective network designs incorporate a combination of convergence techniques

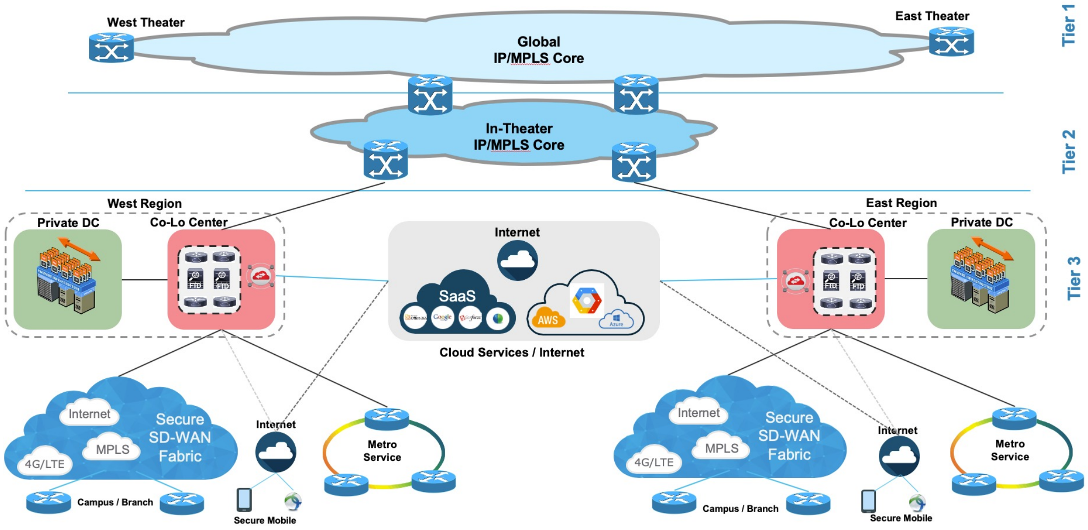

Part 3: WAN Services
Key Takeaways

- Understand the application usage before adding services like QoS or Multicast to the WAN
- QoS should be always included in the initial WAN design deployment
- Leverage federated security cloud proxy and localized stack at the branch in a phased approach for consumption
- Don't look at point solution for automation, rather look and the architecture and then fit the solution.

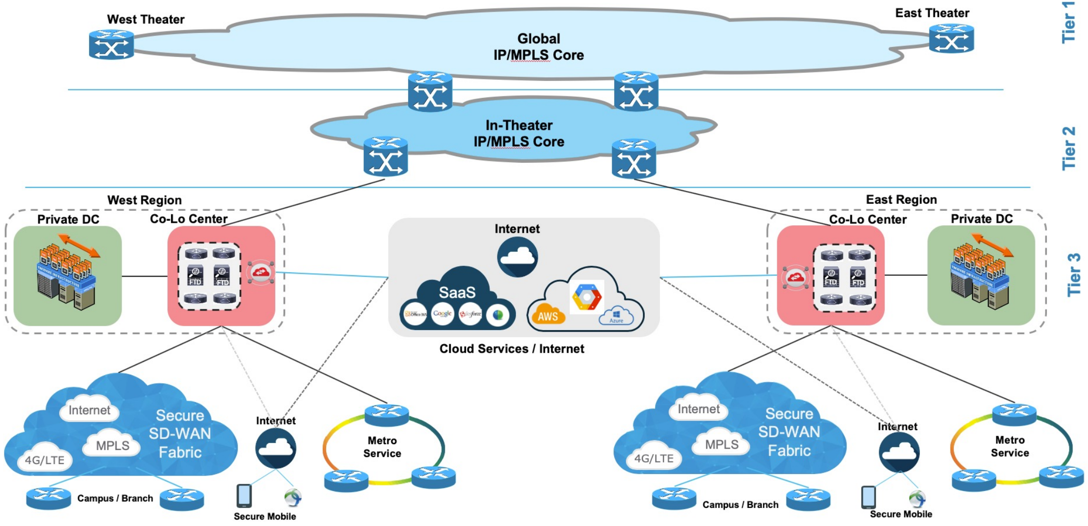

Part 4: L3 Segmentation and Cloud Ready Solutions for the WAN

Key Takeaways

- Make L3 Segmentation a fundamental element in any new WAN designs
- Understand the business and technical criteria for proper next-gen WAN solutions
- Incorporate the Cloud Ready Design fundamentals into all new and existing designs moving forward
- Leverage high-speed encryption (WAN MACsec) where applicable
- Begin to incorporate automation tools into network operations to simplify and error-proof configuration changes
- Keep it simple whenever possible!!!

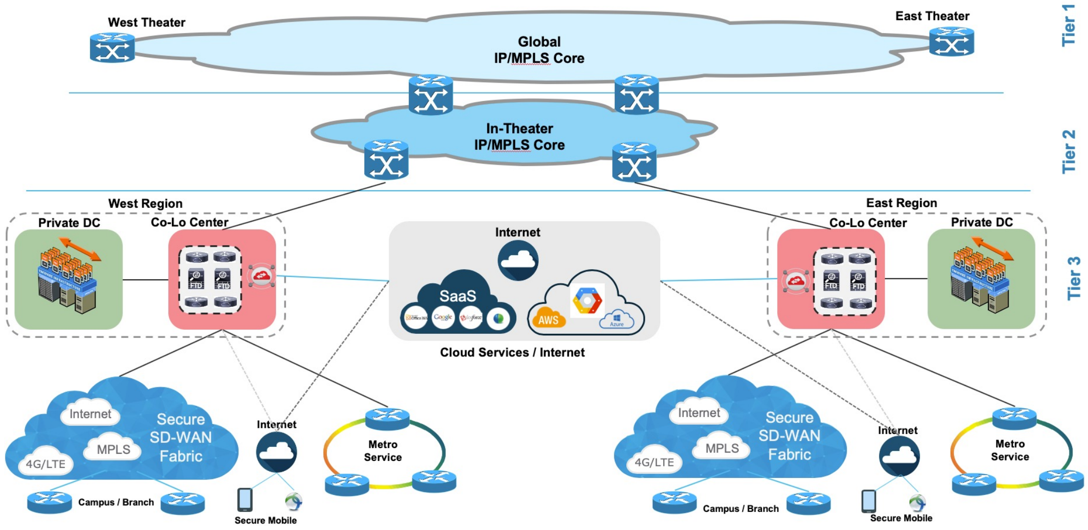

Спасибо! Вопросы?