Cisco live!

What You Make Possible

TOMORROW
starts here.

CISCO

WAN Architectures and Design Principles

BRKRST-2041

TOMORROW starts here.

Housekeeping

- We value your feedback- don't forget to complete your online session evaluations after each session & complete the Overall Conference Evaluation which will be available online
- Visit the World of Solutions
- Please switch off your mobile phones
- Please remember to wear your badge at all times

Agenda

- WAN Technologies & Solutions
  - WAN Transport Technologies
  - WAN Overlay Technologies
  - WAN Optimisation
  - Wide Area Network Quality of Service
- WAN Architecture Design Considerations
  - WAN Design and Best Practices
  - Secure WAN Communication with GETVPN
  - DMVPN Over Internet Deployment
- Summary

Hierarchical Network Design

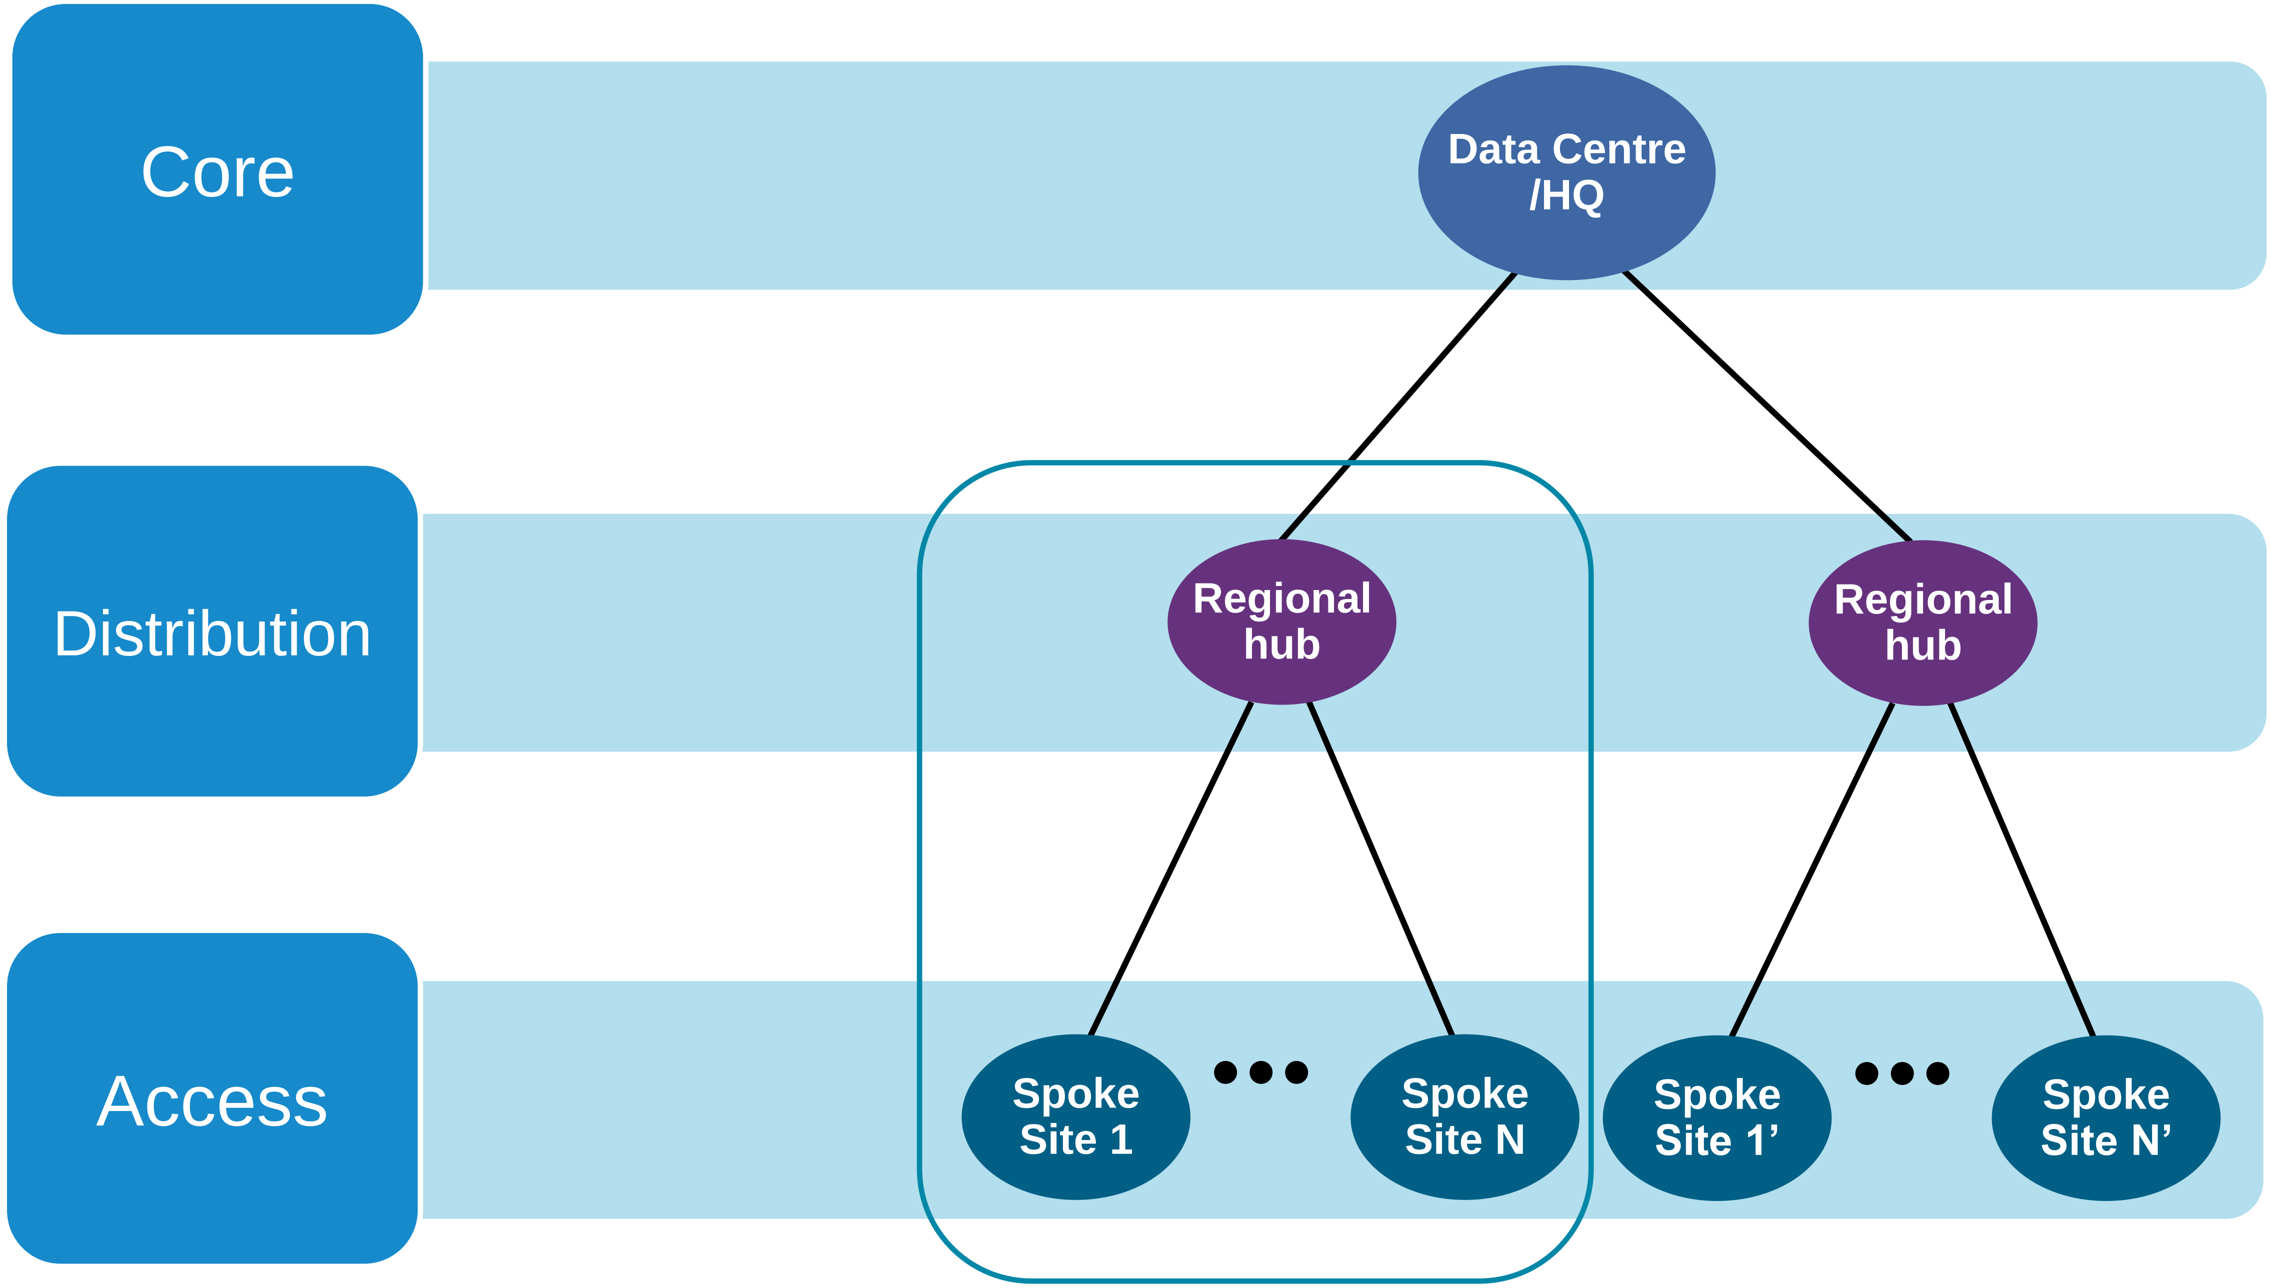

Hierarchical Network Design

- Hierarchical design used to be…
  - – Three routed layers
  - – Core, distribution, access
  - – Only one hierarchical structure end-to-end
- Hierarchical design has become any design that…
  - – Splits the network up into "places," or "regions"
  - – Separates these "regions" by hiding information
  - – Organises these "regions" around a network core
  - – "hub and spoke" at a macro level

MPLS VPN Topology

Definition

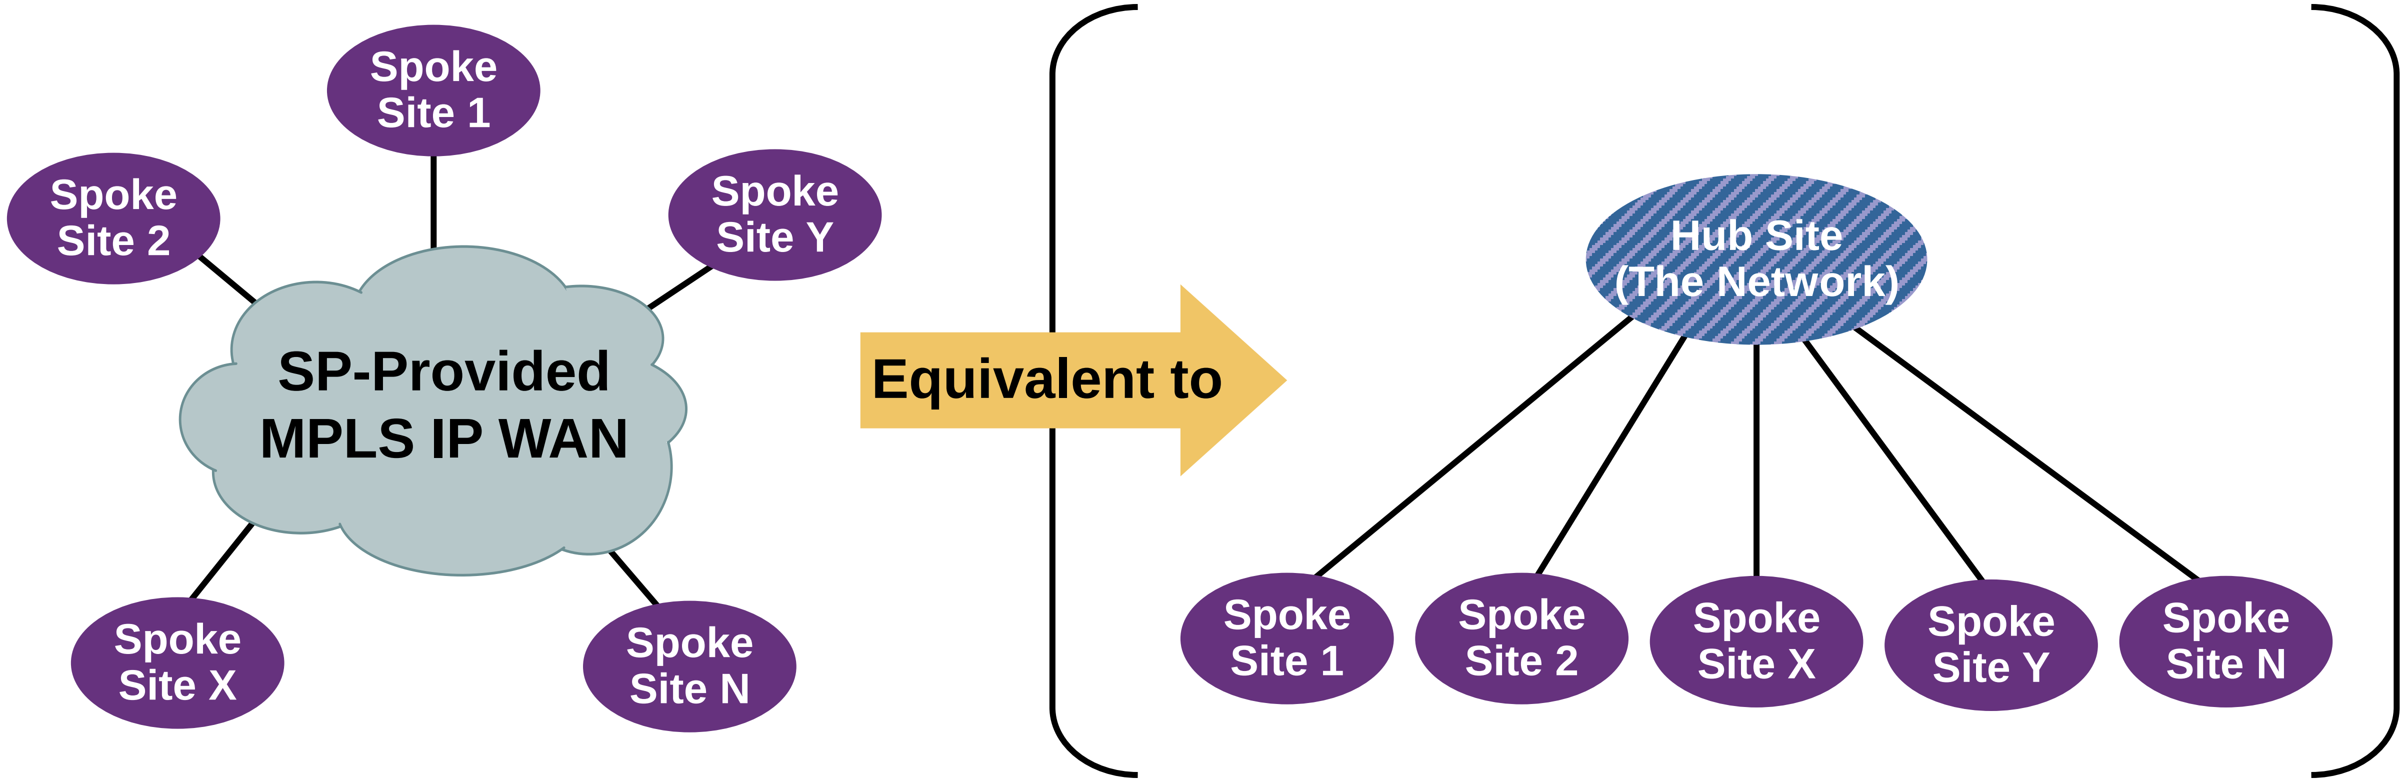

- MPLS WAN is provided by a service provider
- As seen by the enterprise network, every site is one IP “hop” away
- Equivalent to a full mesh, or to a “hubless” hub-and-spoke

MPLS VPN

Layer 3 (L3) Service

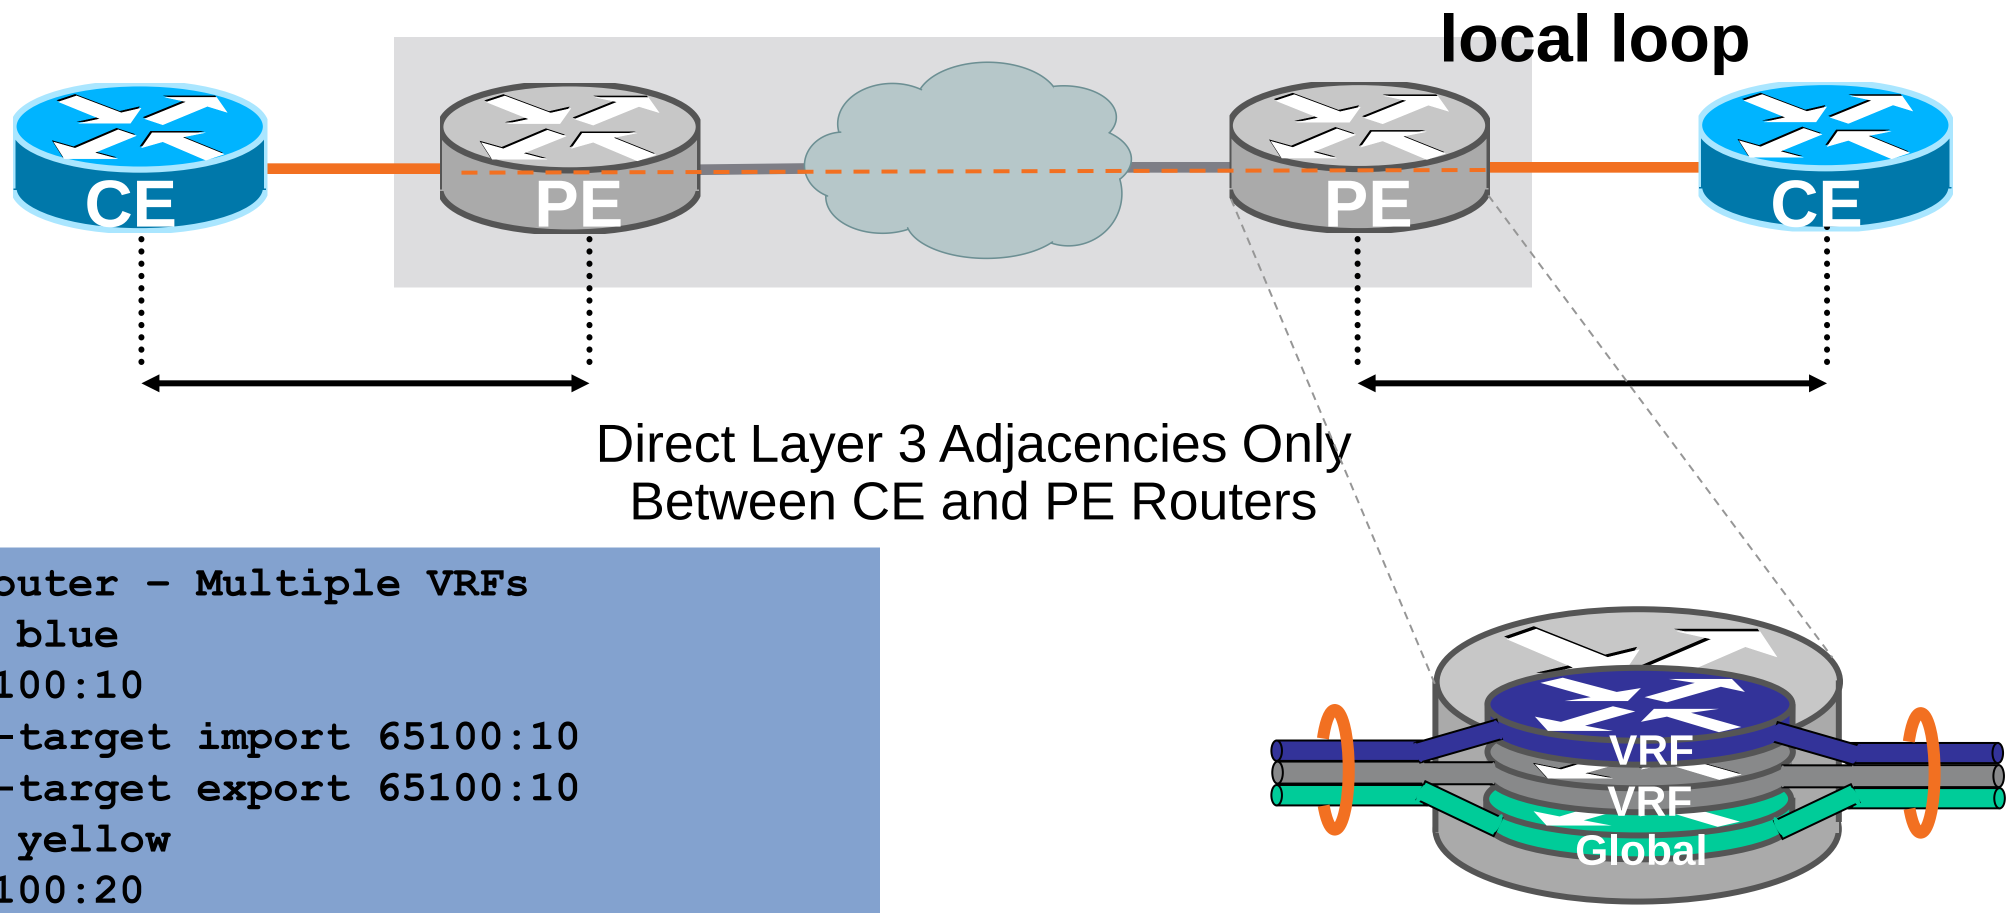

Direct Layer 3 Adjacencies Only
Between CE and PE Routers

```
! PE Router – Multiple VRFs
ip vrf blue
 rd 65100:10
 route-target import 65100:10
 route-target export 65100:10
ip vrf yellow
 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 yellow
```

VRF—Virtual Routing and Forwarding

MPLS VPN Design Trends

- **Single Carrier Designs**:
  - 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 Carrier Designs:**
  - 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.
- **Variants of these designs and site connectivity:**
  - Encryption Overlay (e.g. IPSec, DMVPN, GET VPN, etc.)
  - Sites with On-demand / Permanent backup links

Single Carrier Site Types (Non-Transit)

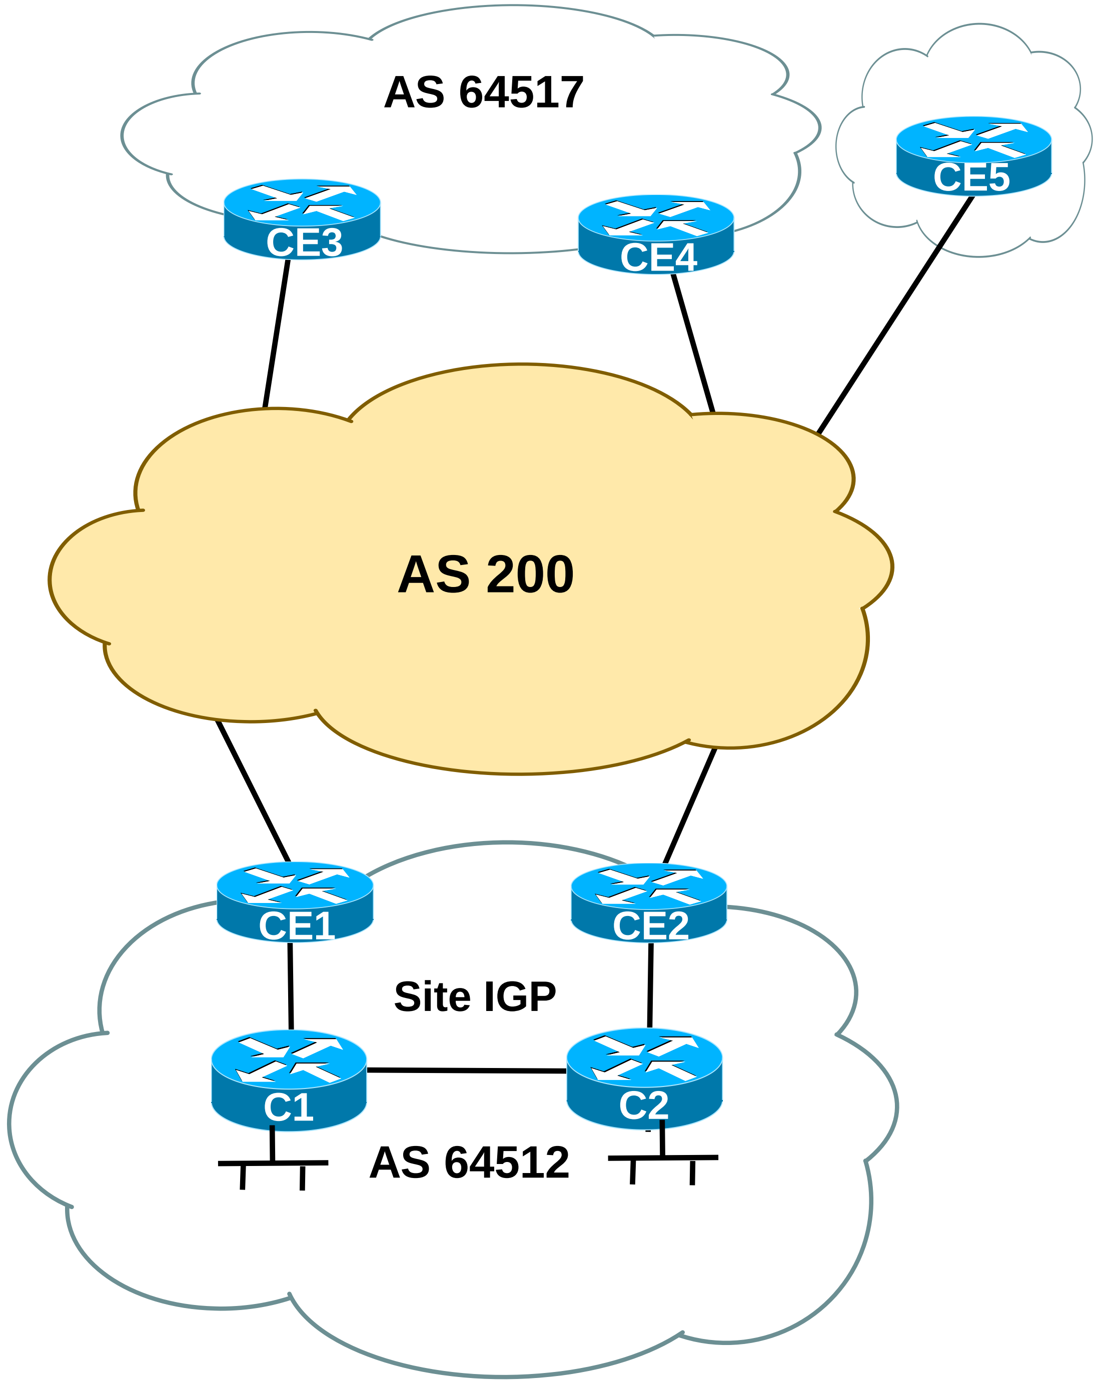

- **Dual Homed Non Transit**
  - Only advertise local prefixes (^$)
  - Typically with Dual CE routers
  - BGP design:
    - eBGP to carrier
    - iBGP between CEs
  - Redistribute cloud learned routes into site IGP
- **Single Homed Non Transit**
  - Advertise local prefixes and optionally use default route.

Dual Carrier: Transit vs. Non Transit

- To guarantee single homed site reachability to a dual homed site experiencing a failure, transit sites had to be elected.
- Transit sites would act as a BGP bridge transiting routes between the two provider clouds.
- To minimise latency costs of transits, transits need to be selected with geographic diversity (e.g. from the East, West and Central US.)

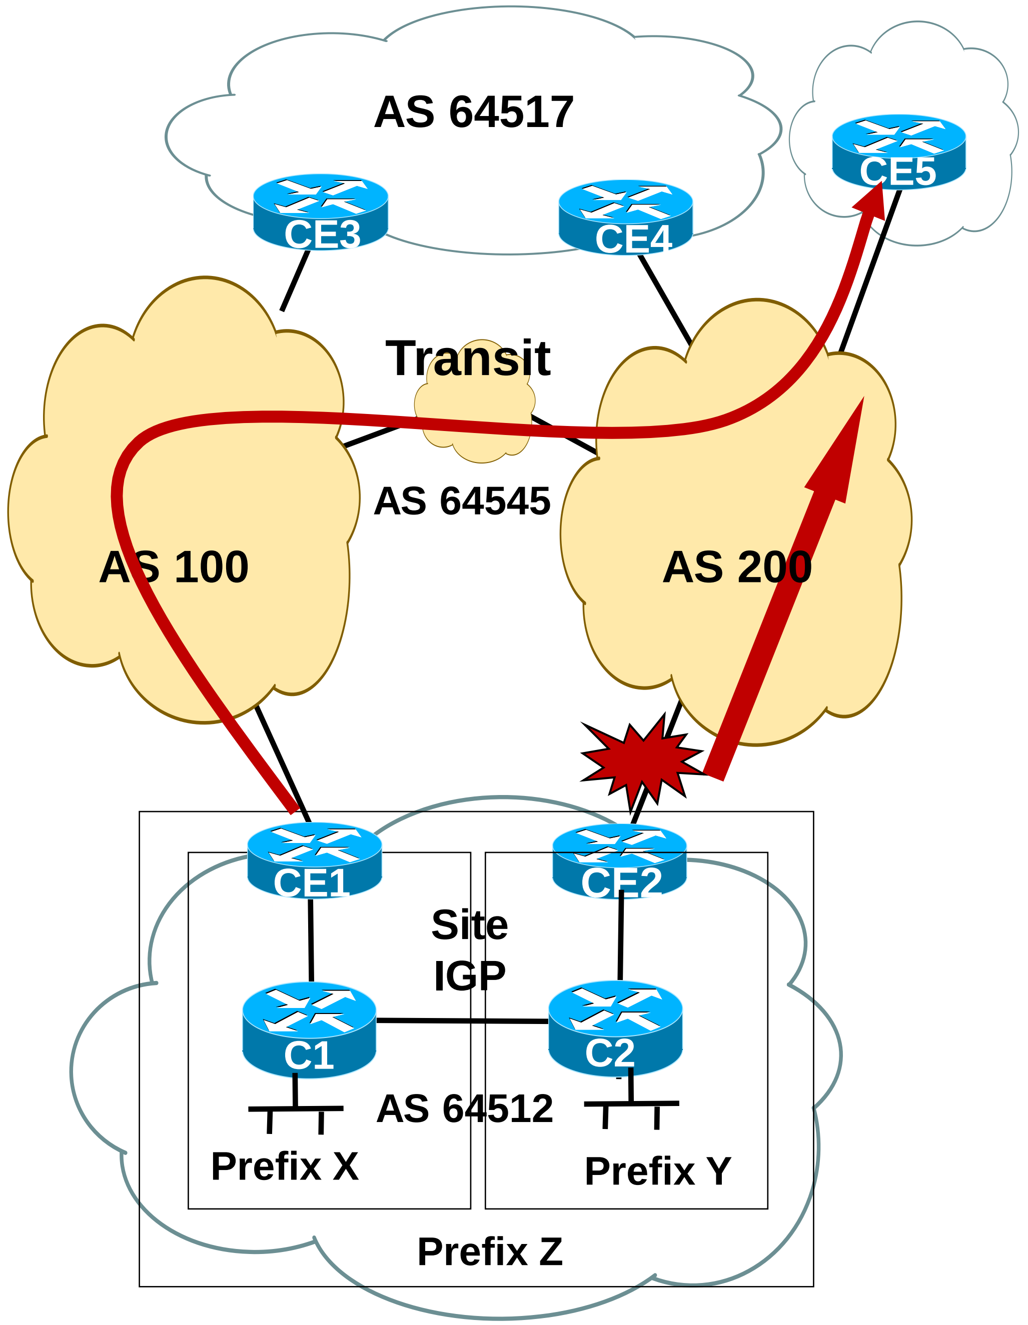

Single vs. Dual Carriers

Single Provider	Dual Providers
Pro: Common QoS support model	Pro: More fault domains
Pro: Only one vendor to "tune"	Pro: More product offerings to business
Pro: Reduced head end circuits	Pro: Ability to leverage vendors for better pricing
Pro: Overall simpler design	Pro: Nice to have a second vendor option
Con: Carrier failure could be catastrophic	Con: Increased Bandwidth "Paying for bandwidth twice"
Con: Do not have another carrier "in your pocket"	Con: Increased overall design complexity
	Con: May be reduced to "common denominator" between carriers

Resiliency Drivers vs. Simplicity

Agenda

- WAN Technologies & Solutions
  - WAN Transport Technologies
  - WAN Overlay Technologies
  - WAN Optimisation
  - Wide Area Network Quality of Service
- WAN Architecture Design Considerations
  - WAN Design and Best Practices
  - Secure WAN Communication with GETVPN
  - DMVPN Over Internet Deployment
- Summary

Tunnelling Technologies

Packet Encapsulation over IP

- IPSec—Encapsulating Security Payload (ESP)
  – Strong encryption
  – IP Unicast only
- Generic Routing Encapsulation (GRE)
  – IP Unicast, Multicast, Broadcast
  – Multiprotocol support
- Layer 2 Tunnelling Protocol—Version 3 (L2TPv3)
  – Layer 2 payloads (Ethernet, Serial,…)
  – Pseudowire capable
- Other Tunnelling Technologies – L3VPNomGRE, LISP, OTV

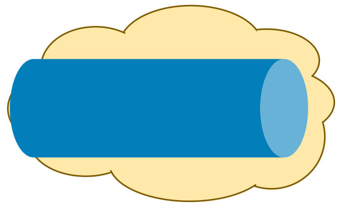

Tunnels

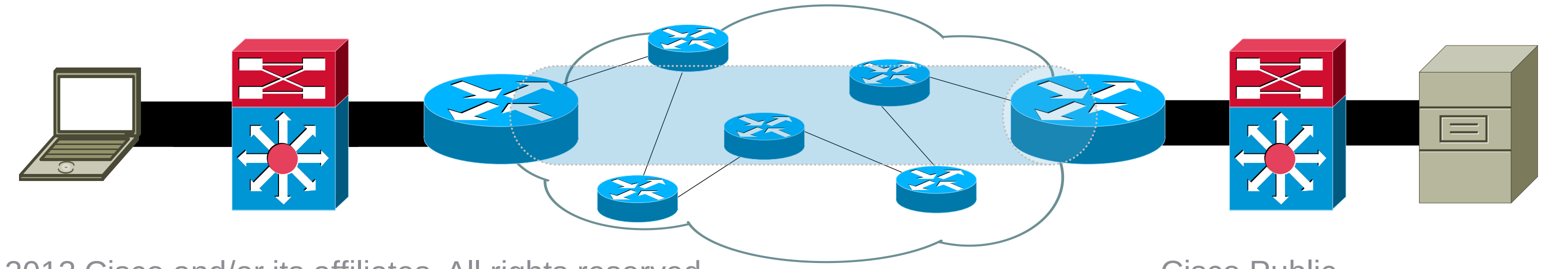

Tunnelling

GRE and IPSec Transport and Tunnel Modes

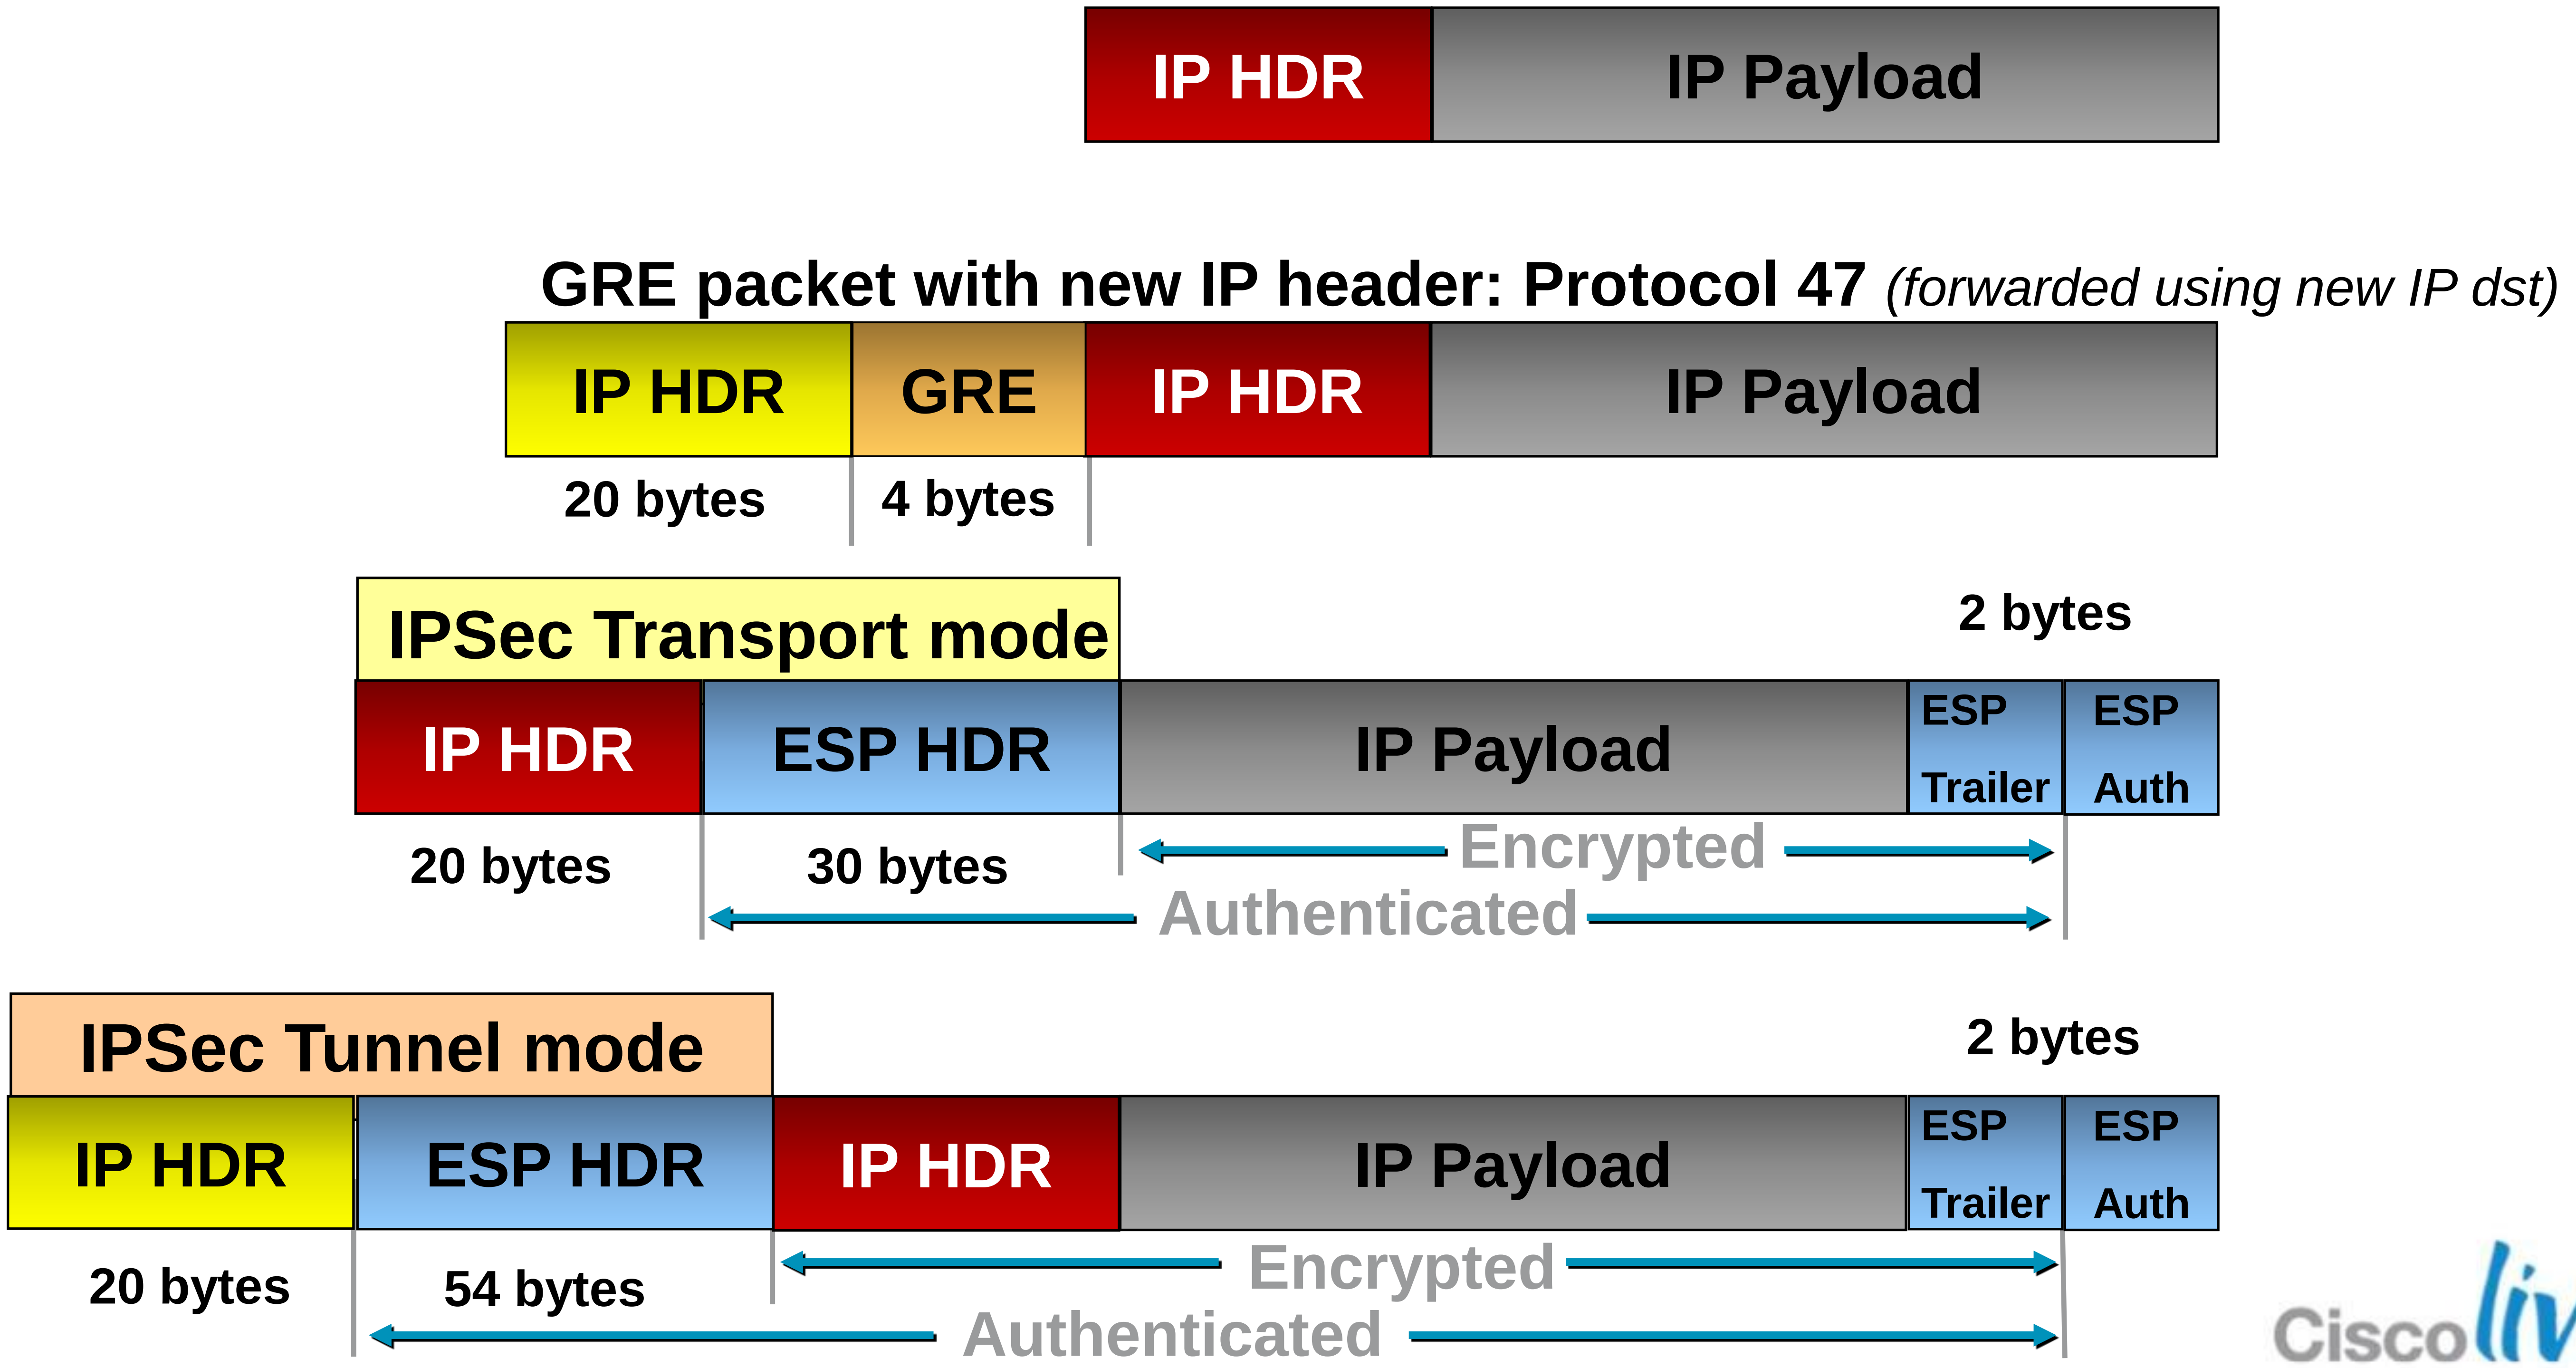

Locator/Identifier Separation Protocol (LISP)

Dynamic Tunnelling Analogous to a DNS but for Network Infrastructure

- DNS resolves IP addresses for URLs

DNS
URL Resolution

- LISP resolves locators for queried identities

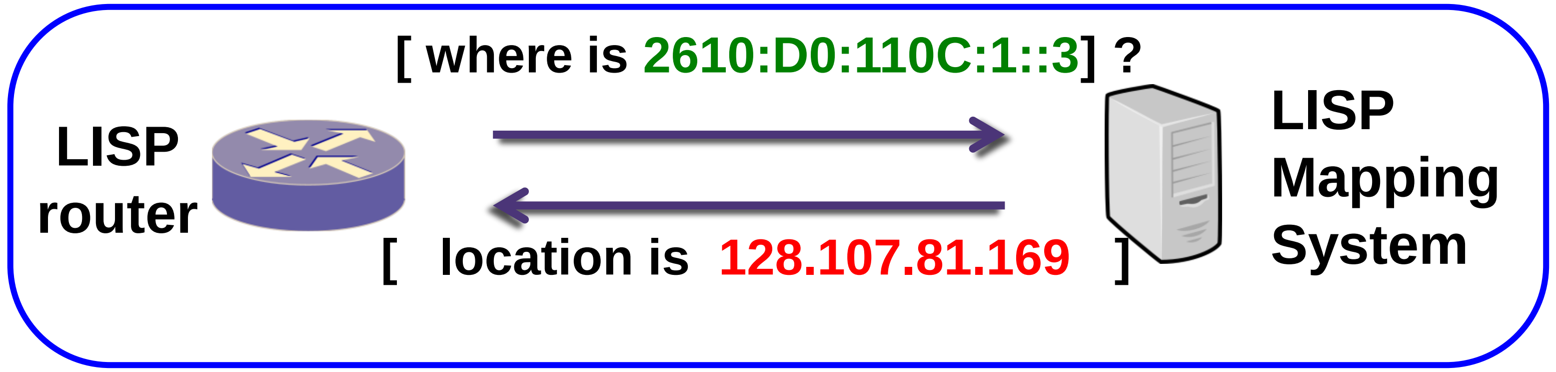

LISP
Identity-to-location
Map Resolution

This Topic Is Covered in Detail in BRKRST-3045

LISP Overview - Terminologies

- **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**

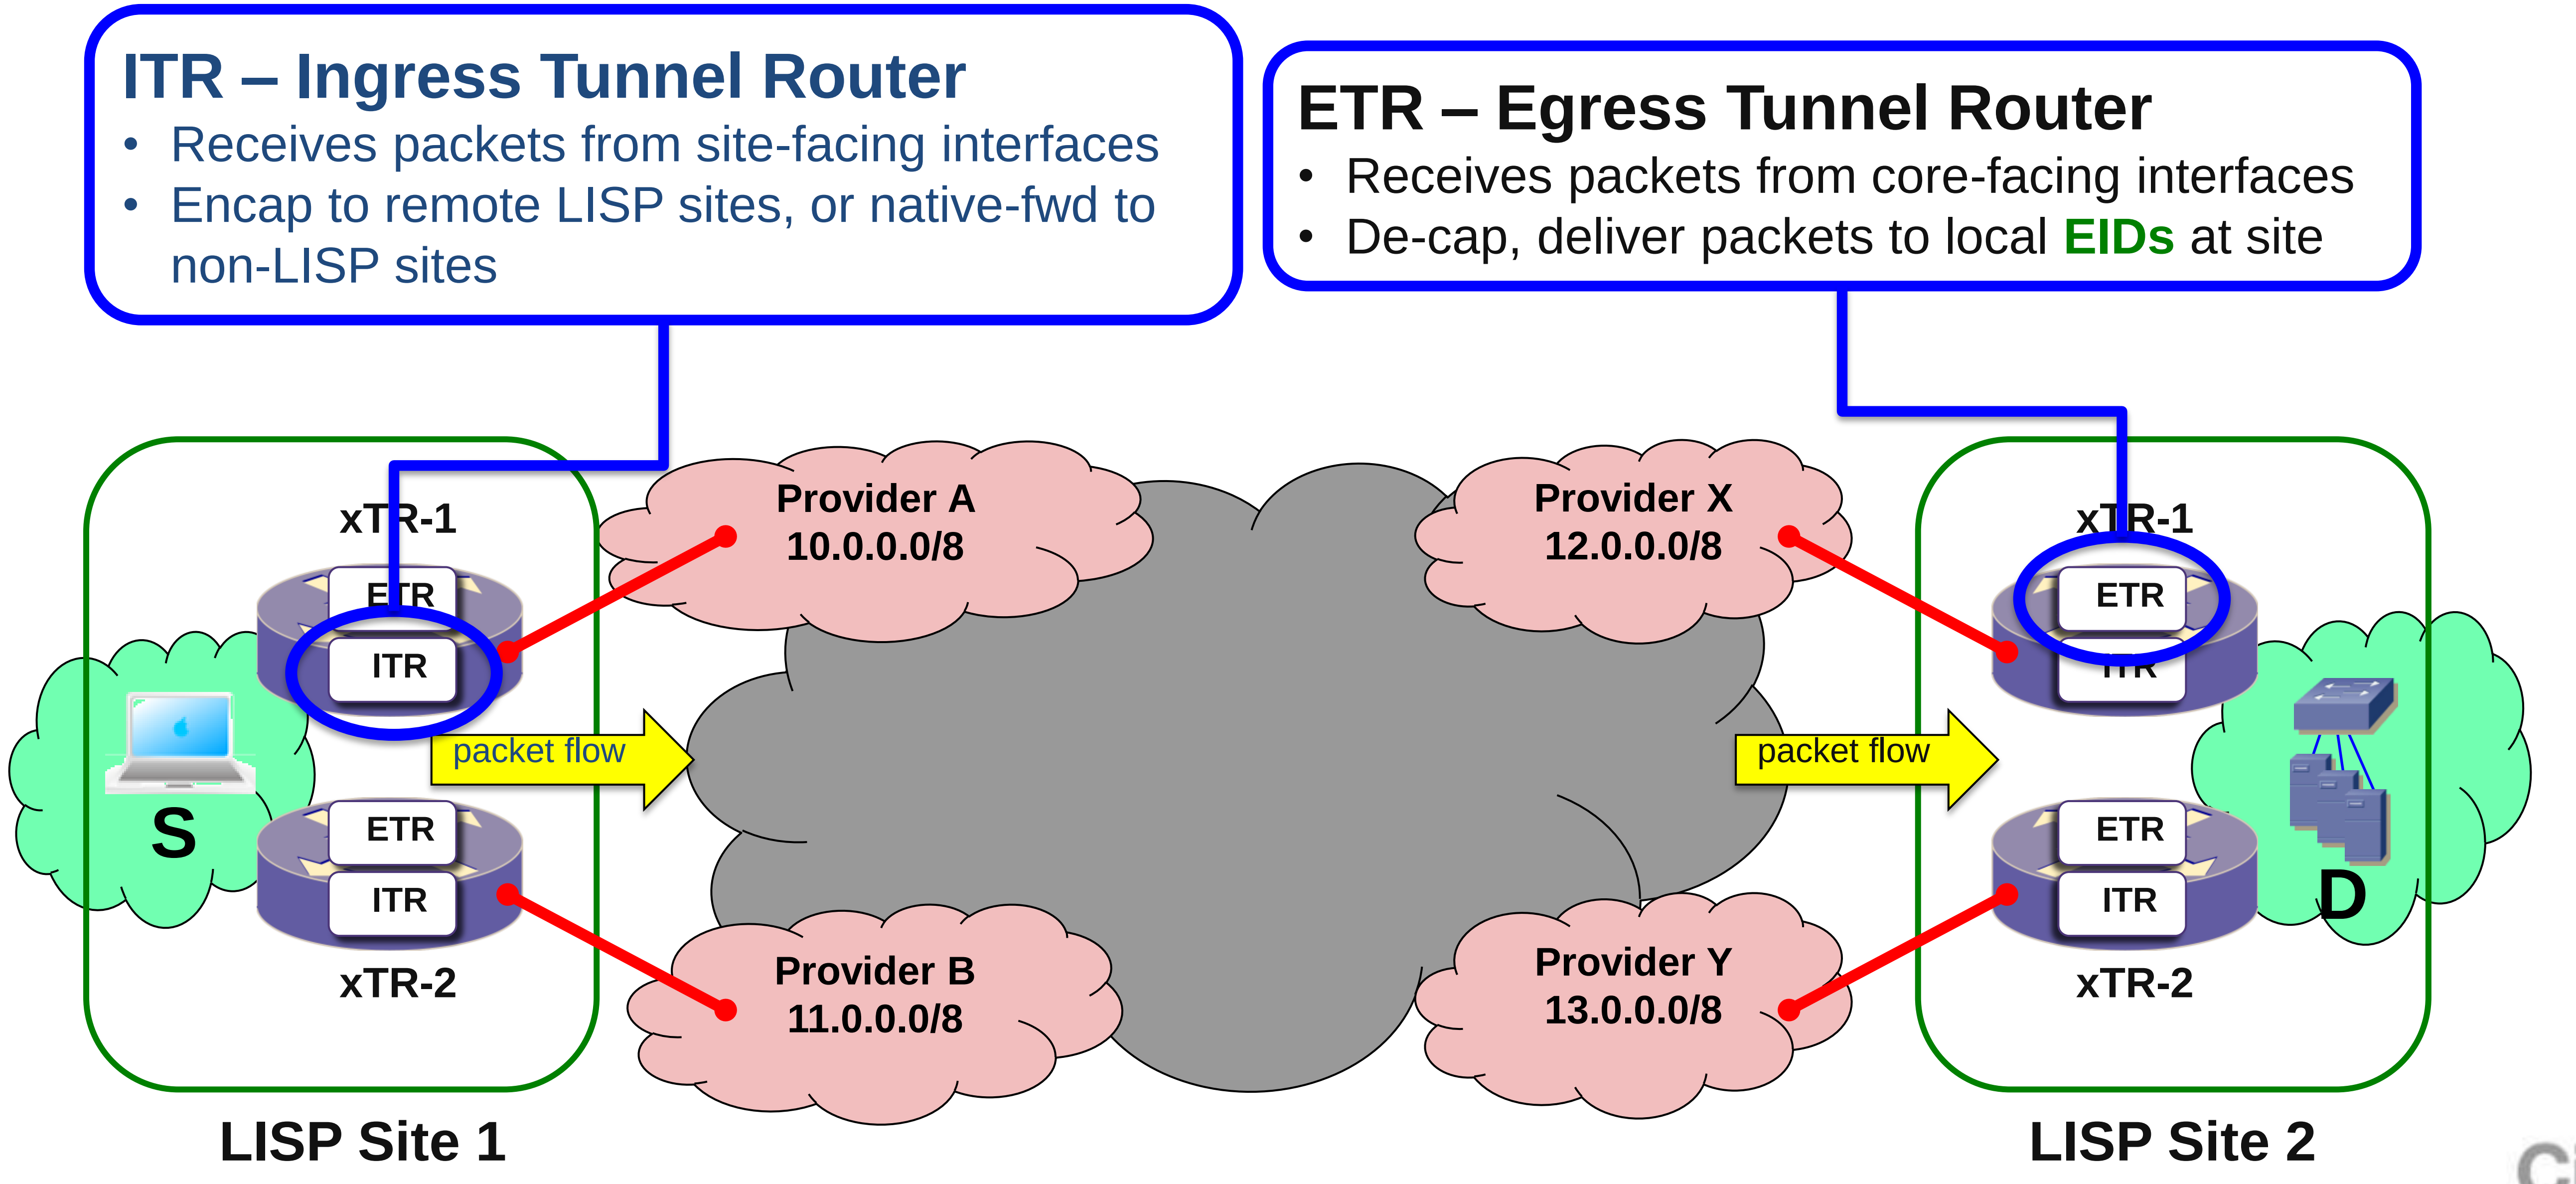

LISP Operation Example

LISP Data Plane - Unicast Packet Forwarding

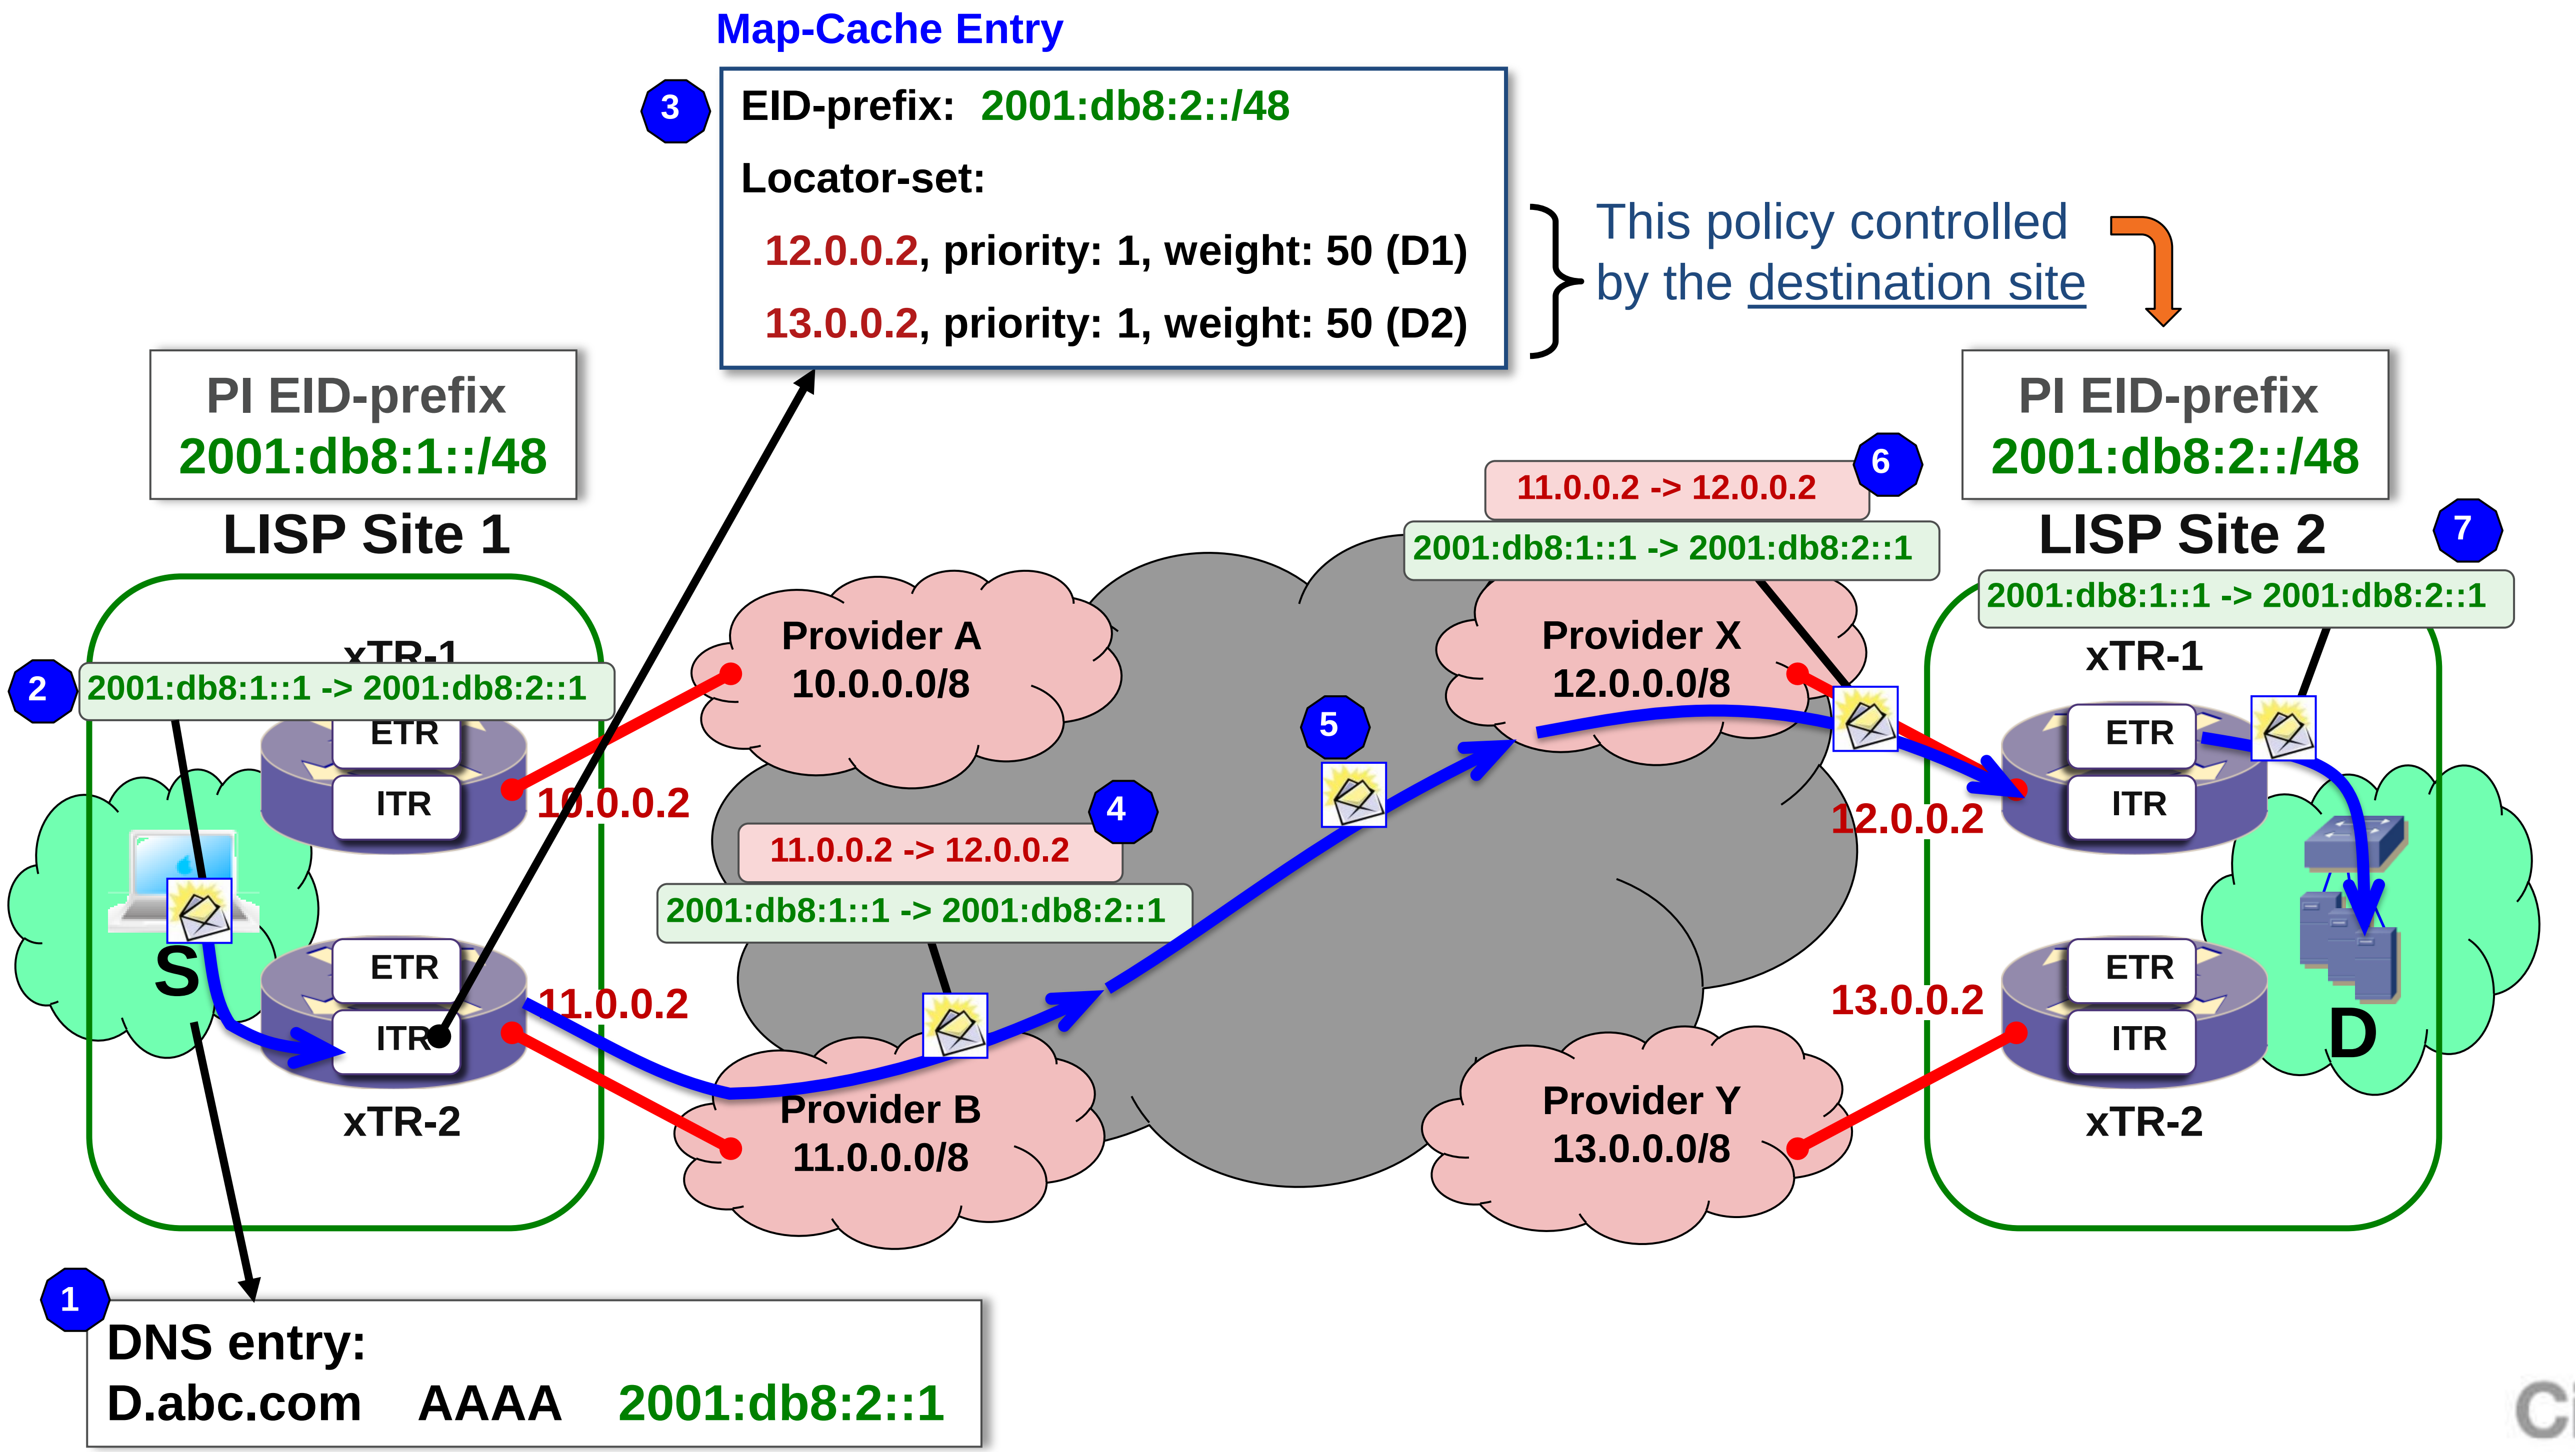

LISP Use Cases

IPv6 Transition

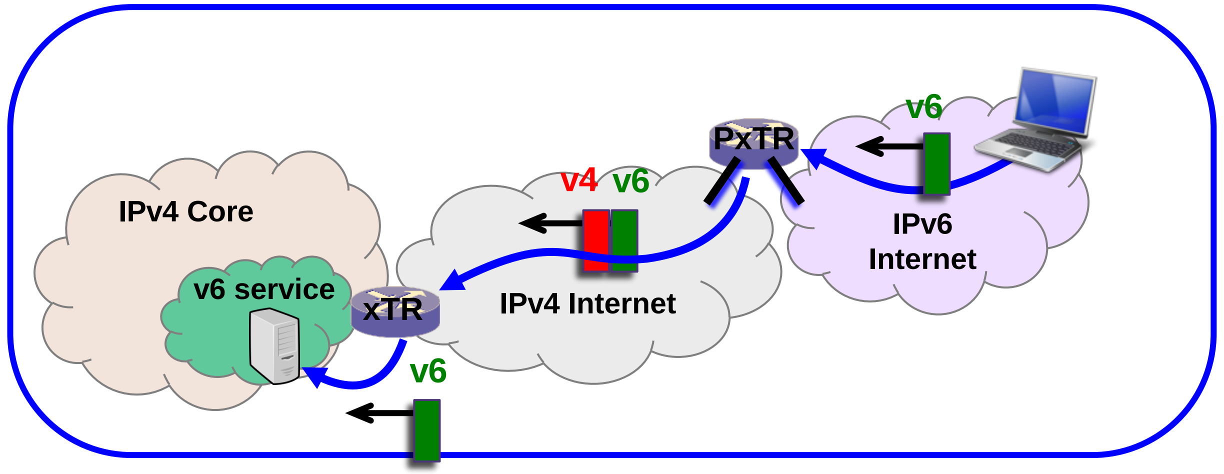

- IPv6-over-IPv4, IPv6-over-IPv6
- IPv4-over-IPv6, IPv4-over-IPv4

Efficient Multi-Homing

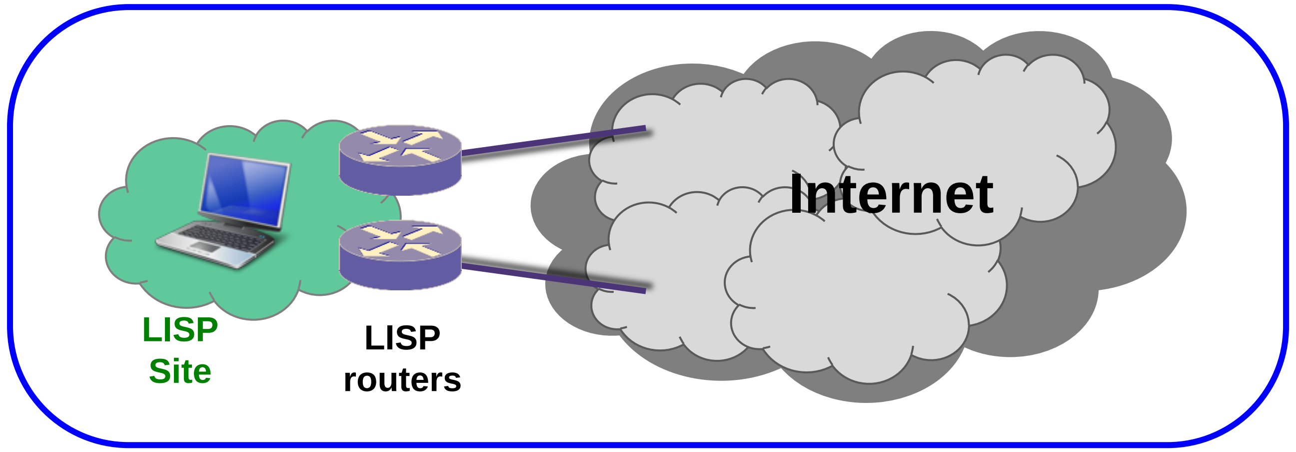

- IP Portability
- Ingress Traffic Engineering Without BGP

Virtualisation/Multi-tenancy

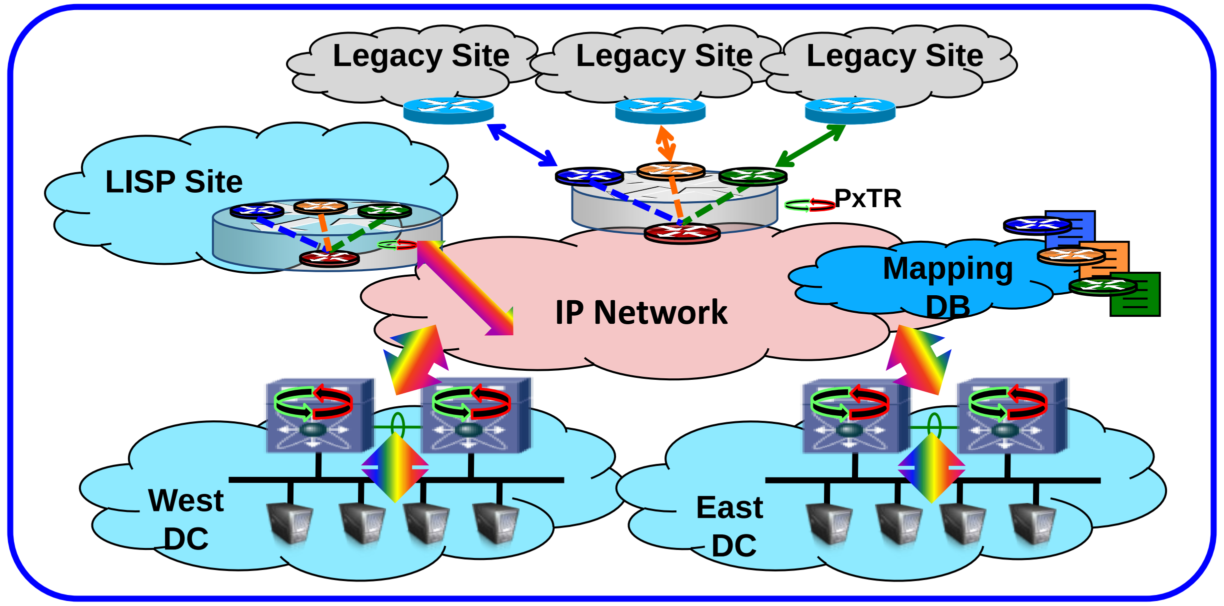

- Large Scale Segmentation

Data Centre/ VM Mobility

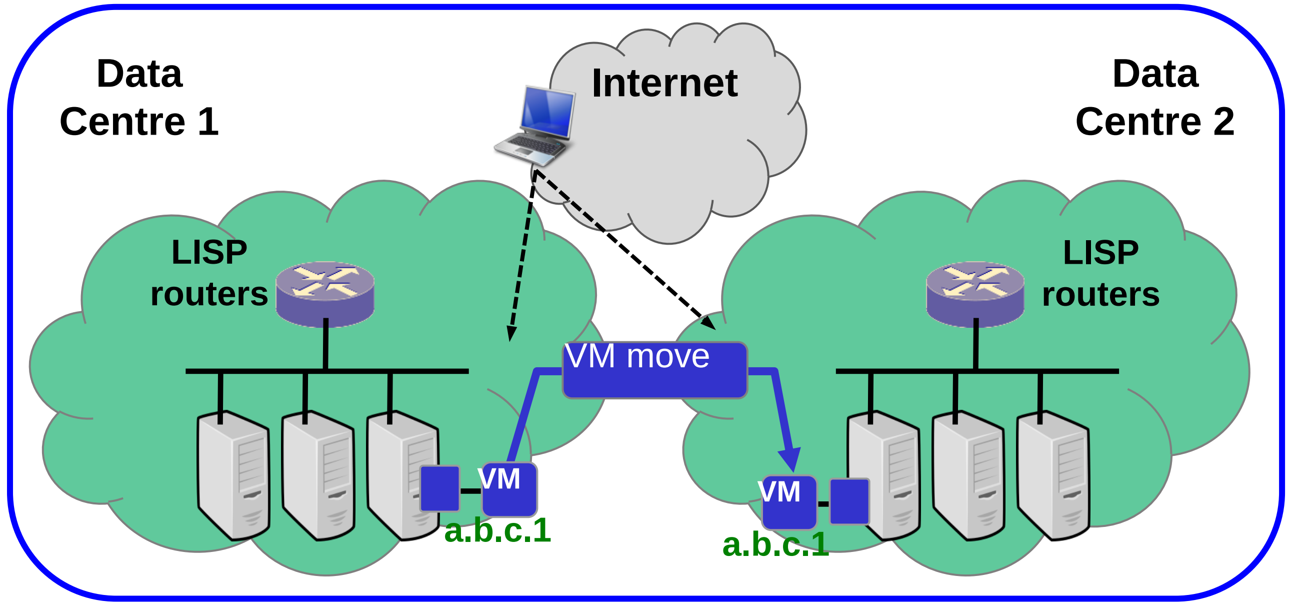

- Cloud / Layer 3 VM Move

VPN Technology

Positioning EzVPN, DMVPN, GETVPN

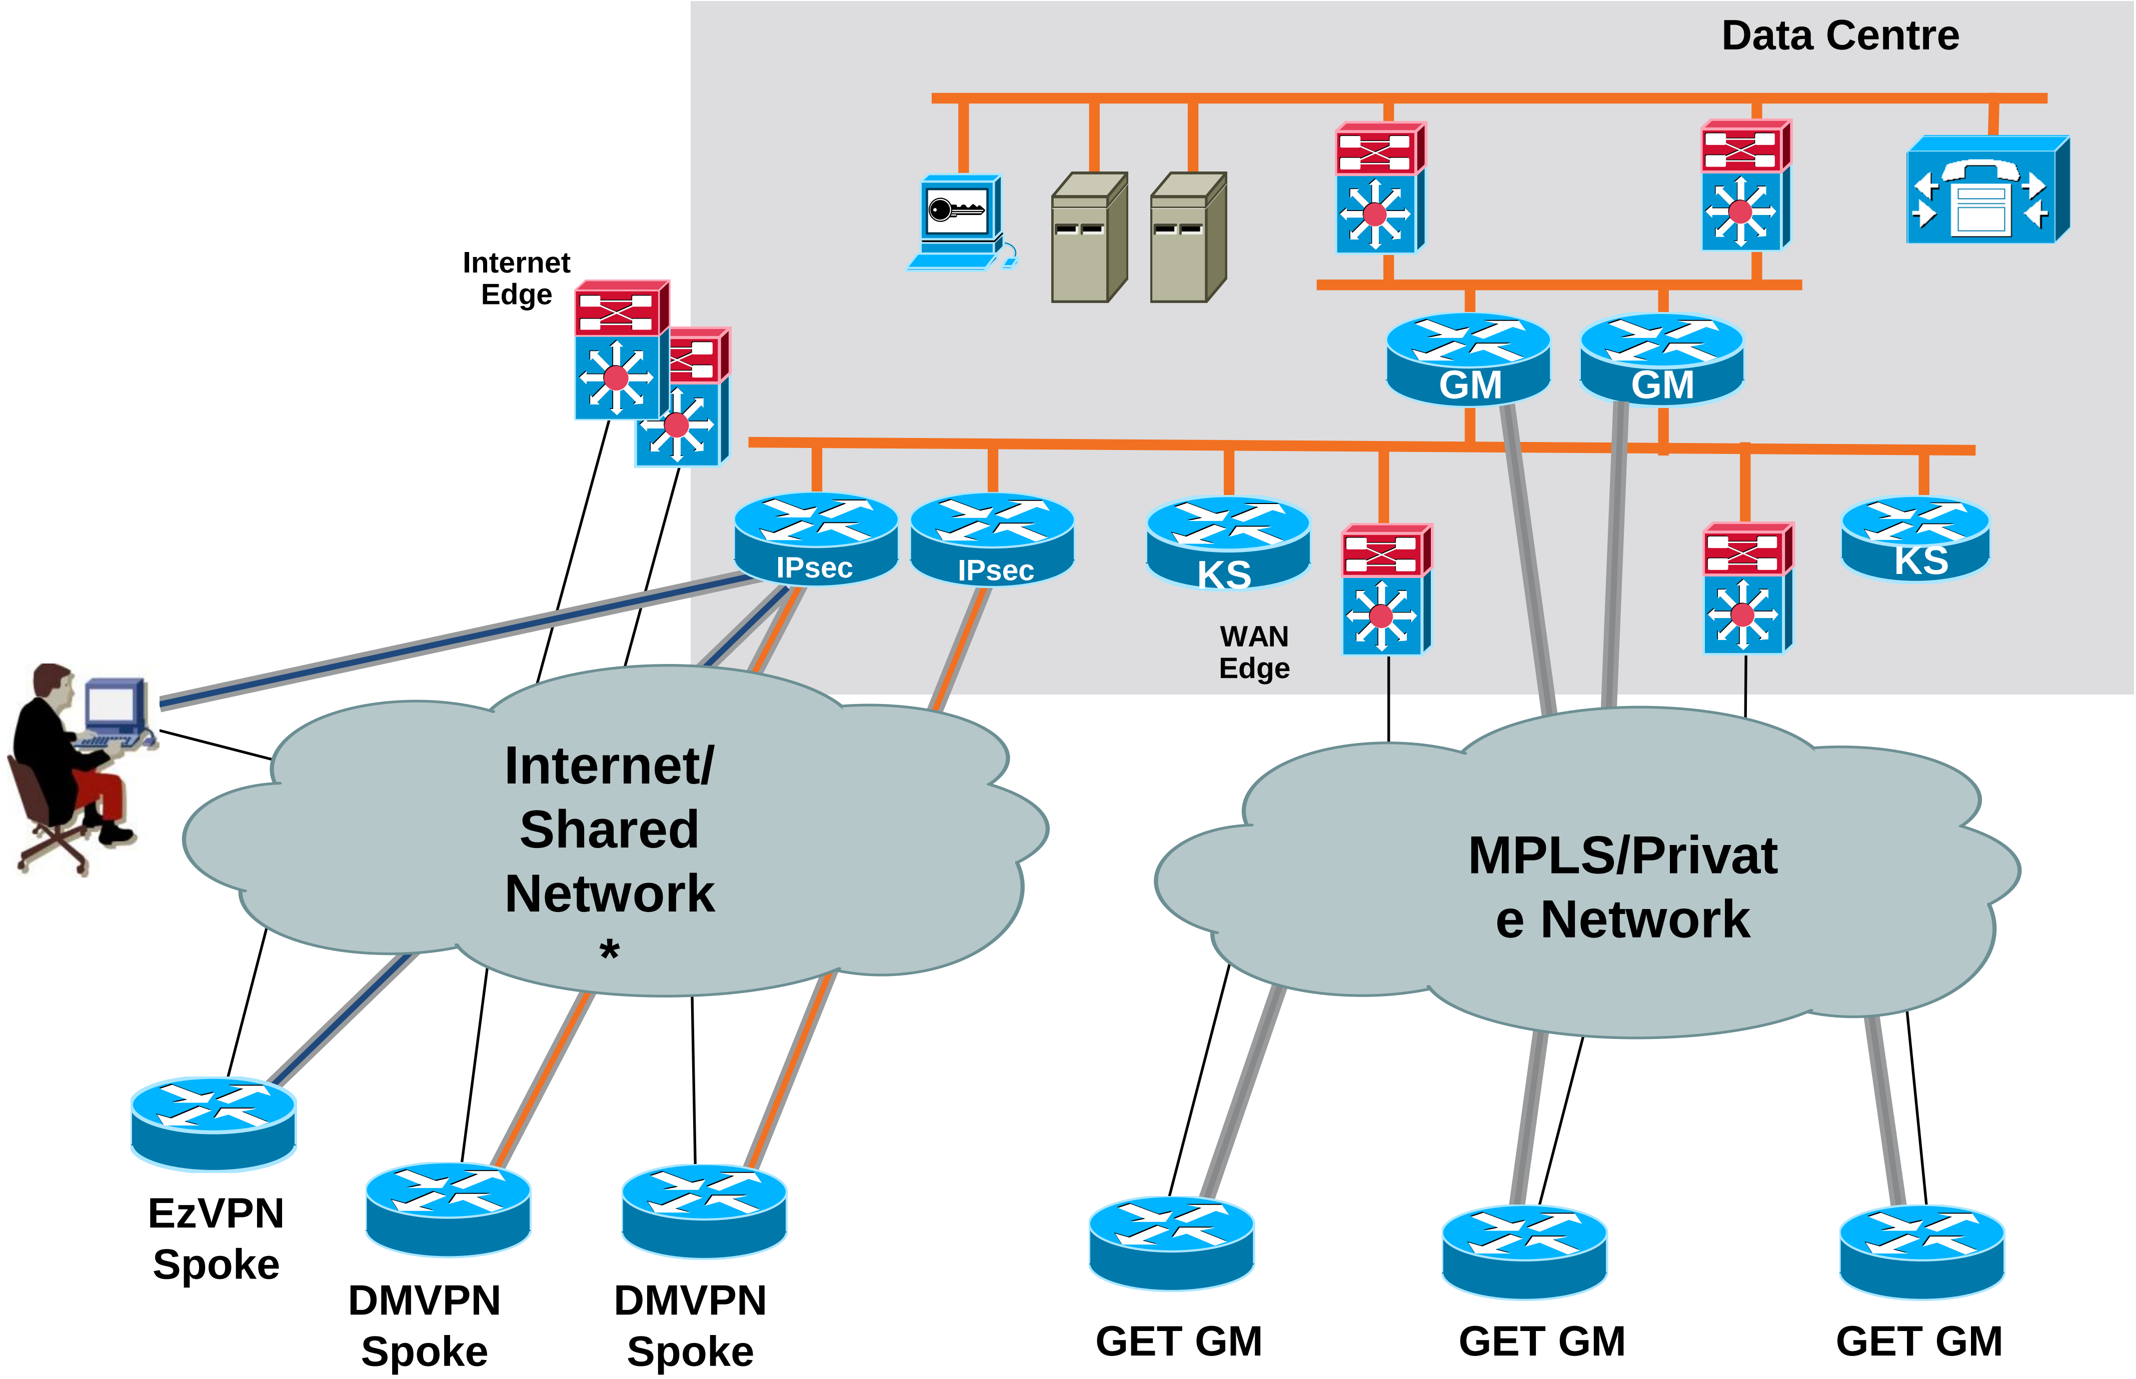

* Note: DMVPN Can Also Be Used on MPLS/Private Network

VPN Technology Comparison

	EzVPN	DMVPN	GETVPN
Infrastructure Network	▪ Public Internet Transport	▪ Private & Public Internet Transport	▪ Private IP Transport
Network Style	▪Hub-Spoke; (Client to Site)	▪ Hub-Spoke and Spoke-to-Spoke; (Site-to-Site)	▪ Any-to-Any; (Site-to-Site)
Routing	▪ Reverse-route Injection	▪ Dynamic routing on tunnels	▪ Dynamic routing on IP WAN
Failover Redundancy	▪ Stateful Hub Crypto Failover	▪ Route Distribution Model	▪ Route Distribution Model + Stateful
Encryption Style	▪ Peer-to-Peer Protection	▪ Peer-to-Peer Protection	▪ Group Protection
IP Multicast	▪ Multicast replication at hub	▪ Multicast replication at hub	▪ Multicast replication in IP WAN network

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

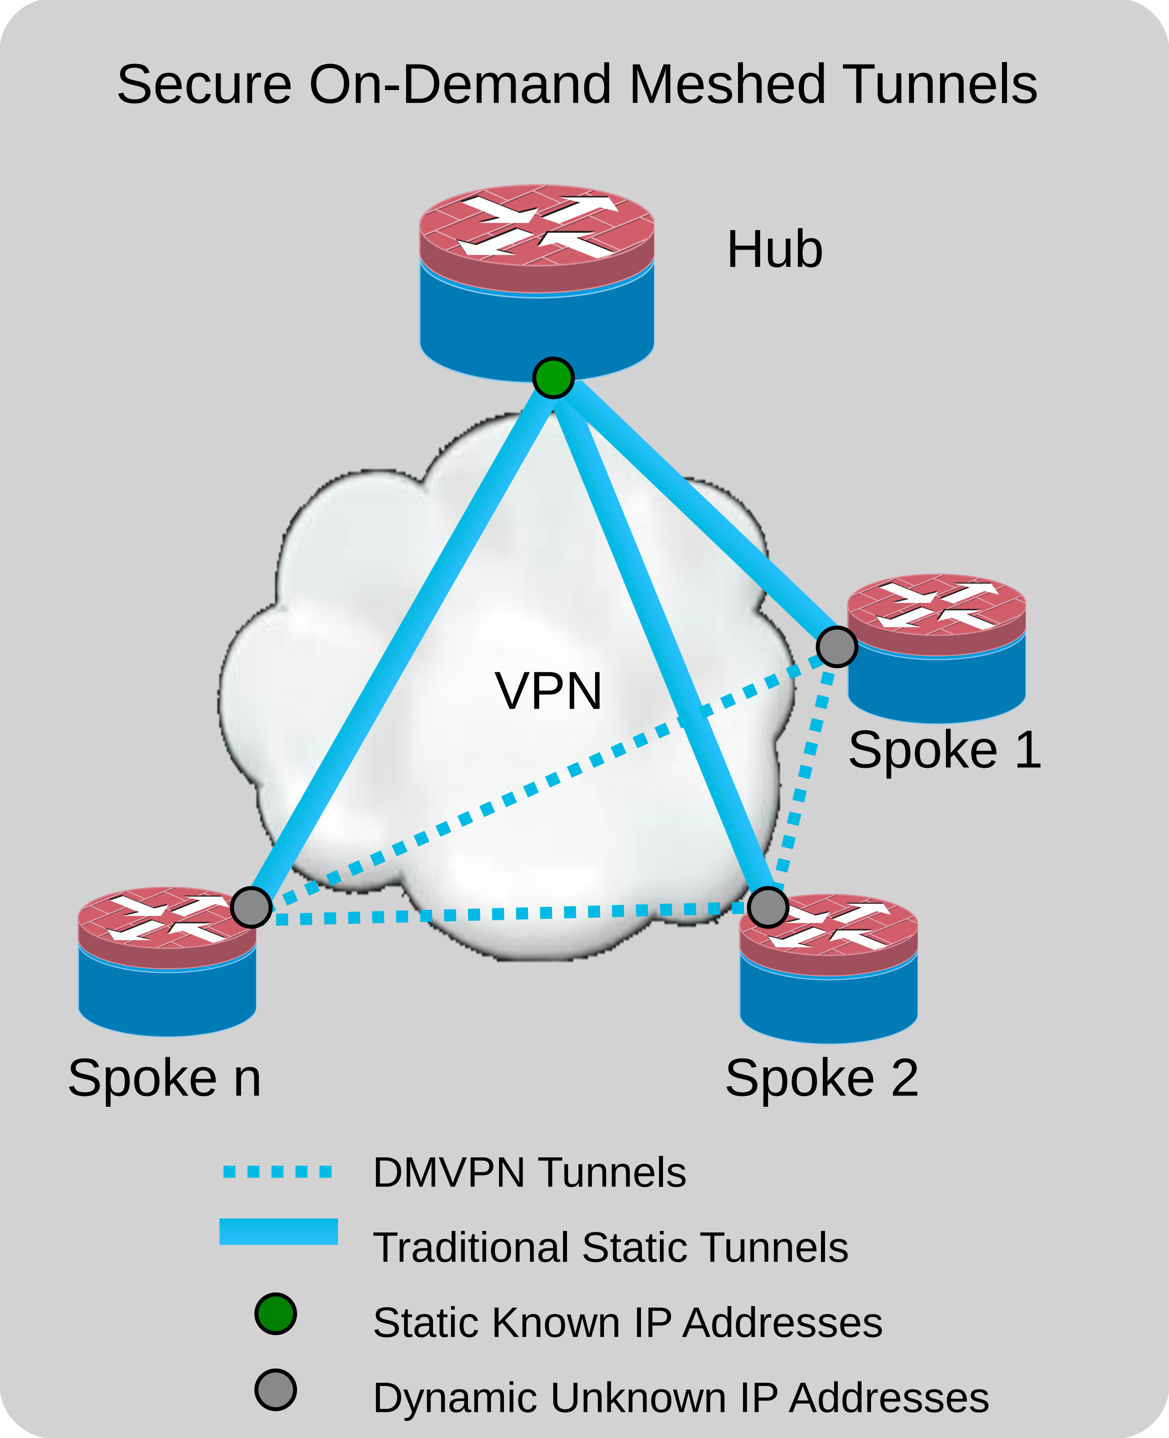

Dynamic Multipoint VPN (DMVPN)

Operational Example

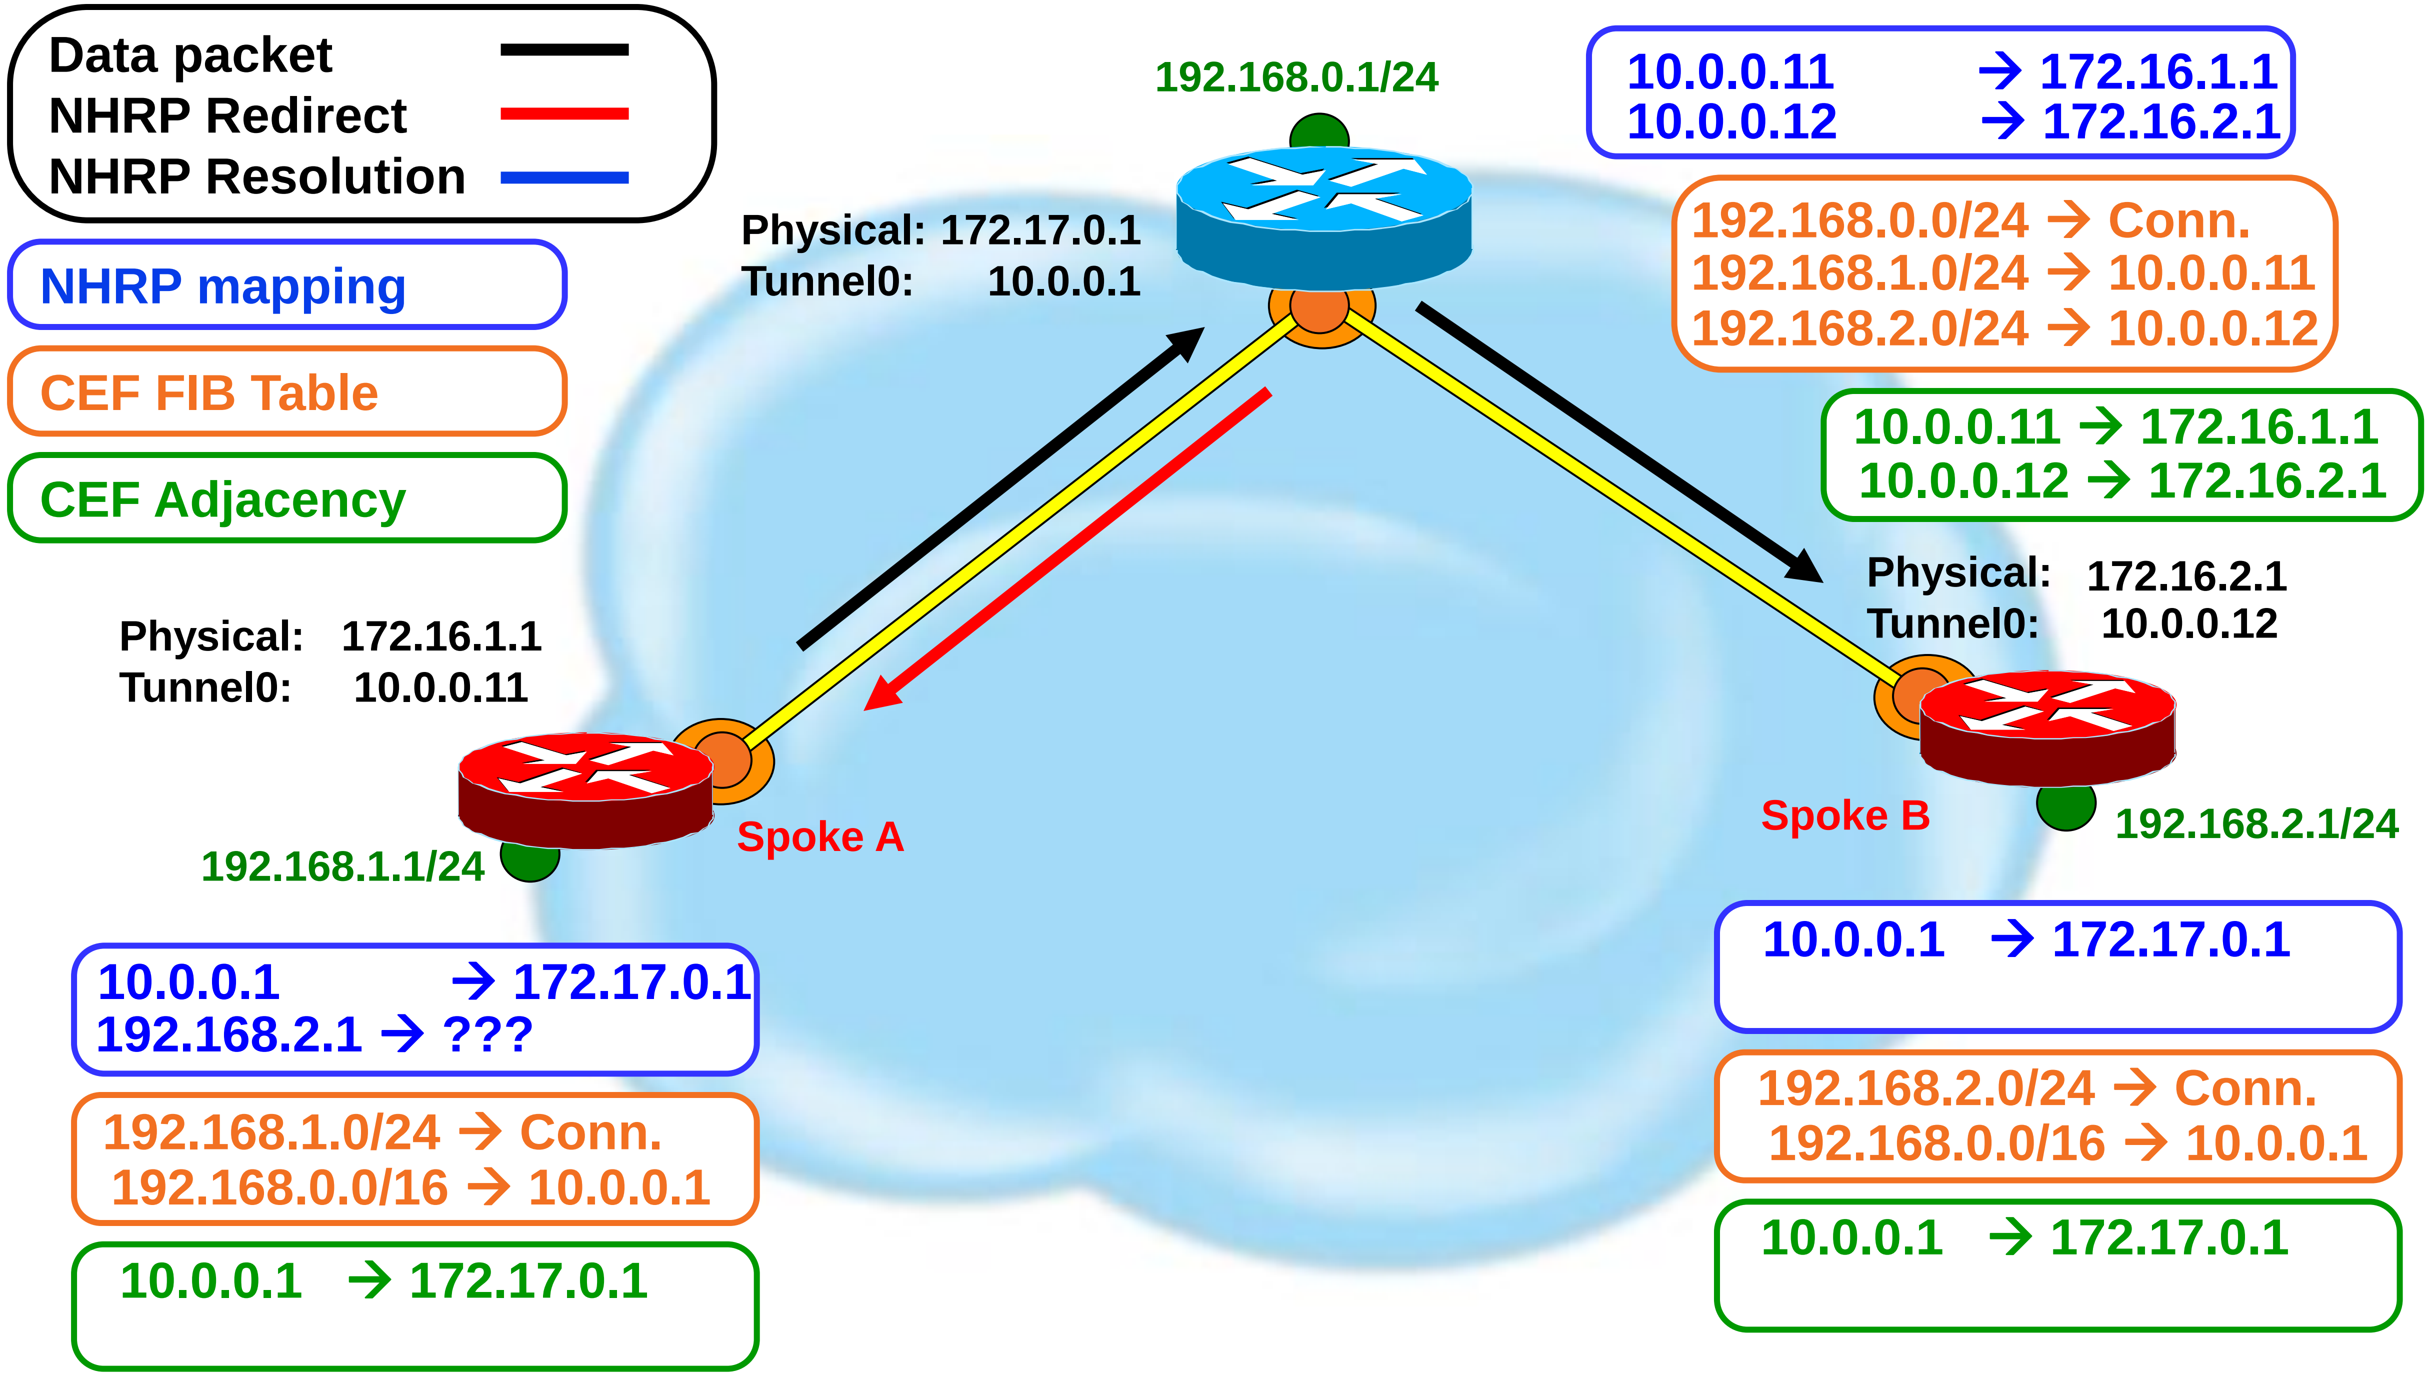

Dynamic Multipoint VPN (DMVPN)

Operational Example (cont.)

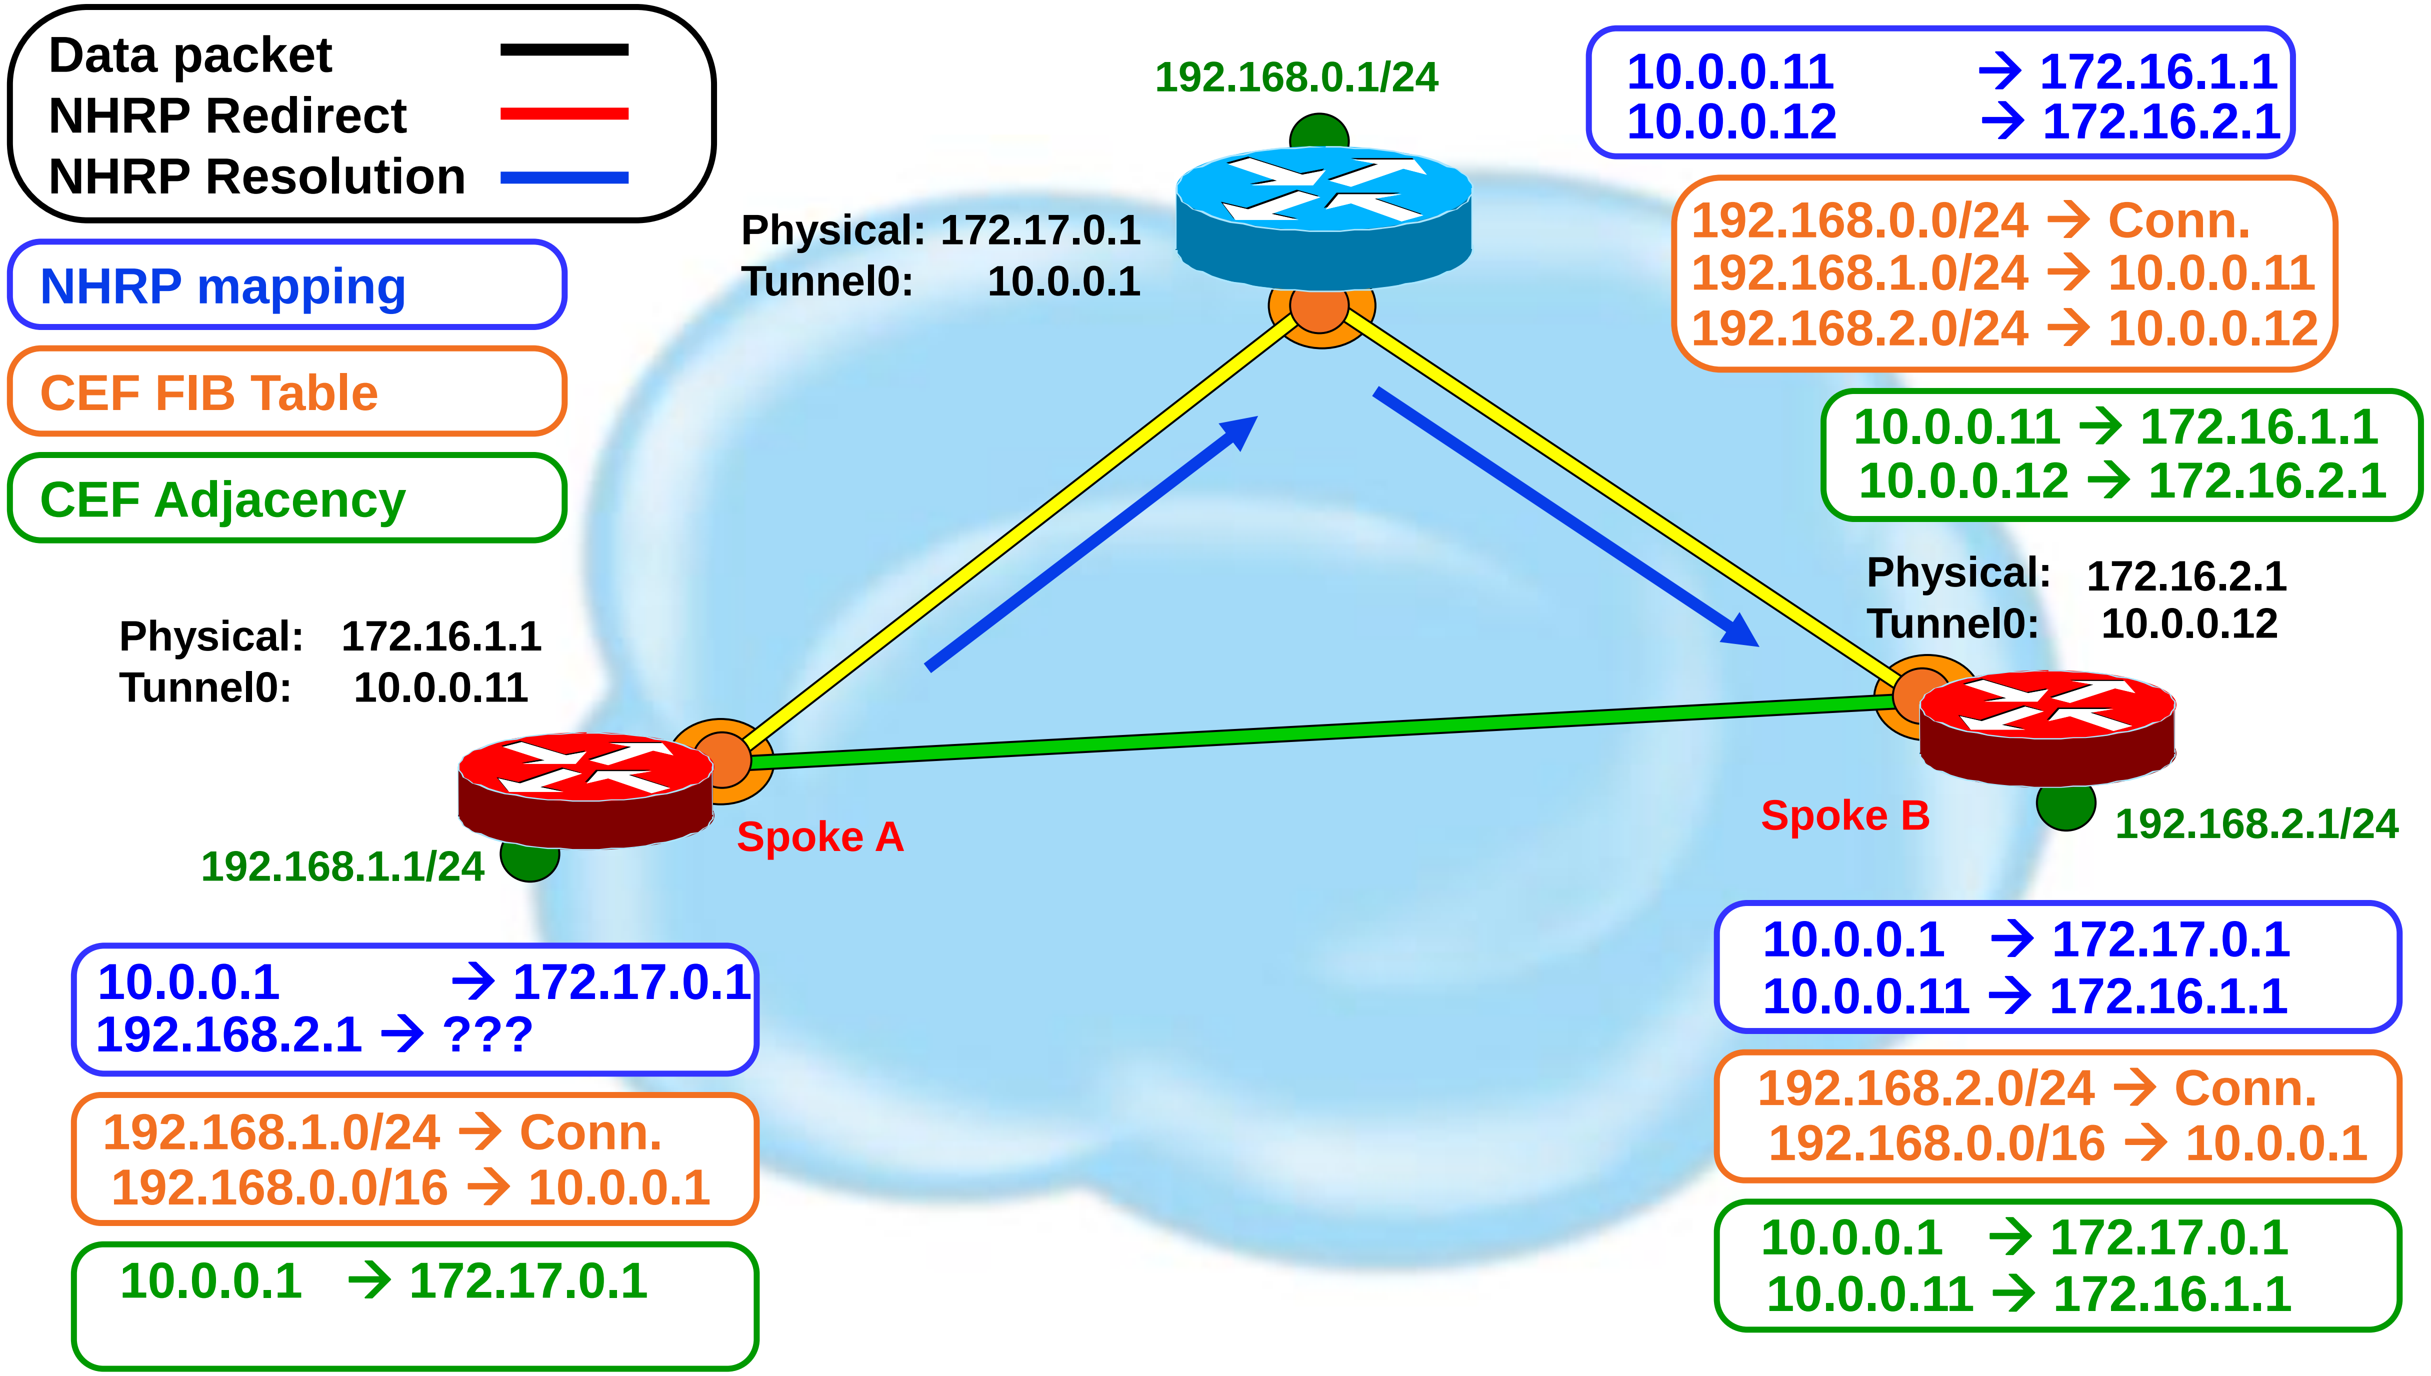

Network Designs

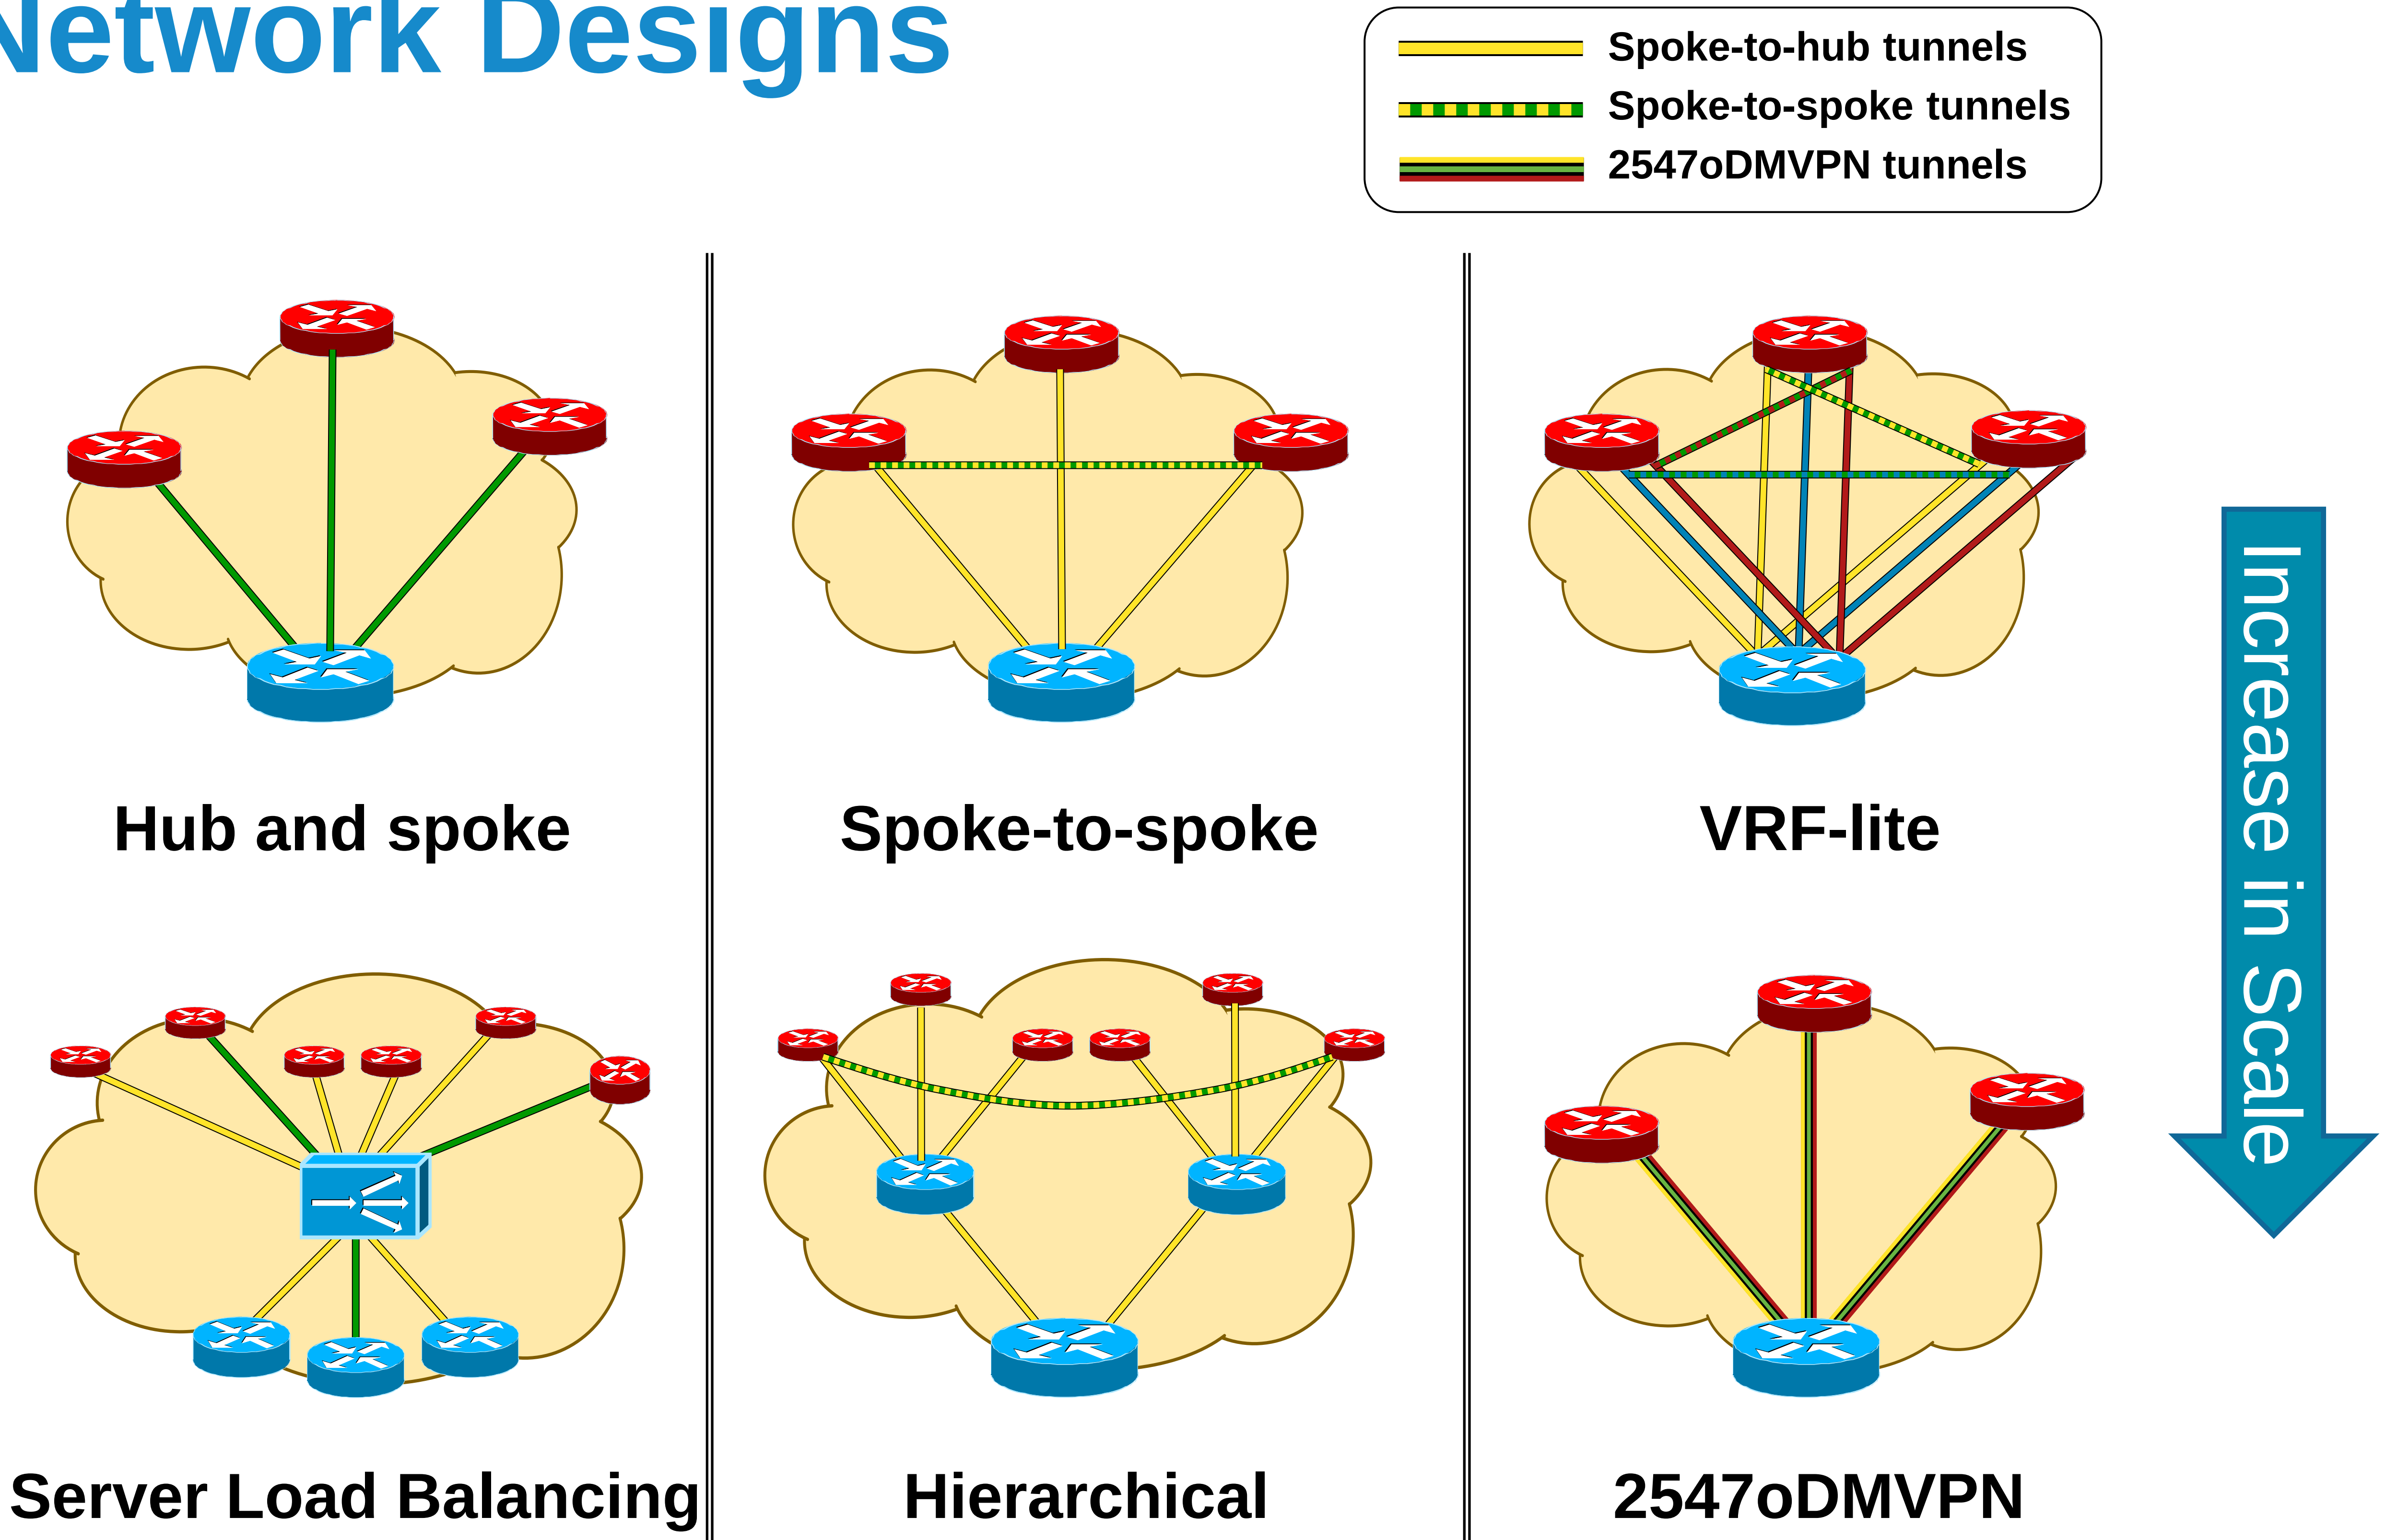

Any-to-Any Encryption

Before and After GETVPN

Public/Private WAN

Before: IPSec P2P Tunnels

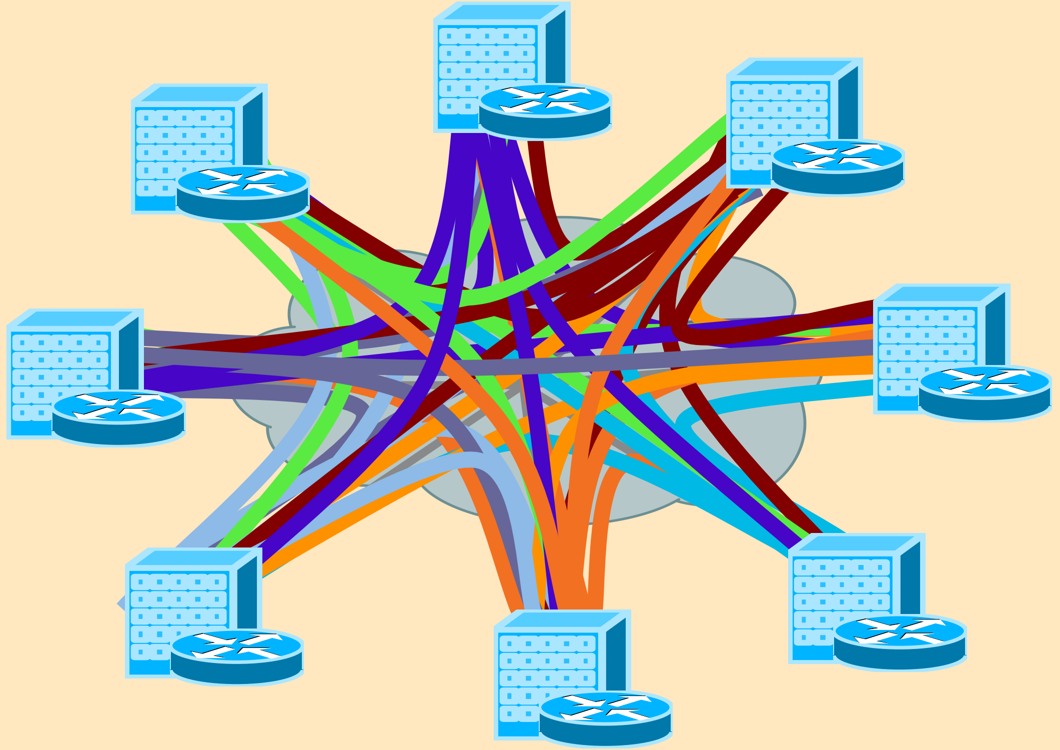

- Scalability—an issue (N^2 problem)
- Overlay routing
- Any-to-any instant connectivity can't be done to scale
- Limited QoS
- Inefficient Multicast replication

Private WAN

After: Tunnel-Less VPN

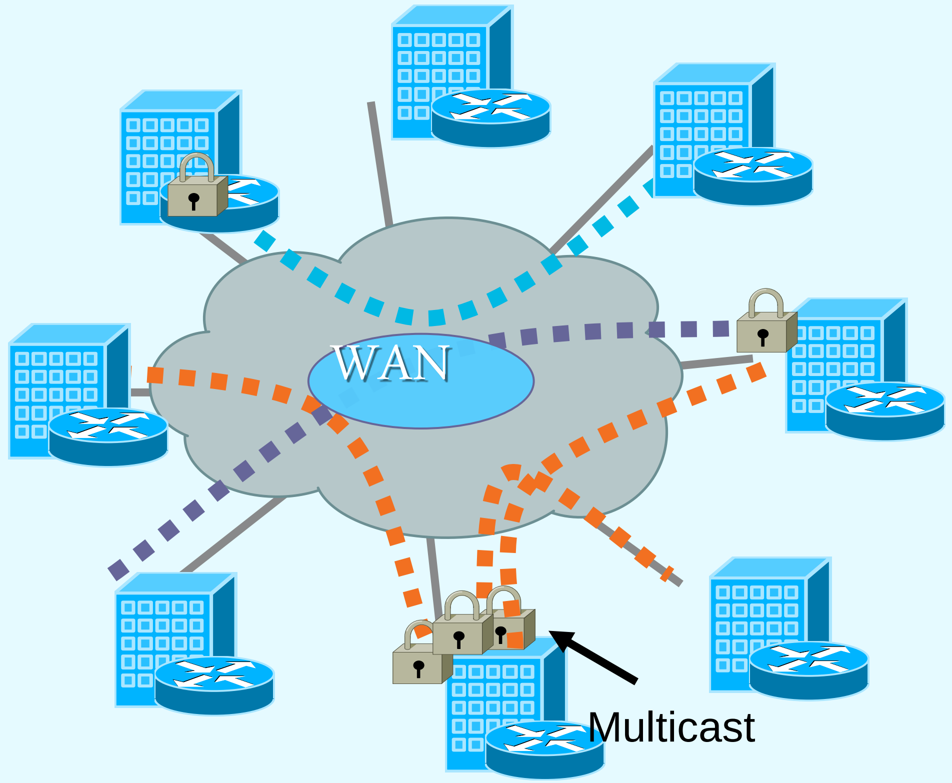

- Scalable architecture for any-to-any connectivity and encryption
- No overlays—native routing
- Any-to-any instant connectivity
- Enhanced QoS
- Efficient Multicast replication

Group Security Functions

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

Routing Member
- Forwarding
- Replication
- Routing

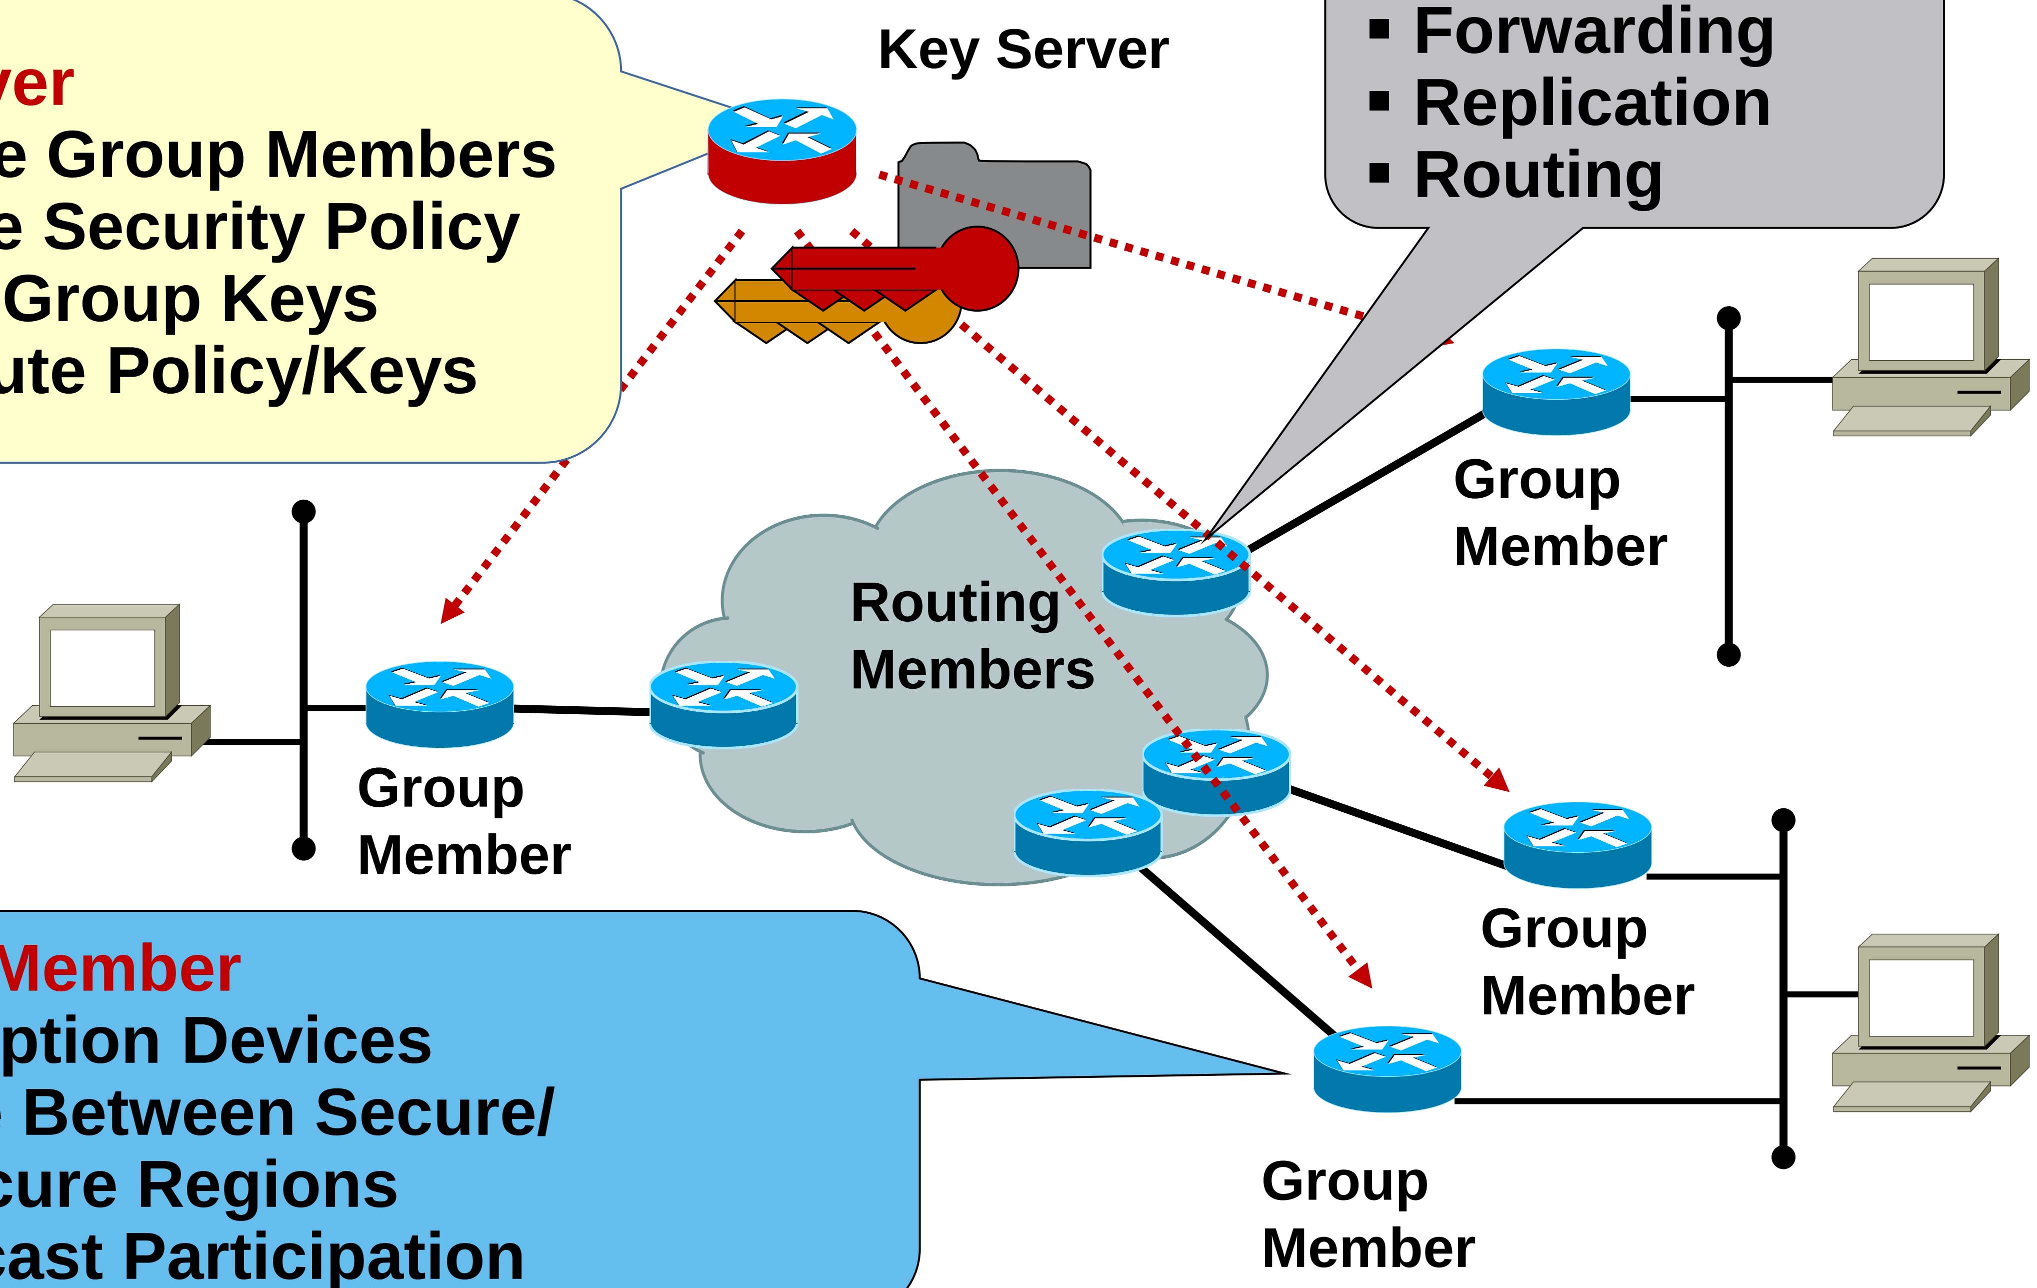

Group Member
- Encryption Devices
- Route Between Secure/ Unsecure Regions
- Multicast Participation

Group Security Elements

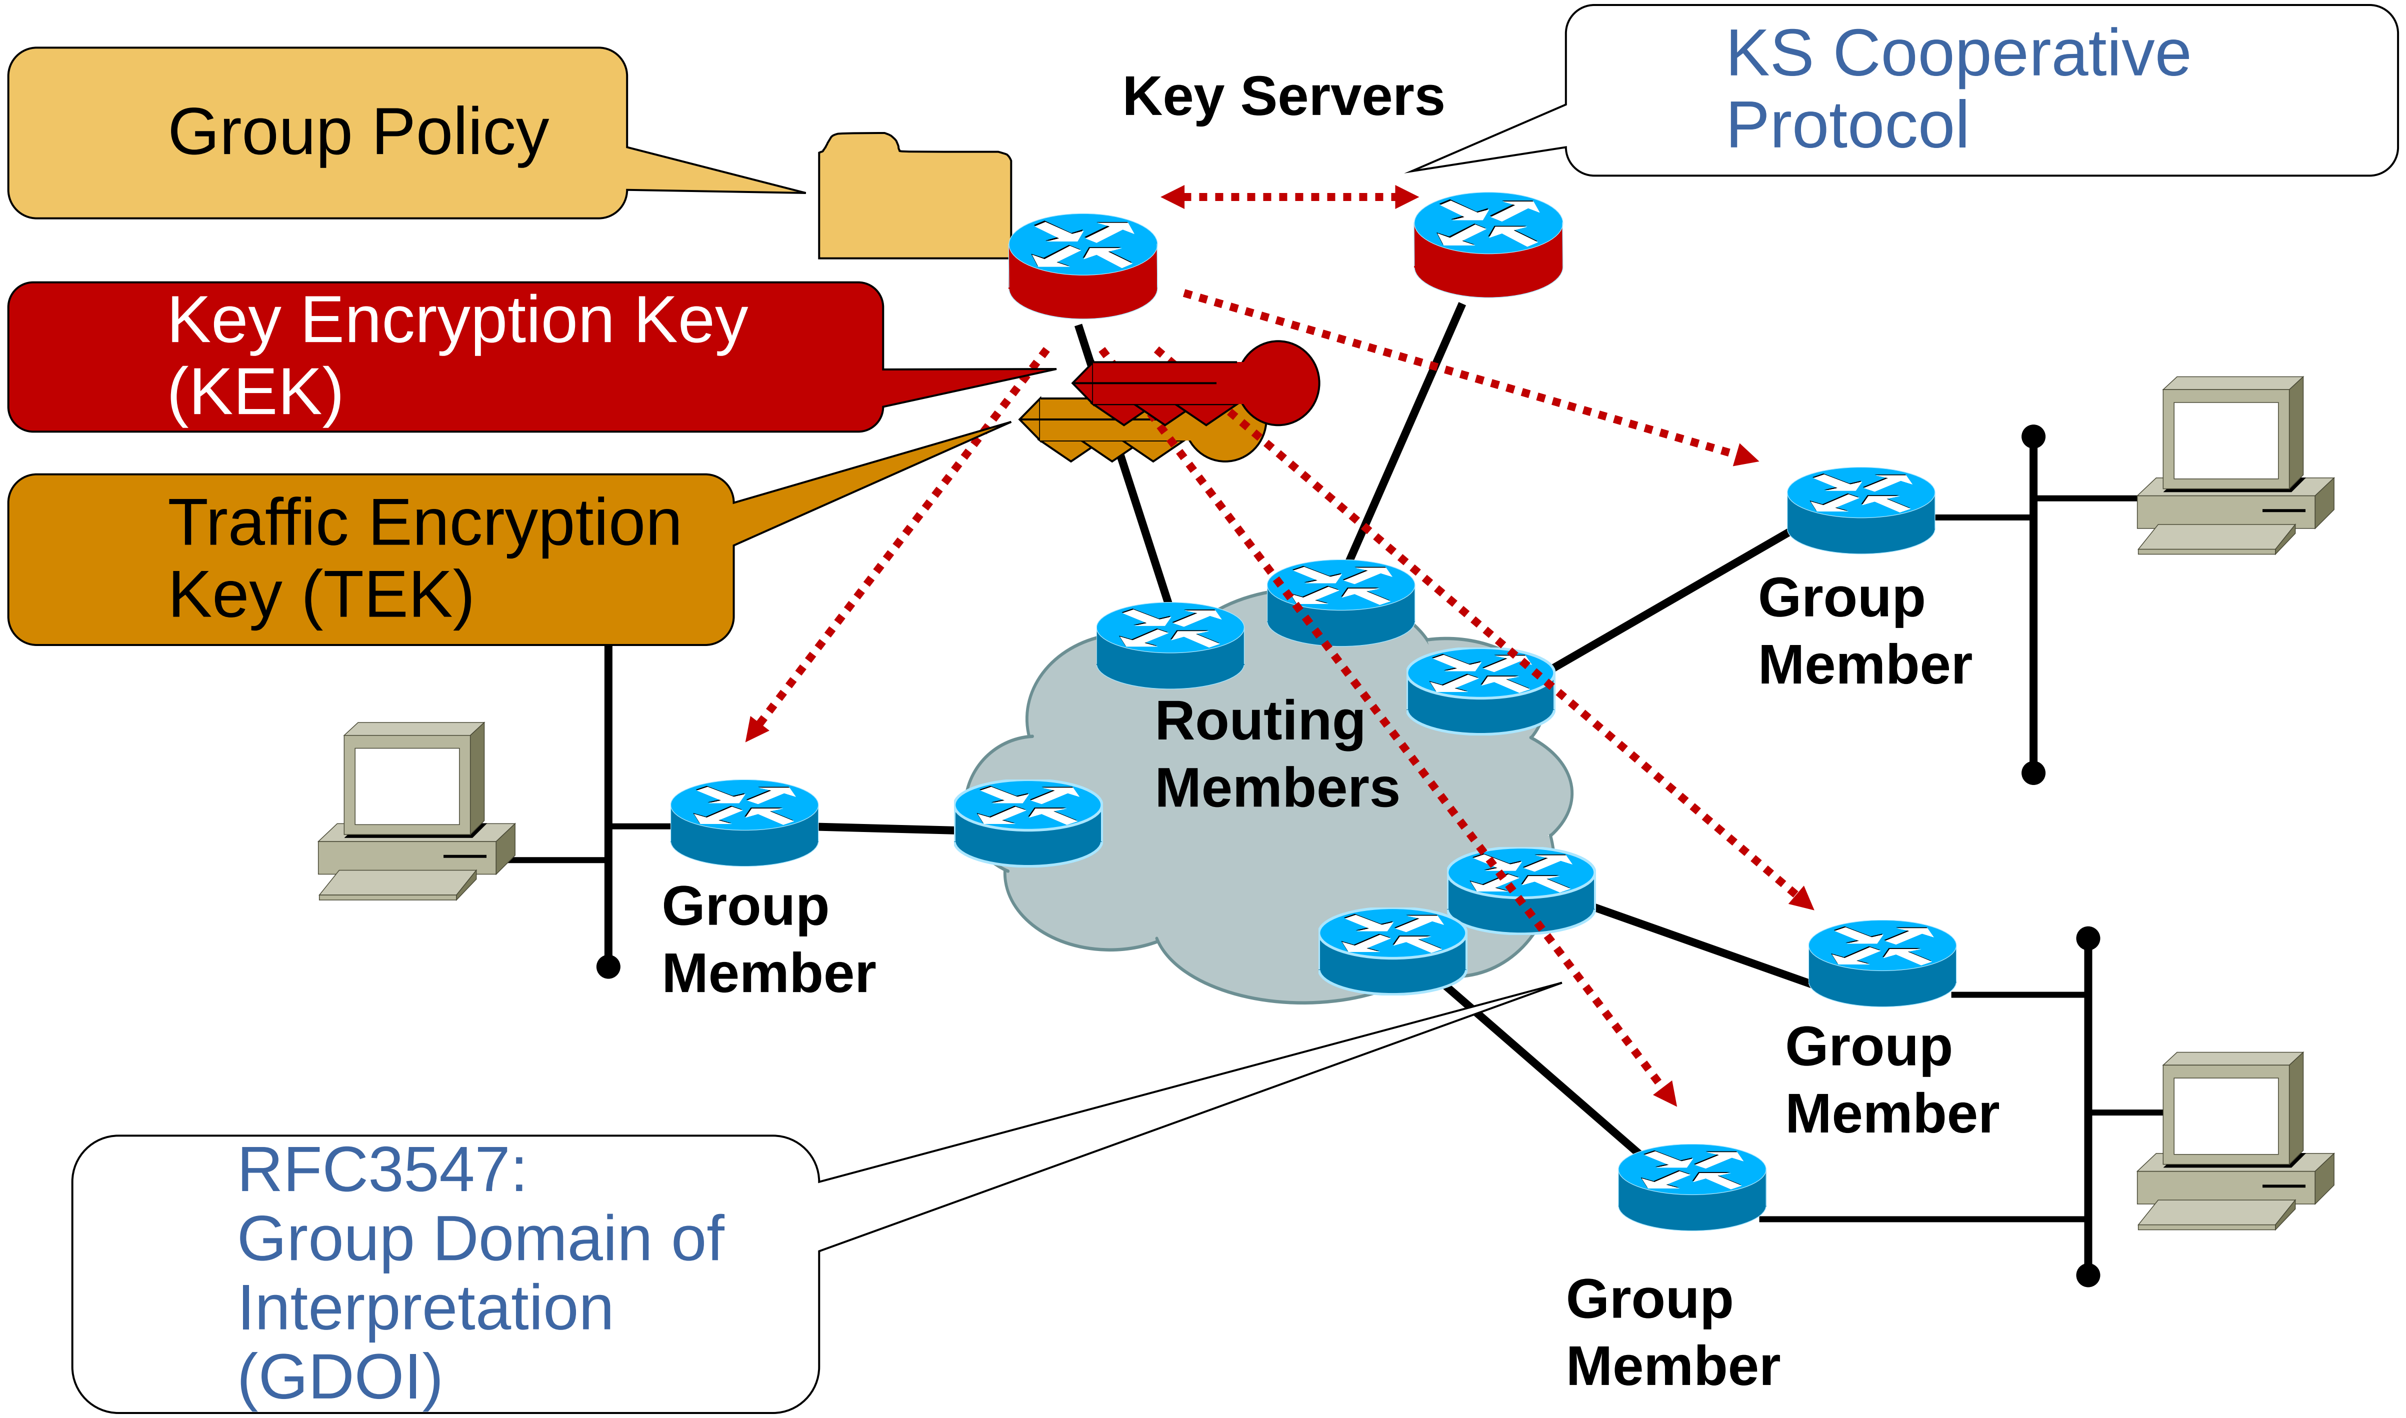

GETVPN - Group Key Technology

Operation Example

- **Step 1**: Group Members (GM) "register" via GDOI (IKE) with the Key Server (KS)
  - KS authenticates and authorises 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"

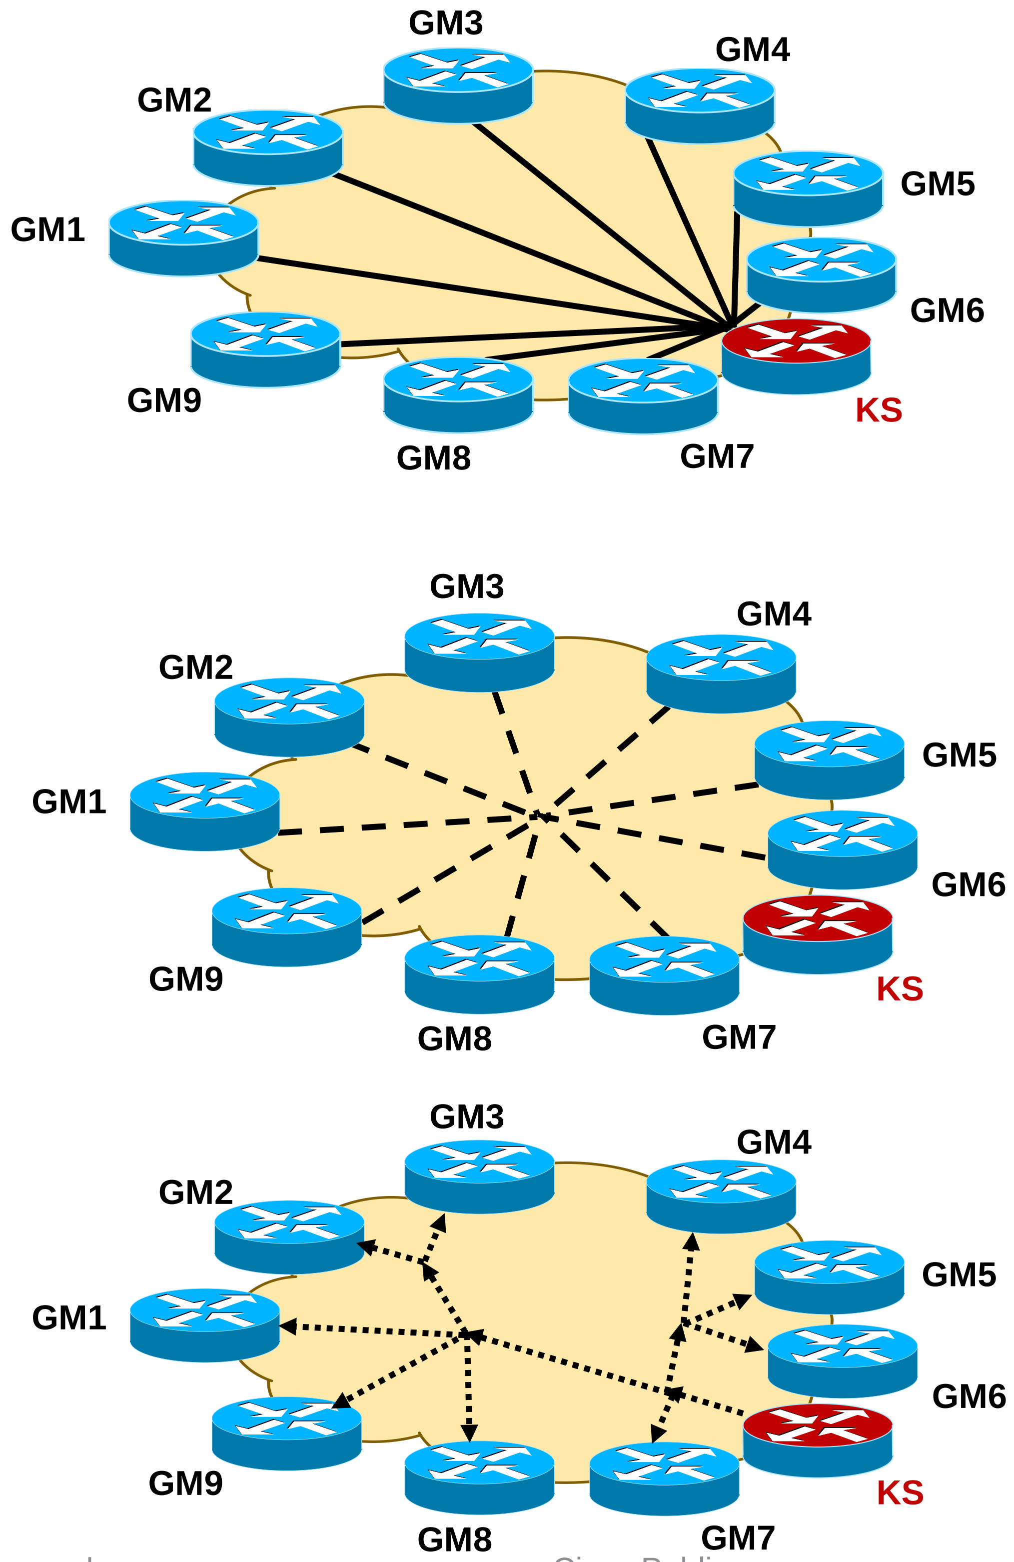

GETVPN Virtualisation Deployment Model

GETVPN Segmented WAN

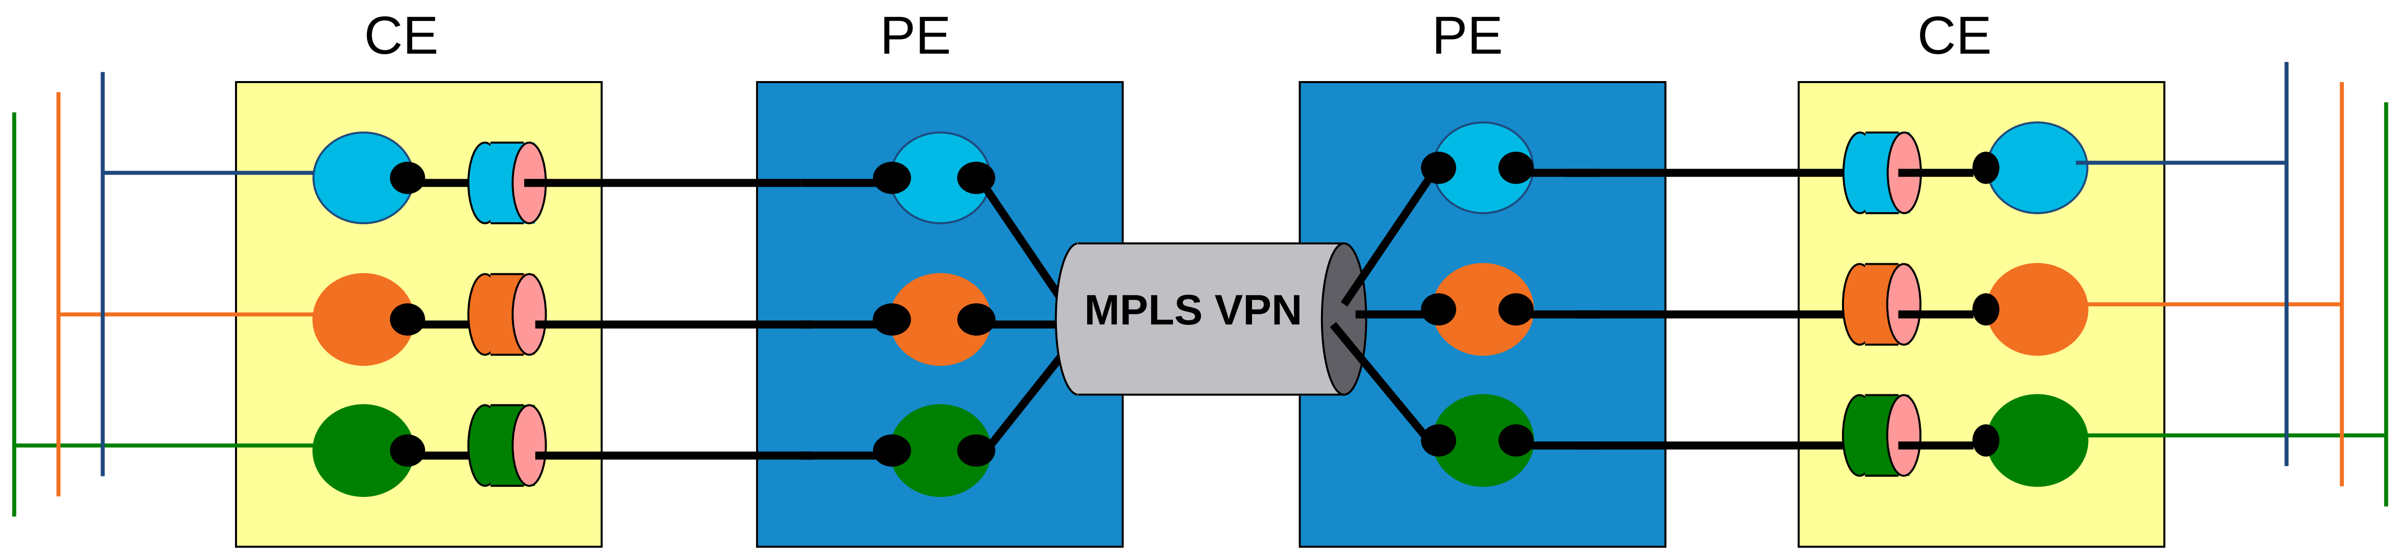

LISP with GETVPN

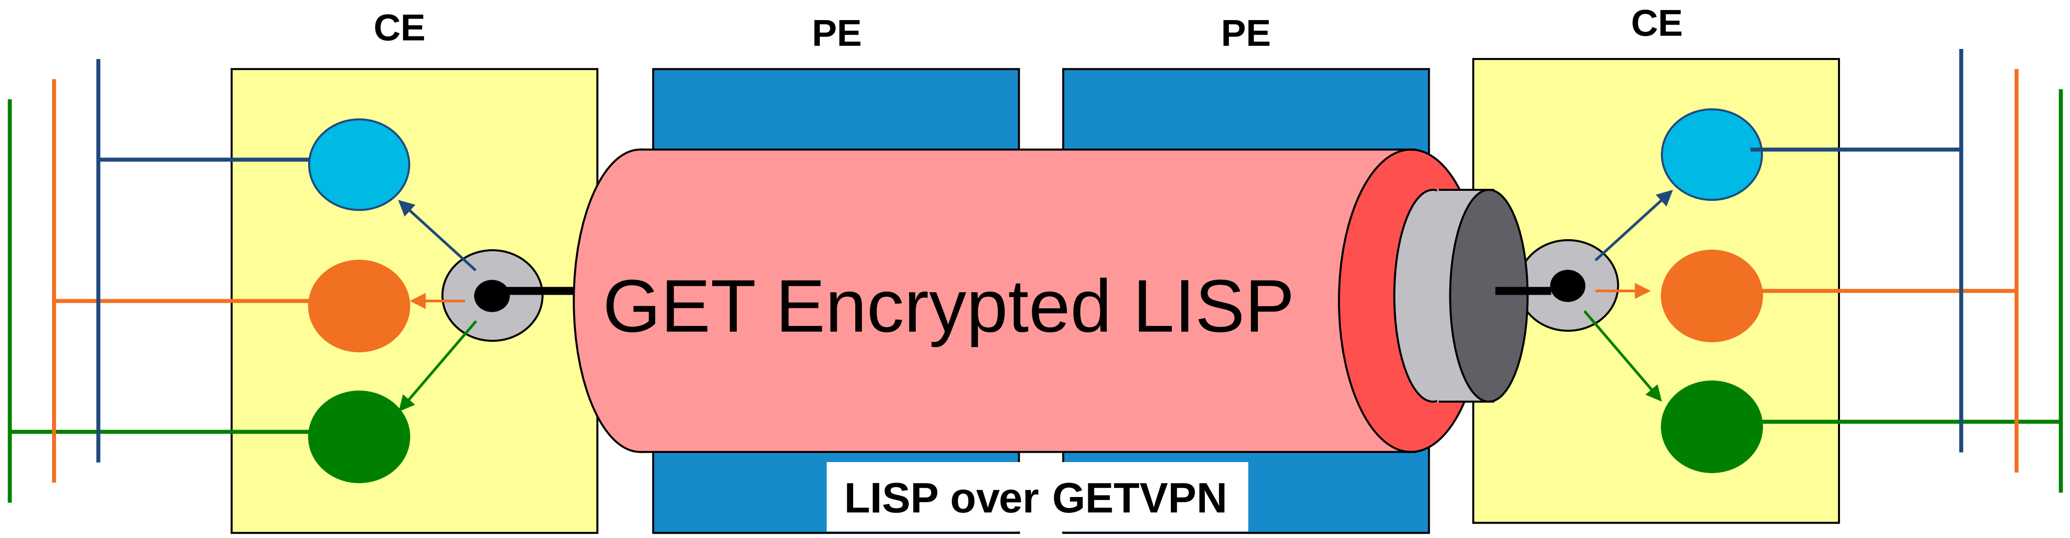

Agenda

- WAN Technologies & Solutions
  - WAN Transport Technologies
  - WAN Overlay Technologies
  - WAN Optimisation
  - Wide Area Network Quality of Service
- WAN Architecture Design Considerations
  - WAN Design and Best Practices
  - Secure WAN Communication with GETVPN
  - DMVPN Over Internet Deployment
- Summary

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

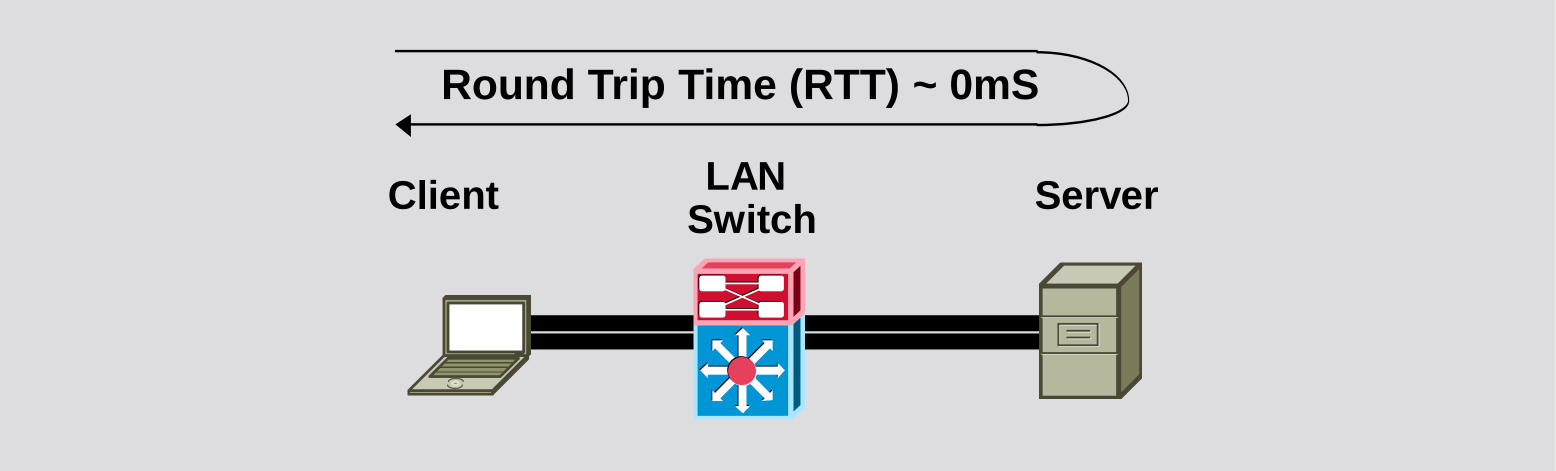

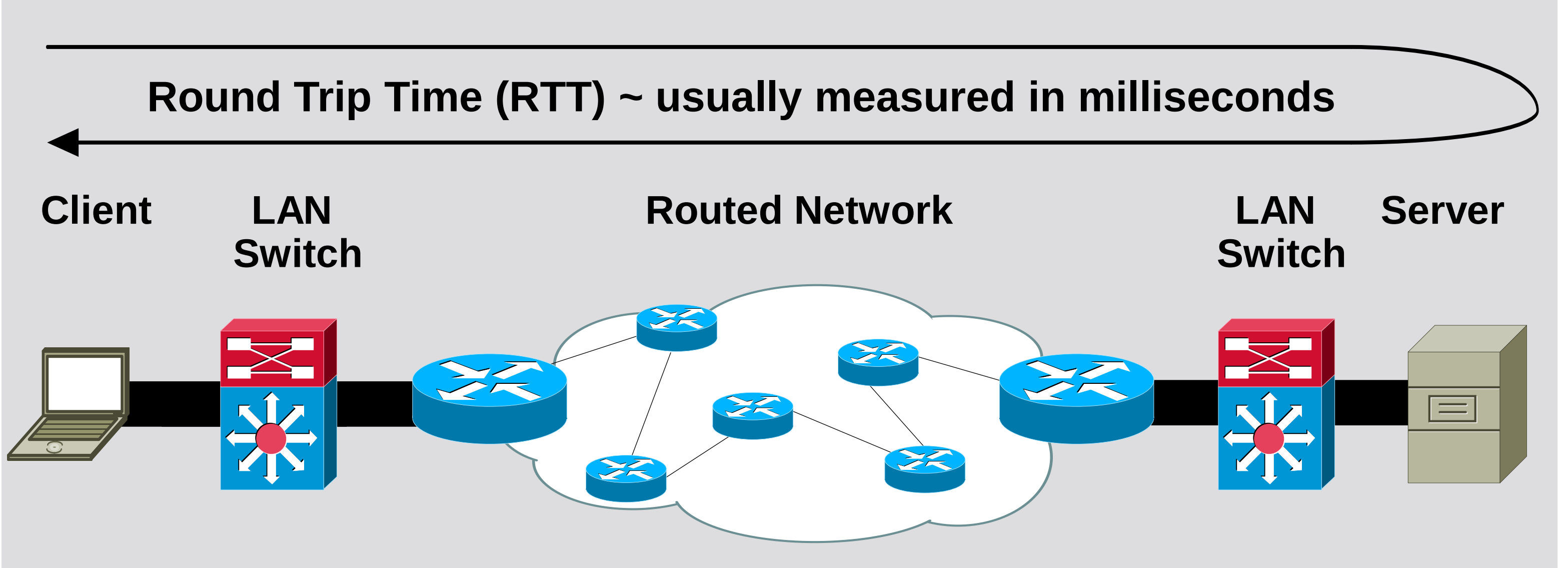

WAN Packet Loss and Latency =
Slow Application Performance =
Keep and manage servers in branch offices ($$$)

TCP Behaviour

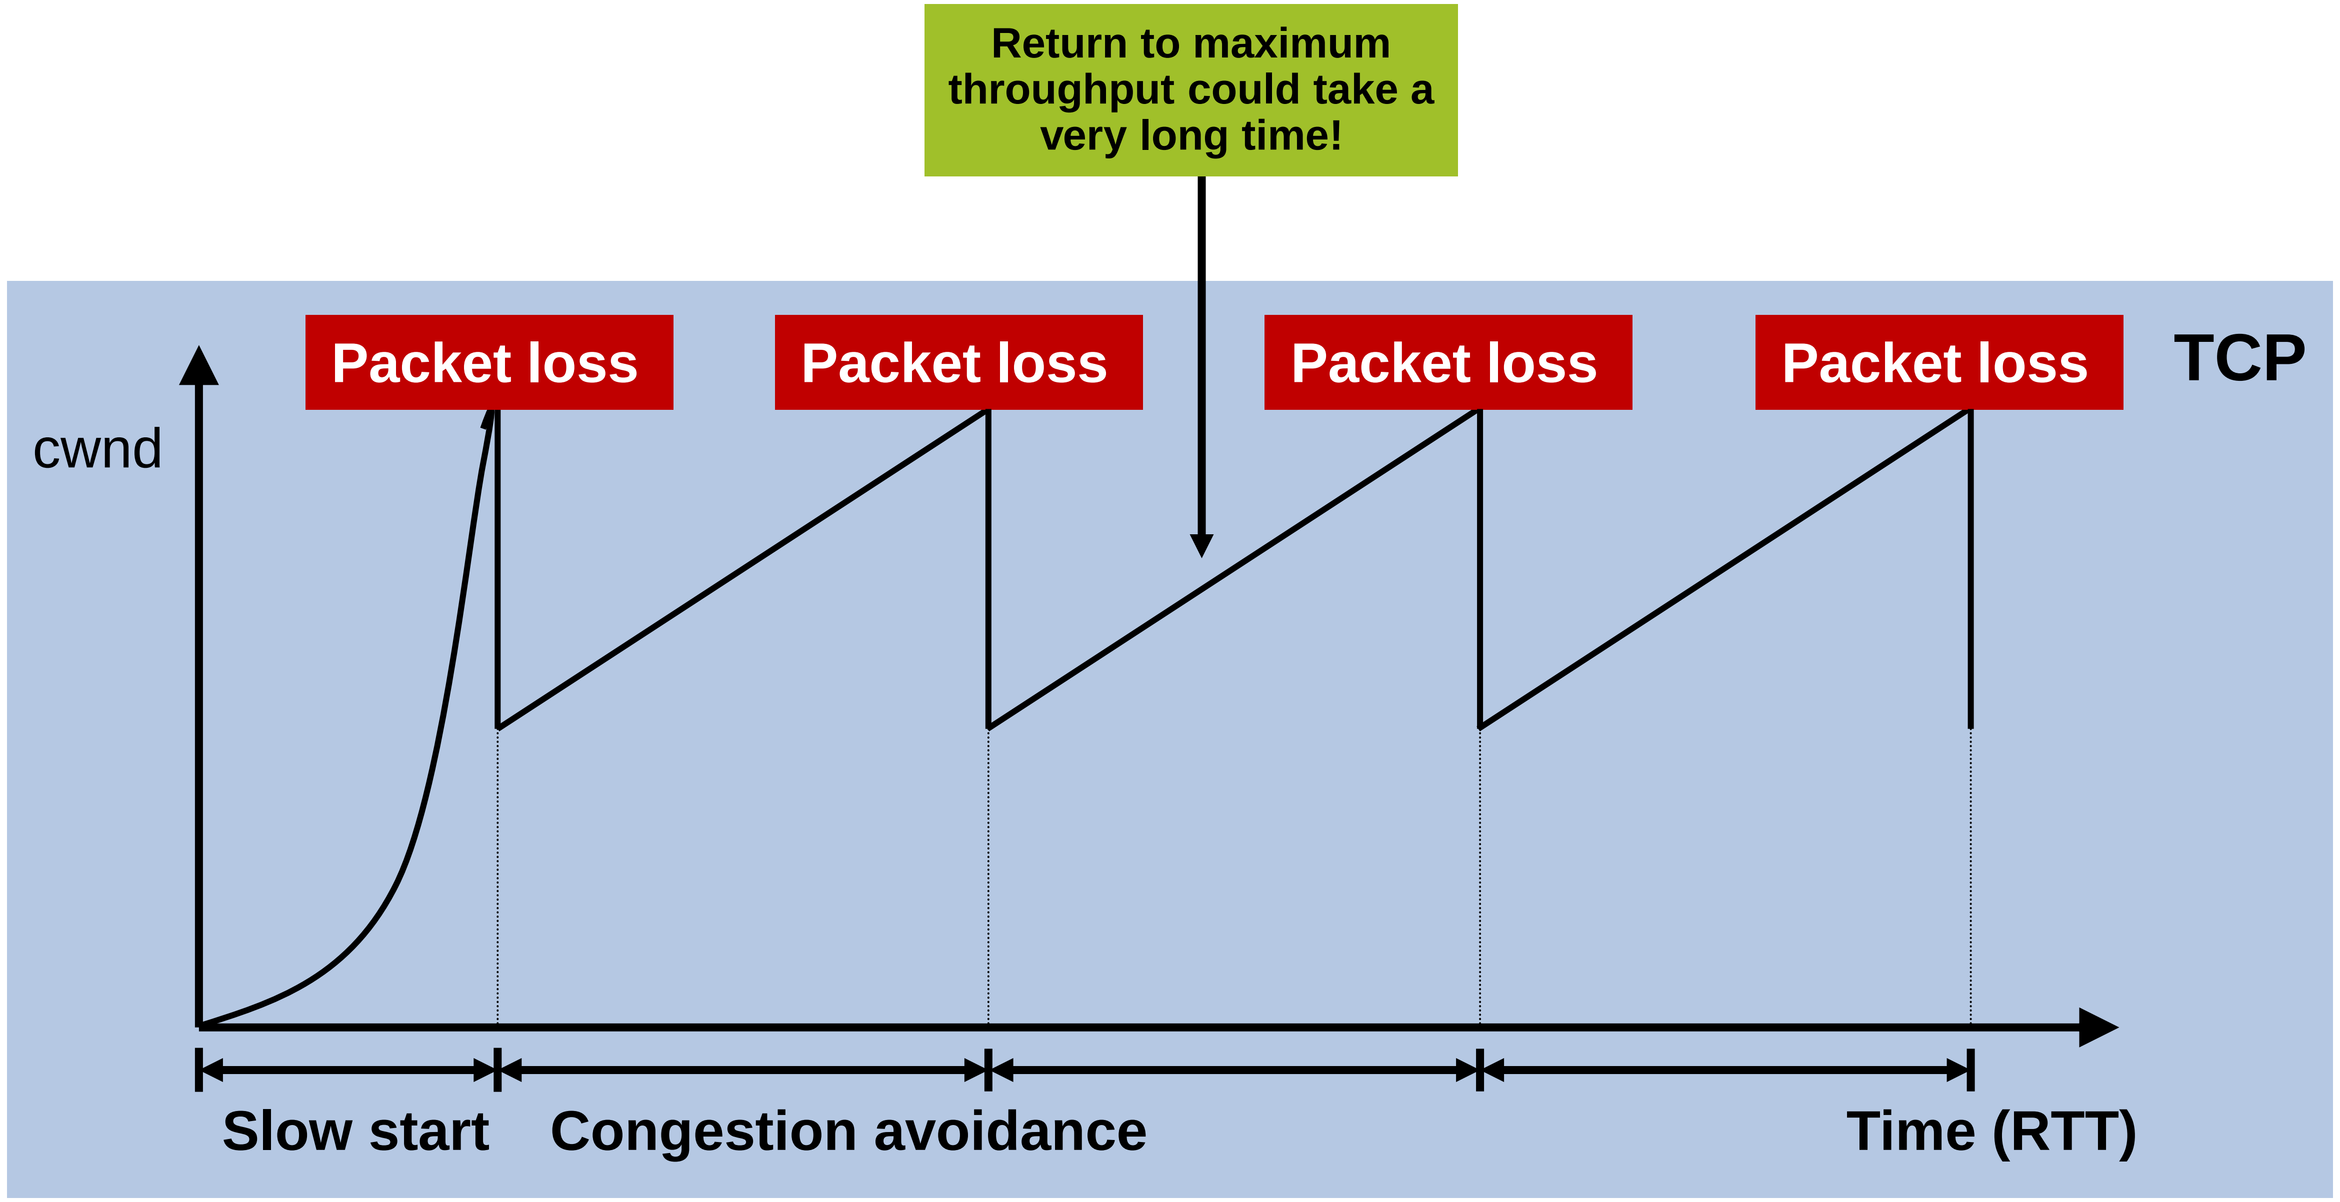

WAAS—TCP Performance Improvement

- Transport Flow Optimisation (TFO) overcomes TCP and WAN bottlenecks
- Shields nodes connections from WAN conditions
  - Clients experience fast acknowledgement
  - Minimise perceived packet loss
  - Eliminate need to use inefficient congestion handling

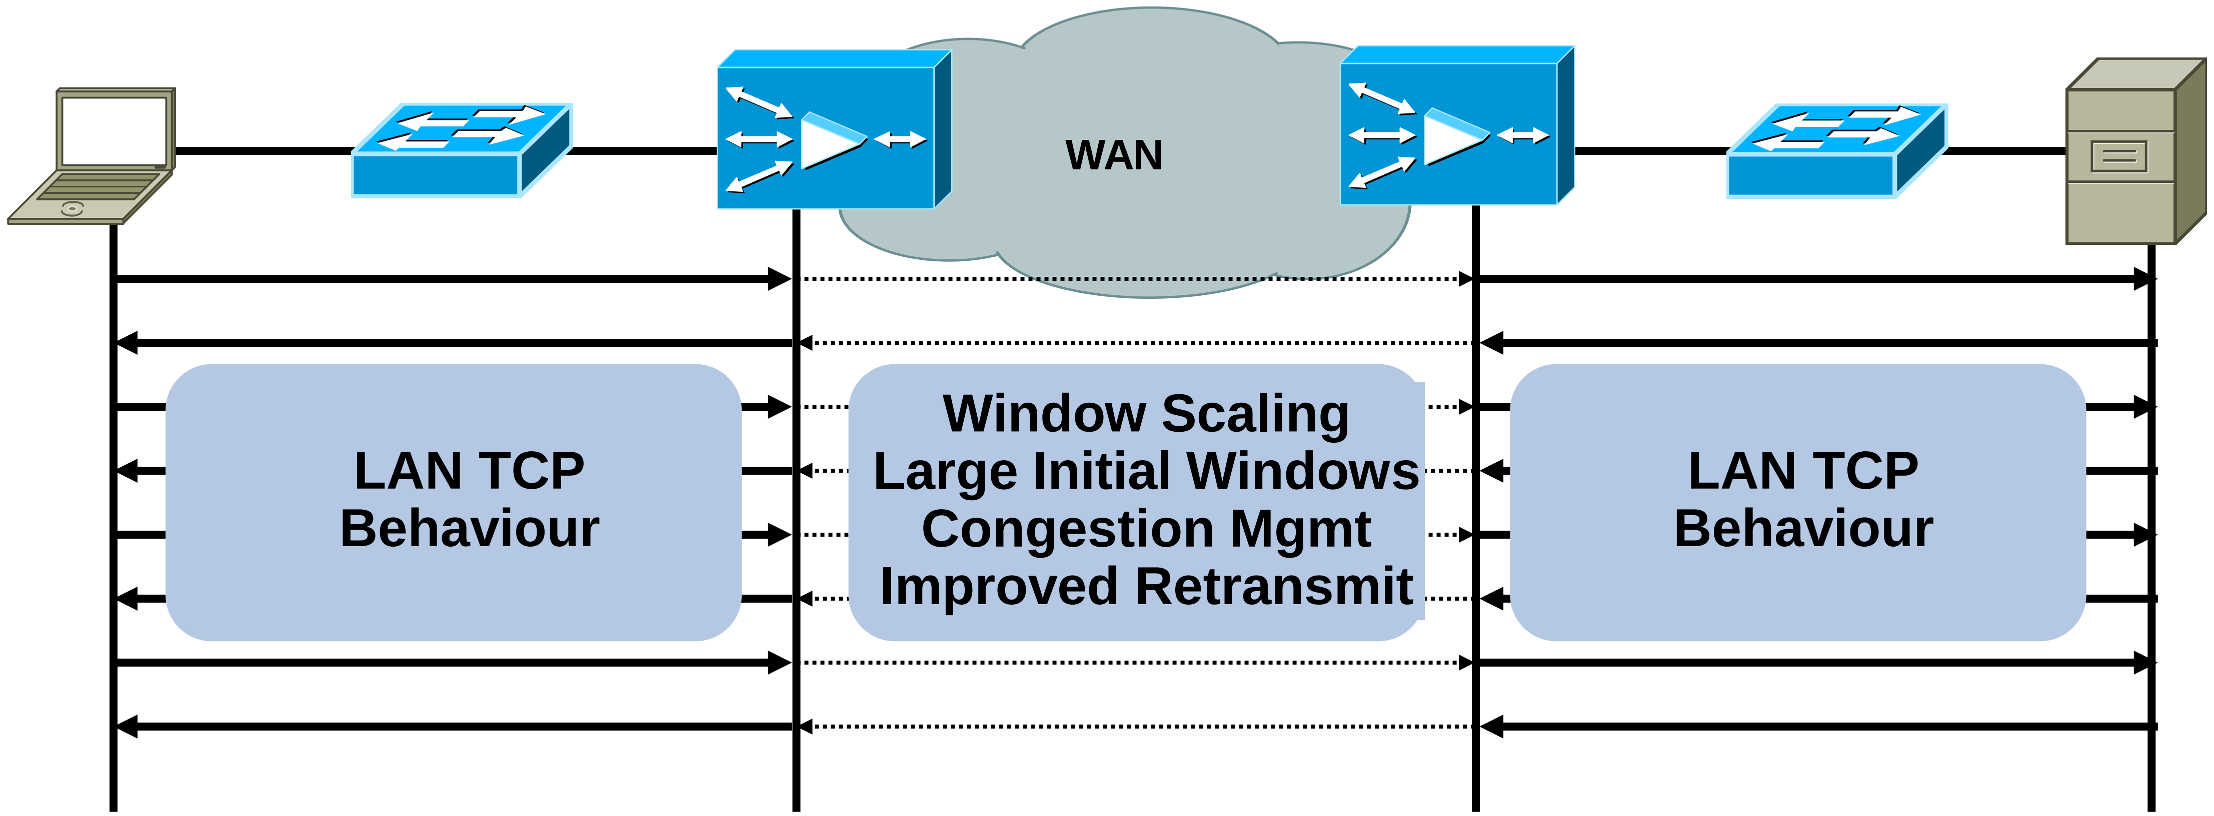

WAAS Overview

DRE and LZ Manage Bandwidth Utilisation

- Data Redundancy Elimination (DRE) provides advanced compression to eliminate redundancy from network flows regardless of application
- LZ compression provides generic compression for all traffic

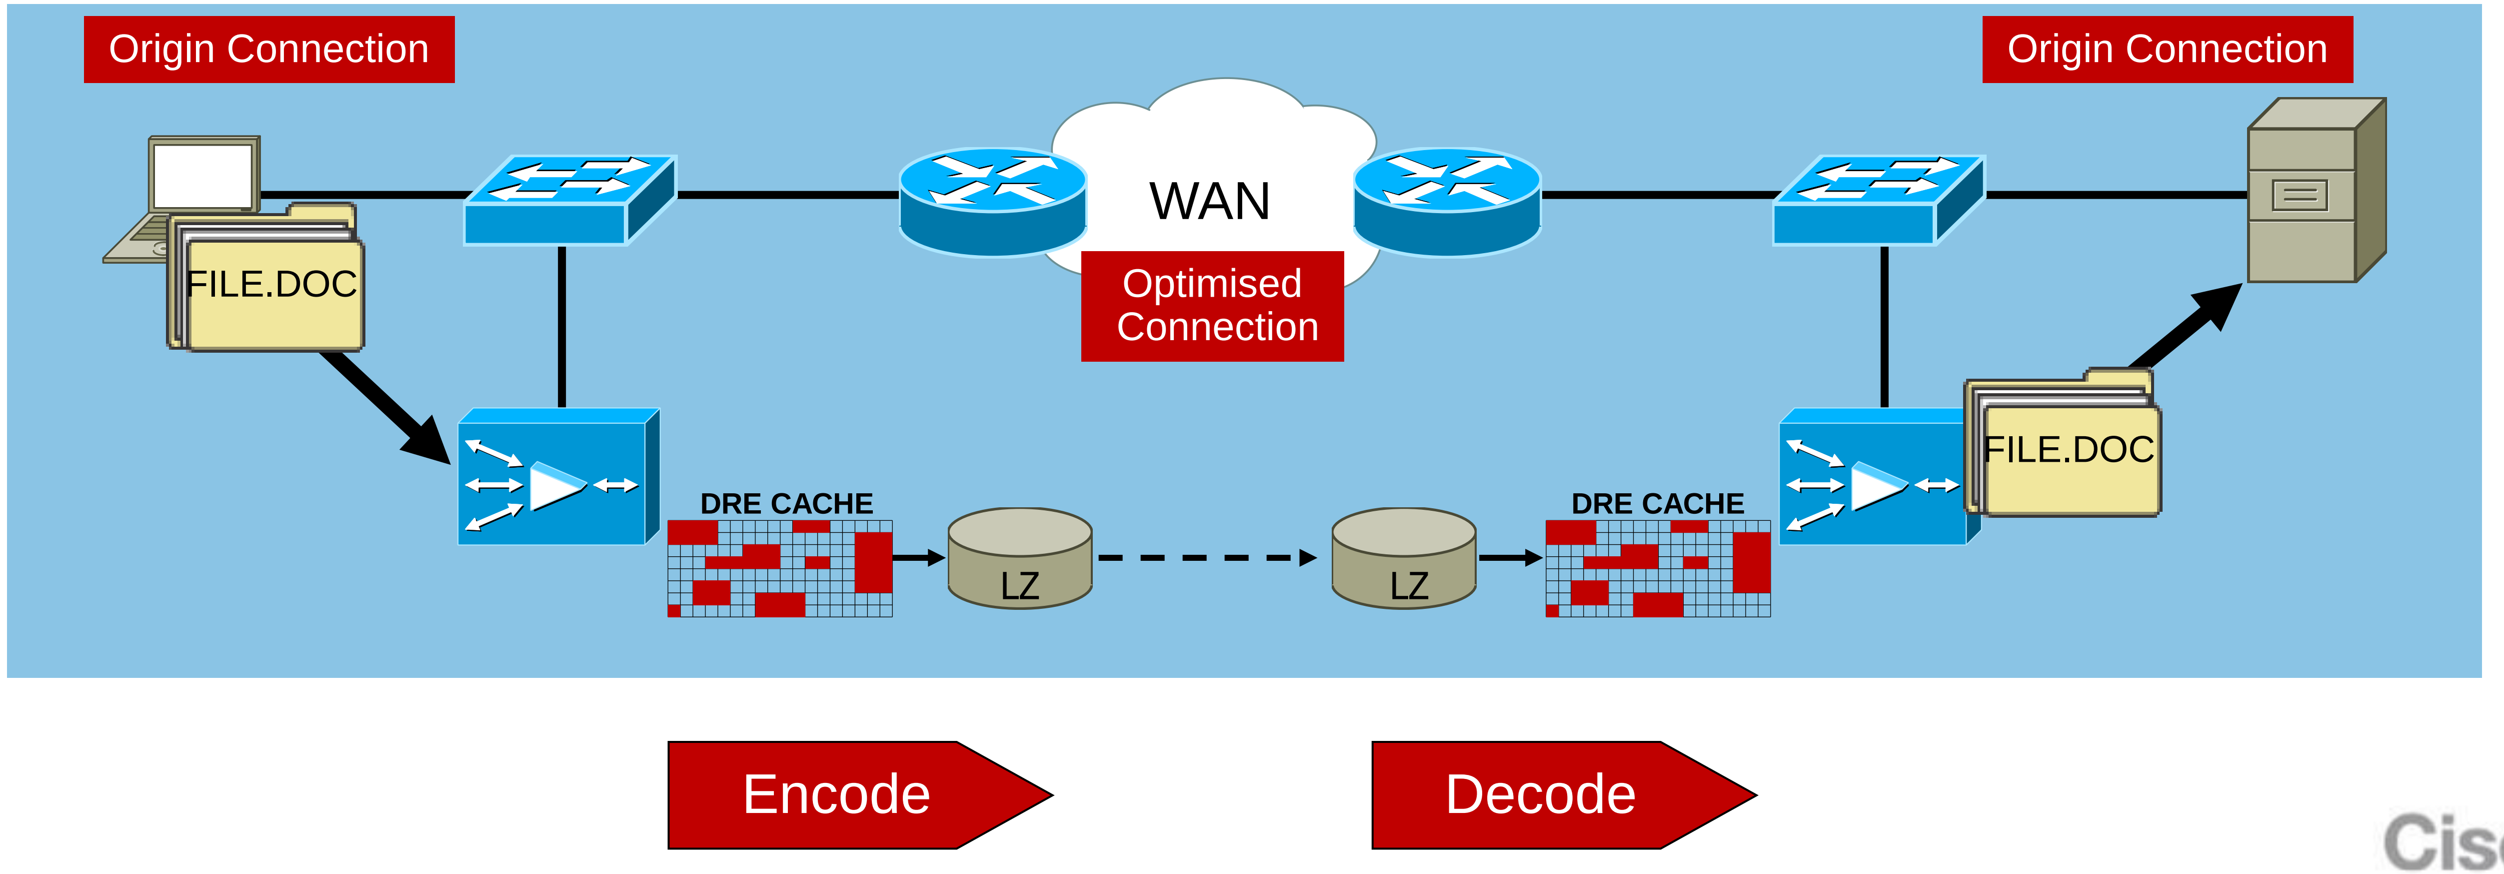

Comparing TCP and Transport

Flow Optimisation

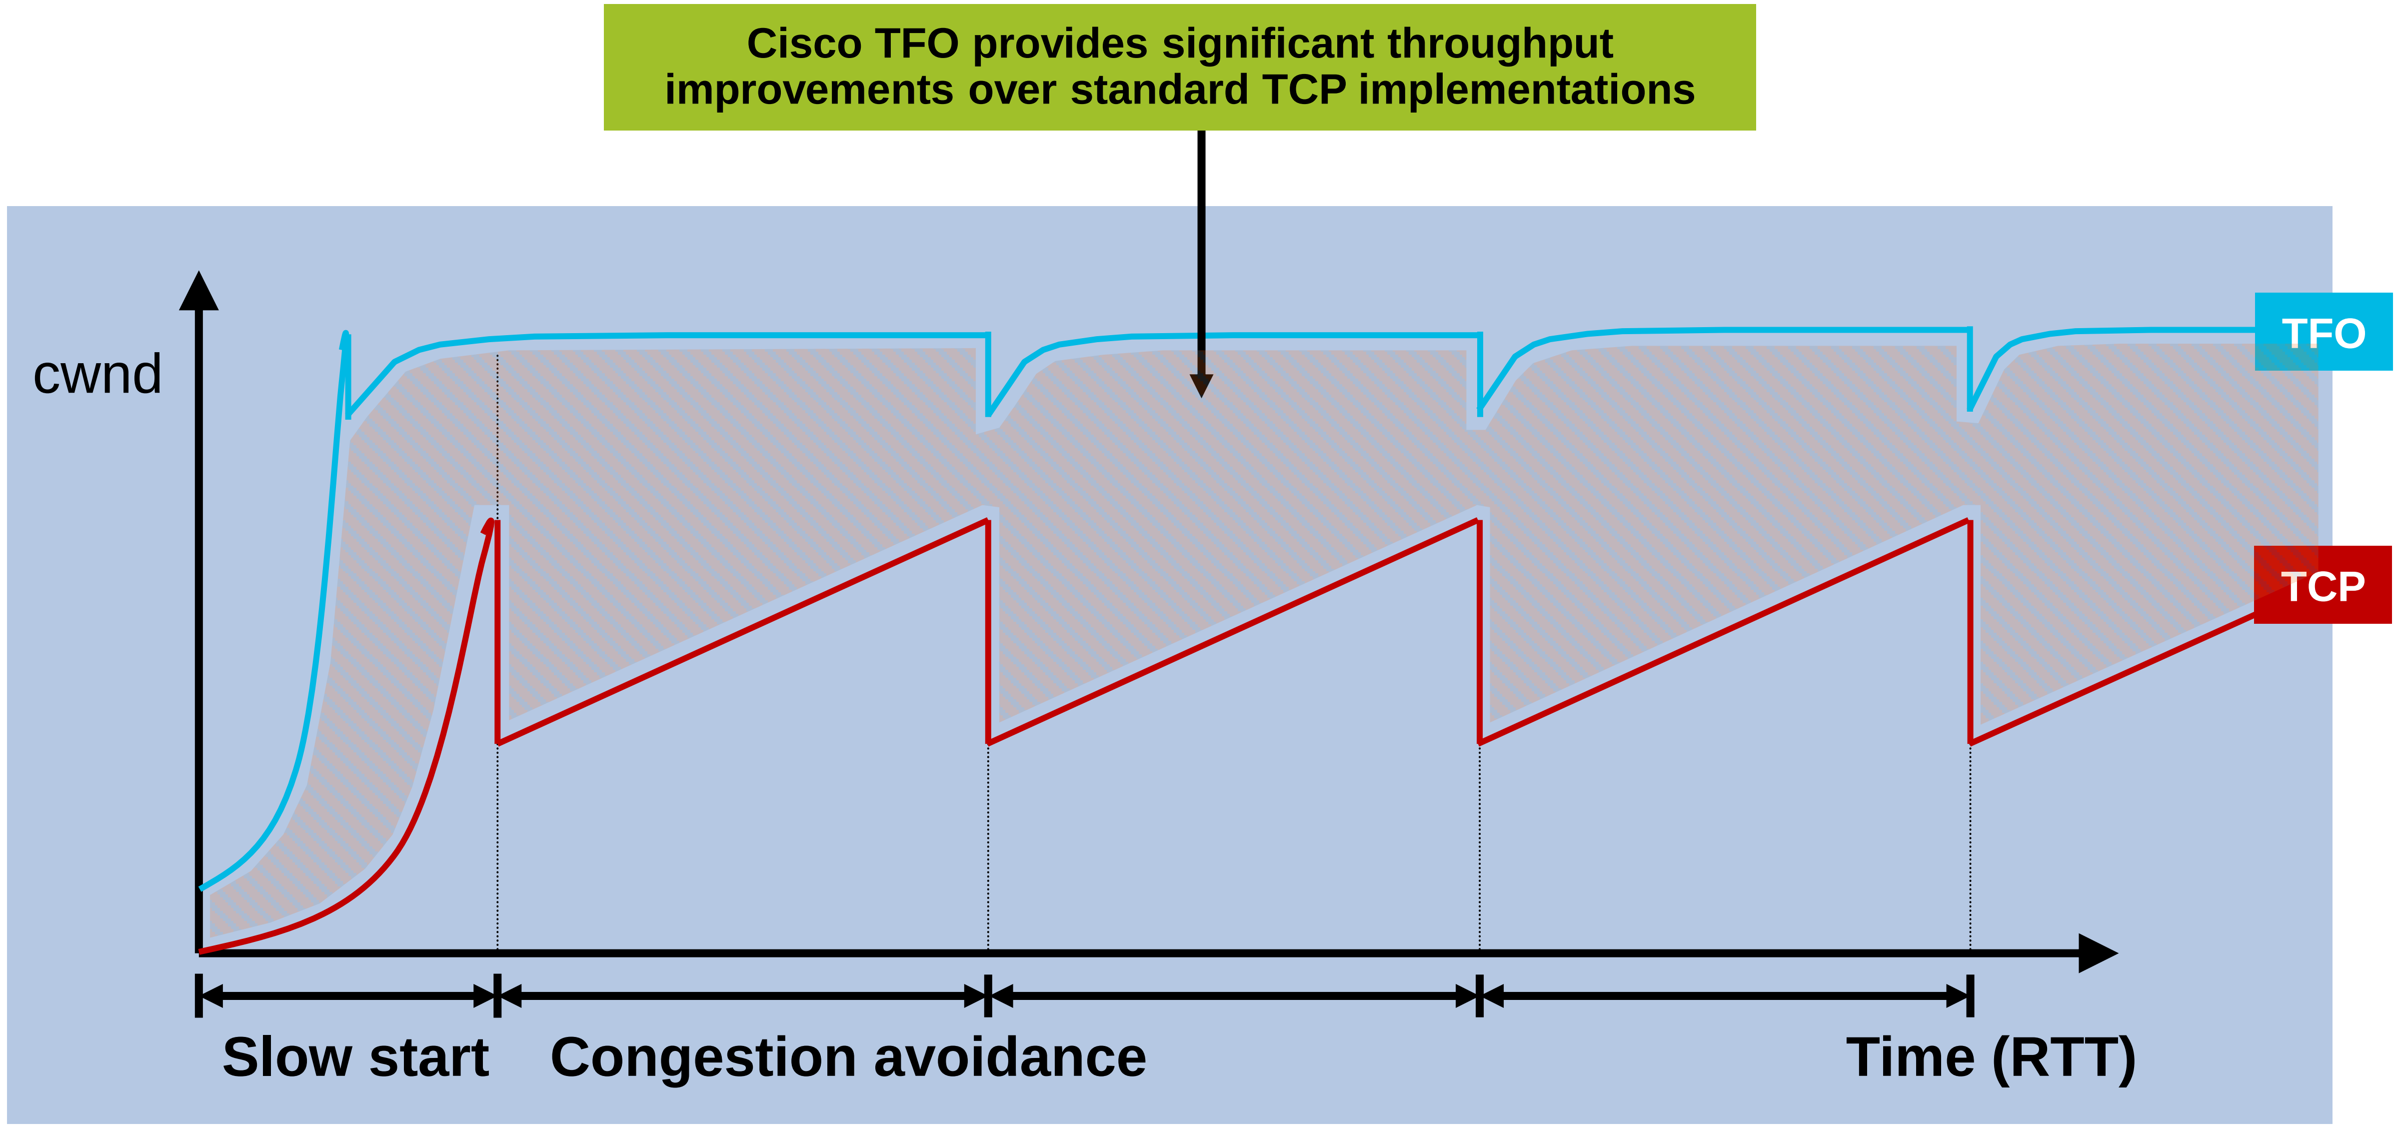

Introducing Cisco WAAS Express

Extend Cisco WAAS product portfolio across ISR G2s

Available Now!
15.1(2)T2

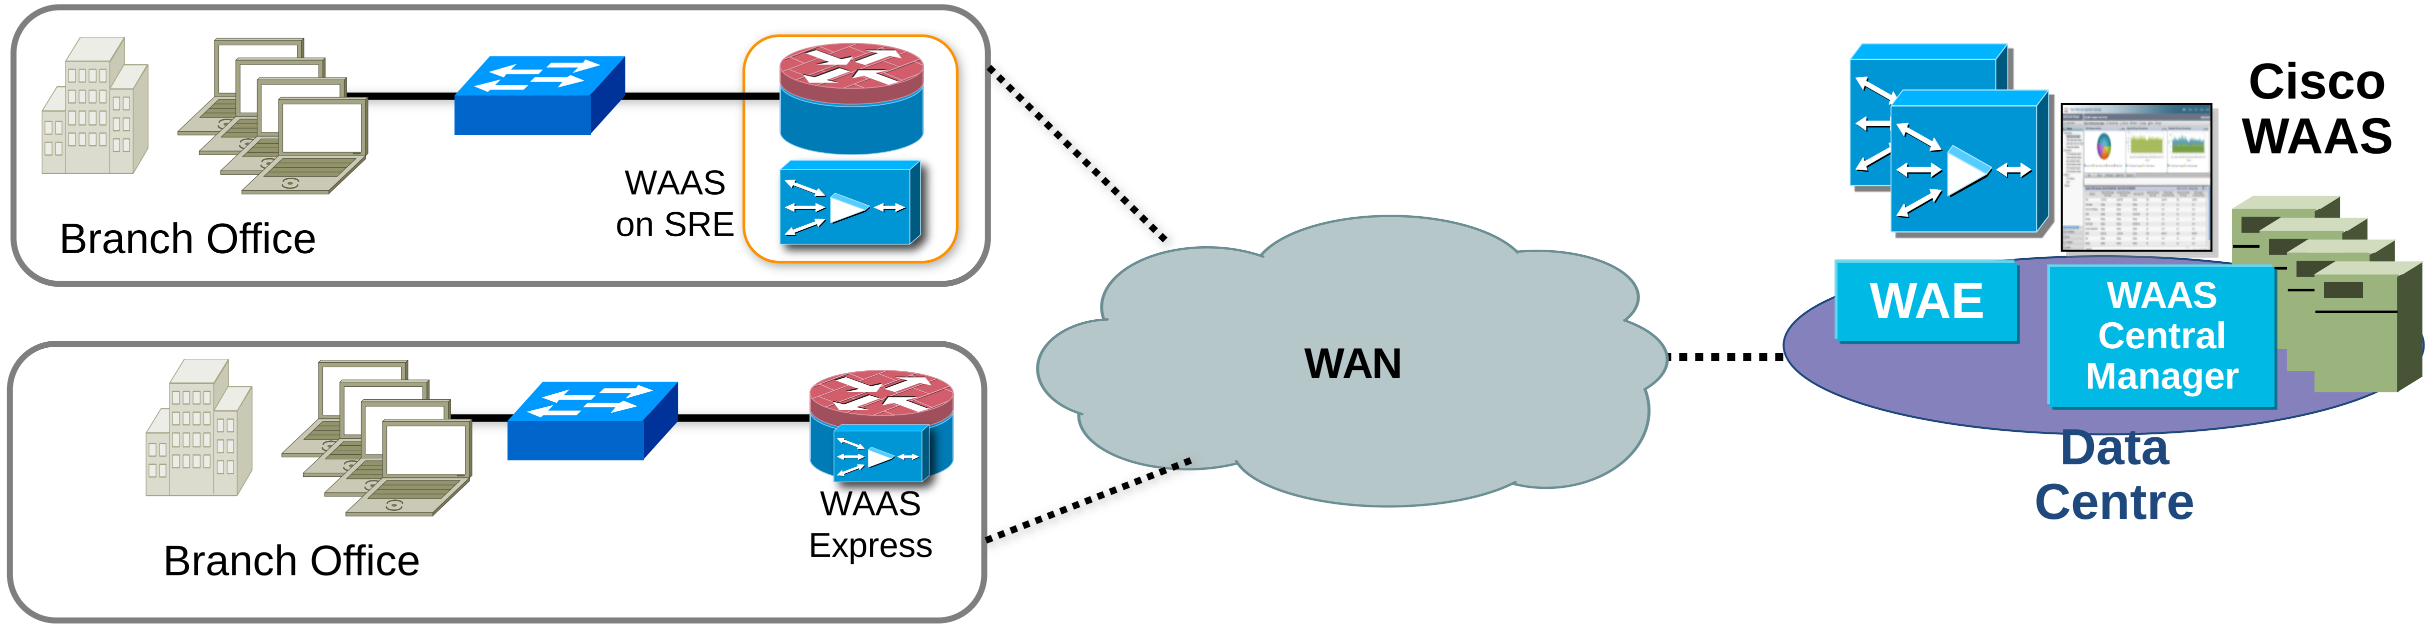

Simple	Cost Effective	Investment Protection
▪ **IOS Based**, Router Integrated WAN Optimisation Solution ▪ Simple **software feature activation** ▪ Network transparency and **integration with IOS** based services	▪ Defer **costly WAN Bandwidth** upgrades ▪ **Reduce truck roll costs** – IOS integrated solution ▪ Capex savings – **Small branch footprint**	▪ Part of Cisco WAAS portfolio – **Leverage existing WAAS deployment** ▪ Easy **migration to WAAS on SRE** as business needs grow ▪ Integrated policy **provisioning, monitoring, and, reporting**

WAAS/WAAS Express Feature Comparison

Features	WAAS Express	Cisco WAAS hardware (version 4.2.1)
Auto-discovery of end nodes	Supported	Supported
TFO (Transport Optimisation)	Supported	Supported
Compression	Supported	Supported
DRE (Data Redundancy Elimination)	- Memory based. - Non-persistent cache	- Disk based. -Persistent cache.
BIC-TCP	Supported	Supported
WAAS Central Manager	Cisco WAAS Version 4.3.1+	Supported
Application Optimisers	Supported*** 15.2(3)T	Supported
Caching	Not Supported	Supported

*** IOS 15.2(3)T Apr 2012, HTTP, SSL/HTTPS, CIFS Application Optimisation

Integrated Branch-WAN Services

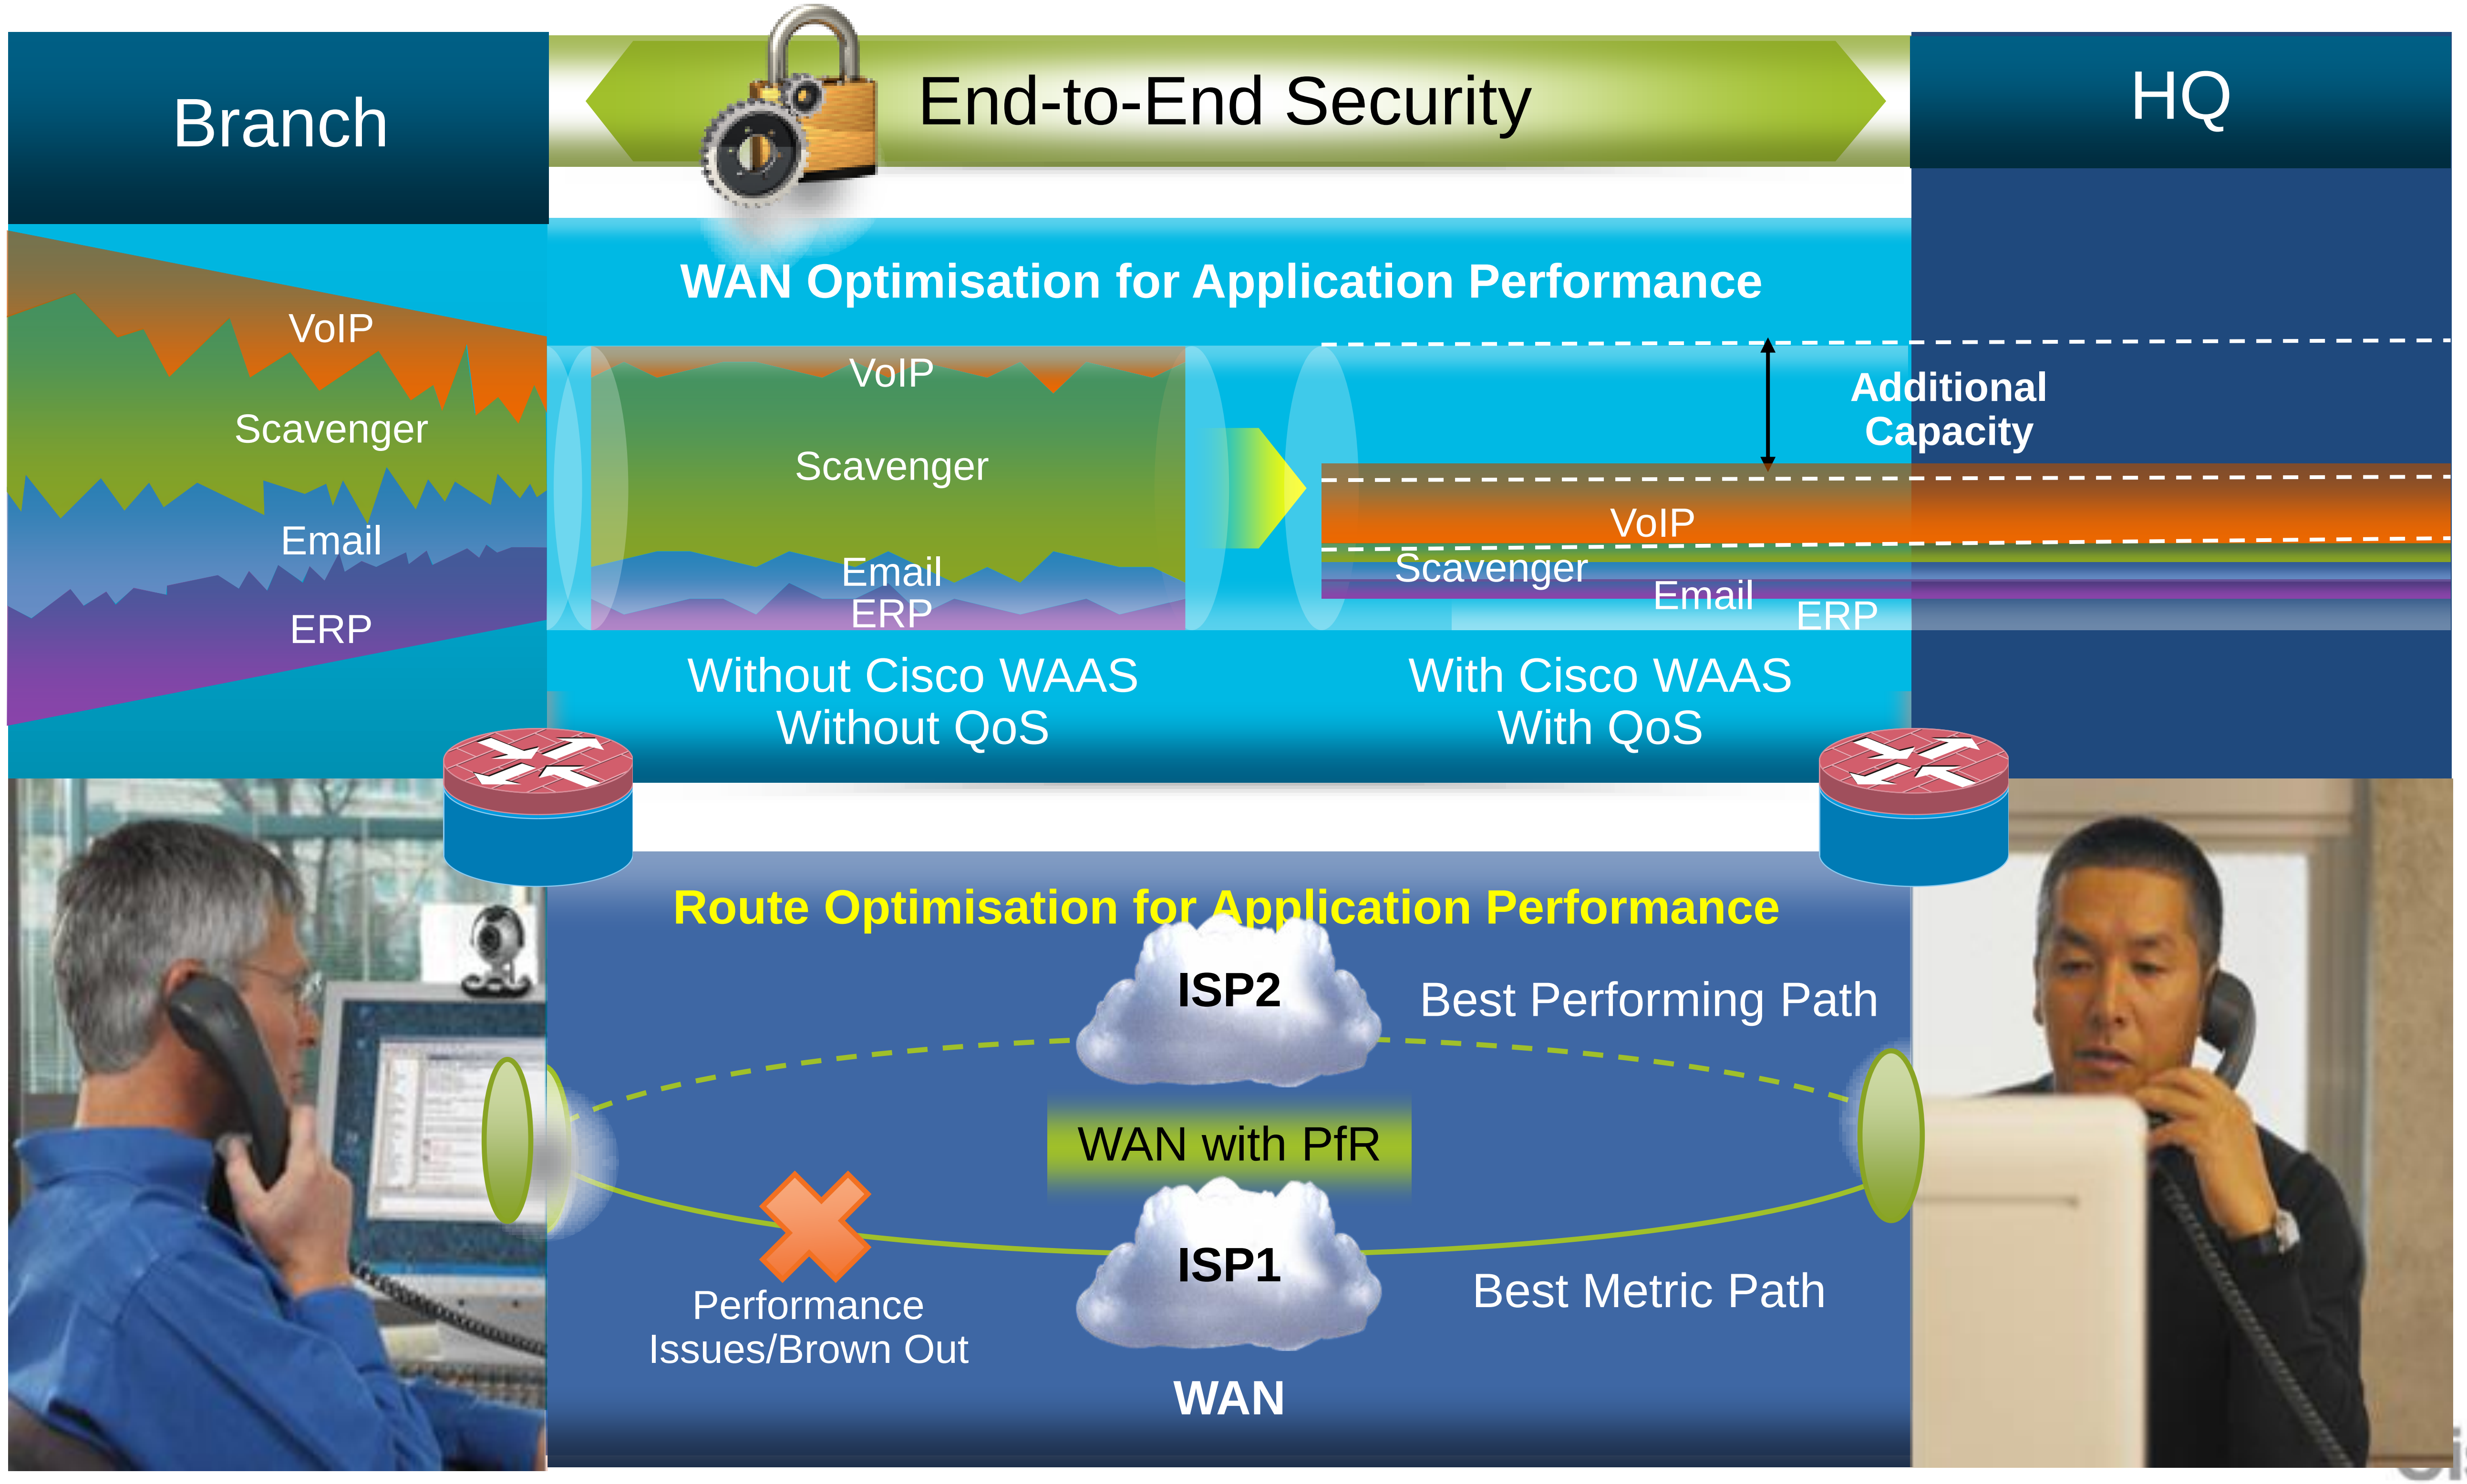

Agenda

- WAN Technologies & Solutions
  - WAN Transport Technologies
  - WAN Overlay Technologies
  - WAN Optimisation
  - Wide Area Network Quality of Service
- WAN Architecture Design Considerations
  - WAN Design and Best Practices
  - Secure WAN Communication with GETVPN
  - DMVPN Over Internet Deployment
- Summary

Quality of Service Operations

How Does It Work and Essential Elements

Classification and Marking
IDENTIFY & PRIORITIZE

Queuing and Dropping
MANAGE & SORT

Post-Queuing Operations
PROCESS & SEND

- **Classification and Marking:**
  - The first element to a QoS policy is to classify/identify the traffic that is to be treated differently. Following classification, marking tools can set an attribute of a frame or packet to a specific value.
- **Policing:**
  - Determine whether packets are conforming to administratively-defined traffic rates and take action accordingly. Such action could include marking, remarking or dropping a packet.
- **Scheduling (including Queuing and Dropping):**
  - Scheduling tools determine how a frame/packet exits a device. Queuing algorithms are activated only when a device is experiencing congestion and are deactivated when the congestion clears.

Enabling QoS in the WAN

Traffic Profiles and Requirements

Voice

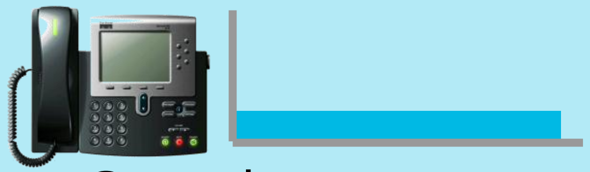

- Smooth
- Benign
- Drop sensitive
- Delay sensitive
- UDP priority

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

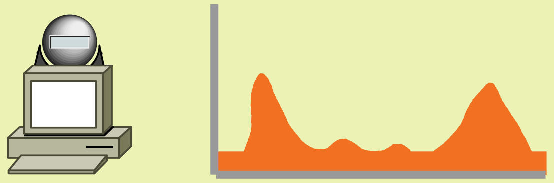

- Bursty
- Greedy
- Drop sensitive
- Delay sensitive
- UDP priority

SD/VC has the Same Requirements as VoIP, but Has Radically Different Traffic Patterns (BW Varies Greatly)

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

One-Way Requirements

Telepresence

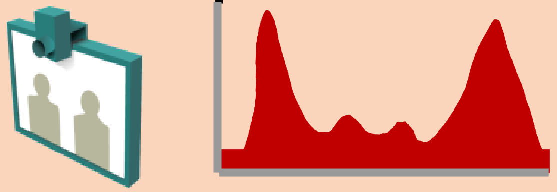

- Bursty
- Drop sensitive
- Delay sensitive
- Jitter sensitive
- UDP priority

HD/VC has Tighter Requirements than VoIP in terms of 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

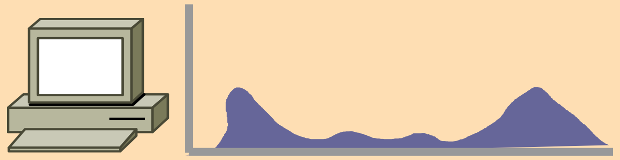

- Smooth/bursty
- Benign/greedy
- Drop insensitive
- Delay insensitive
- TCP retransmits

Traffic patterns for Data Vary Among Applications

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

Scheduling Tools

LLQ/CBWFQ Subsystems

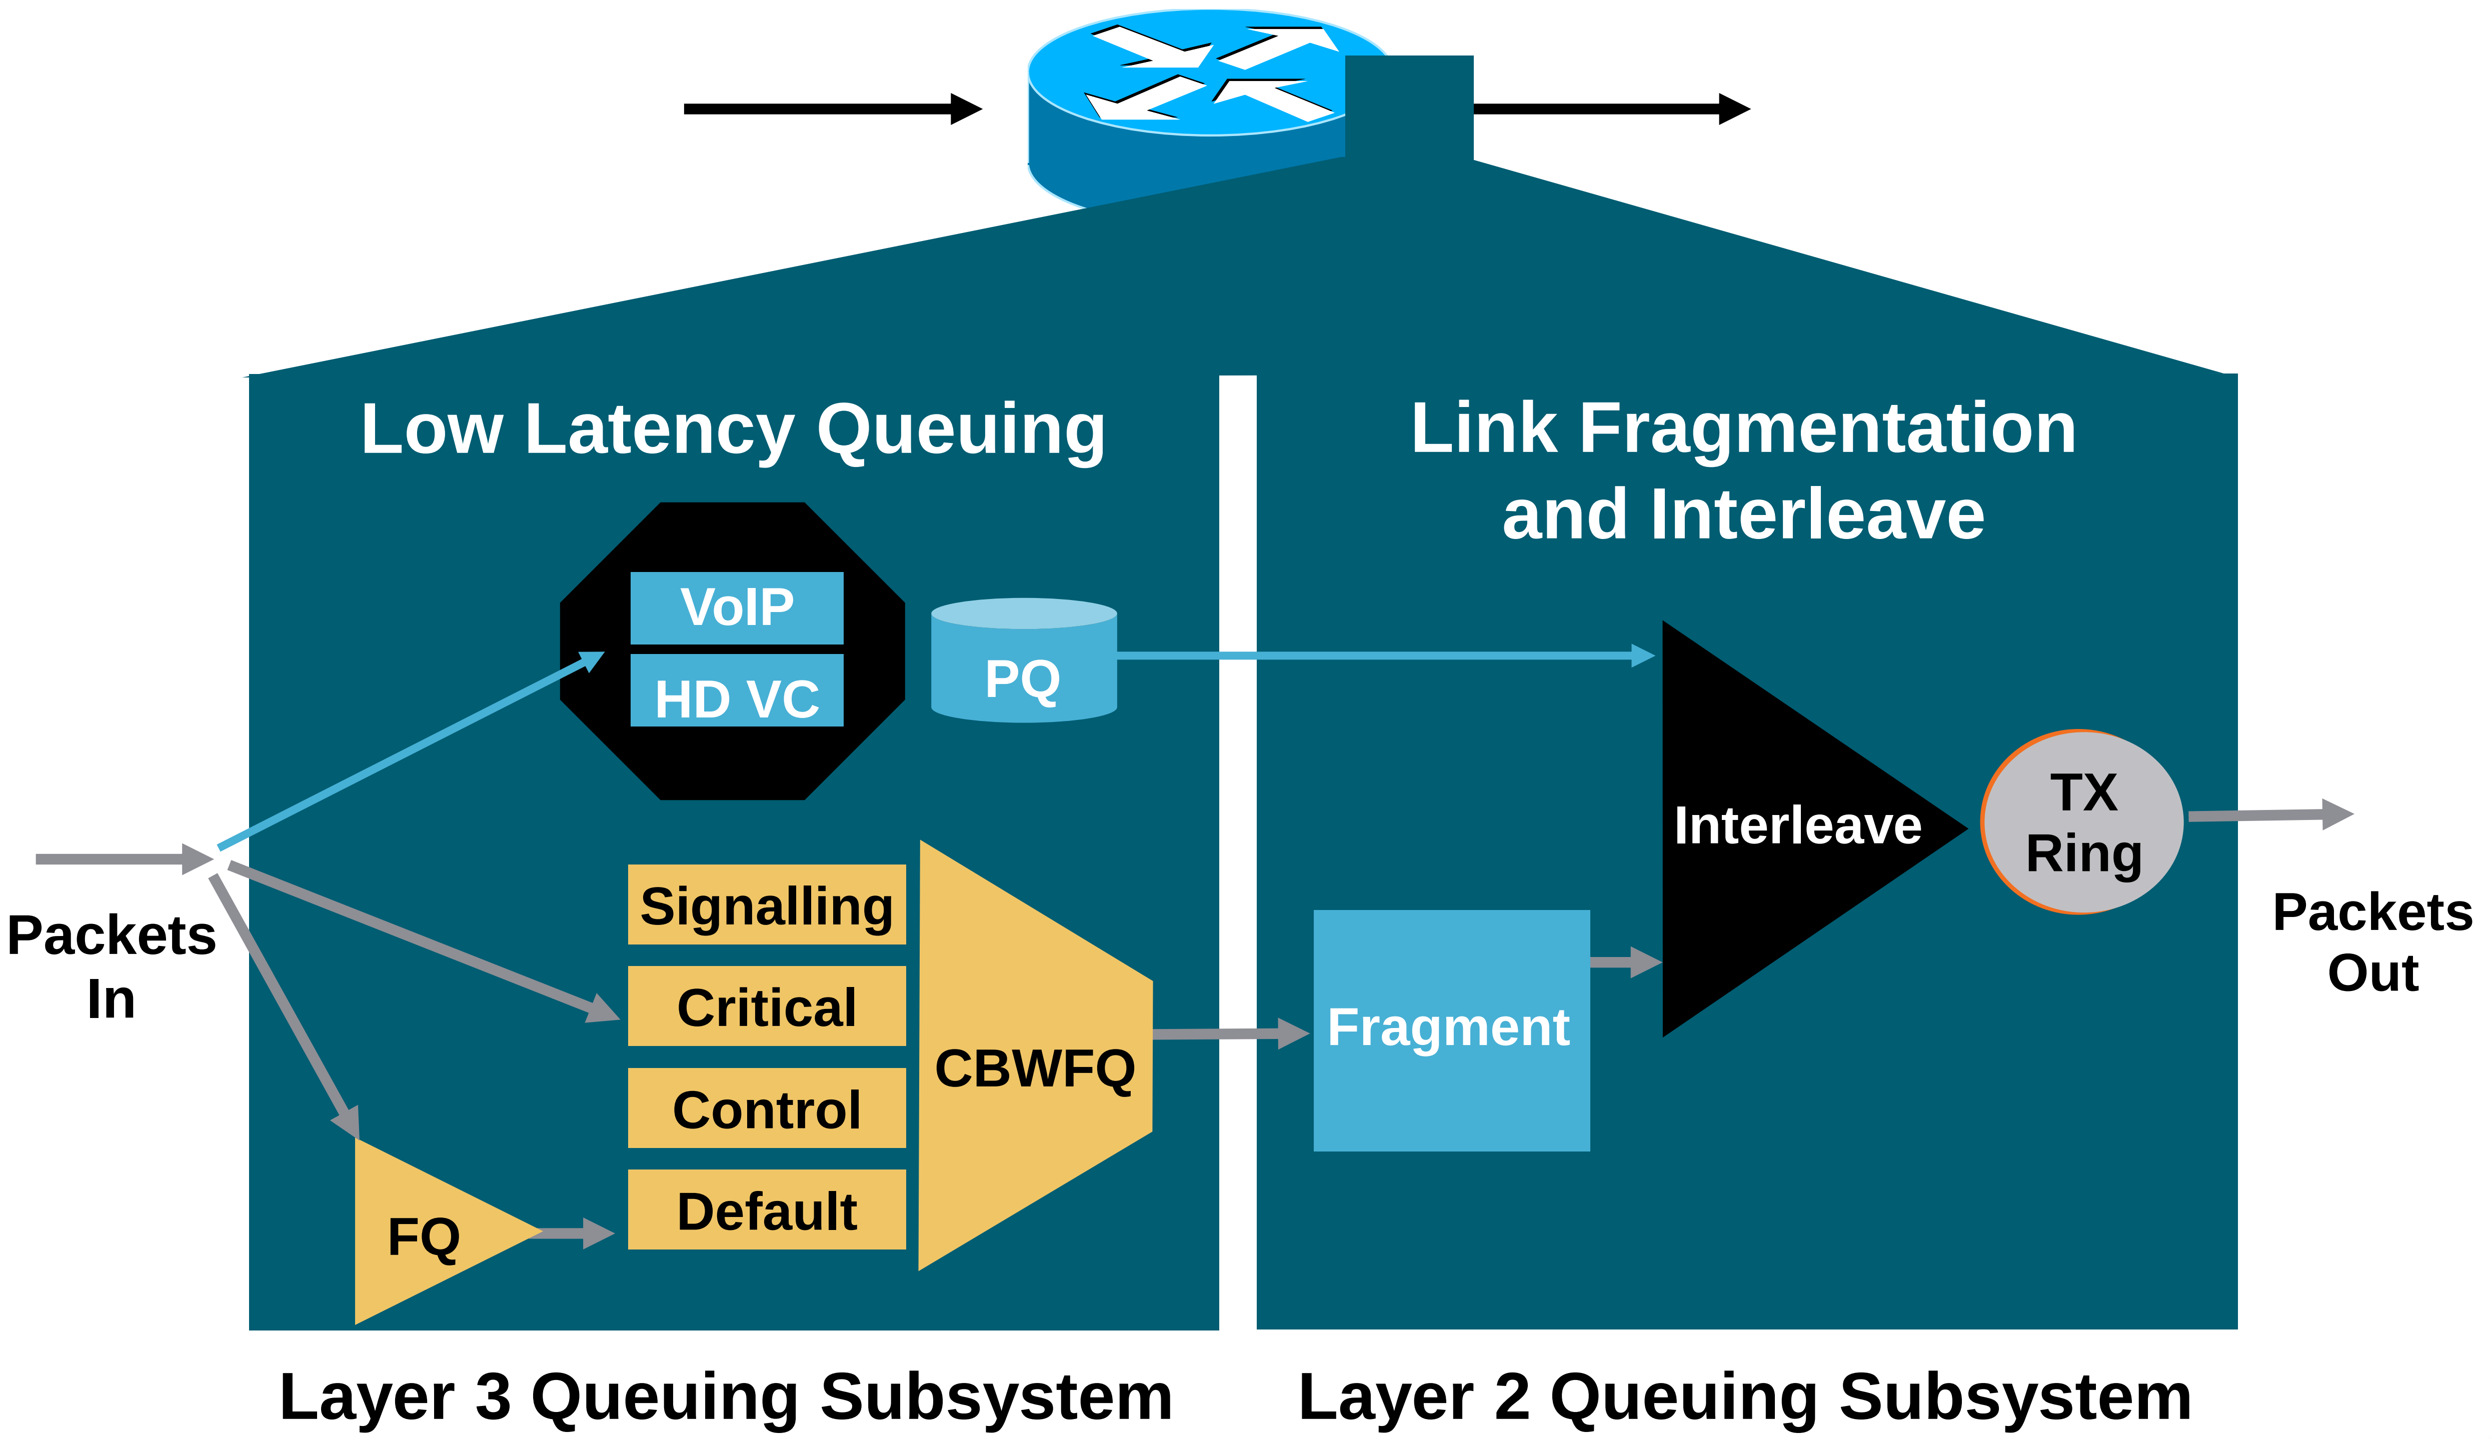

Traffic Shaping

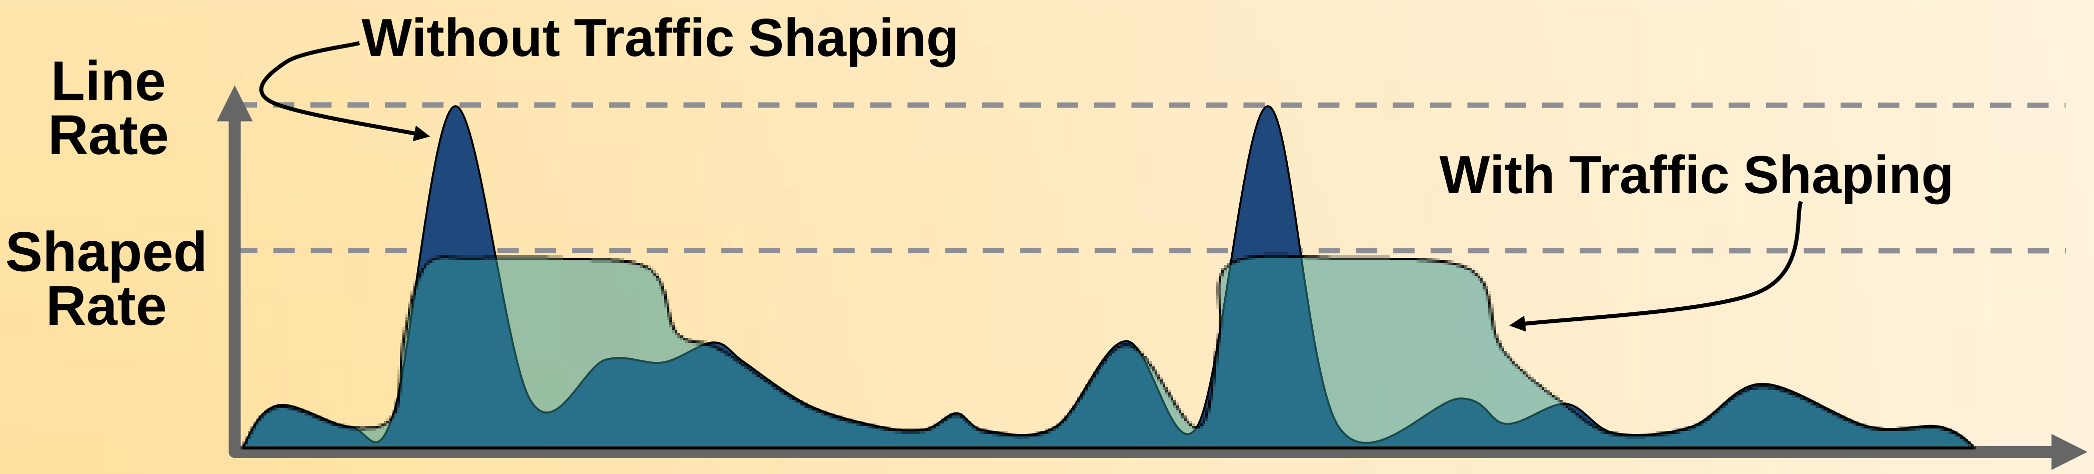

Traffic Shaping Limits the Transmit Rate to a Value Lower Than Line Rate

- Policers typically drop traffic
- Shapers typically delay excess traffic, smoothing bursts and preventing unnecessary drops
- Very common with Ethernet WAN, as well as Non-Broadcast Multiple-Access (NBMA) network topologies such as Frame-Relay and ATM

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

Policy-map CHILD
   class Voice
     police cir percent 10
     priority level 1
   class Video
     police cir percent 20
     priority level 2
   class Control
     bandwidth remaining ration 1
   class class-default
    bandwidth remaining ratio 9

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

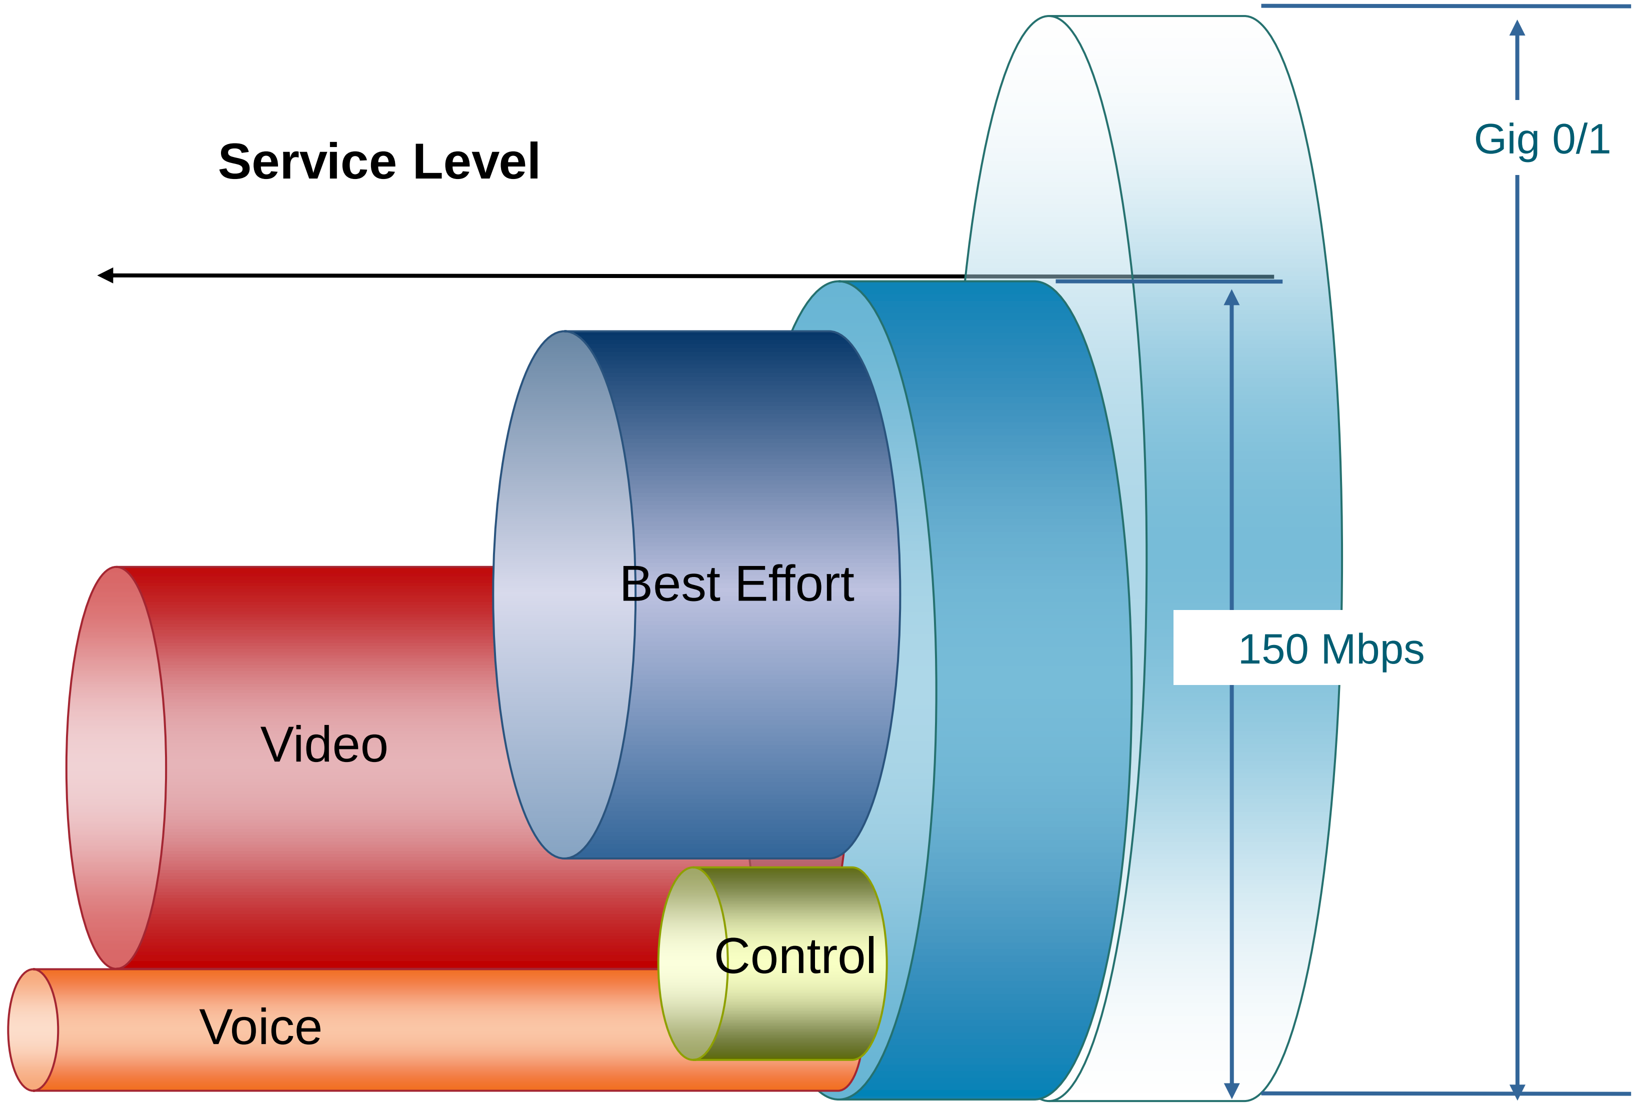

MPLS VPN QoS Considerations

MPLS VPN Port QoS Roles

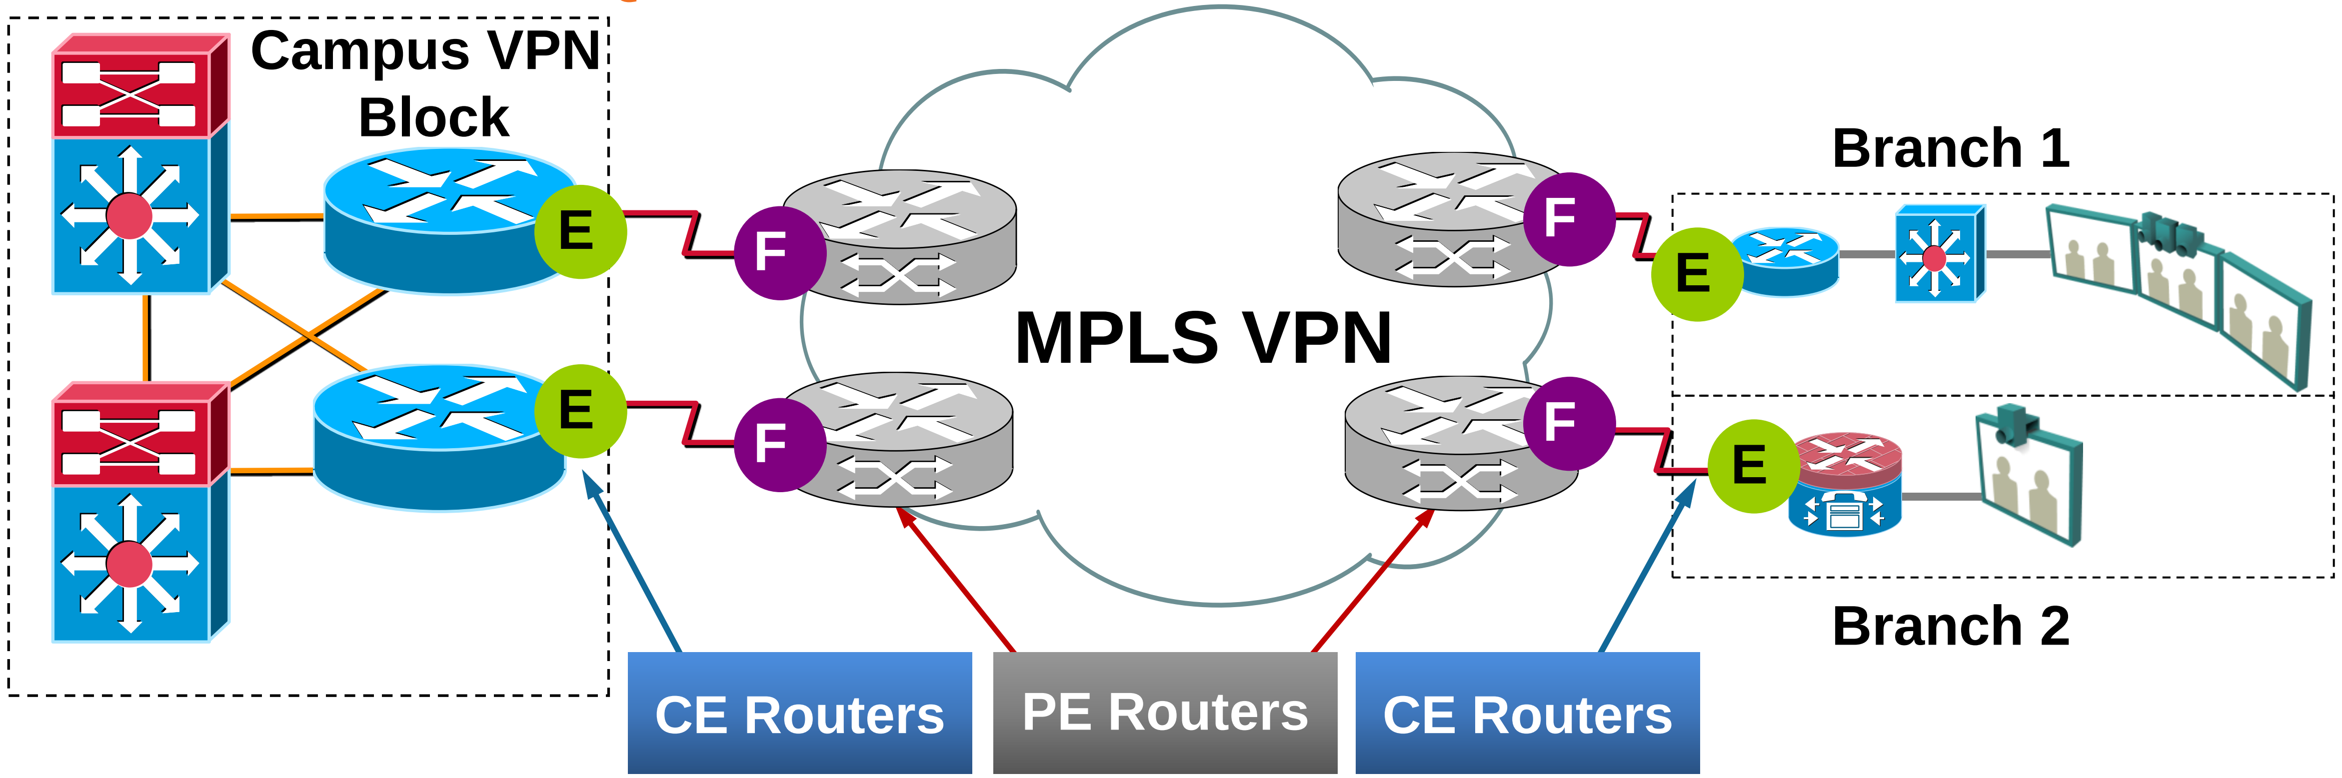

Enterprise Subscriber (Unmanaged CE Routers)

E Outbound Policies:	Inbound Policies:
HQoS Shaper (if required)	
≤ 33% of BW { + LLQ for VoIP (EF)	Trust DSCP
+ LLQ or CBWFQ for RT-Interactive (CS4) }	
+ Remark RTI (if necessary)	+ Restore RT-Interactive to CS4 (if necessary)
+ CBWFQ for Signalling (CS3)	
+ Remark Signalling (if necessary)	+ Restore Signalling to CS3 (if necessary)

Service Provider:

F Outbound Policies:	Inbound Policies:
+ LLQ for Real-Time	Trust DSCP
+ CBWFQ for Critical Data	Police on a per-Class Basis

GRE/IPSec QoS Consideration

ToS Byte Preservation

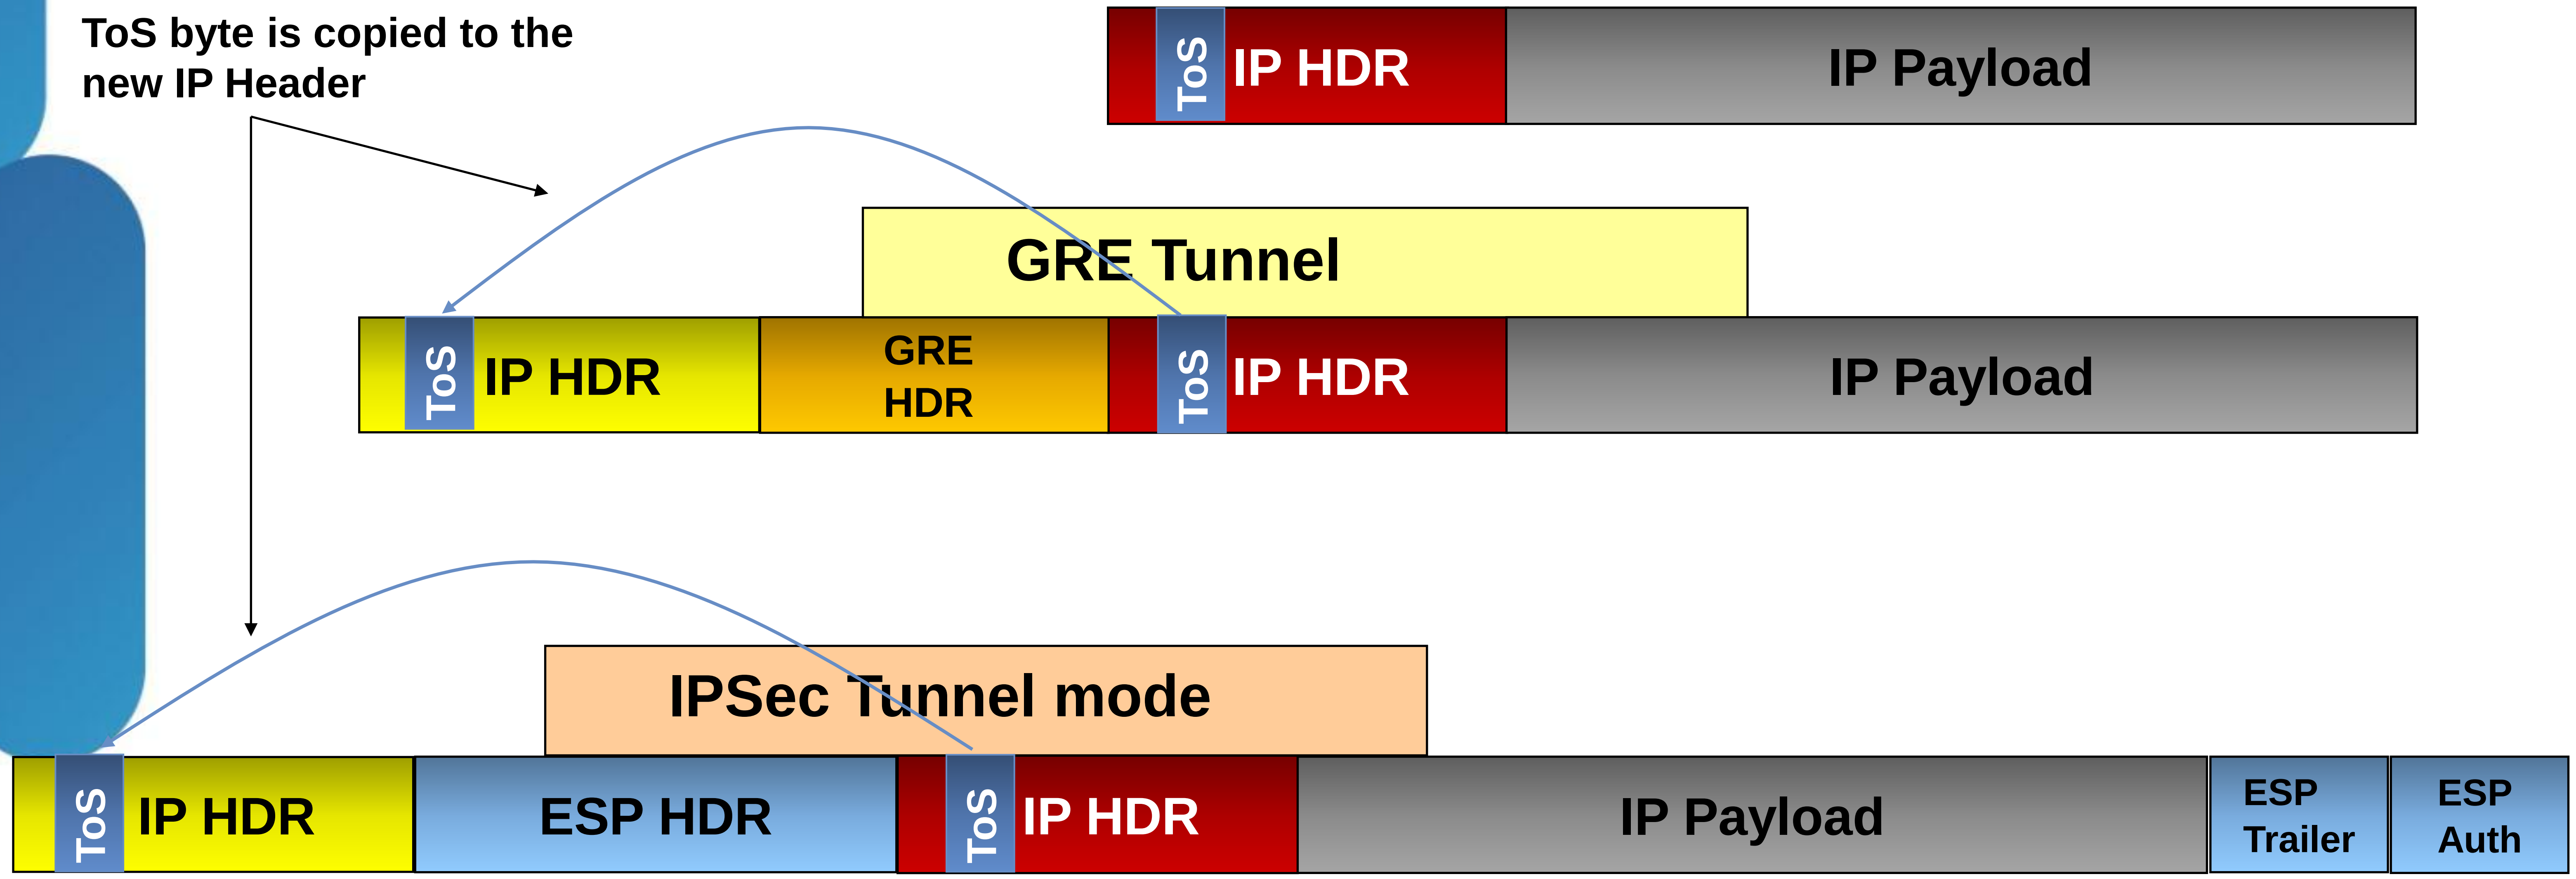

GRE/IPSec Network QoS Design

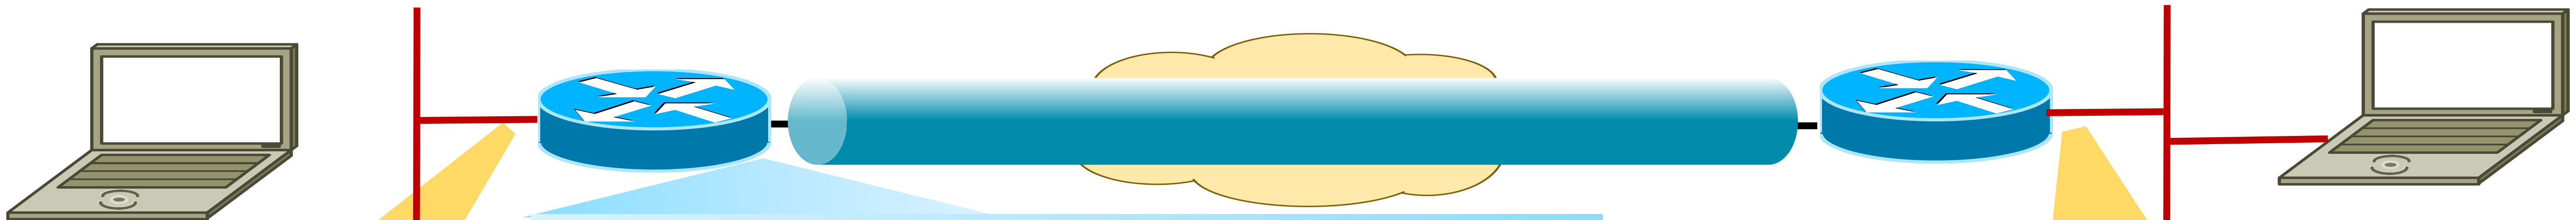

DSCP AF41

Packet Initially Marked to DSCP AF41

DSCP AF41	DSCP CS5
DSCP AF41	DSCP AF41
By Default ToS Values is Copied To IPSec Header	Top-Most ToS is Rewrote on egress

DSCP AF41

Packet decapsulated To reveal the original ToS Byte

Direction of Packet Flow

```
policy-map WAN-SP-CLASS-OUTPUT
 class VOICE
  priority percent 10
 class VIDEO-INTERACTIVE
  priority percent 23
  set ip dscp cs5
 class NETWORK-MGMT
  bandwidth percent 5
  service-policy MARK-BGP
 class class-default
  bandwidth percent 25
  random-detect
!
policy-map Int-Gig-Agg-HE
 class class-default
  shape average 1000000000
  service-policy WAN-Out
```

Remarks the DSCP value on the encrypted/encapsulated header on egress interface

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.
- 7200/ISR G1/G2 – 12.4(22)T or later
- ASR1000 – IOS XE RLS 3.6

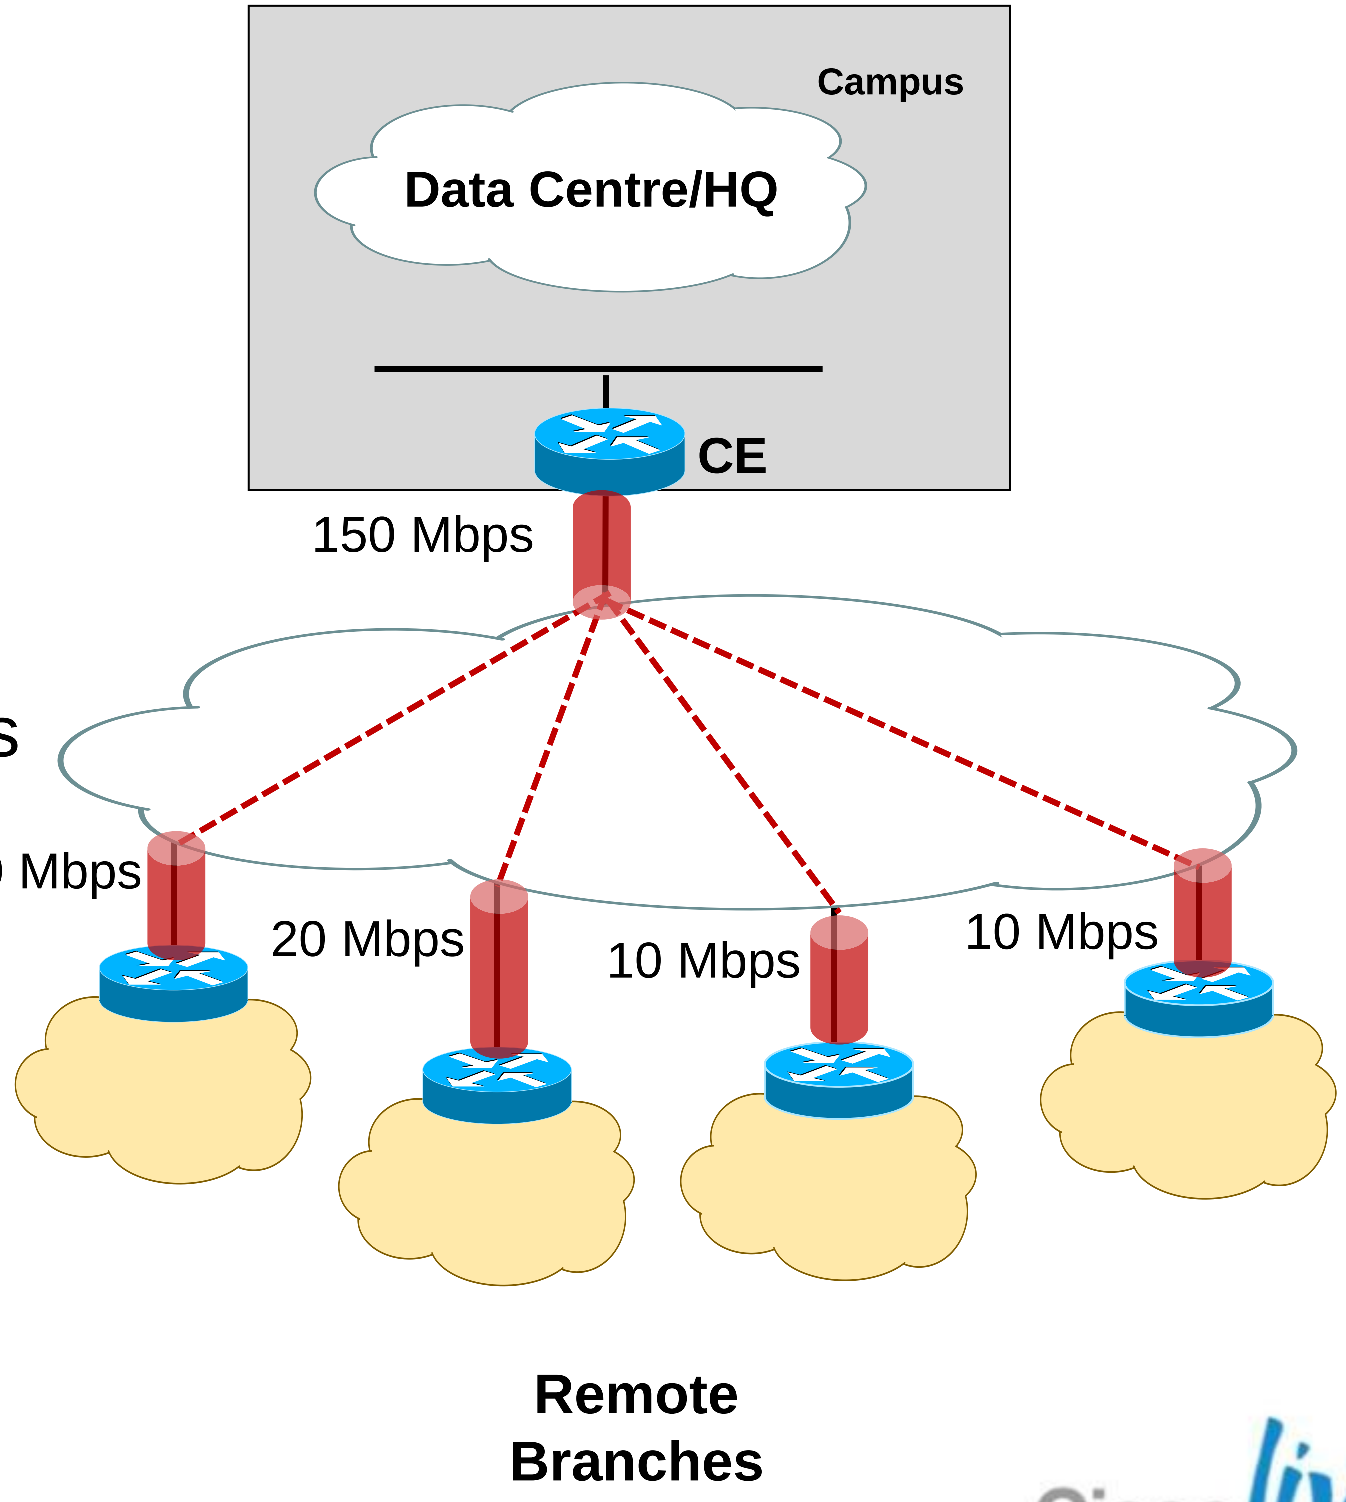

IOS Configuration Reference for Per-Tunnel QoS for DMVPN:
http://www.cisco.com/en/US/docs/ios/sec_secure_connectivity/configuration/guide/sec_per_tunnel_qos.html

Per-tunnel QoS

Configurations

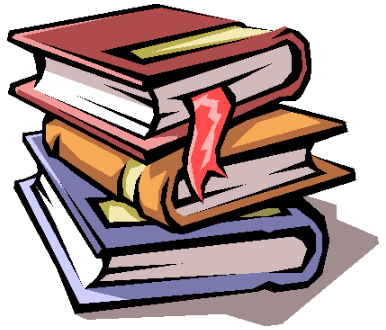

Hub

```
class-map match-all typeA_voice
  match access-group 100
class-map match-all typeB_voice
  match access-group 100
class-map match-all typeA_Routing
  match ip precedence 6
class-map match-all typeB_Routing
  match ip precedence 6

policy-map typeA
  class typeA_voice
    priority 1000
  class typeA_Routing
    bandwidth percent 20

policy-map typeB
  class typeB_voice
    priority percent 20
  class typeB_Routing
    bandwidth percent 10

policy-map typeA_parent
  class class-default
    shape average 3000000
     service-policy typeA

policy-map typeB_parent
  class class-default
    shape average 2000000
     service-policy typeB
```

Hub (cont)

```
interface Tunnel0
  ip address 10.0.0.1 255.255.255.0
  …
  ip nhrp map group typeA service-policy output typeA_parent
  ip nhrp map group typeB service-policy output typeB_parent
  …
  ip nhrp redirect
  no ip split-horizon eigrp 100
  ip summary-address eigrp 100 192.168.0.0 255.255.192.0 5
  …
```

Spoke1

```
interface Tunnel0
  ip address 10.0.0.11 255.255.255.0
  …
  ip nhrp group typeA
  ip nhrp map multicast 172.17.0.1
  ip nhrp map 10.0.0.1 172.17.0.1
  ip nhrp nhs 10.0.0.1
  …
```

Spoke2

```
interface Tunnel0
  ip address 10.0.0.12 255.255.255.0
  …
  ip nhrp group typeB
  ip nhrp map multicast 172.17.0.1
  ip nhrp map 10.0.0.1 172.17.0.1
  ip nhrp nhs 10.0.0.1
  …
```

Spoke3

```
interface Tunnel0
  ip address 10.0.0.13 255.255.255.0
  …
  ip nhrp group typeA
  ip nhrp map multicast 172.17.0.1
  ip nhrp map 10.0.0.1 172.17.0.1
  ip nhrp nhs 10.0.0.1
  …
```

Agenda

- WAN Technologies & Solutions
  - WAN Transport Technologies
  - WAN Overlay Technologies
  - WAN Optimisation
  - Wide Area Network Quality of Service
- WAN Architecture Design Considerations
  - WAN Design and Best Practices
  - Secure WAN Communication with GETVPN
  - DMVPN Over Internet Deployment
- Summary

Smart Business Architecture

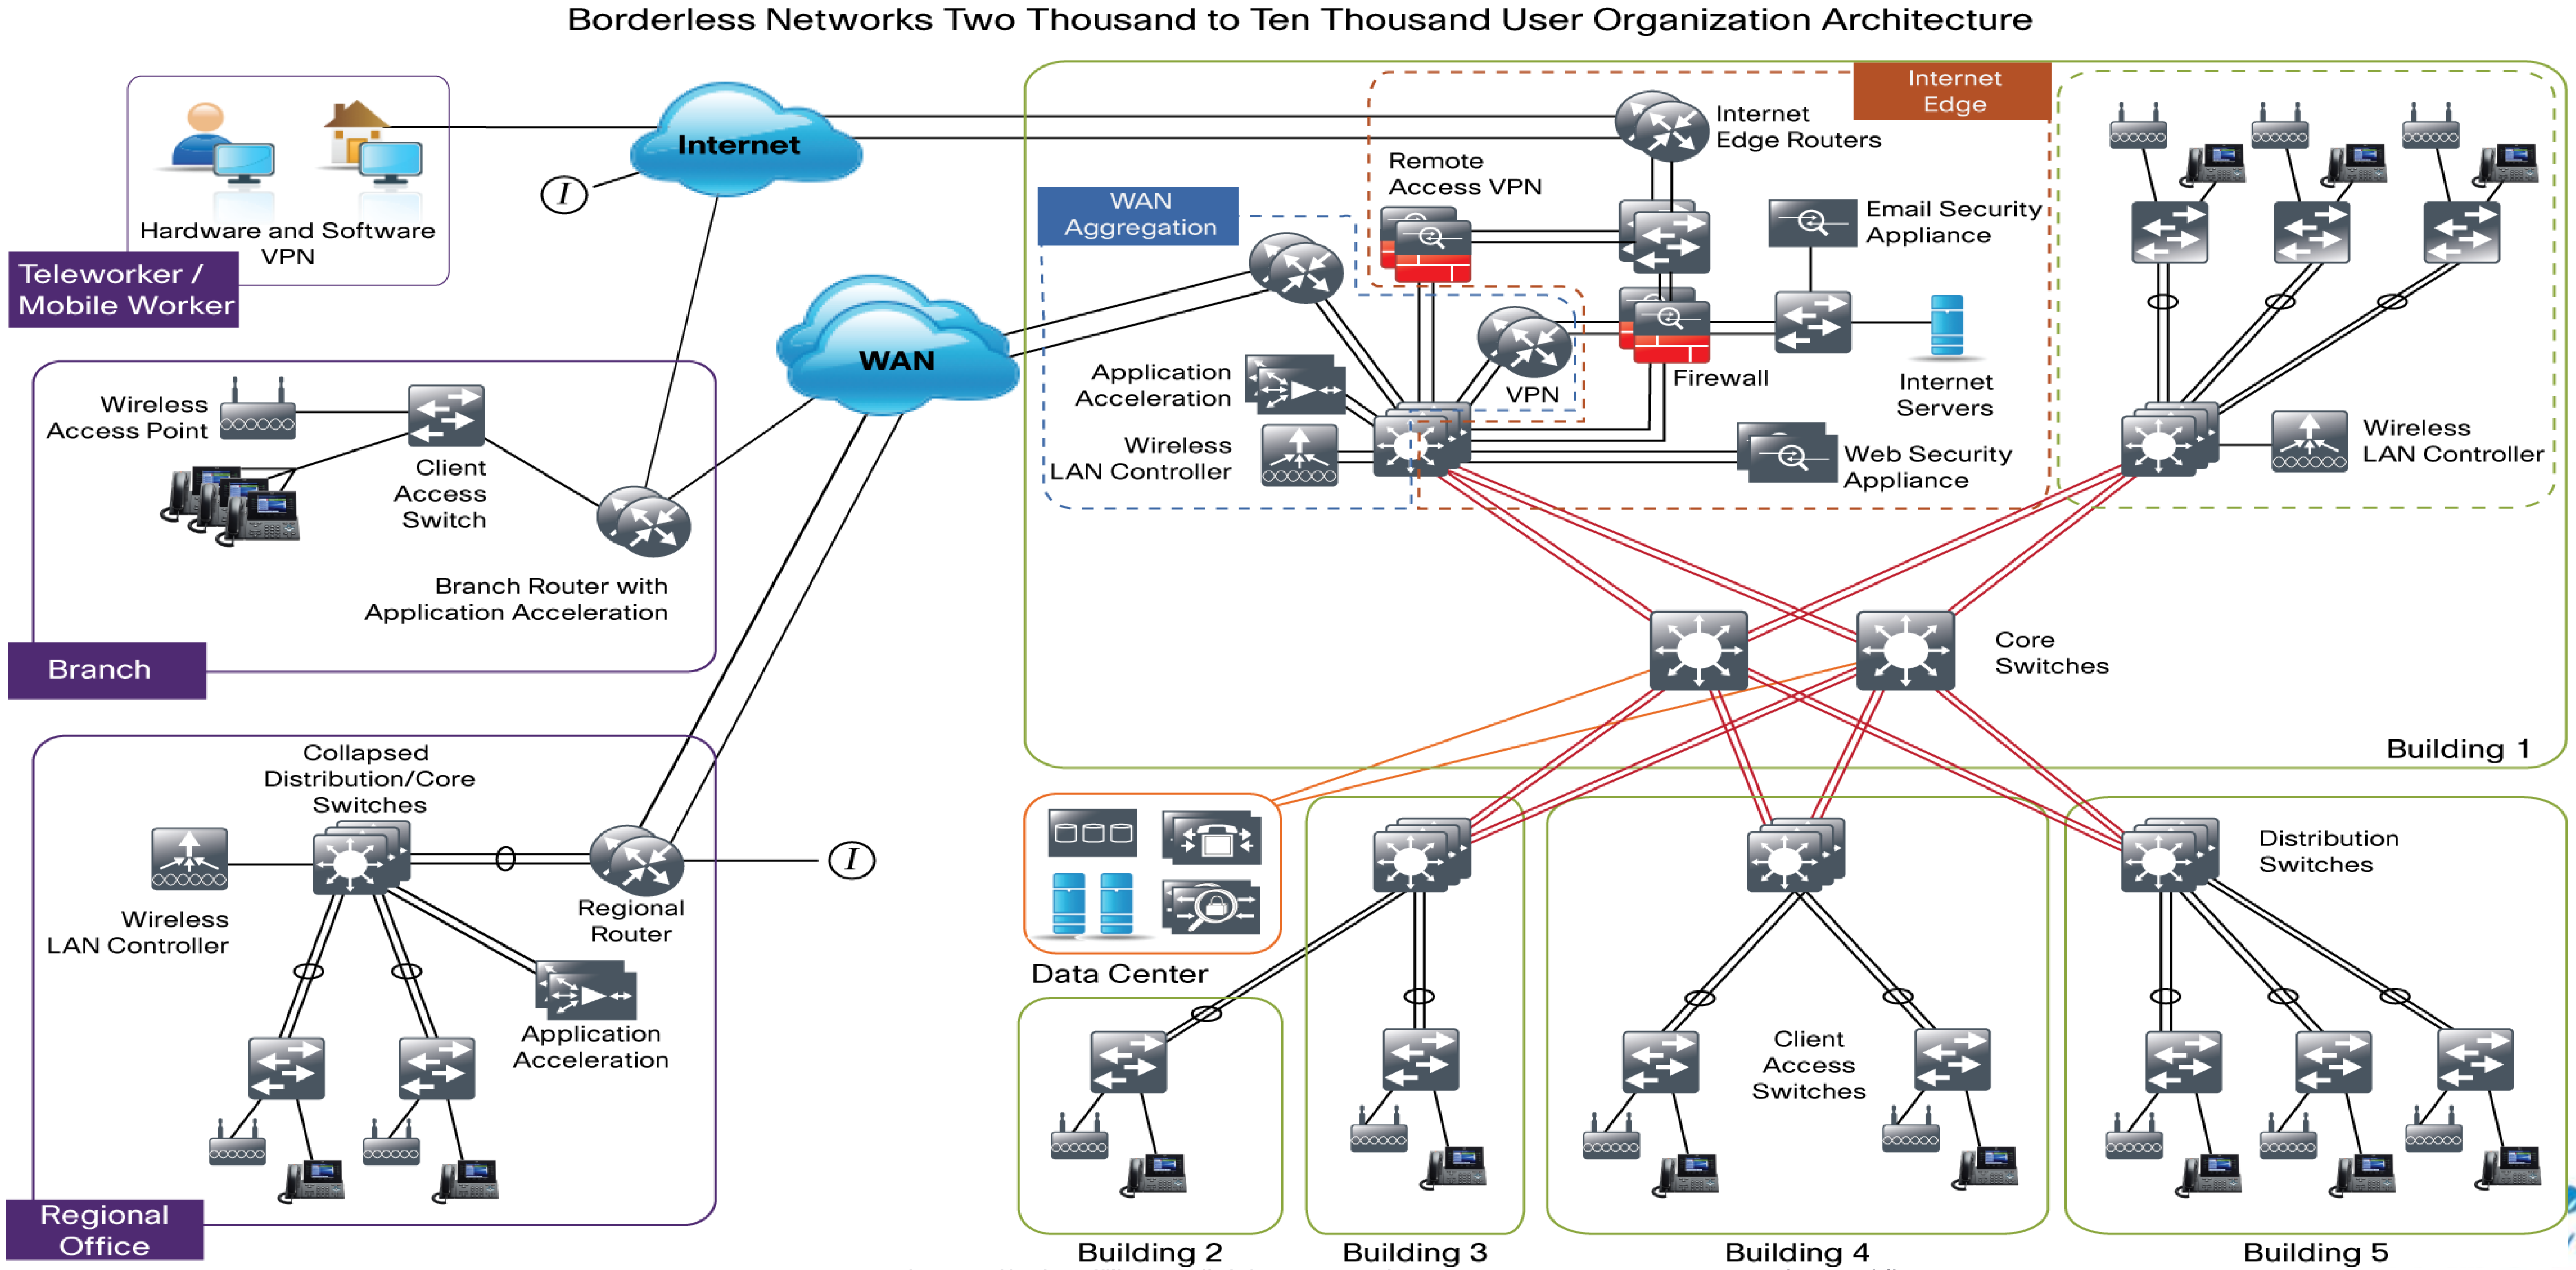

High Performance WAN Headend

Data Centre/ Campus

WAN Services/ Distribution

WAN Edge

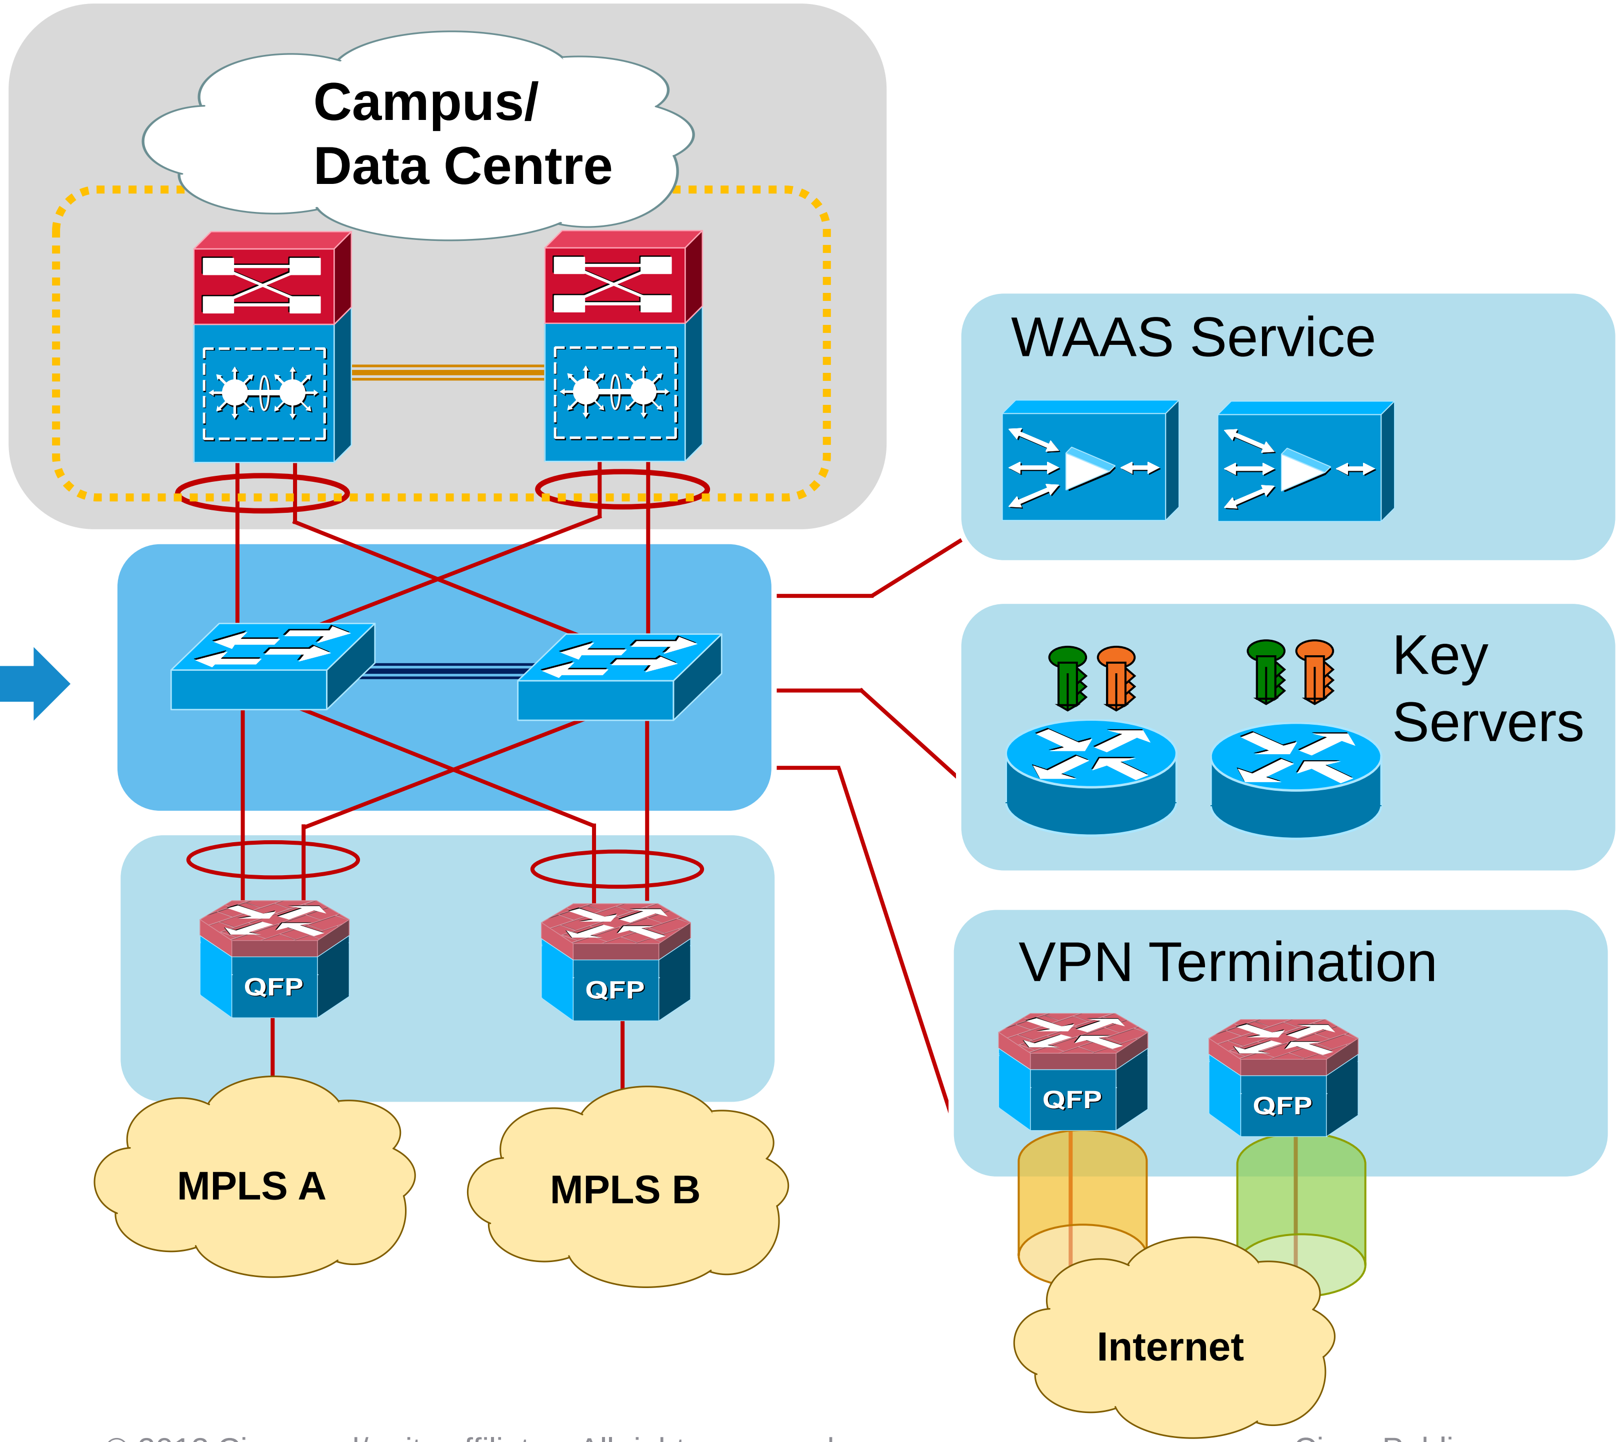

Remote Branch

Transport & Redundancy Options

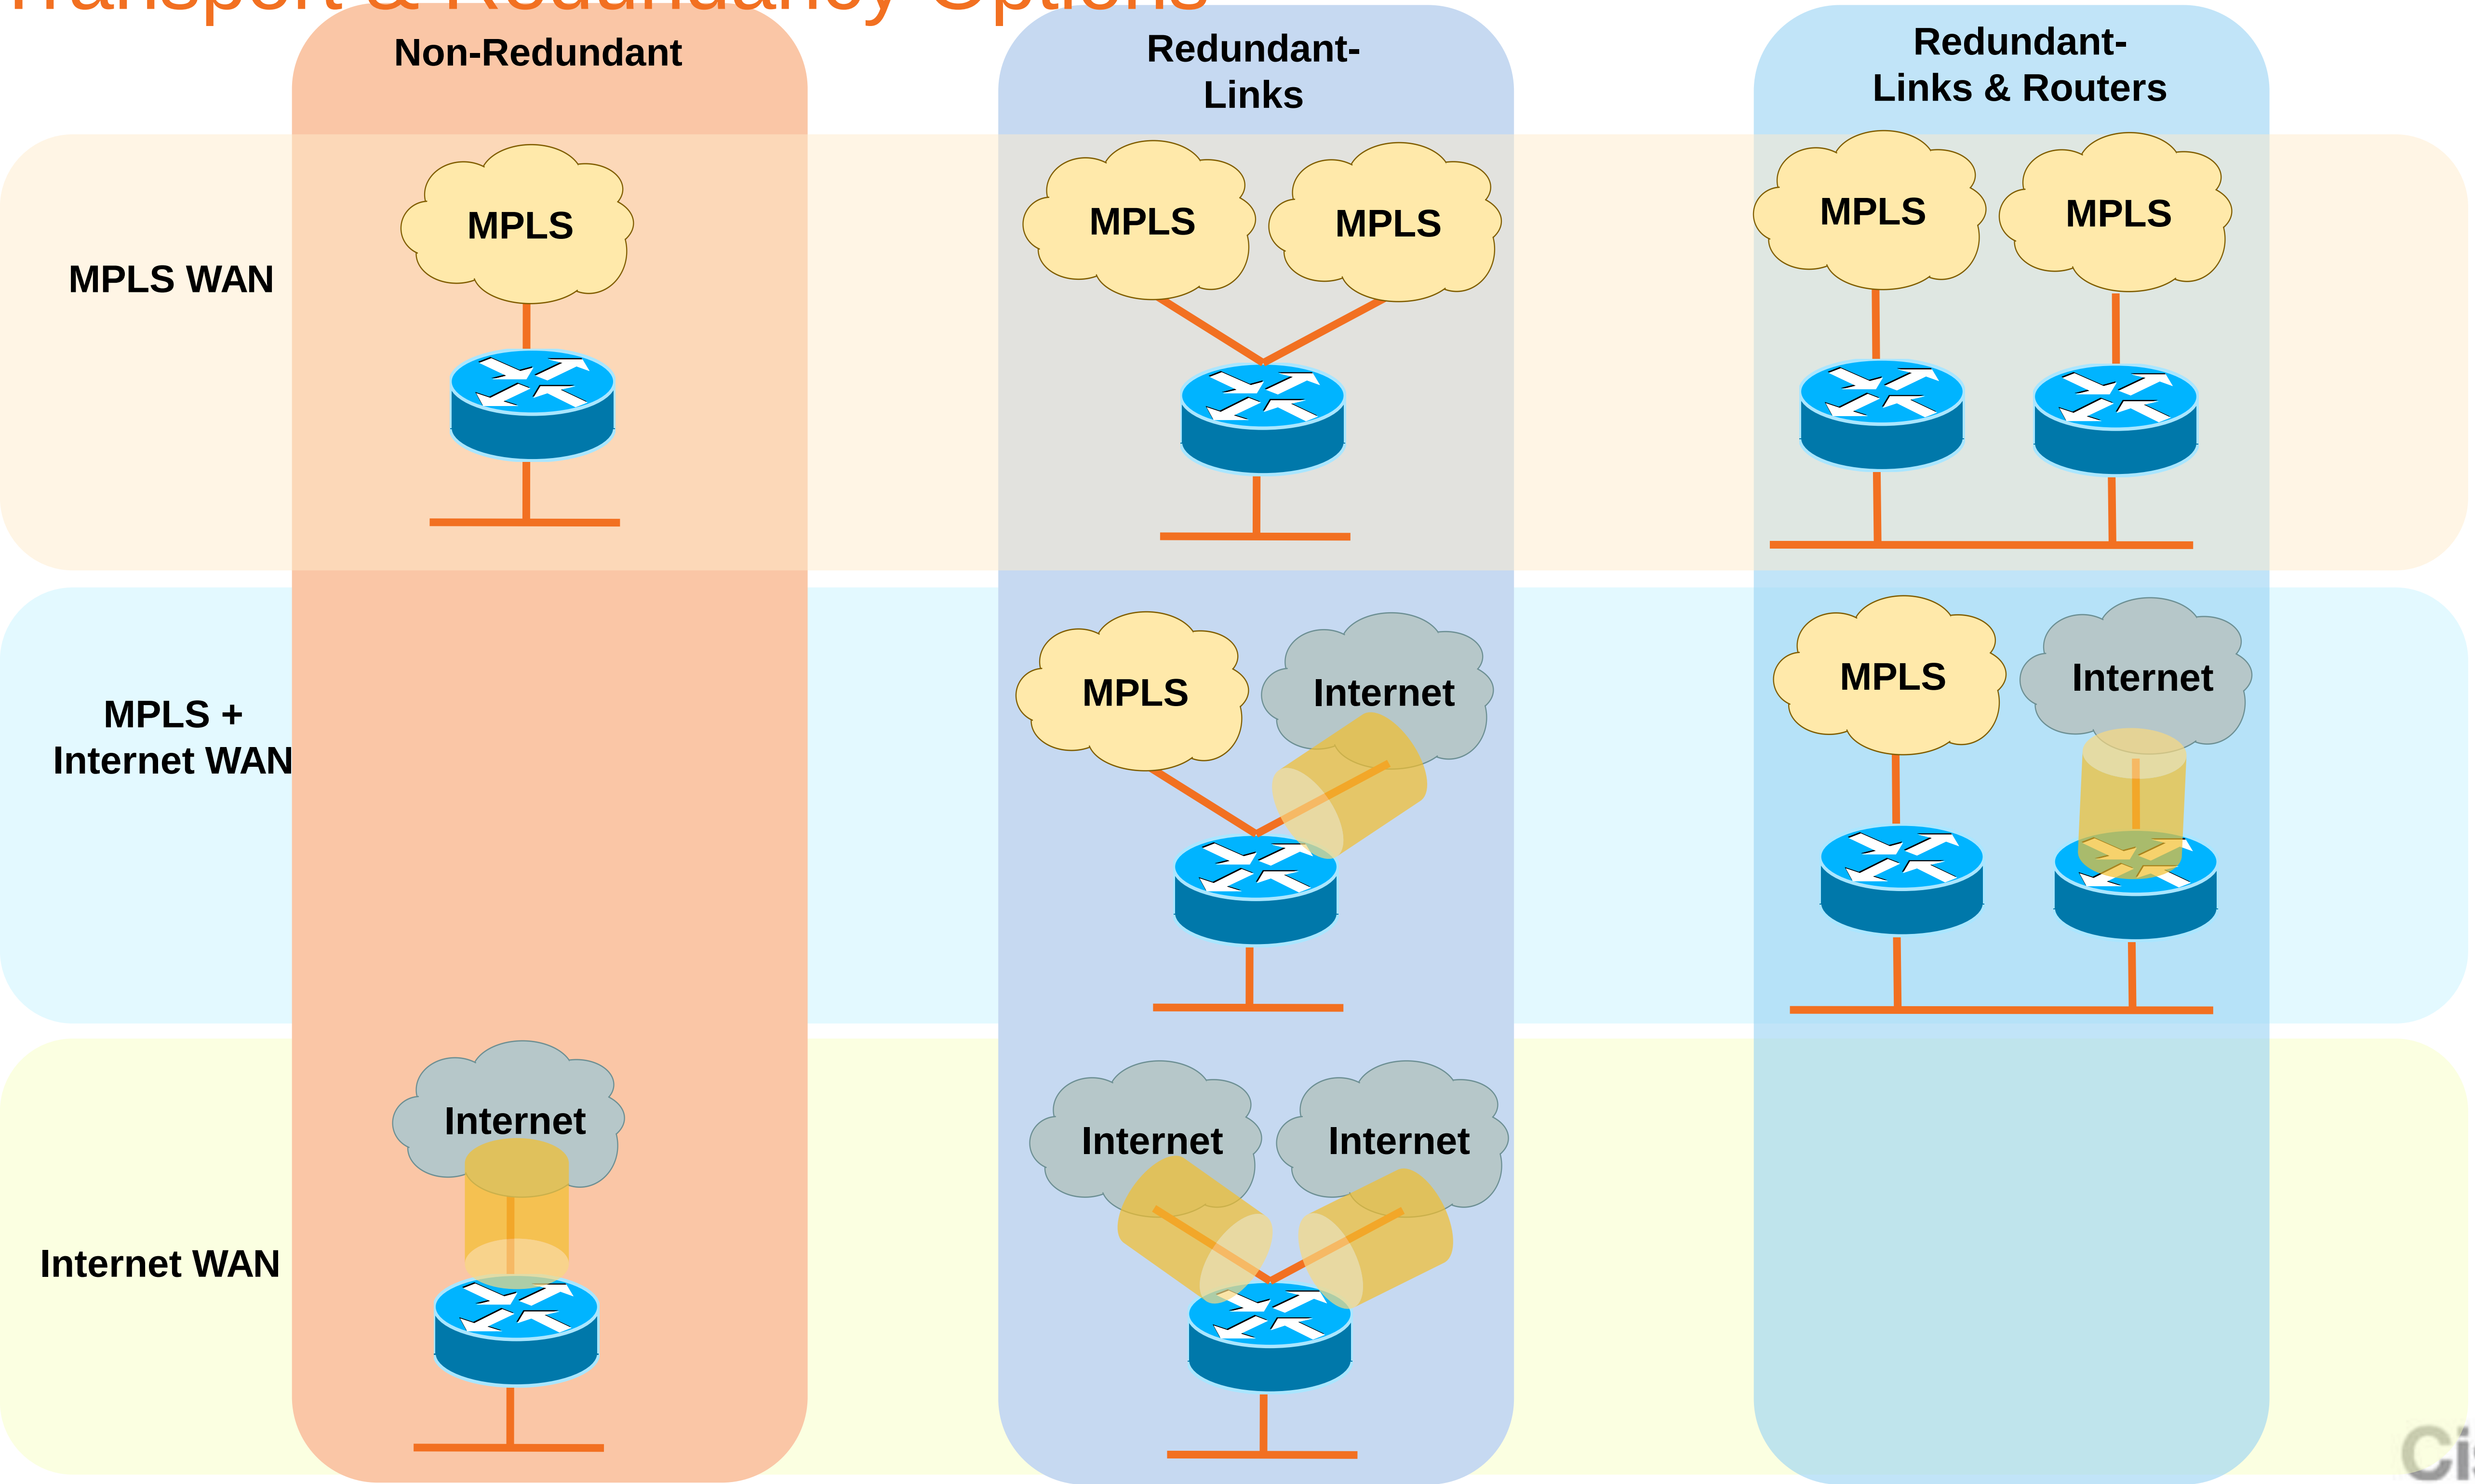

Routing Topology at Hub Location

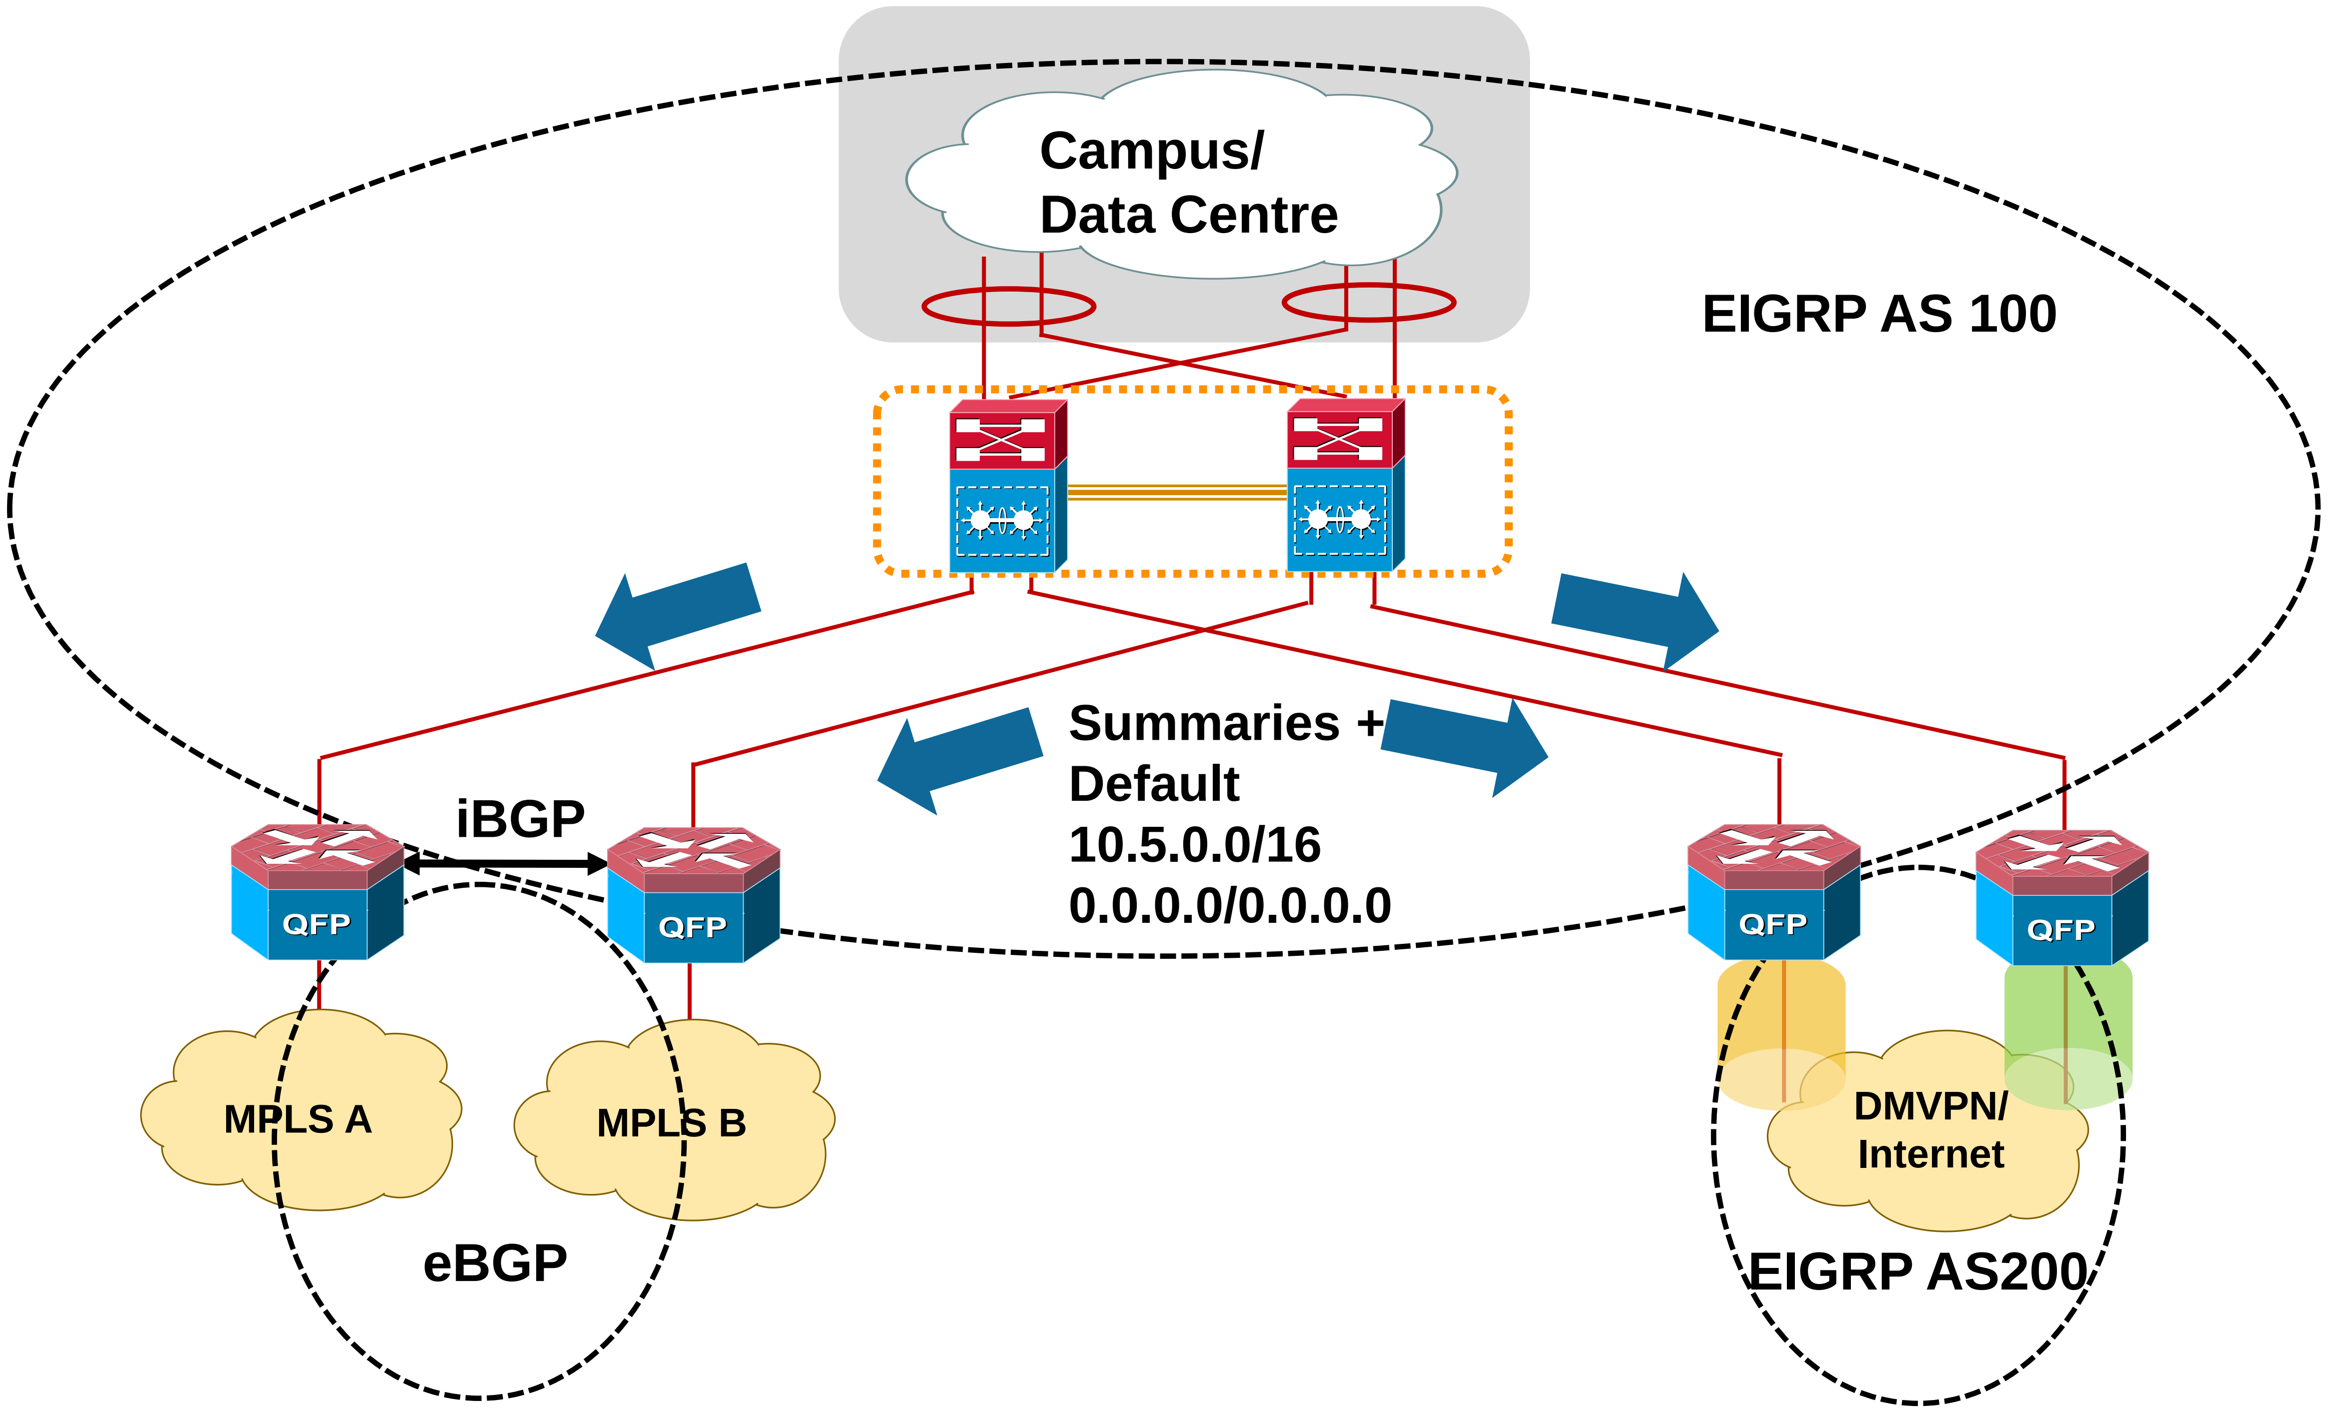

WAN Edge

Connection Methods Compared

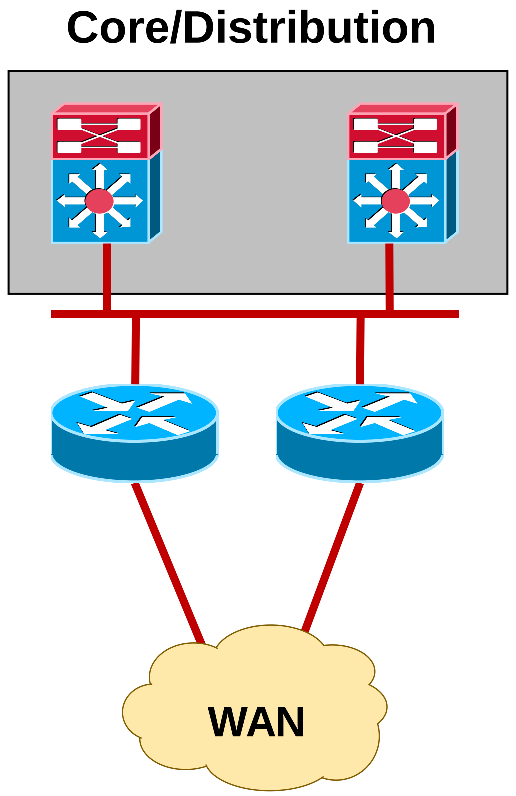

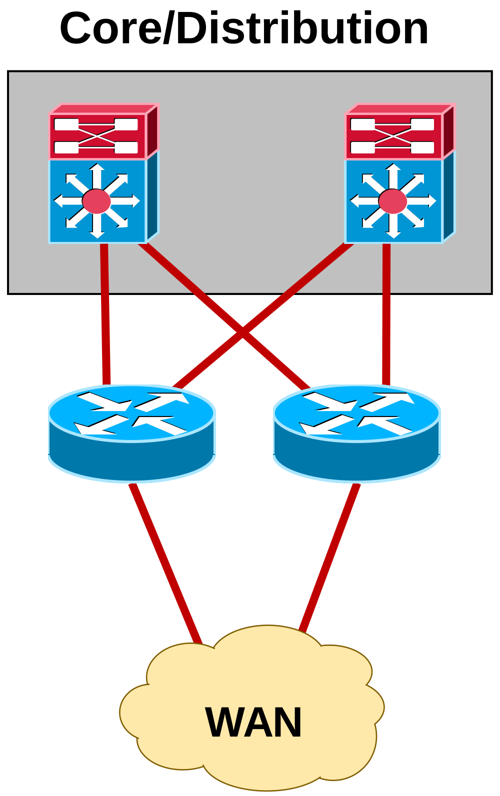

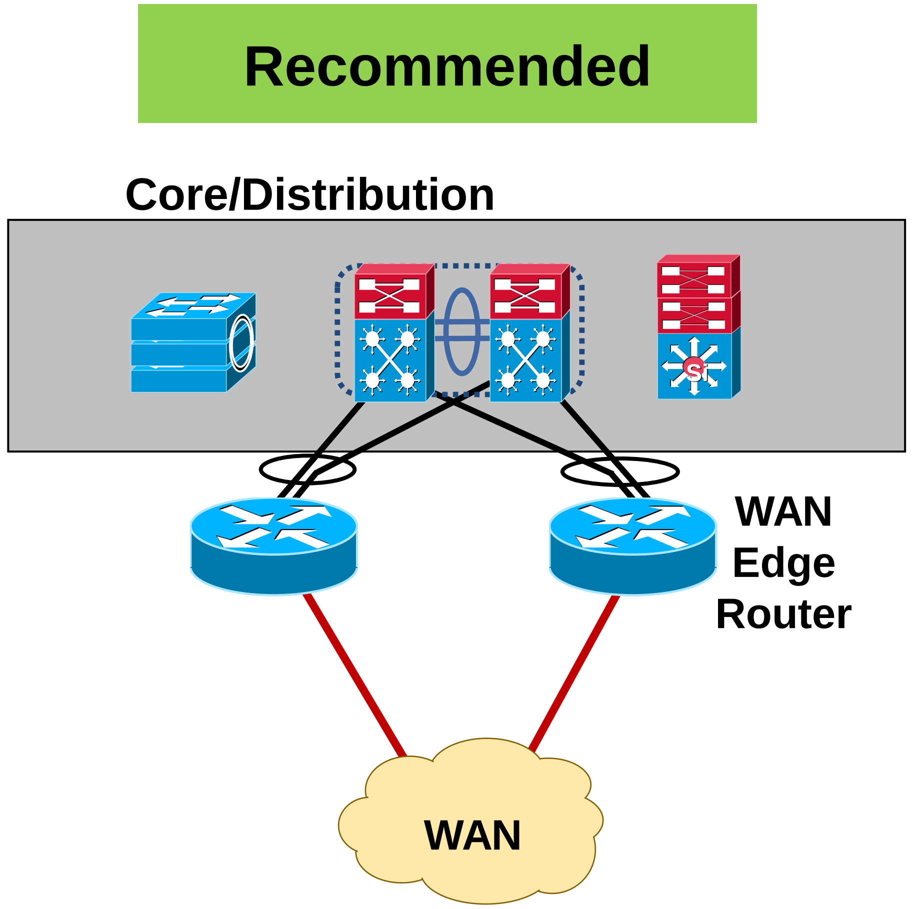

- All:
  - No static routes
  - No FHRPs

- Single Logical Control Plane
- Port-Channel for H/A

Optimise Convergence and Redundancy

Multichassis EtherChannel

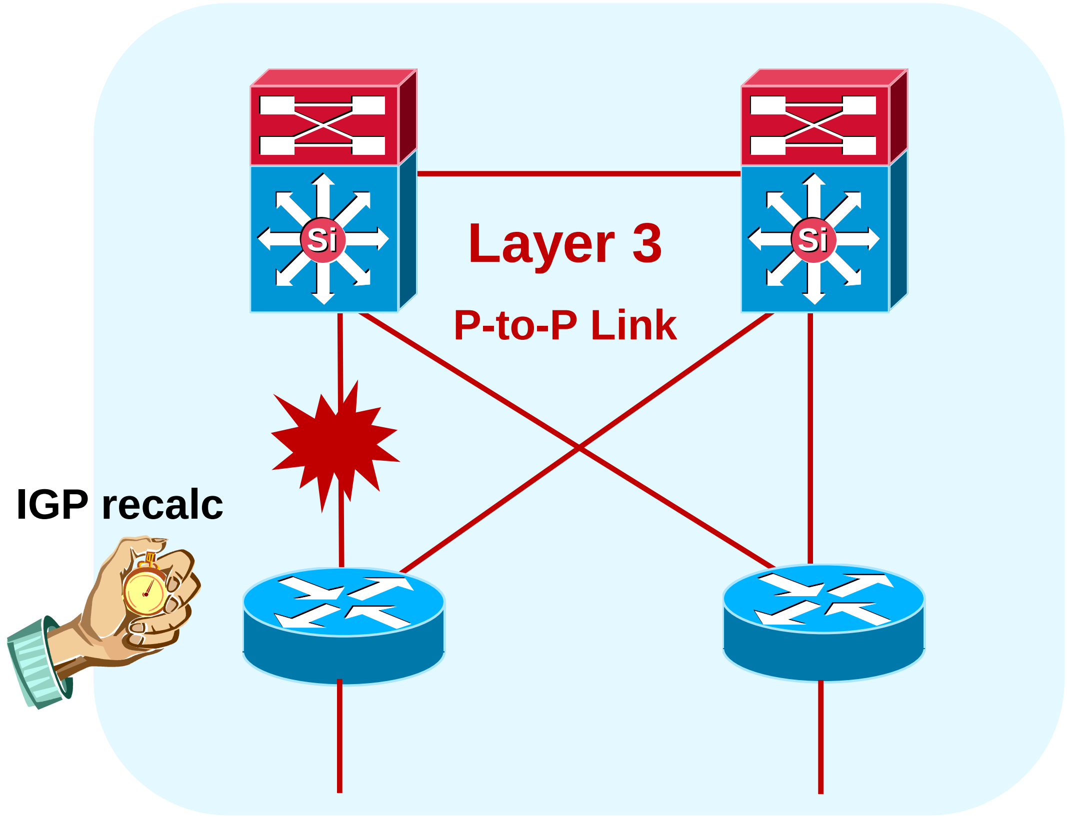

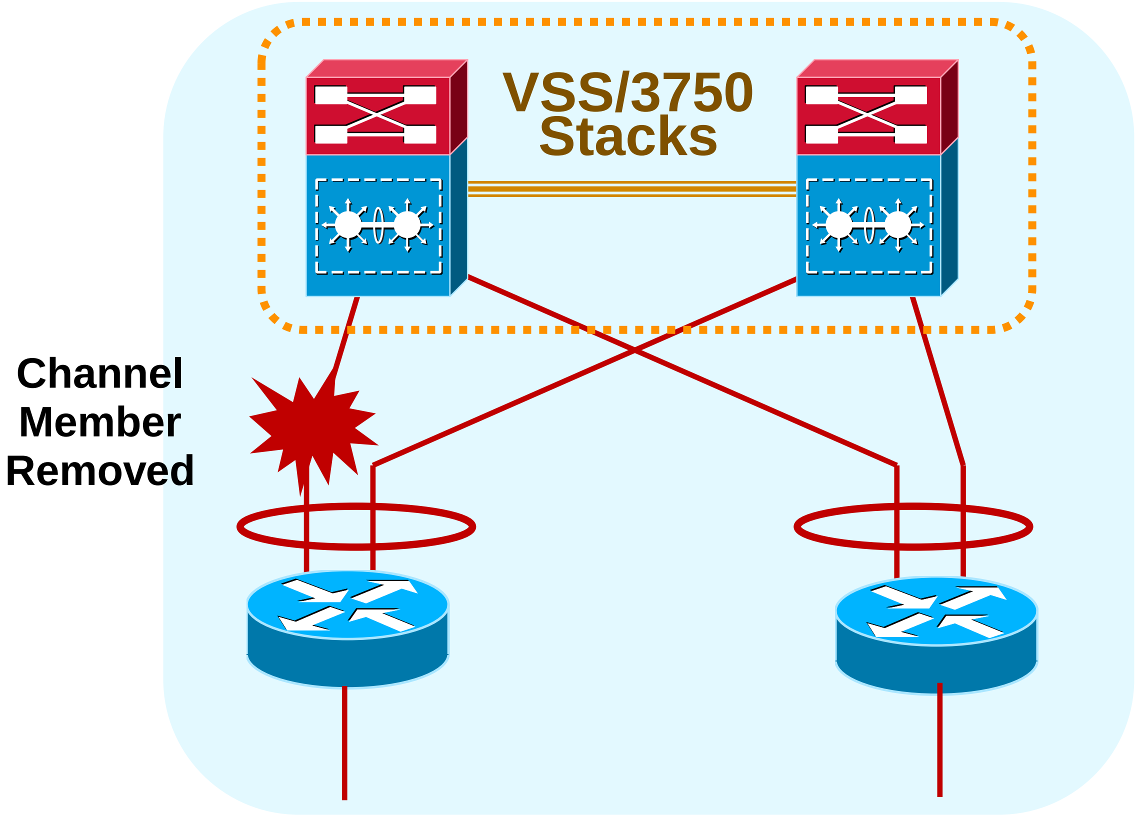

- Link redundancy achieved through redundant L3 paths
- Flow based load-balancing through CEF forwarding across
- Routing protocol reconvergence when uplink failed
- Convergence time may depends on routing protocol used and the size of routing entries

- Provide Link Redundancy and reduce peering complexity
- Tune L3/L4 load-balancing hash to achieve maximum utilisation
- No L3 reconvergence required when member link failed
- No individual flow can go faster than the speed of an individual member of the link

Link Recovery Comparison

ECMP vs. Multichassis EtherChannel

- ECMP convergence is dependent on the number of routes
- MEC convergence is consistent, independent of the number of routes

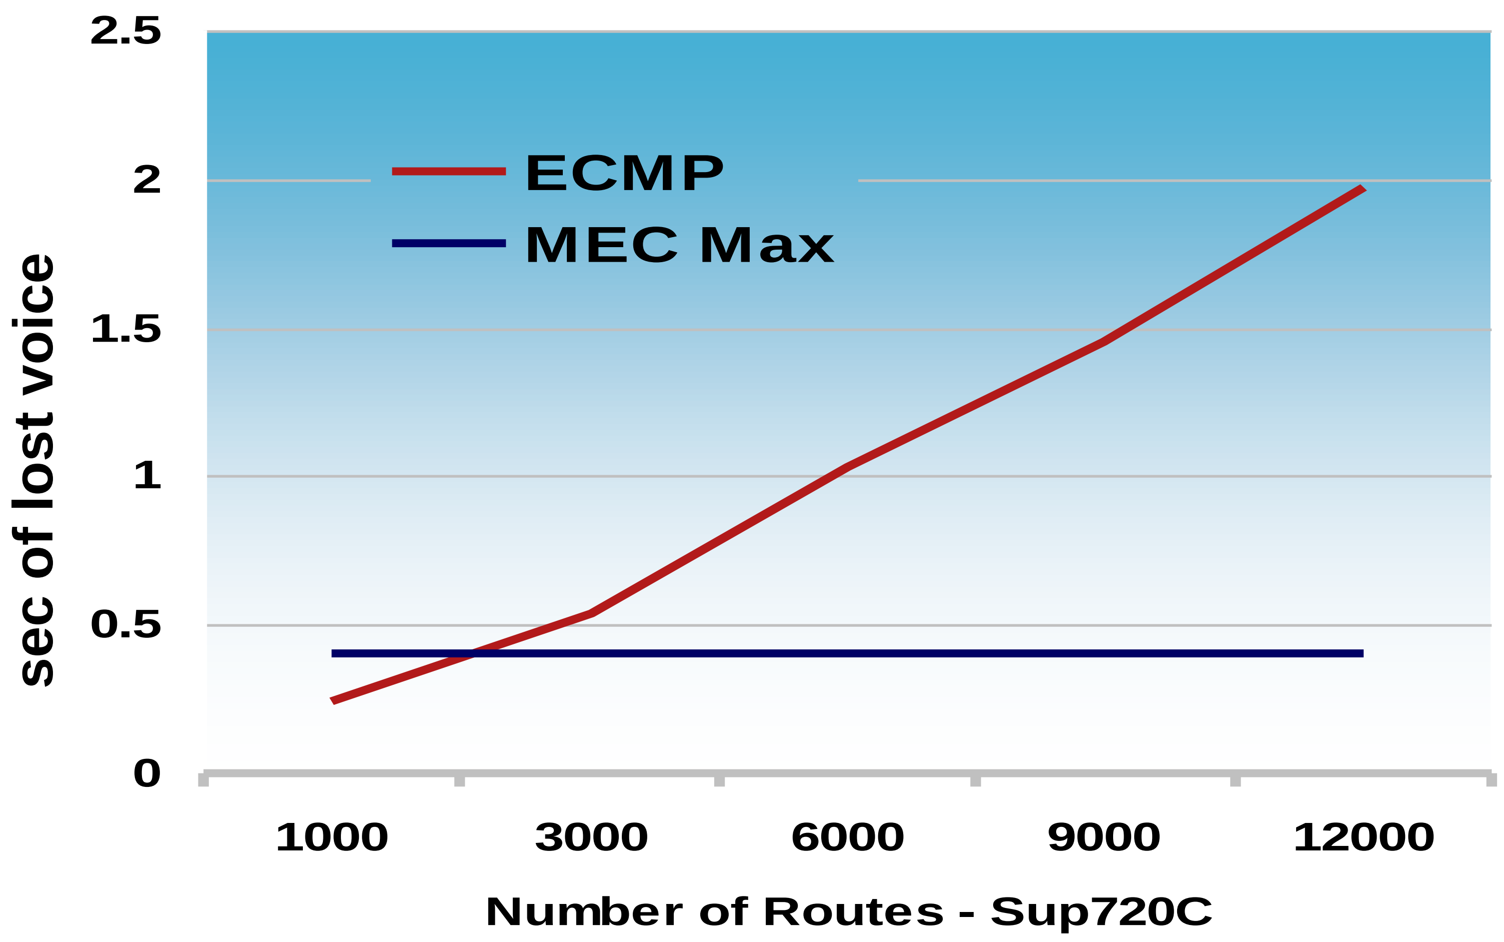

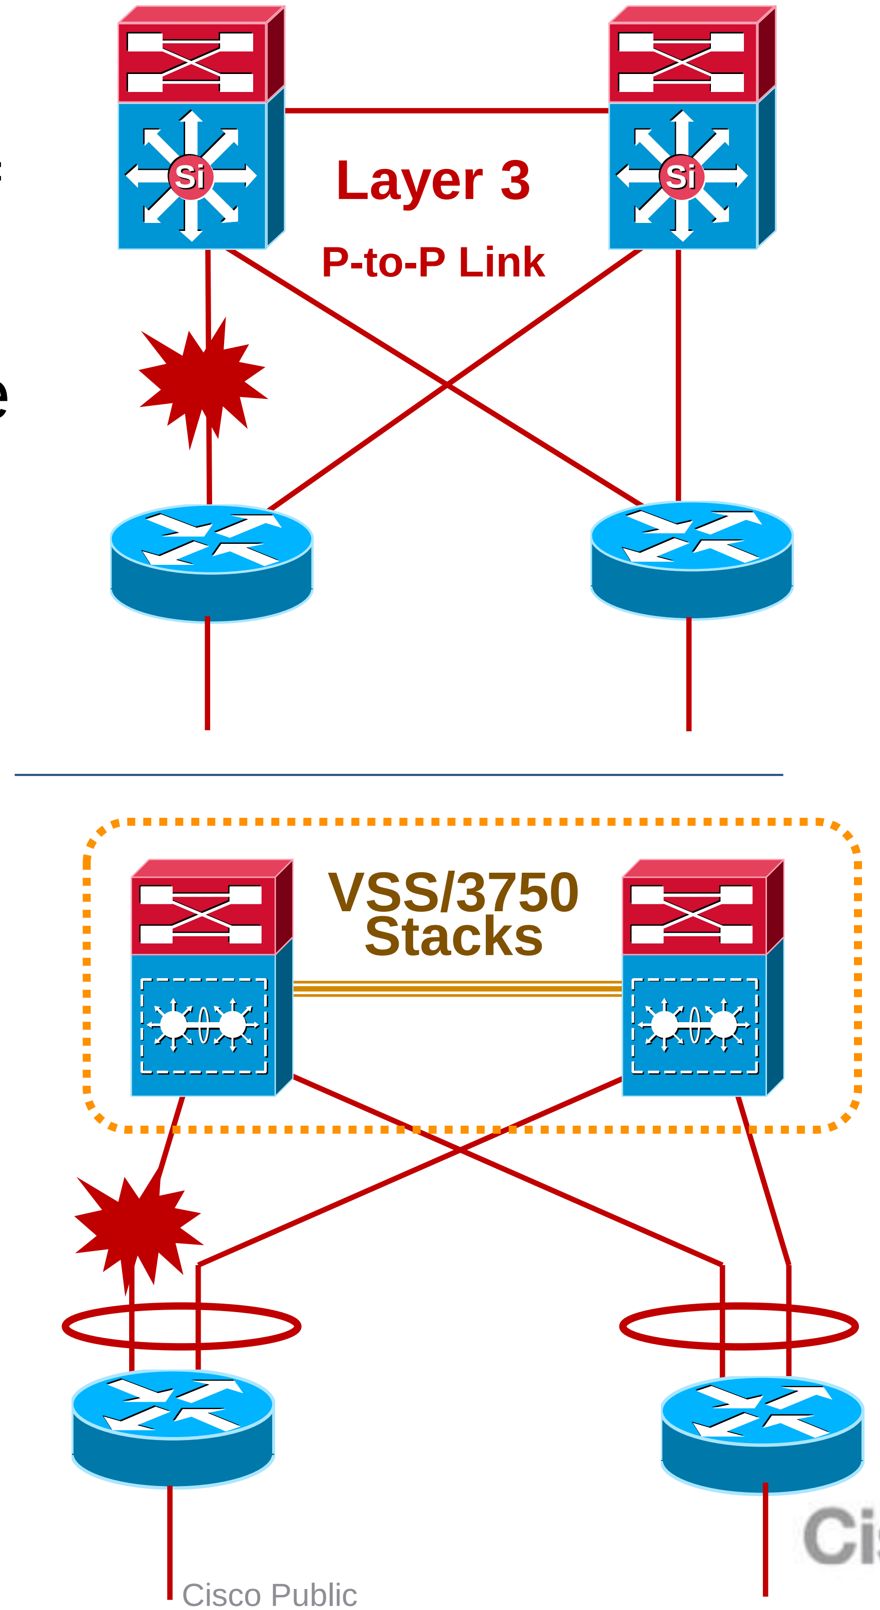

Best Practice —

Summarise at Service Distribution

- It is important to force summarisation at the distribution towards WAN Edge and towards campus & data centre
- Summarisation limit the number of peers an EIGRP router must query (minimise SIA) or the number of LSAs an OSPF peer must process

```
interface Port-channel1
 description Interface to MPLS-A-CE
 no switchport
 ip address 10.4.128.1 255.255.255.252
 ip pim sparse-mode
 ip summary-address eigrp 100 10.5.0.0 255.255.0.0
```

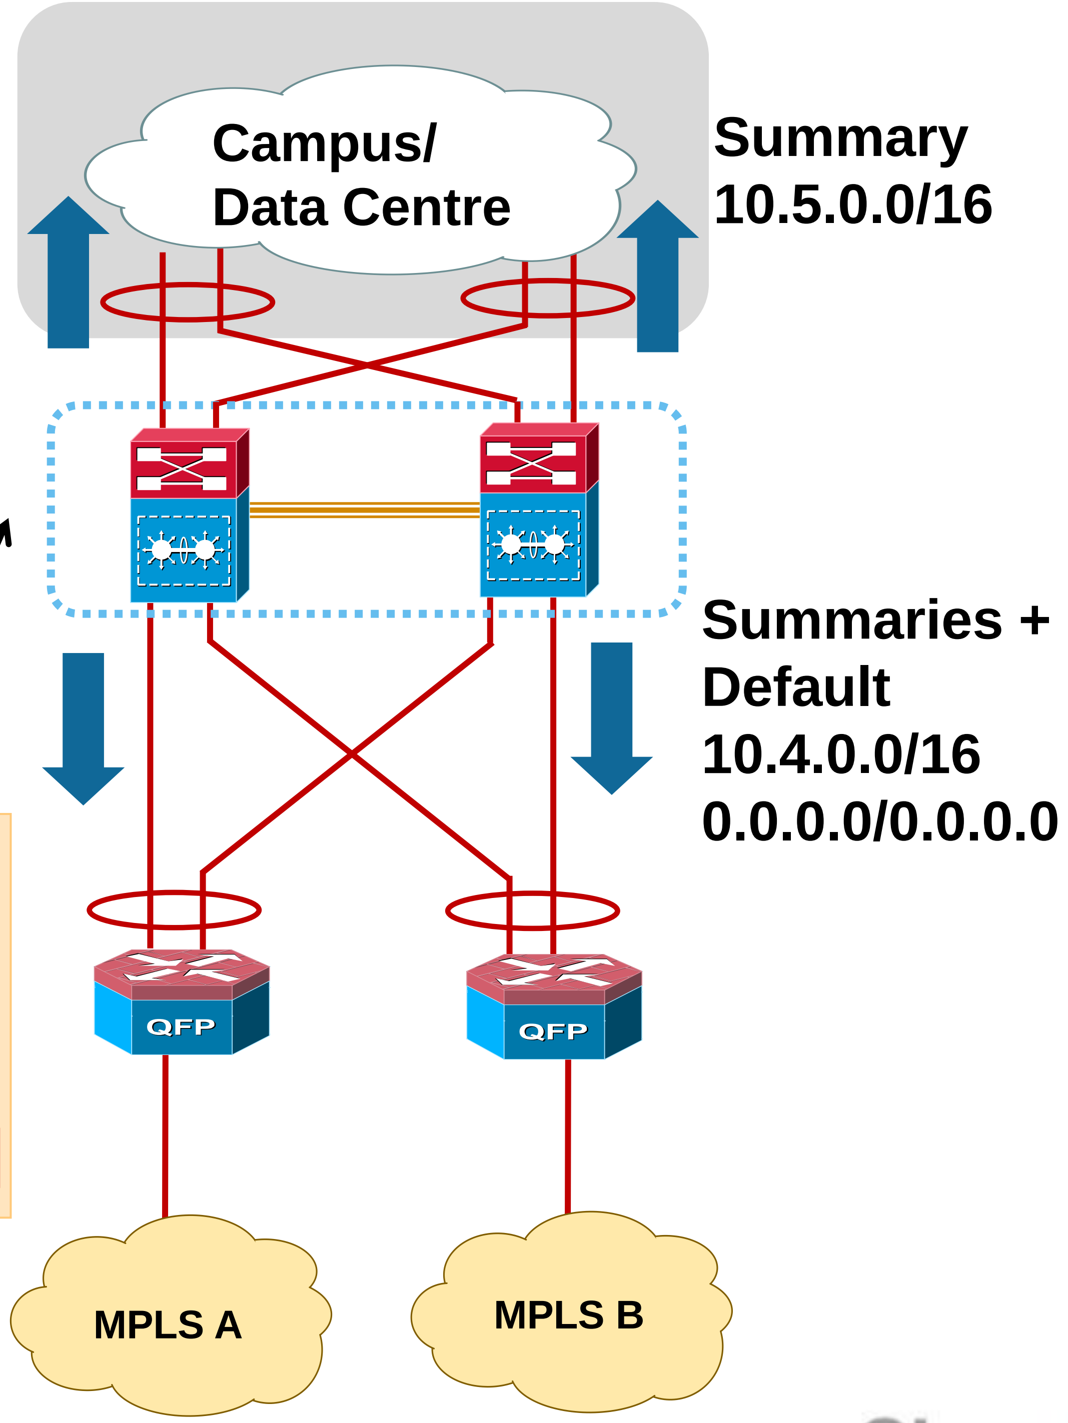

Best Practice –

Preventing Routing Loops with Route Tag and Filter

- Mutual route redistribution between protocols can cause routing loops without preventative measures
- Use route-map to set tags and then redistribute based on the tags
- Routes are implicitly tagged when distributed from eBGP to EIGRP/OSPF with carrier AS
- Use route-map to block re-learning of WAN routes via the distribution layer (already known via iBGP)

```
router eigrp 100
 distribute-list route-map BLOCK-TAGGED-ROUTES in
 default-metric [BW] 100 255 1 1500
 redistribute bgp 65500

route-map BLOCK-TAGGED-ROUTES deny 10
 match tag 65401 65402

route-map BLOCK-TAGGED-ROUTES permit 20
```

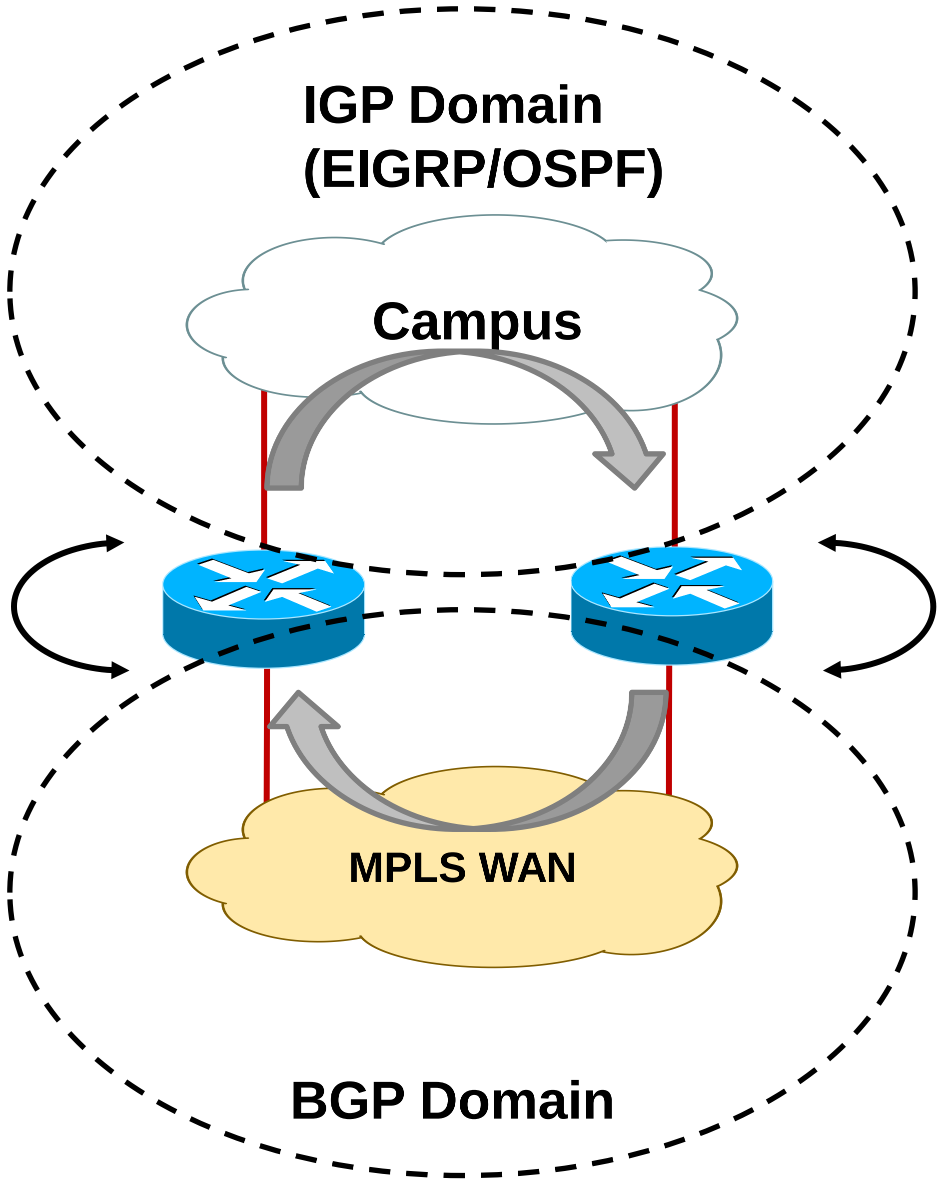

Dual Carriers with BGP as CE-PE Protocol

Use iBGP for Intelligent Path Selection

- Run iBGP between the CE routers to exchange prefixes associated with each carrier
- CE routers will use only BGP path selection information to select both the primary and secondary preferences for any destinations announced by the IGP and BGP
- Use IGP (OSPF/EIGRP) for prefix re-advertisement will result in equal-cost paths at remote-site

```
bn-br200-3945-1# sh ip bgp 10.5.128.0/21
BGP routing table entry for 10.5.128.0/21, version 71
Paths: (2 available, best #2, table default, RIB-failure(17))
  Not advertised to any peer
  65401 65402, (aggregated by 65511 10.5.128.254)
    10.4.142.26 from 10.4.142.26 (192.168.100.3)
      Origin IGP, localpref 100, valid, external, atomic-
aggregate
  65402, (aggregated by 65511 10.5.128.254)
    10.4.143.26 (metric 51456) from 10.5.0.10 (10.5.0.253)
      Origin IGP, metric 0, localpref 100, valid, internal,
atomic-aggregate, best
```

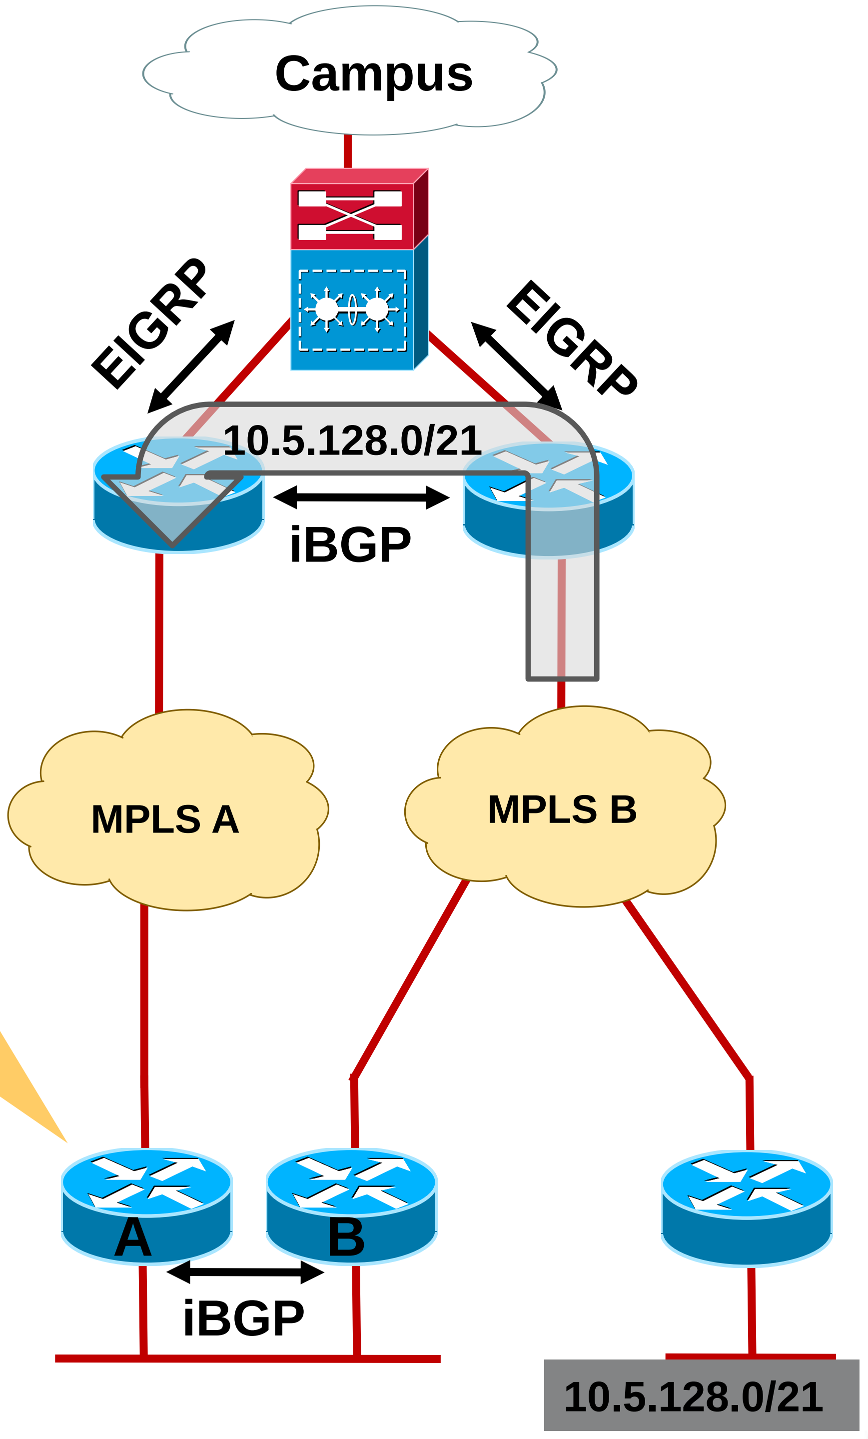

Best Practice - Implement AS-Path Filter

Prevent Branch Site Becoming Transit Network

- Dual carrier sites can unintentionally become transit network during network failure event and causing network congestion due to transit traffic
- Design the network so that transit path between two carriers only occurs at sites with enough bandwidth
- Implement AS-Path filter to allow only locally originated routes to be advertised on the outbound updates for branches that should not be transit

```
router bgp 65511
 neighbor 10.4.142.26 route-map NO-TRANSIT-AS
out
!
ip as-path access-list 10 permit ^$
!
route-map NO-TRANSIT-AS permit 10
 match as-path 10
```

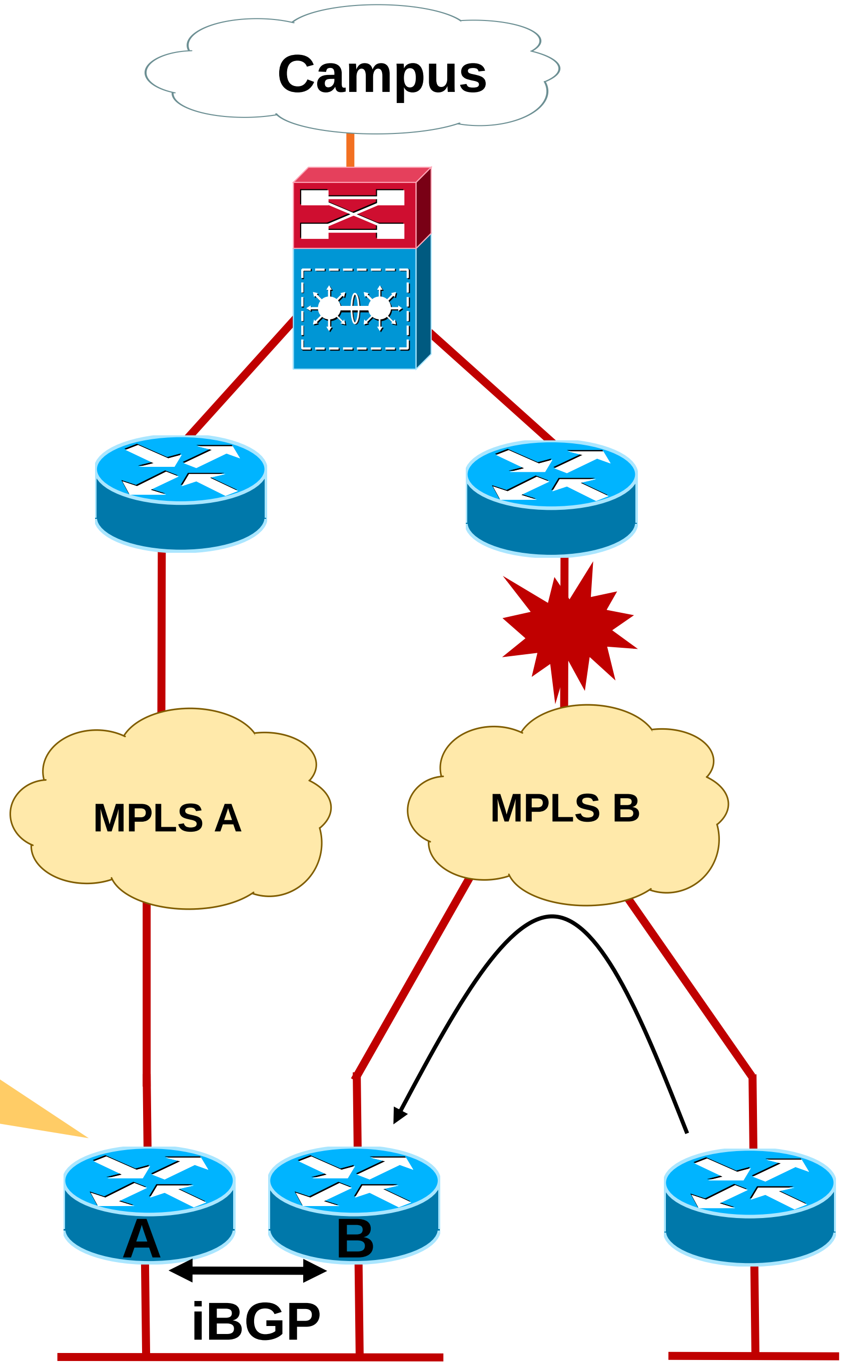

EIGRP Metric Calculation - Review

- EIGRP Composite Metric

EIGRP Metric = $256*([K_1*Bw + K_2*Bw/(256-Load) + K_3*Delay]*[K_5/(Reliability + K_4)])$

Bandwidth [Bw] (minimum along path)
Delay (aggregate)
Load (1-255)
Reliability (1-255)
MTU (minimum along path)

- For default behaviuor (K1=K3=1), the formula metric is following:
  metric = bandwidth + delay
- EIGRP uses the following formula to scale the bandwidth & delay
  *bandwidth = (10000000/bandwidth(i)) * 256*
  *delay = delay(i) *256*

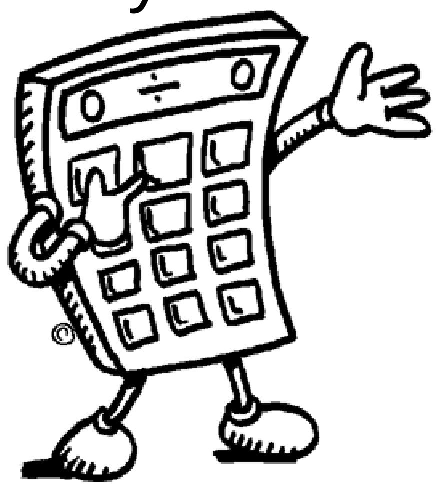

Best Practice – Use Delay Parameter to Influence EIGRP Path Selection

- EIGRP uses the minimum bandwidth along the path and the total delay to compute routing metrics
- Does anything else use these values?
  - EIGRP also uses interface Bandwidth parameter to avoid congestion by pacing routing updates (default is 50% of bandwidth)
  - Interface Bandwidth parameter is also used for QoS policy calculation
  - Performance Routing (PfR) leverages Bandwidth parameter for traffic load sharing
- **Delay parameter should always be used to influence EIGRP routing decision**

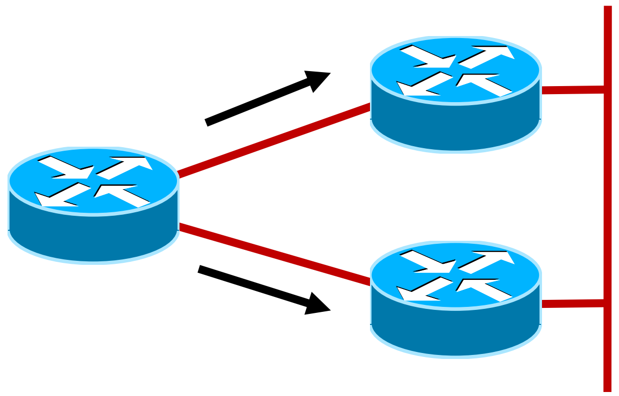

MPLS + Internet WAN

Prefer the MPLS Path over Internet

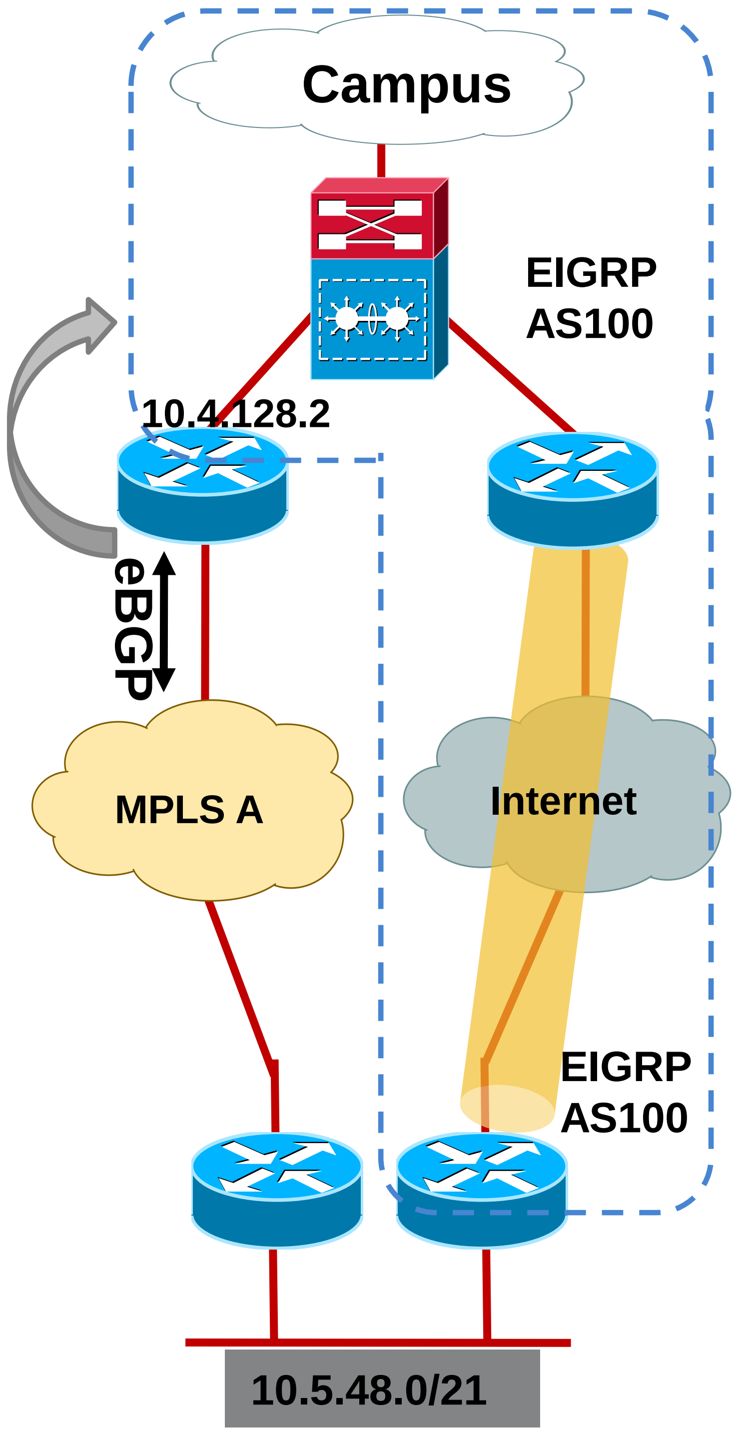

- eBGP routes are redistributed into EIGRP 100 as external routes with default Admin Distance 170
- Running same EIGRP AS for both campus and DMVPN network would result in Internet path preferred over MPLS path
- Multiple EIGRP AS processes can be used to provide control of the routing
  - EIGRP 100 is used in campus location
    EIGRP 200 over DMVPN tunnels
  - Routes from EIGRP 200 redistributed into EIGRP 100 appear as external route (distance = 170)
- Routes from both WAN sources are equal-cost paths. To prefer MPLS path over DMVPN use eigrp delay to modify path preference

MPLS + Internet WAN

Use Autonomous System for IGP Path Differentiation

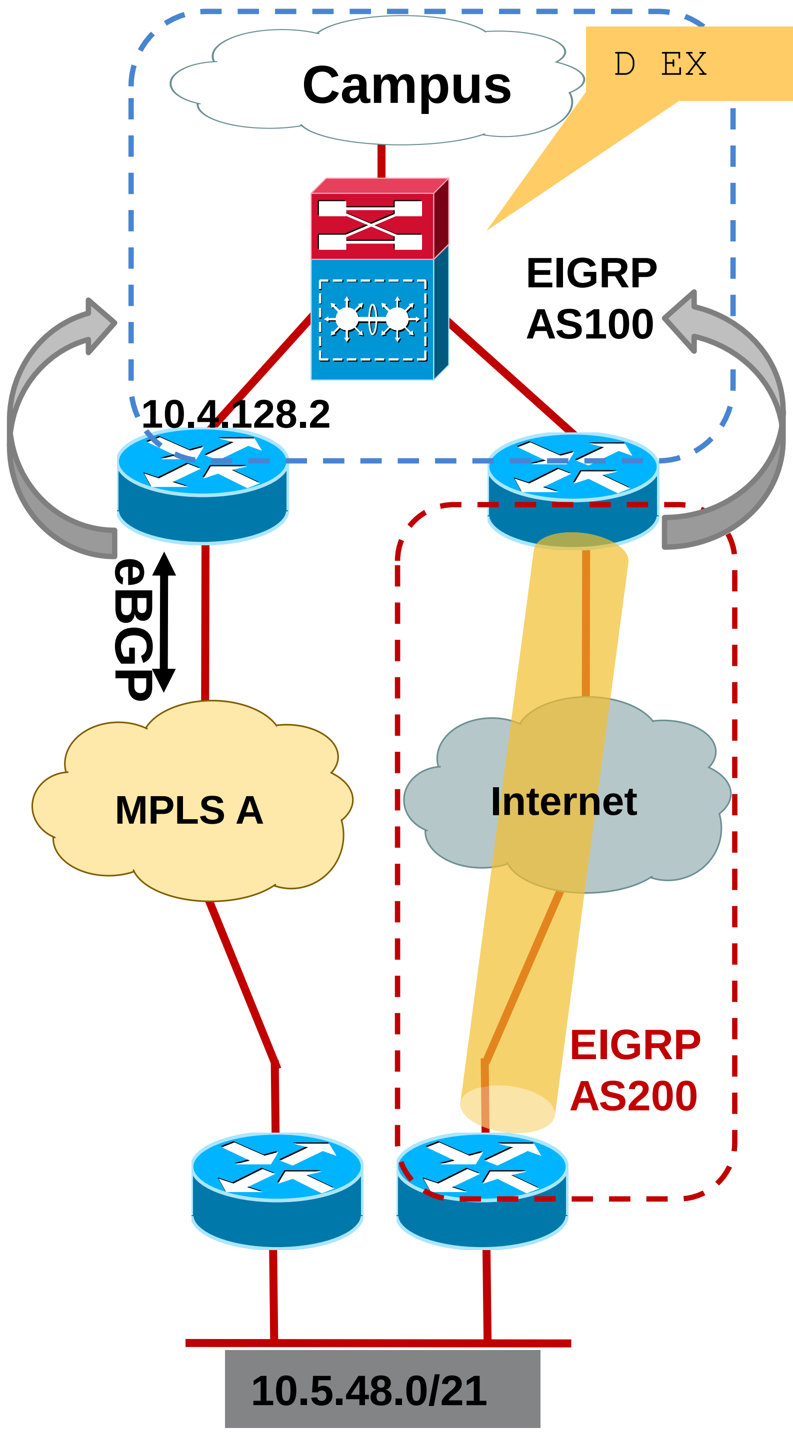

- eBGP routes are redistributed into EIGRP 100 as external routes with default Admin Distance 170
- Running same EIGRP AS for both campus and DMVPN network would result in Internet path preferred over MPLS path
- Multiple EIGRP AS processes can be used to provide control of the routing
  - EIGRP 100 is used in campus location EIGRP 200 over DMVPN tunnels
  - Routes from EIGRP 200 redistributed into EIGRP 100 appear as external route (distance = 170)
- Routes from both WAN sources are equal-cost paths. To prefer MPLS path over DMVPN use eigrp delay to modify path preference

```
MPLS CE router#

router eigrp 100
 default-metric 1000000 10 255 1 1500
```

MPLS VPN BGP Path with IGP Backdoor Path

- eBGP as the PE-CE Routing Protocol
- MPLS VPN as preferred path learned via eBGP
- Secondary path via backdoor IGP link (EIGRP or OSPF) over tunneled connection (DMVPN over Internet)
- Default configuration the failover to backup path works as expected

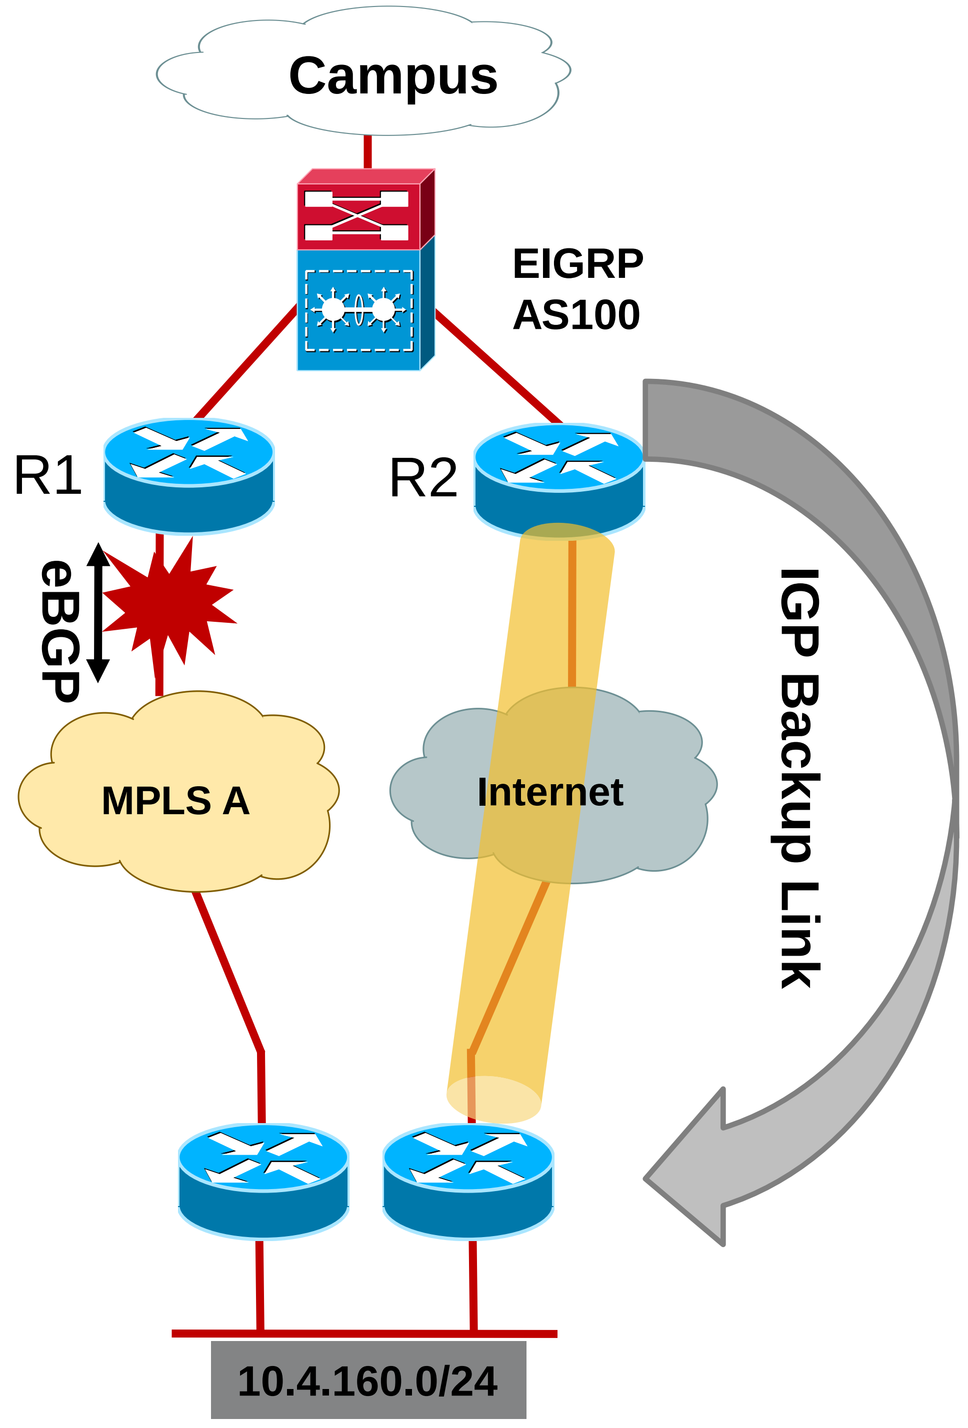

MPLS VPN BGP Path with IGP Backdoor Path

- After link restore, MPLS CE router receives BGP advertisement for remote-site route.
- Does BGP route get (re)installed in the route table?

```
D EX 10.4.160.0/24 [170/3584]....
```

```
R1# show ip route
B    10.4.144.0/24 [20/0] via 10.4.142.2, 01:30:06
B    10.4.145.0/24 [20/0] via 10.4.142.2, 01:30:06
```

```
B    10.4.160.0/24 [20/0]....
```

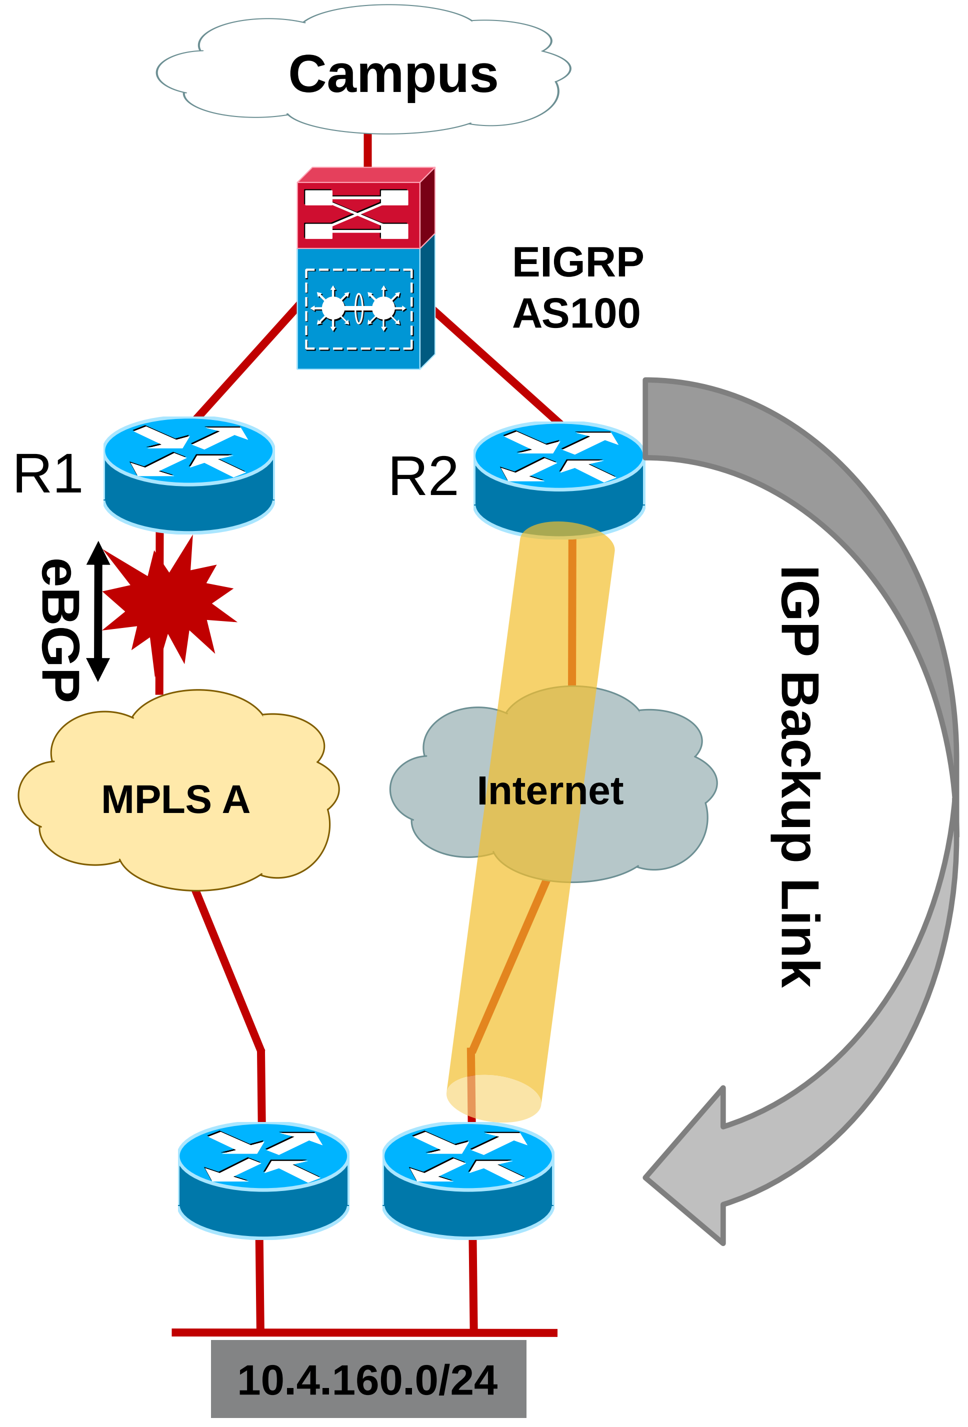

BGP Route Selection Criteria

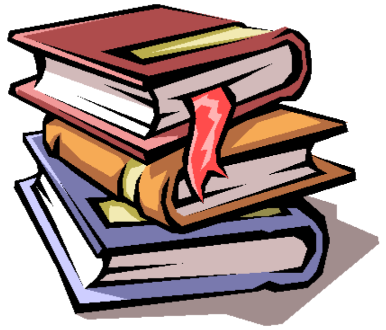

BGP Prefers Path with:

1. Highest Weight
2. Highest Local PREF
3. Locally originated via network or aggregate BGP
4. Shortest AS_PATH
5. Lowest Origin type
   IGP>EGP>INCOMPLETE
6. Lowest MED
7. eBGP over iBGP paths
8. Lowest IGP metric to BGP next hop

BGP Prefers Path with Highest Weight

- Routes redistributed into BGP are considered locally originated and get a default weight of 32768
- The eBGP learned prefix has default weight of 0
- Path with *highest* weight is selected

```
ASR1004-1#show ip bgp 10.4.160.0 255.255.255.0
BGP routing table entry for 10.4.160.0/24, version 22
Paths: (3 available, best #3, table default)
  Advertised to update-groups:
     4          5
  65401 65401
    10.4.142.2 from 10.4.142.2 (192.168.100.3)
      Origin IGP, localpref 200, valid, external
Local
    10.4.128.1 from 0.0.0.0 (10.4.142.1)
      Origin incomplete, metric 26883072, localpref 100, weight 32768, valid, sourced, best
```

Prefer the eBGP Path over IGP

Set the eBGP weight > 32768

- To resolve this issue set the weights on route learned via eBGP peer higher than 32768

neighbor 10.4.142.2 weight 35000

```
ASR1004-1#show ip bgp 10.4.160.0 255.255.255.0
BGP routing table entry for 10.4.160.0/24, version 22
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  65401 65401
    10.4.142.2 from 10.4.142.2 (192.168.100.3)
      Origin IGP, metric 0, localpref 100, weight 35000, valid, external, best
```

```
ASR1004-1#show ip route
....
B    10.4.160.0/24 [20/0] via 10.4.142.2, 05:00:06
```

Agenda

- WAN Technologies & Solutions
  - WAN Transport Technologies
  - WAN Overlay Technologies
  - WAN Optimisation
  - Wide Area Network Quality of Service
- WAN Architecture Design Considerations
  - WAN Design and Best Practices
  - Secure WAN Communication with GETVPN
  - DMVPN Over Internet Deployment
- Summary

GETVPN Topology

COOP Key Server

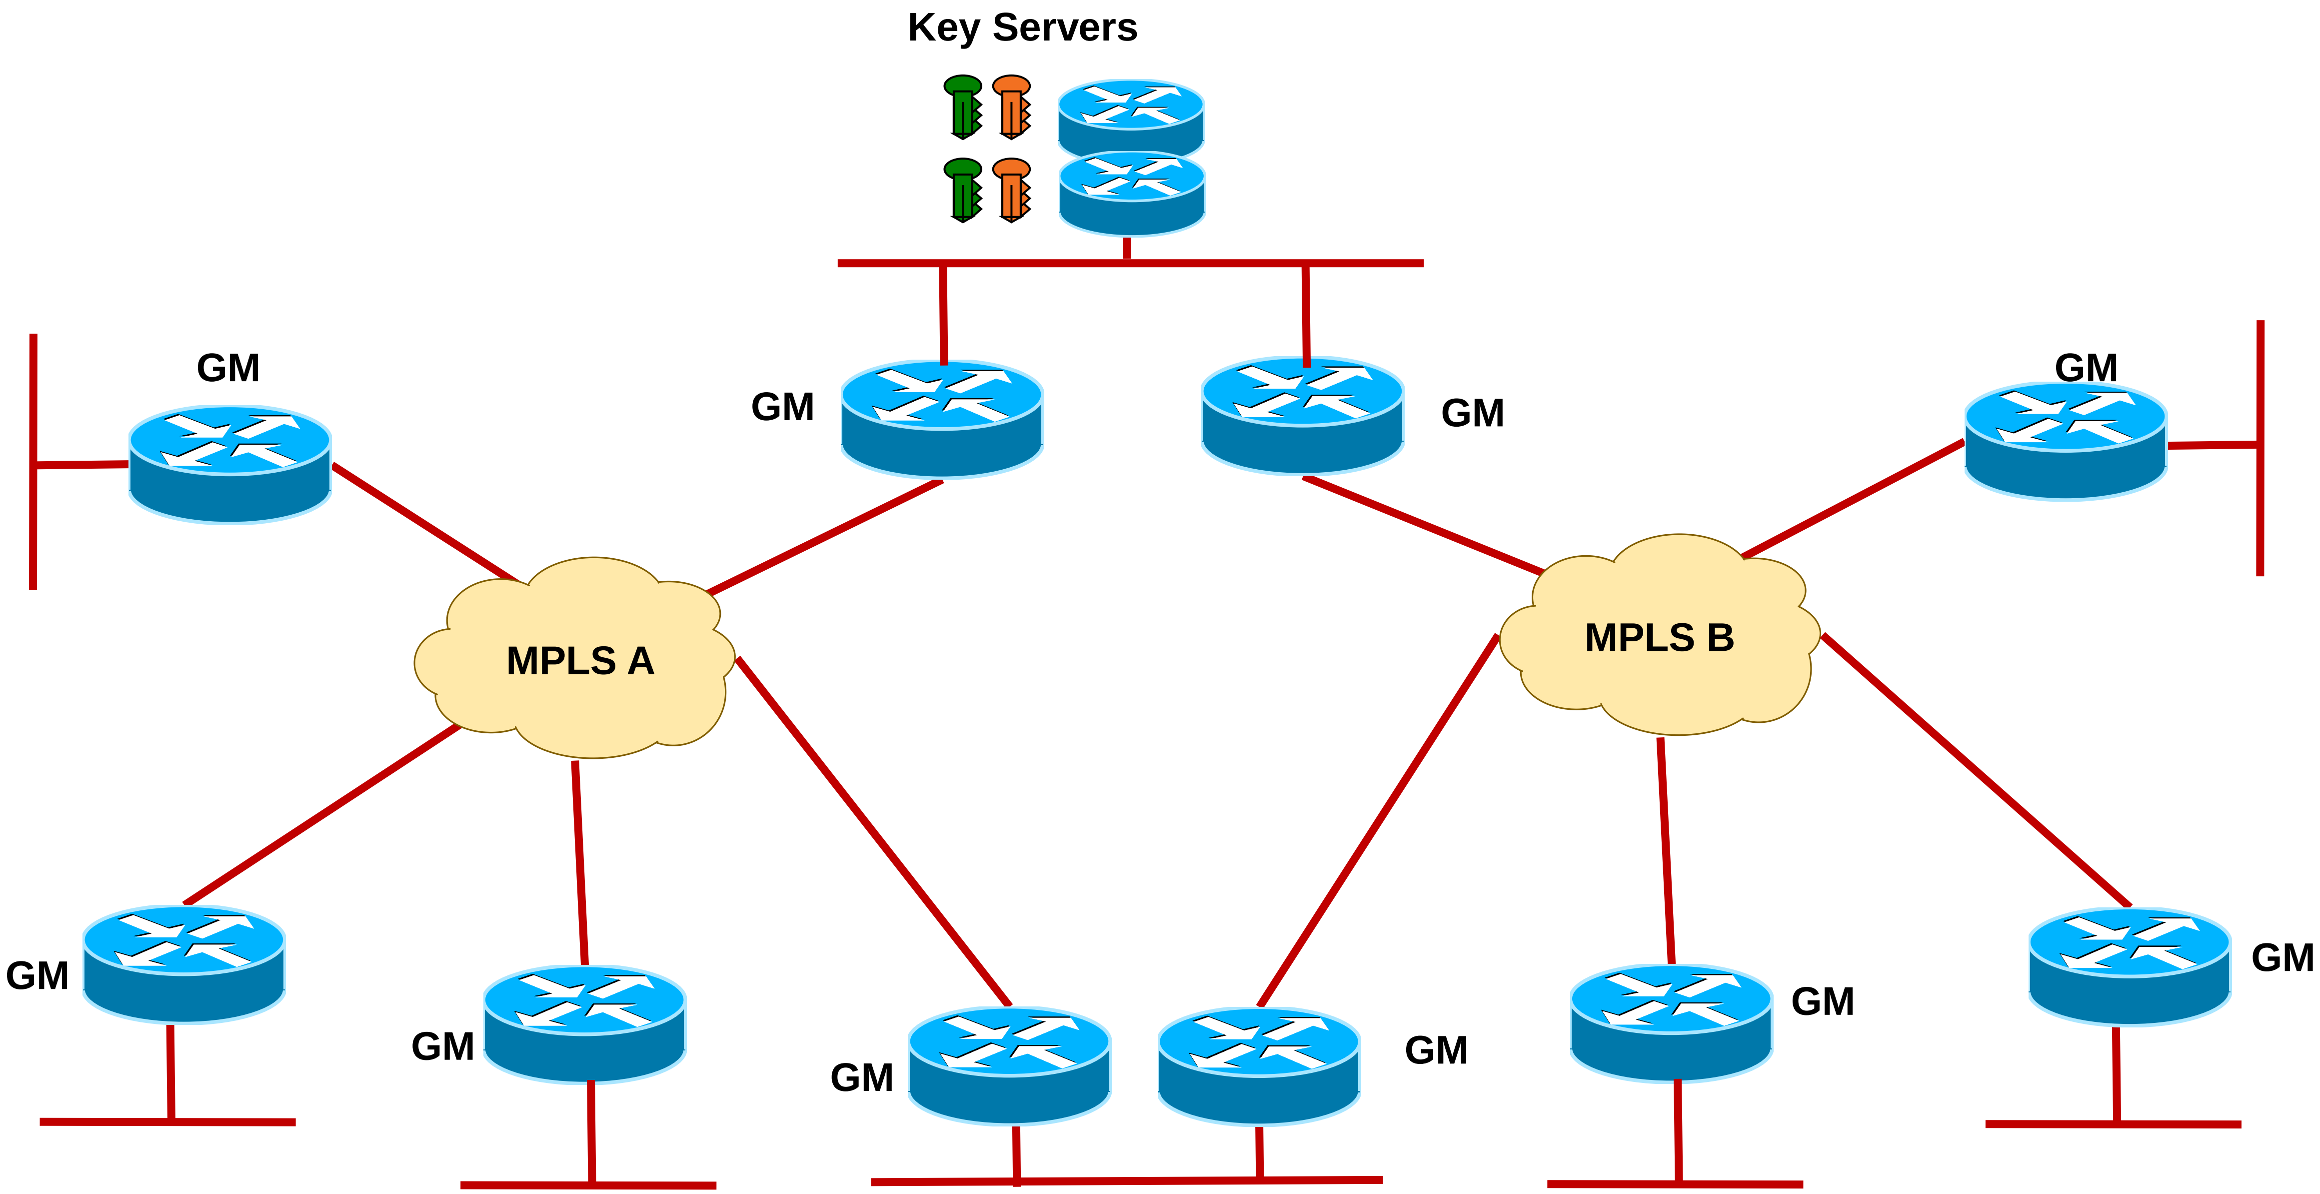

Best Practice - High Availability with Cooperative Key Servers

- Two or more KSs known as COOP KSs manage a common set of keys and security policies for GETVPN group members
- Group members can register to any one of the available KSs
- Cooperative KSs periodically exchange and synchronise group's database, policy and keys
- Primary KS is responsible to generate and distribute group keys

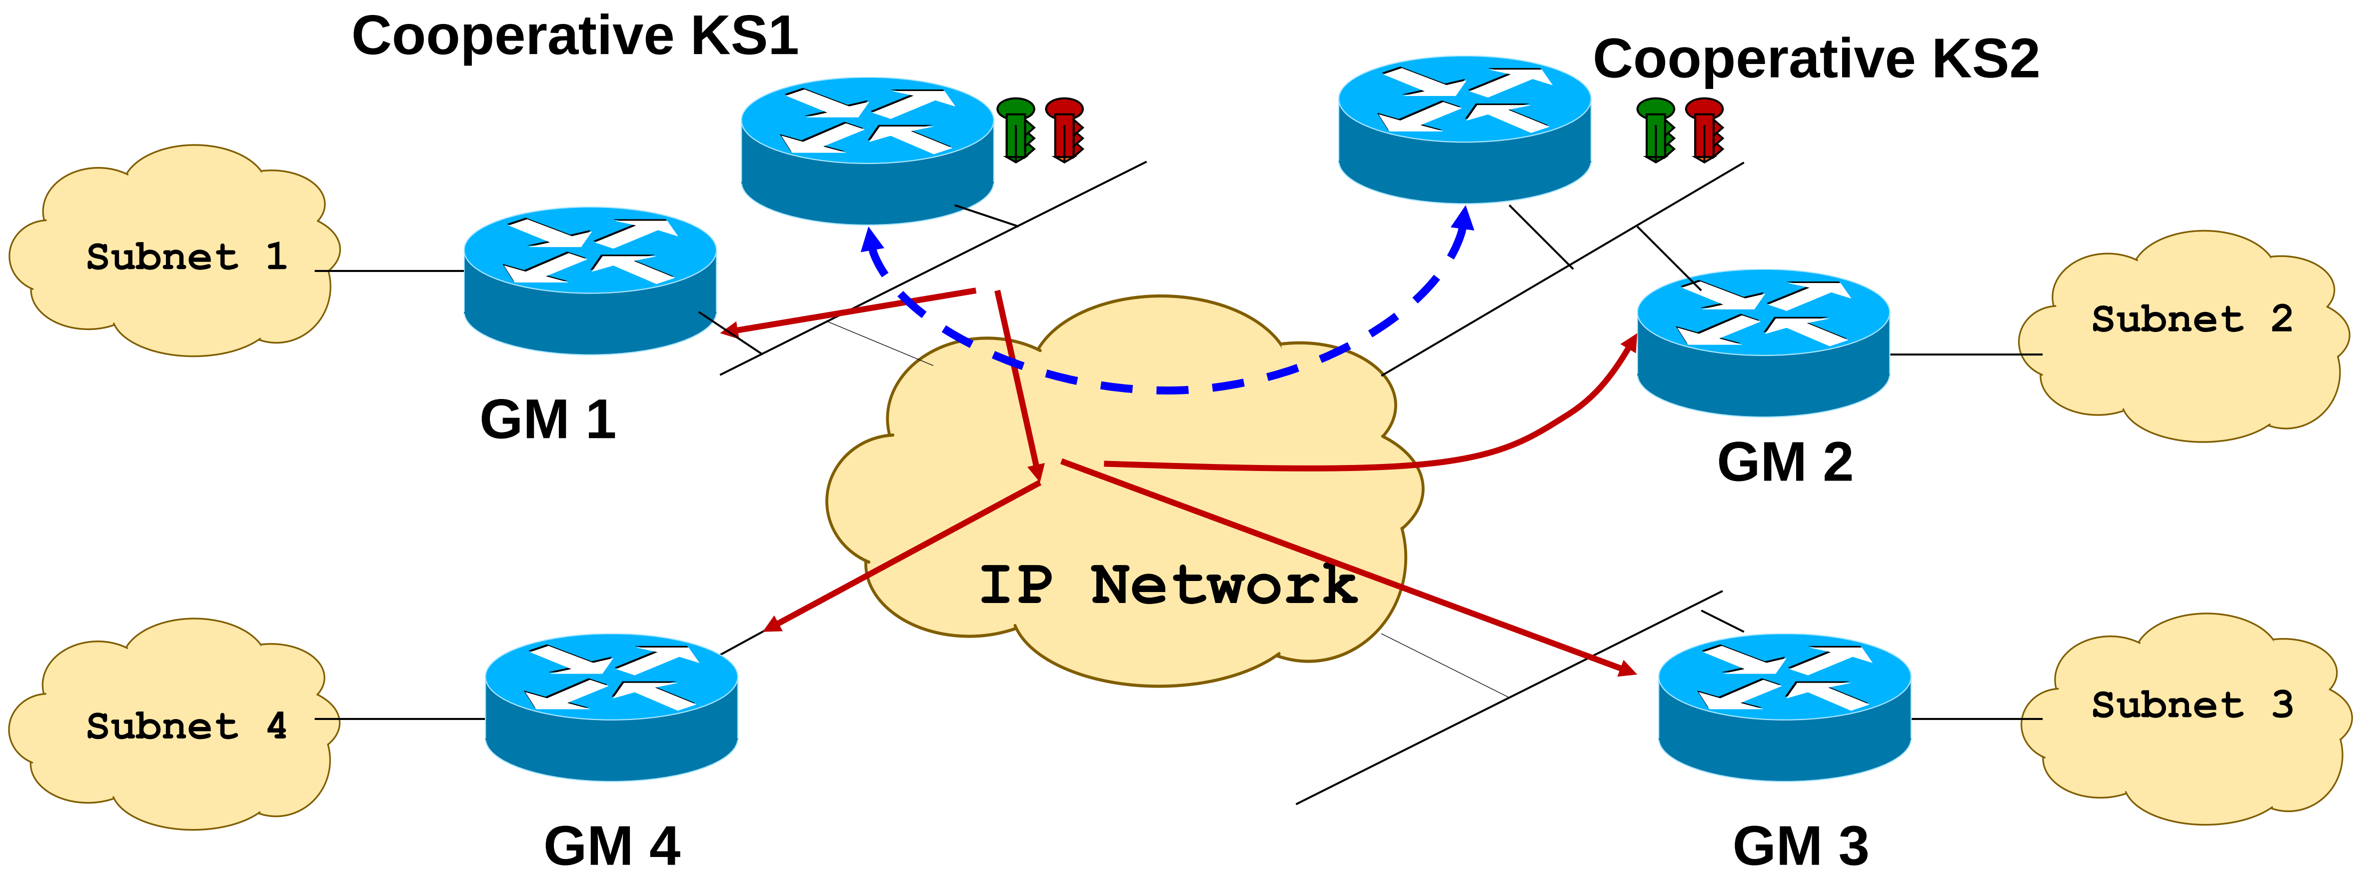

Transition from Clear-text to GETVPN

SA Receive-Only Method

- Goal
  - –Incrementally deploy infrastructure without encryption
  - –Immediate transition to encryption controlled by KS
- Method
  - –Deploy KS with Receive-only SA's (don't encrypt, allow decryption)
  - –Deploy GM throughout infrastructure and monitor rekey processes
  - –Transition KS to Normal SA (encrypt, decrypt)
- Assessment
  - –Pro: Simple transition to network-wide encryption
  - –Con: Correct policies imperative
  - –Con: Deferred encryption until all CE are capable of GM functions

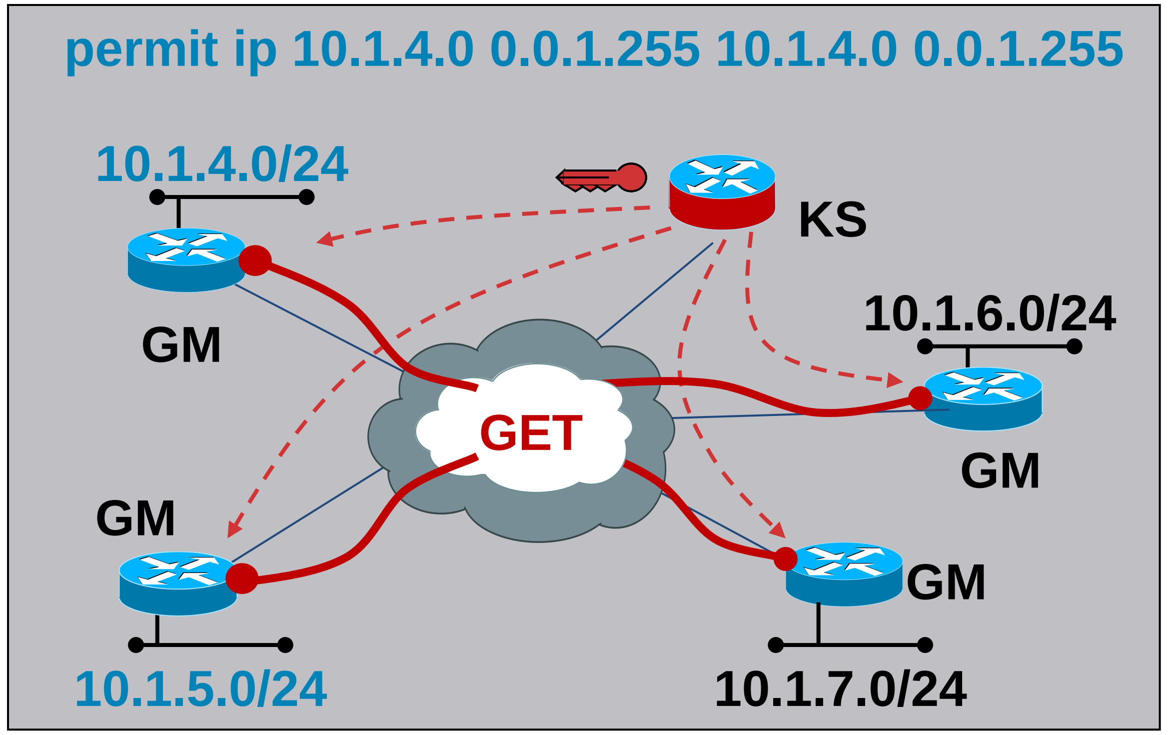

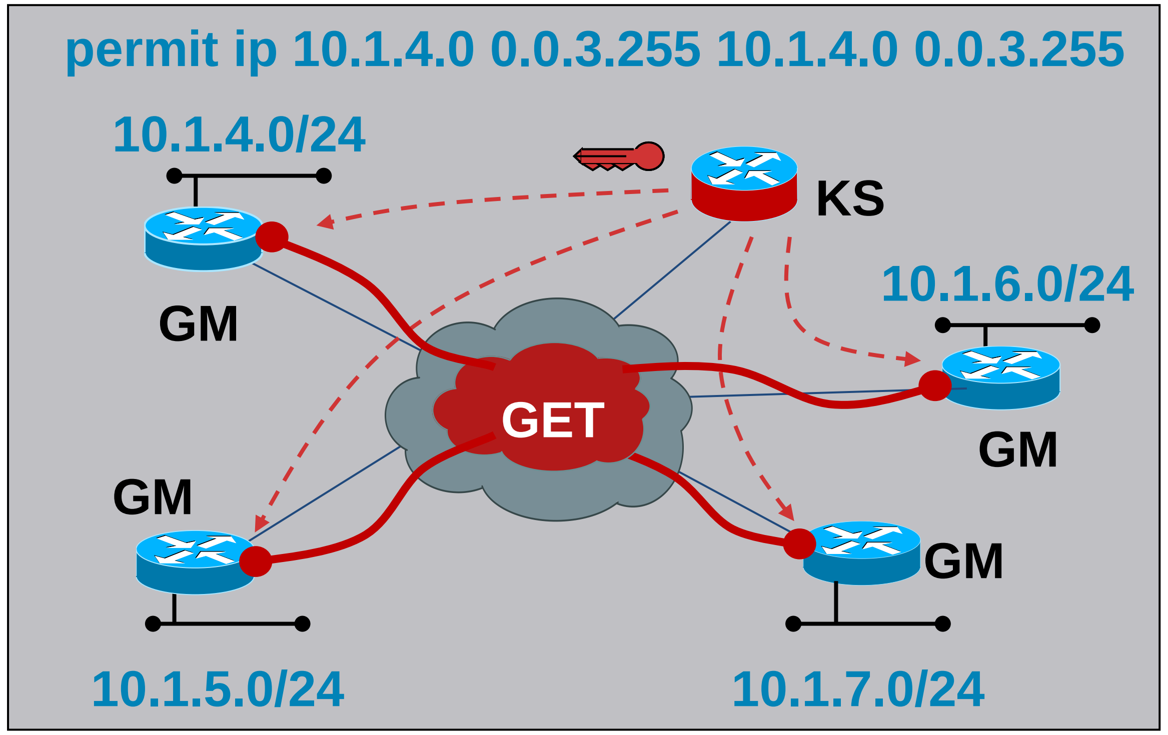

Group Member

Secured Group Member Interface

```
interface Serial0/0
ip address 192.168.1.14 255.255.255.252
  crypto map svn                            <- WAN ENCRYPTION
  access-group pack-filter out              <- ALLOW IPsec and Control
```

Packet filter (after encryption)

```
ip access-list extended pack-filter
   permit esp any any                             <- ALLOW IPsec
   permit ip host 192.168.1.14 host 192.168.1.13  <- ALLOW ROUTE ADJACENCY
   permit tcp host 192.168.1.14 eq ssh any        <- ALLOW SECURE SHELL
```

Crypto Map Association to Group Security

```
crypto map svn 10 gdoi<- GROUP CRYPTO MAP ENTRY
  set group secure-wan                      <- GROUP MEMBERSHIP
  match address control_plane               <- LOCAL POLICY (EXCLUDE)
```

Group Member Policy Exceptions

```
ip access-list extended control_plane          <- CONTROL PLANE PROTOCOLS
   deny ip host 192.168.1.14 host 192.168.1.13 <- PE-CE LINK (BGP, ICMP)
   deny tcp host 192.168.1.14 eq ssh any       <- MANAGEMENT SECURE SHELL
```

Group Member Association

```
crypto gdoi group secure-wan                <- GROUP ENCRYPTION
   identity number 3333                     <- MEMBER'S GROUP IDENTITY
   server address ipv4 <ks1_address>        <- KS ADDRESS TO REGISTER
   server address ipv4 <ks2_address>        <- ALTERNATE KS REGISTRATION
```

Key Server

```
crypto gdoi group secure-wan
     identity number 3333                          <- GROUP ID
     server local                                  <- KEY SERVER
     rekey address ipv4 102                        <- REKEY ADDRESSES REKEY
     rekey retransmit 40 number 3                  <- REKEY RETRANSMITS
     rekey authentication mypubkey rsa my_rsa      <- KS MSG AUTHENTICATION
     saipsec 10                                    <- SECURITY ASSOCIATION
     profile gdoi-p                                <- CRYPTO ATTRIBUTES SELECTION
     match address ipv4ipsec-policy                <- ENCRYPTION POLICY
     no replay                                     <- NO ANTI-REPLAY
     address ipv4 <ks_address>                     <- KS ADDRESS
```

Rekey Profile (needed for multicast rekey only)

```
access-list 102 permit any host 239.192.1.1   <- REKEY SOURCE / DESTINATION
```

Encryption IPsec Proxy ID's (mandatory)

```
ip access-list extended ipsec-policy                         <- ENCRYPTION POLICY
   deny udp any eq 848 any eq 848                            <- ALLOW GDOI
   permit ip 10.0.0.0 0.255.255.255 10.0.0.0 0.255.255.255   <- UNICAST
   permit ip 10.0.0.0 0.255.255.255 232.0.0.0 0.255.255.255 <- MULTICAST
```

Agenda

- WAN Technologies & Solutions
  - WAN Transport Technologies
  - WAN Overlay Technologies
  - WAN Optimisation
  - Wide Area Network Quality of Service
- WAN Architecture Design Considerations
  - WAN Design and Best Practices
  - Secure WAN Communication with GETVPN
  - DMVPN Over Internet Deployment
- Summary

DMVPN Deployment over Internet

Multiple Default Routes for VPN Headend

- VPN Headend has a default route to ASA firewall's VPN-DMZ interface to reach Internet
- Remote site policy requires centralised Internet access
- Enable EIGRP between VPN headend & Campus core to propagate default to remote
- Static default (admin dist=0) remains active,
- VPN-DMZ is wrong firewall interface for user traffic
- Adjust admin distance so EIGRP route installed (to core)
- VPN tunnel drops

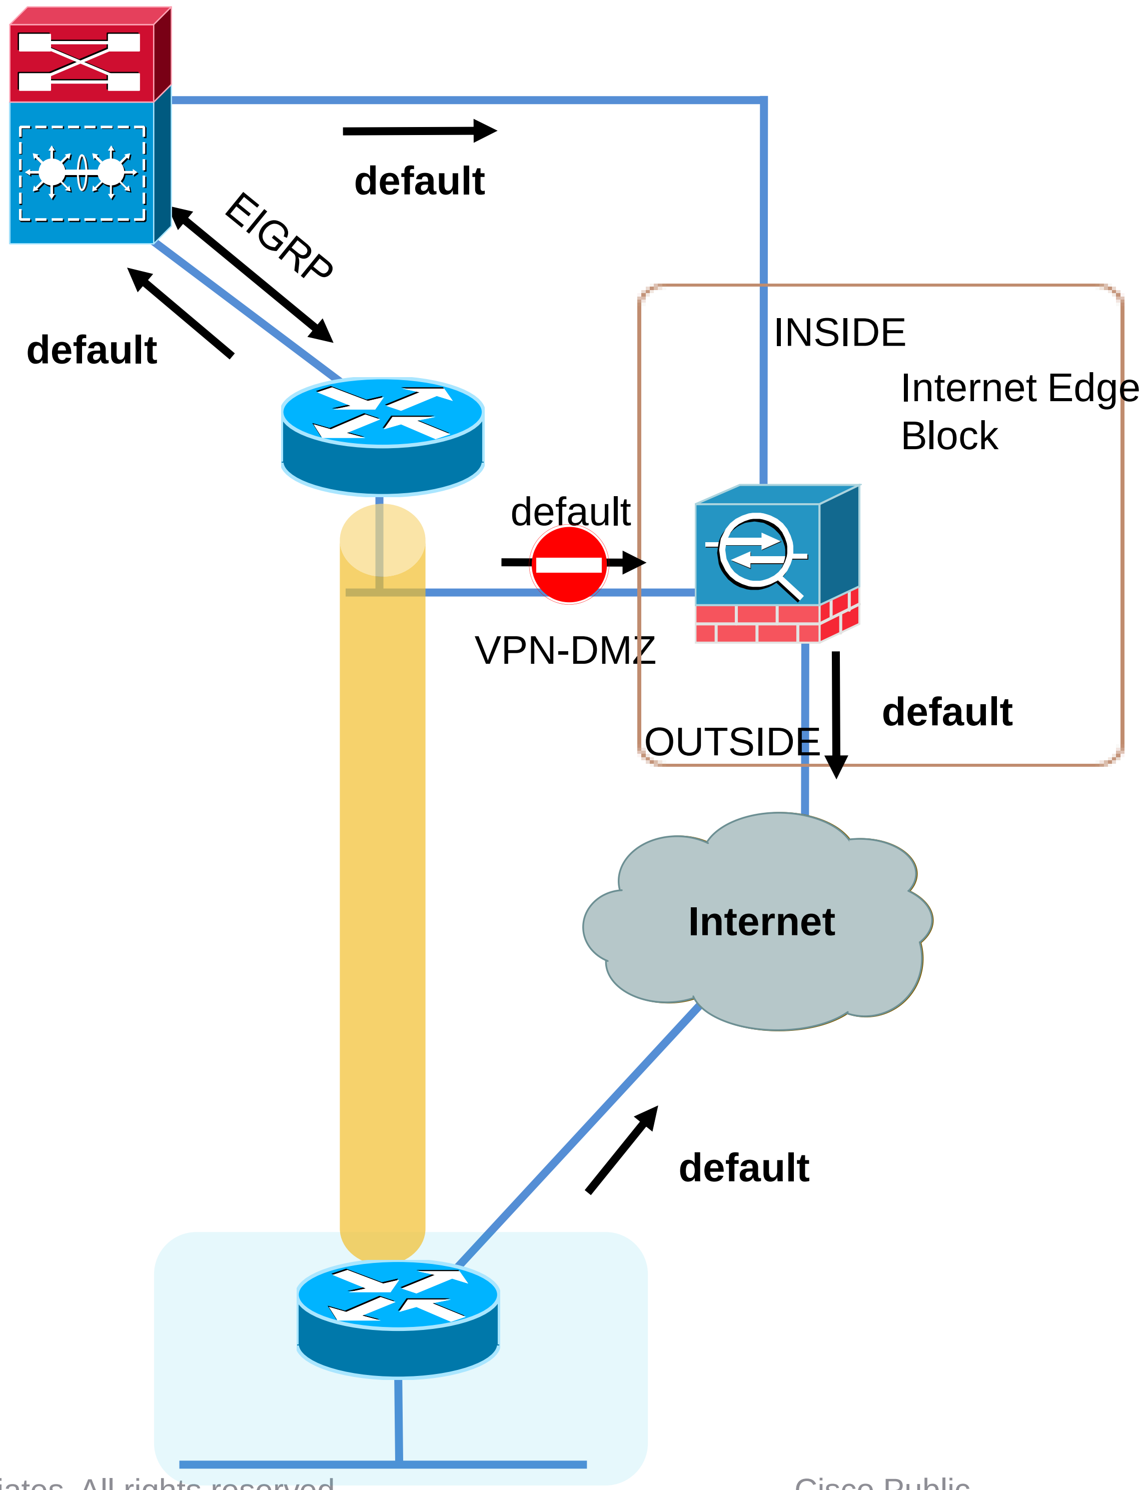

DMVPN Deployment over Internet

- Enable FVRF with DMVPN to separate out the two default routes
- The RED-VRF contains the default route to VPN-DMZ Interface needed for Tunnel Establishment
- A 2nd default route exist on the Global Routing Table used by the user data traffic to reach Internet
- To prevent split tunnelling the default route is advertised to spokes via Tunnel
- Spoke's tunnel drops due to 2nd default route conflict with the one learned from ISP

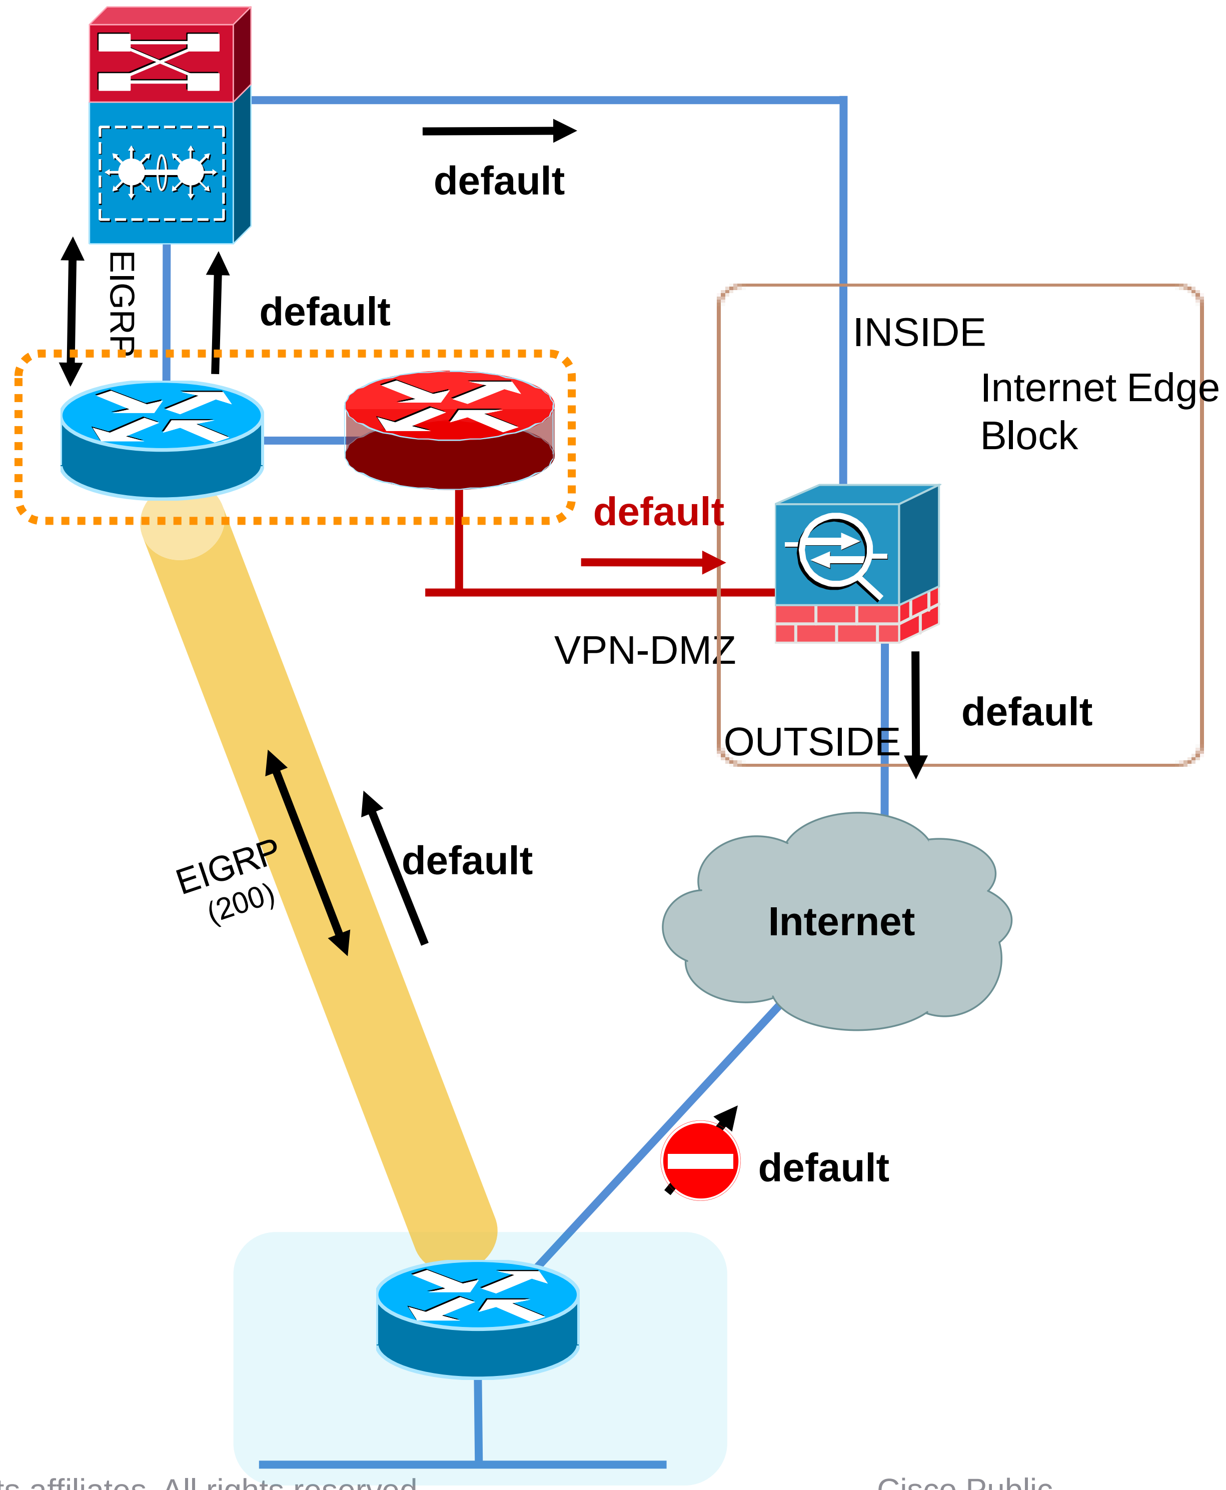

Best Practice – VRF-aware DMVPN

Keeping the Default Routes in Separate VRFs

No Split Tunnelling at Branch location

- Enable FVRF DMVPN on the Spokes
- Allow the ISP learned Default Route in the RED-VRF and used for tunnel establishment
- Global VRF contains Default Route learned via tunnel. User data traffic follow Tunnel to INSIDE interface on firewall
- Allow for consistency for implementing corporate security policy for all users

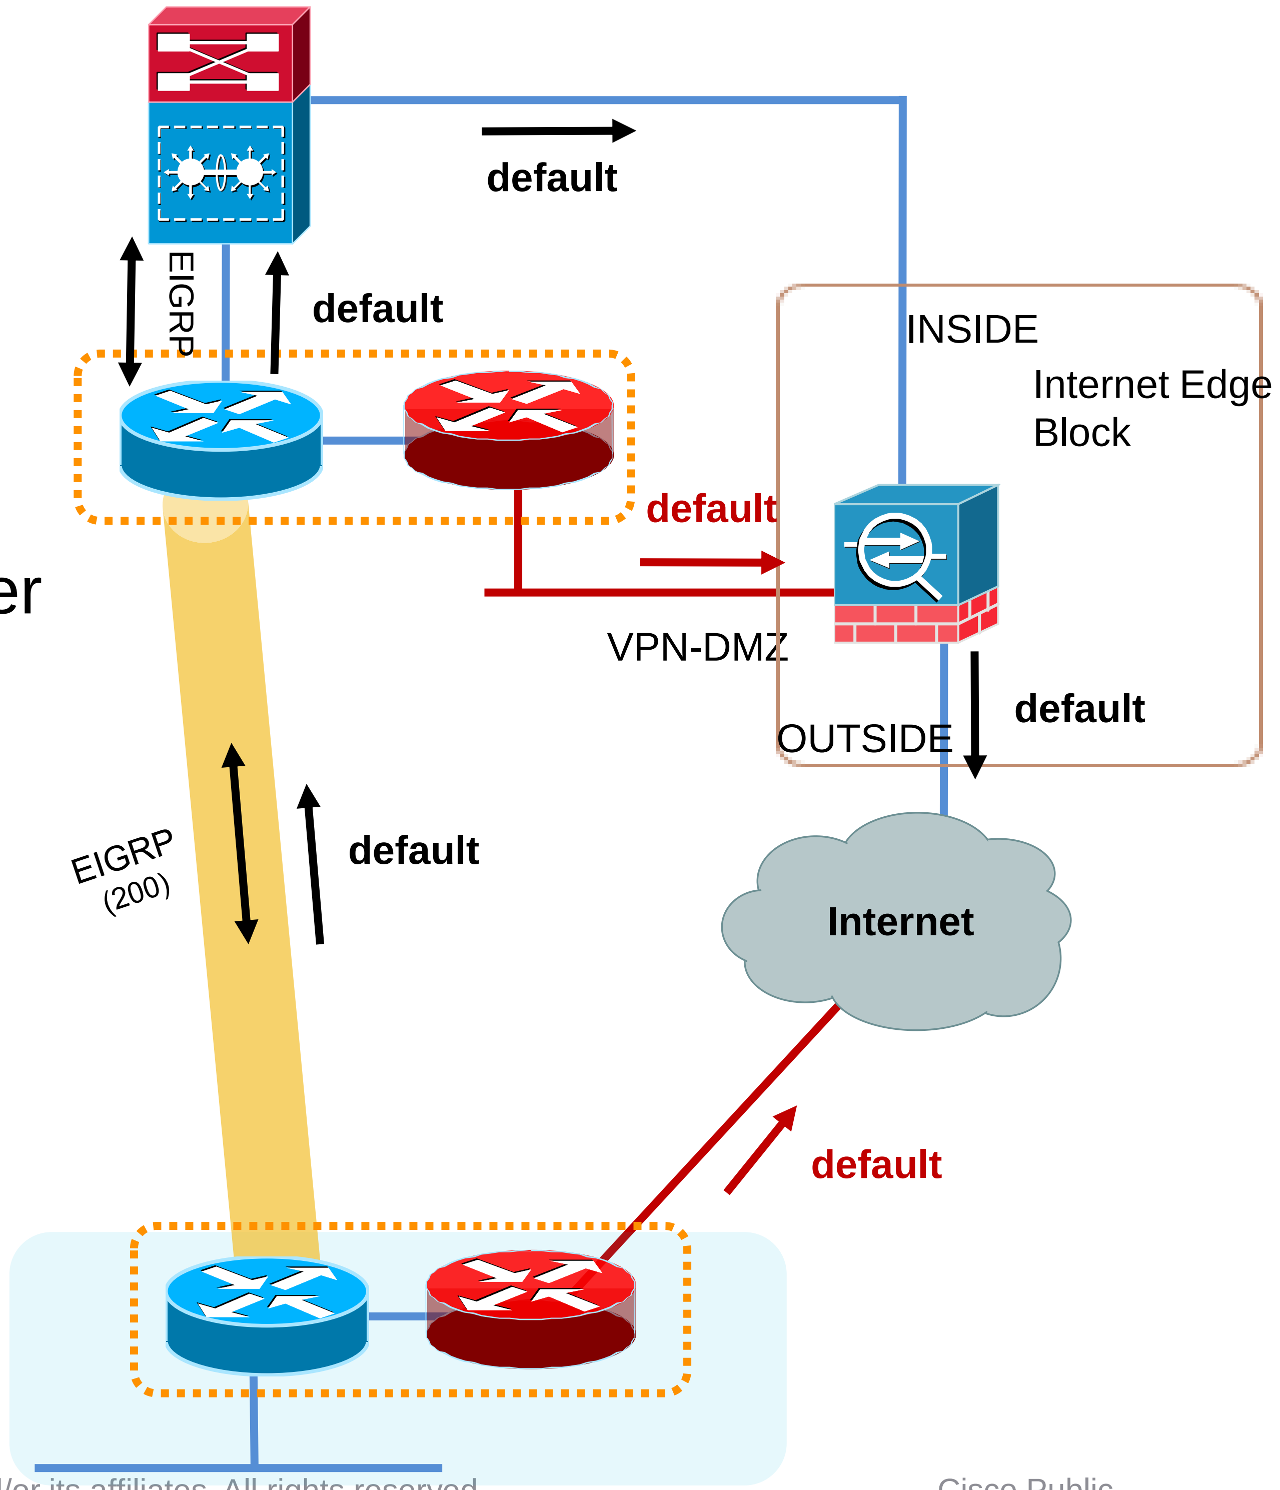

DMVPN and FVRF

Dual Default Routes - Packet Flow

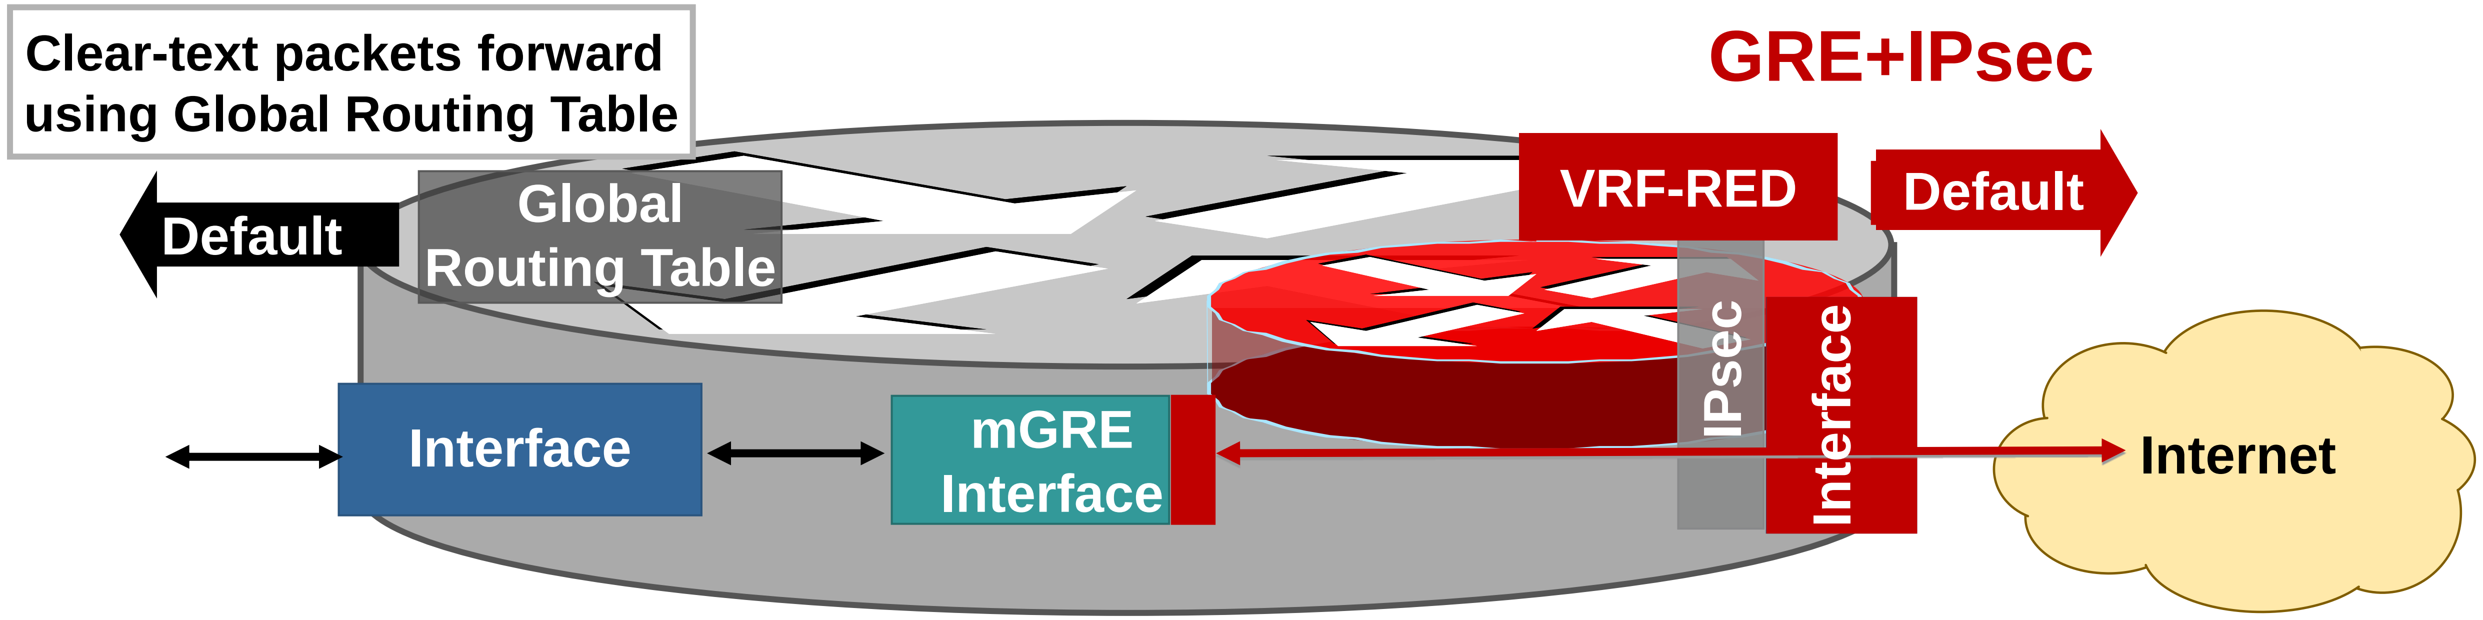

- Based on incoming interface, the IPsec packet is directly associated with VRF
- After decryption the GRE packet is assigned to GRE tunnel in the VRF
- GRE decapsulated clear-text packets forwarded using Global Routing table
- Two routing tables – one global (default) routing table and a separate routing table for VRF

DMVPN and FVRF

Dual Default Routes — Show IP Route Outputs

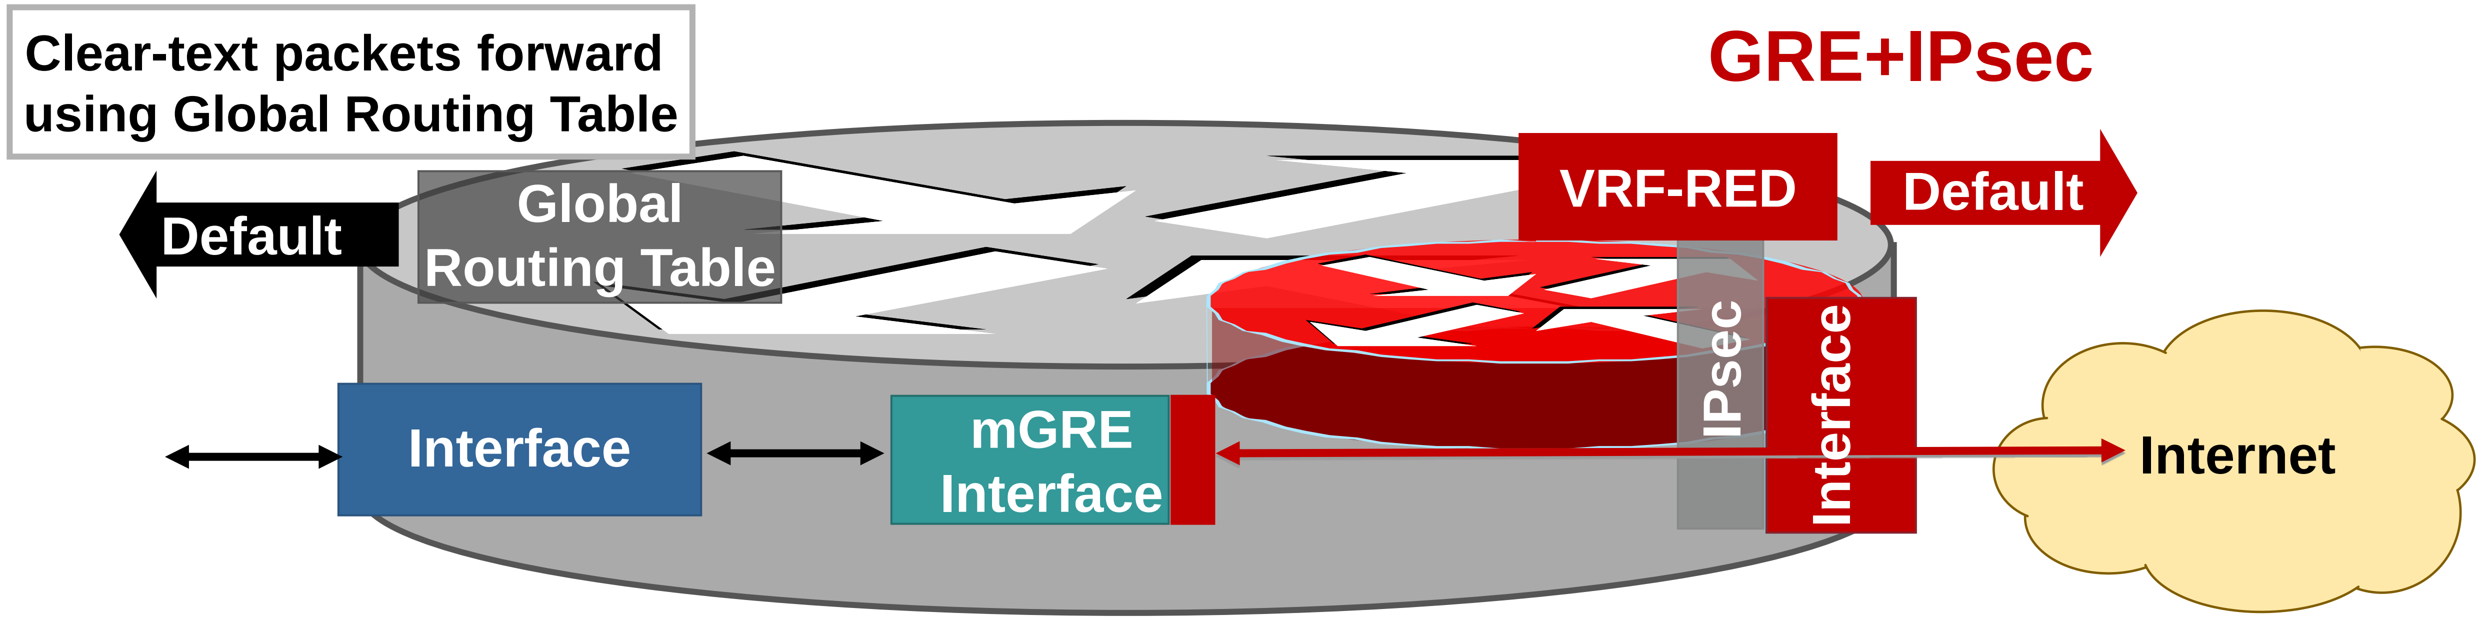


```
bn-vpn-7206-1#sh ip route
Gateway of last resort is 10.4.128.17 to network 0.0.0.0

D*EX  0.0.0.0/0 [170/3328] via 10.4.128.17, 2d22h, Port-channel3
....
```

```
bn-vpn-7206-1#sh ip route vrf RED
Gateway of last resort is 10.4.128.35 to network 0.0.0.0

S*    0.0.0.0/0 [1/0] via 10.4.128.35
....
```

DMVPN and FVRF

Configuration Example

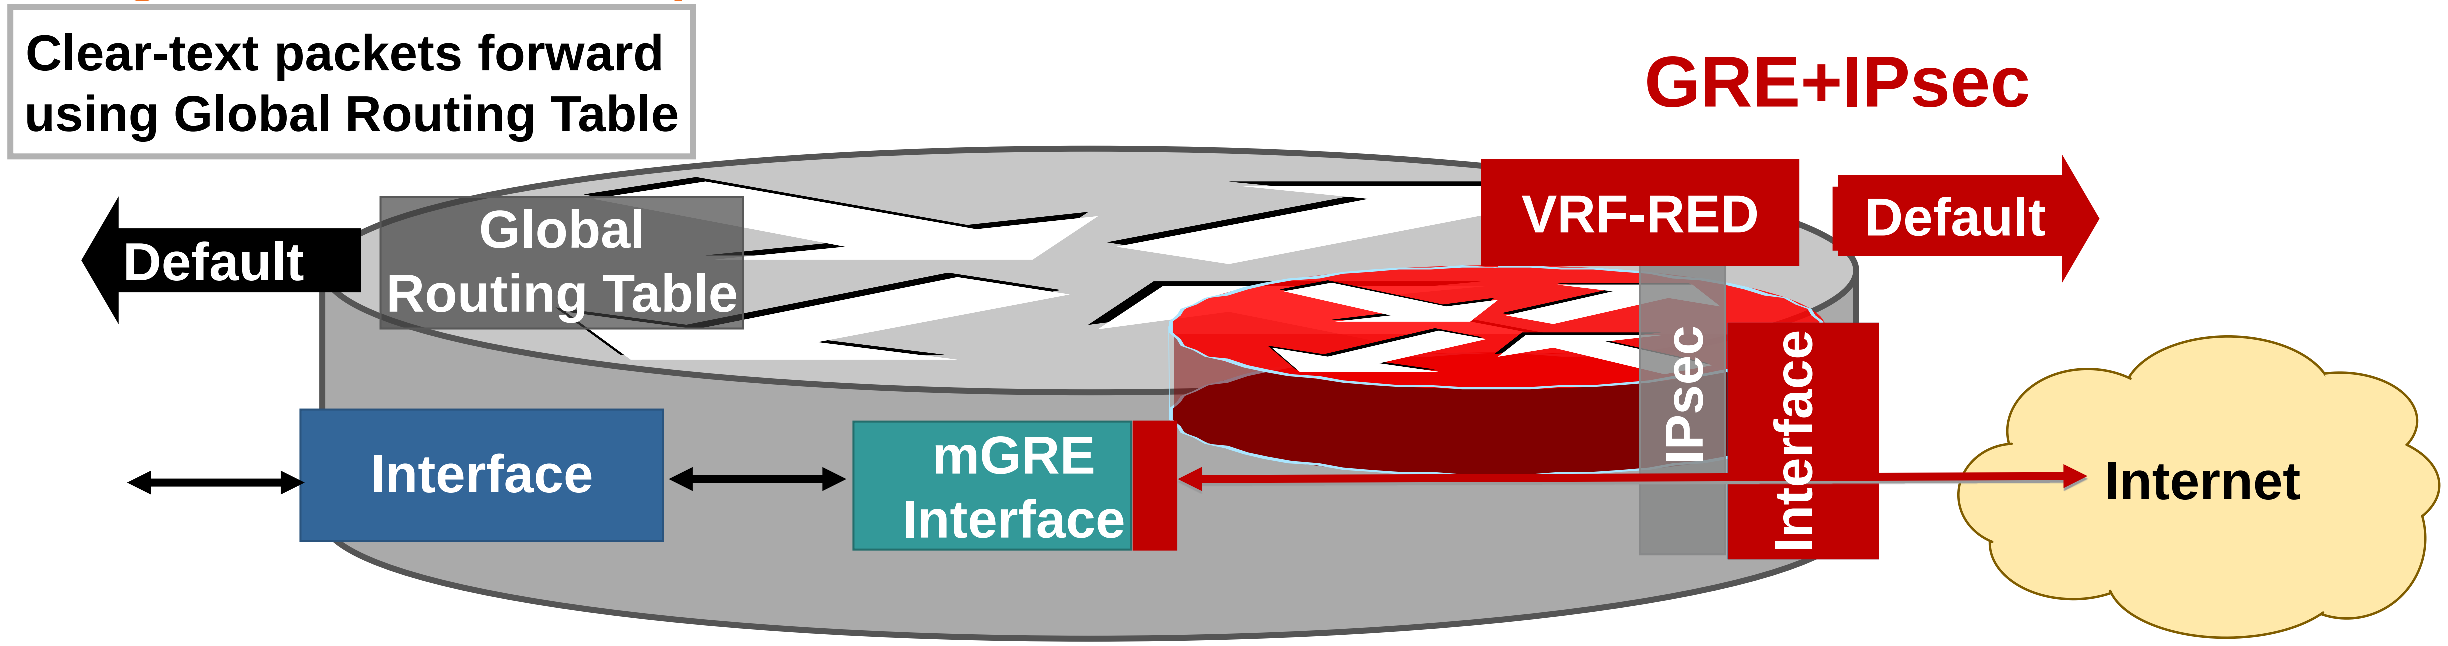


```
ip vrf RED
 rd 65512:1
!
crypto keyring DMVPN-KEYRING vrf RED
  pre-shared-key address 0.0.0.0 0.0.0.0 key cisco123
!
!
crypto isakmp policy 10
 encr aes 256
 authentication pre-share
 group 2
!
crypto isakmp keepalive 30 5
!
crypto isakmp profile FVRF-ISAKMP-RED
   keyring DMVPN-KEYRING
   match identity address 0.0.0.0 RED
!
```

```
interface GigabitEthernet0/1
 ip vrf forwarding RED
 ip address dhcp
!
interface Tunnel10
 ip address 10.4.132.201 255.255.254.0
 ….
 tunnel mode gre multipoint
 tunnel vrf RED
 tunnel protection ipsec profile DMVPN-PROFILE
!
router eigrp 200
 network 10.4.132.0 0.0.0.255
 network 10.4.163.0 0.0.0.127
 eigrp router-id 10.4.132.201
```

Best Practices — Enable Dead Peer Detection (DPD)

Informational RFC 3706

- Dead Peer Detection (DPD) is a mechanism for detecting unreachable IKE peers
- Each peer's DPD state is independent of the others
- Without DPD spoke routers will continue to encrypt traffic using old SPI which would be dropped at the hub. May take up to 60 minutes for spokes to reconverge
- Use ISAKMP keepalives on spokes

```
crypto isakmp keepalives <initial>
  <retry>
```

  – ISAKMP invalid-SPI-recovery is not useful with DMVPN
  – ISAKMP keepalive timeout should be greater than routing protocol hellos
- Not recommended for Hub routers – may cause an increase of CPU overhead with large number of peers

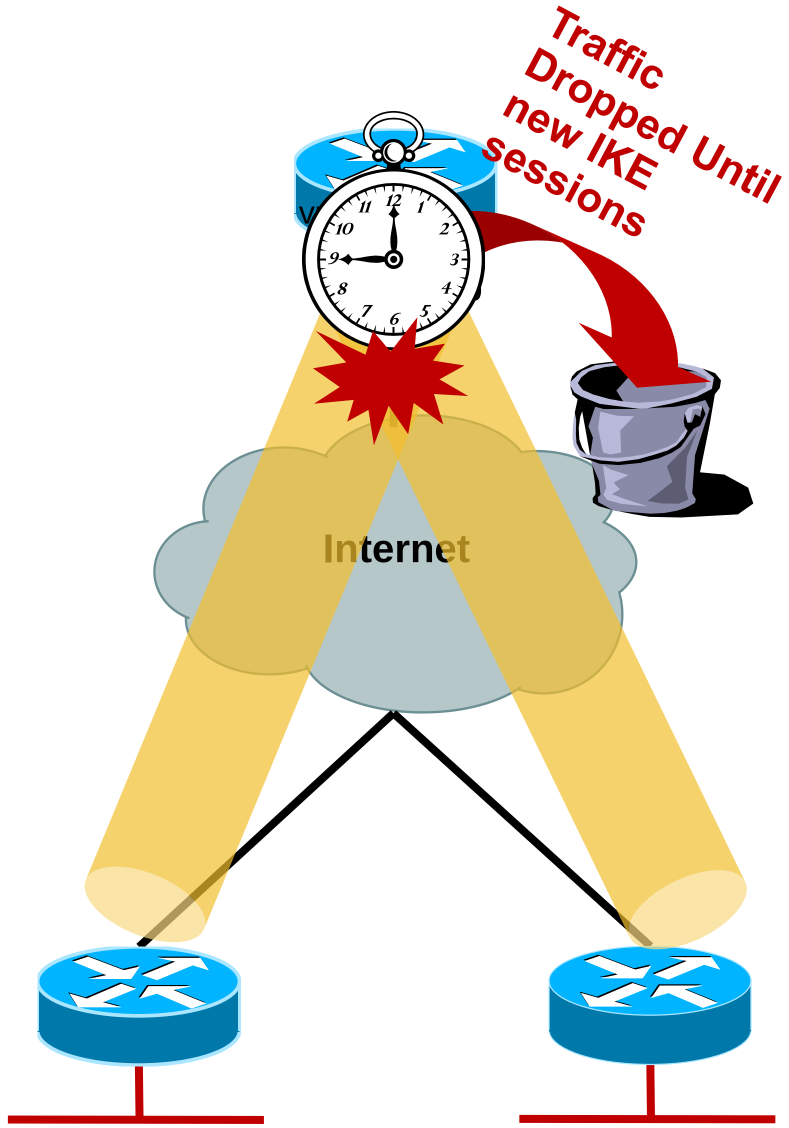

DMVPN Internet Deployment

Dynamic IP Address Assignment on the Spokes

- Spokes are receiving dynamic address assignment from the ISP
- Spoke reboots and receive a new IP address from the ISP, VPN session is established but no traffic passes
- Following error message appears on the spoke

```
"%NHRP-3-PAKREPLY: Receive Registration
Reply packet with error - unique address
registered already(14)"
```

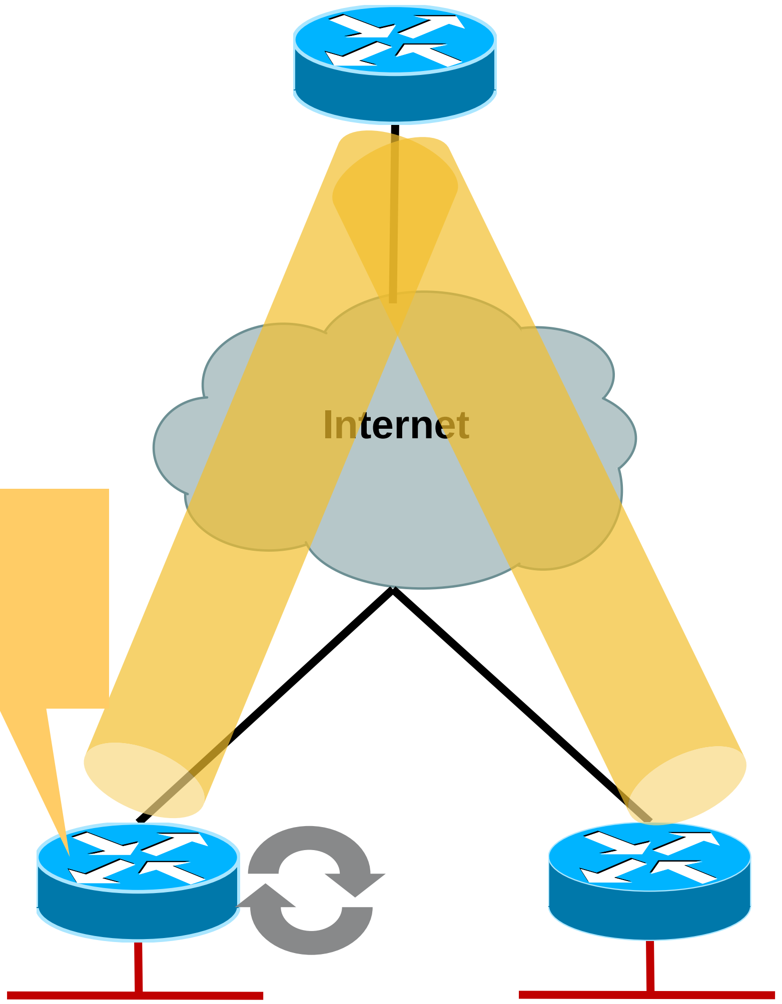

- Hub router (NHS) reject registration attempts for the same private address that uses a different NBMA address
- To resolve this issue, configure following command on spoke routers –
  ip nhrp registration no-unique

Best Practices —

Avoid Fragmentation with IPSec VPN

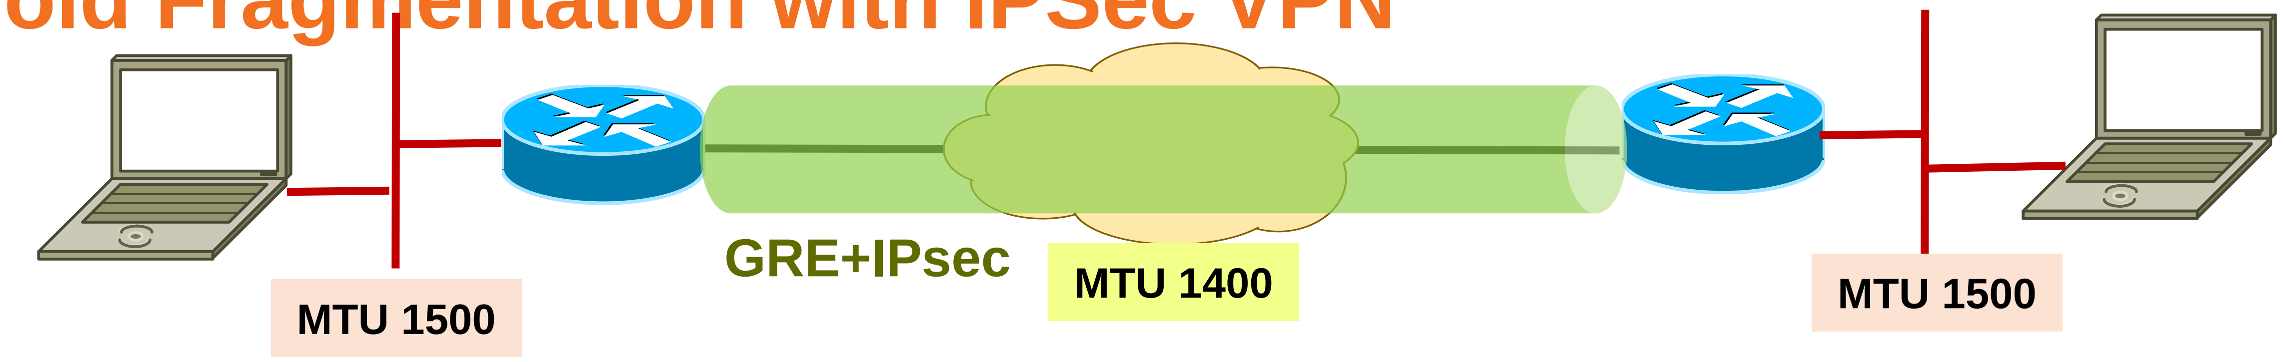

Tunnel Setting (AES256+SHA)	Minimum MTU	Recommended MTU
GRE/IPSec (Tunnel Mode)	1414 bytes	1400 bytes
GRE/IPSec (Transport Mode)	1434 bytes	1400 bytes

- IP fragmentation will cause CPU and memory overhead and resulting in lowering throughput performance
- When one fragment of a datagram is dropped, the entire original IP datagram will have to be resent
- Use '`mode transport`' on transform-set
  - NHRP needs for NAT support and saves 20 bytes
- Avoid MTU issues with the following best practices
  - `ip mtu 1400`
  - `ip tcp adjust-mss 1360`

Best Practices — Multicast over DMVPN

- By default router uses OIL to correlate multicast group join to interface
- This causes problem when hub is connected to multiple spokes over NBMA network
- Any spoke that leaves a multicast group would case all the spokes to be pruned off the multicast group
- Enable PIM NBMA mode under tunnel interface on hubs and spokes

```
ip pim nbma-mode
```

  - Allows the router to track multicast joins based on IP address instead of interface
  - Applies only to PIM sparse-mode
- Router treats NBMA network as a collection of point-to-point circuits, allowing remote sites to be pruned off traffic flows

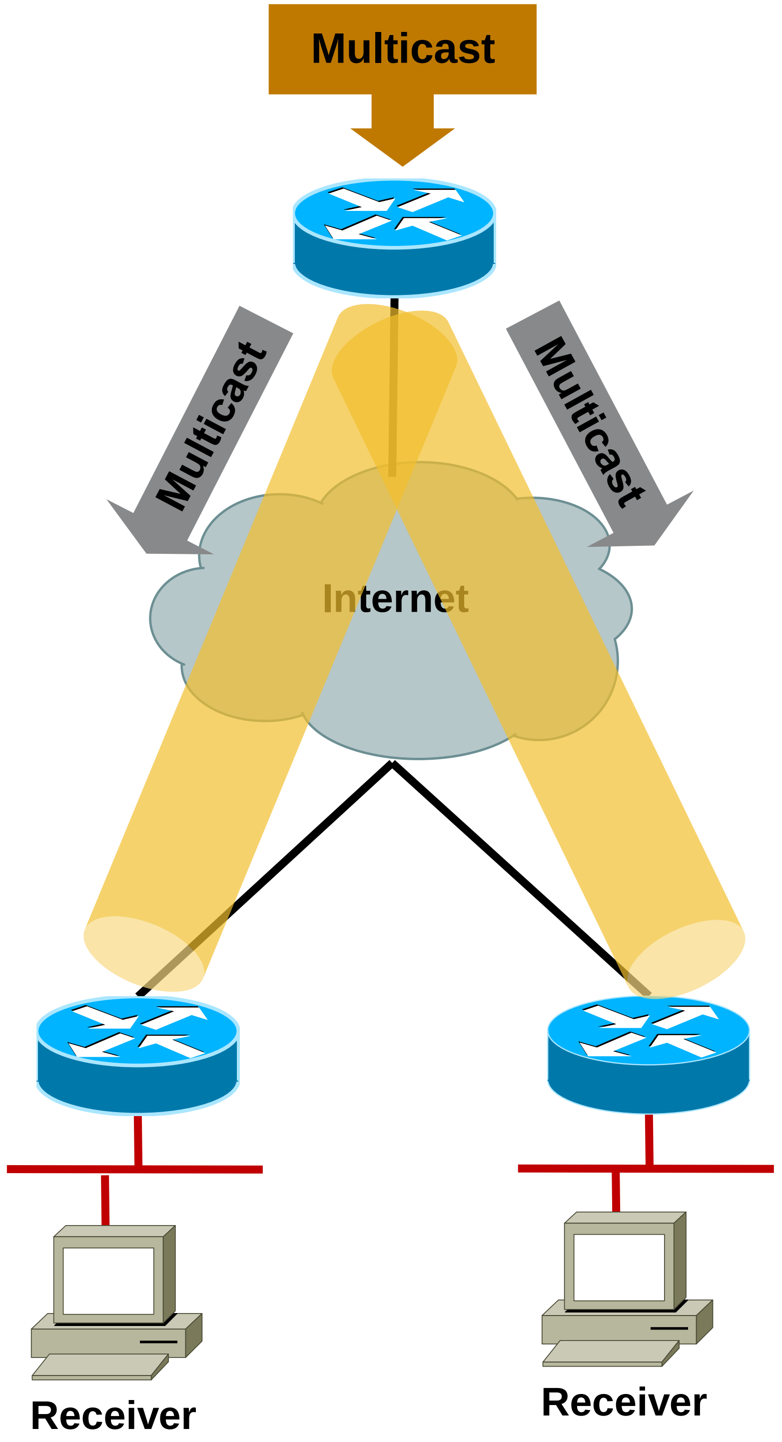

Best Practices — Multicast over DMVPN

- By default router uses OIL to correlate multicast group join to interface
- This causes problem when hub is connected to multiple spokes over NBMA network
- Any spoke that leaves a multicast group would case all the spokes to be pruned off the multicast group
- Enable PIM NBMA mode under tunnel interface on hubs and spokes

```
ip pim nbma-mode
```

  - Allows the router to track multicast joins based on IP address instead of interface
  - Applies only to PIM sparse-mode
- Router treats NBMA network as a collection of point-to-point circuits, allowing remote sites to be pruned off traffic flows

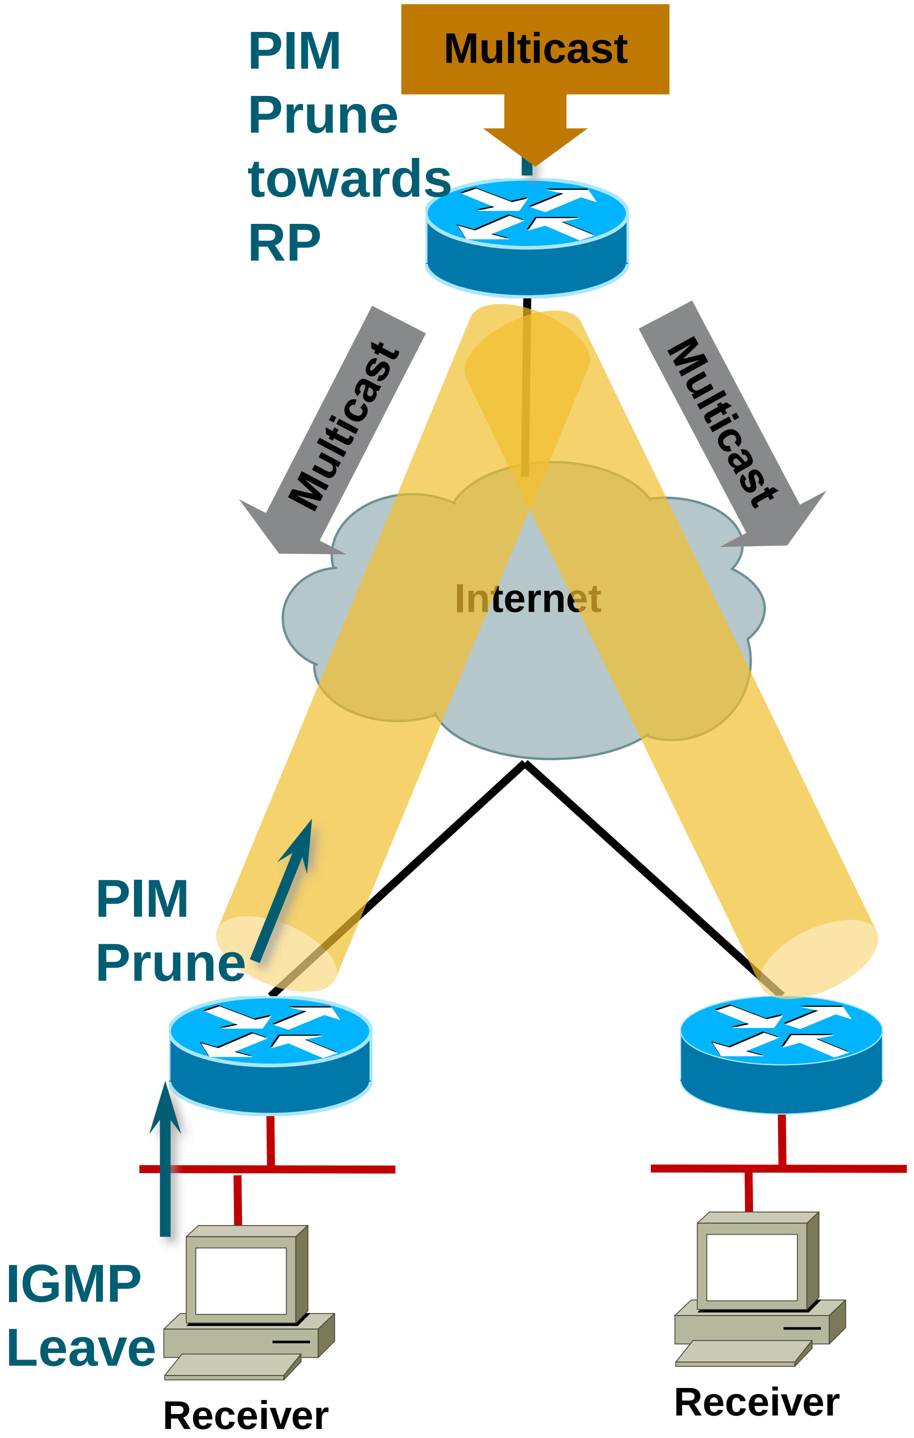

Deploying WCCP with DMVPN Phase 3

- DMVPN deployments with WCCP, WCCP intercept is configured on the tunnels
- Any packet traveling from spoke-to-spoke, on reaching the tunnel, is intercepted by WCCP and sent to the WAE
- This breaks the NHRP condition to send the redirect.
- No dynamic tunnels are established

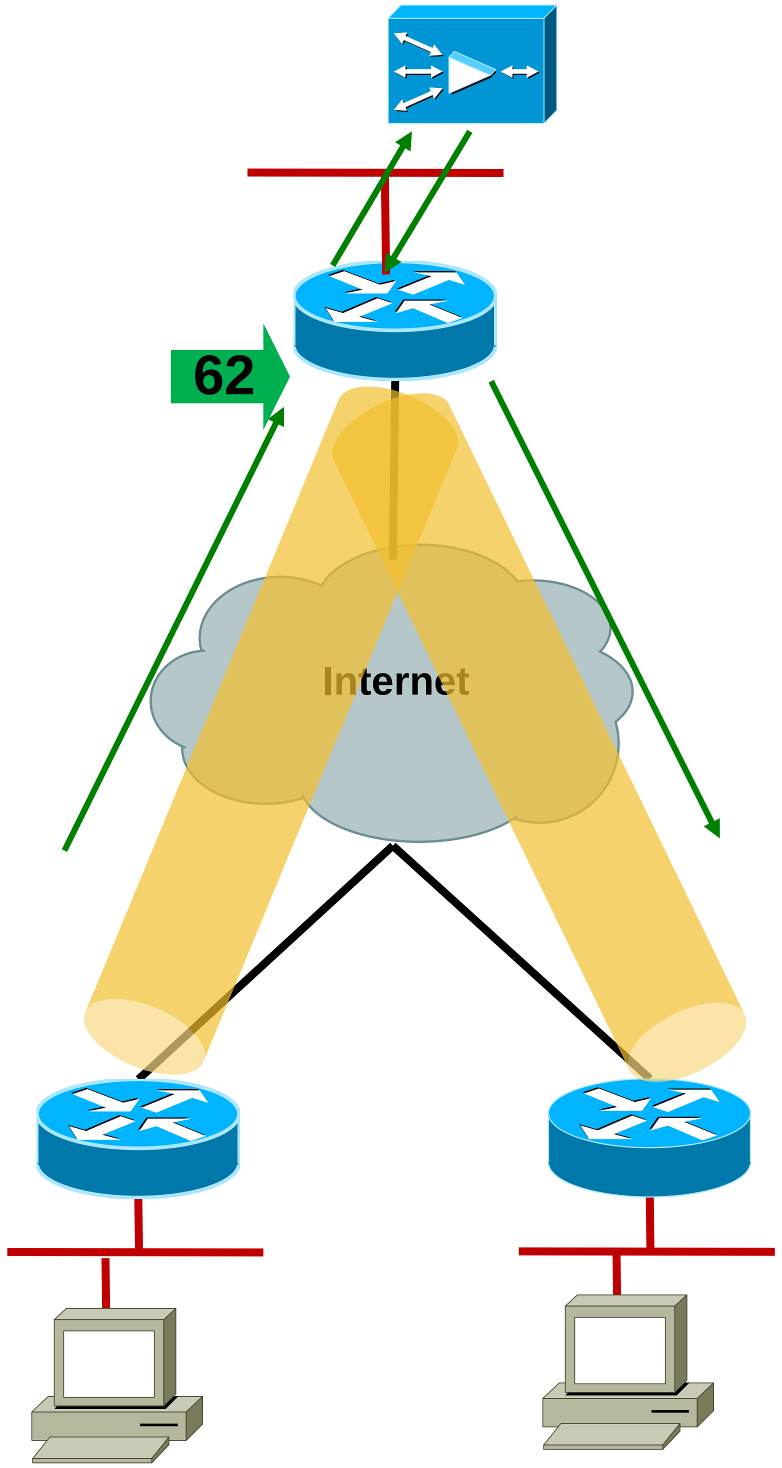

Deploying WCCP with DMVPN Phase 3

- Remove the WCCP intercept on the tunnel interface on the hub and configure it on its LAN interface.
  - `ip wccp 62 redirect out`
- Initial spoke-to-spoke traffic hairpin through hub without being intercepted by WCCP
- Hub creates NHRP redirect message to spoke allows for dynamic spoke-to-spoke tunnel setup

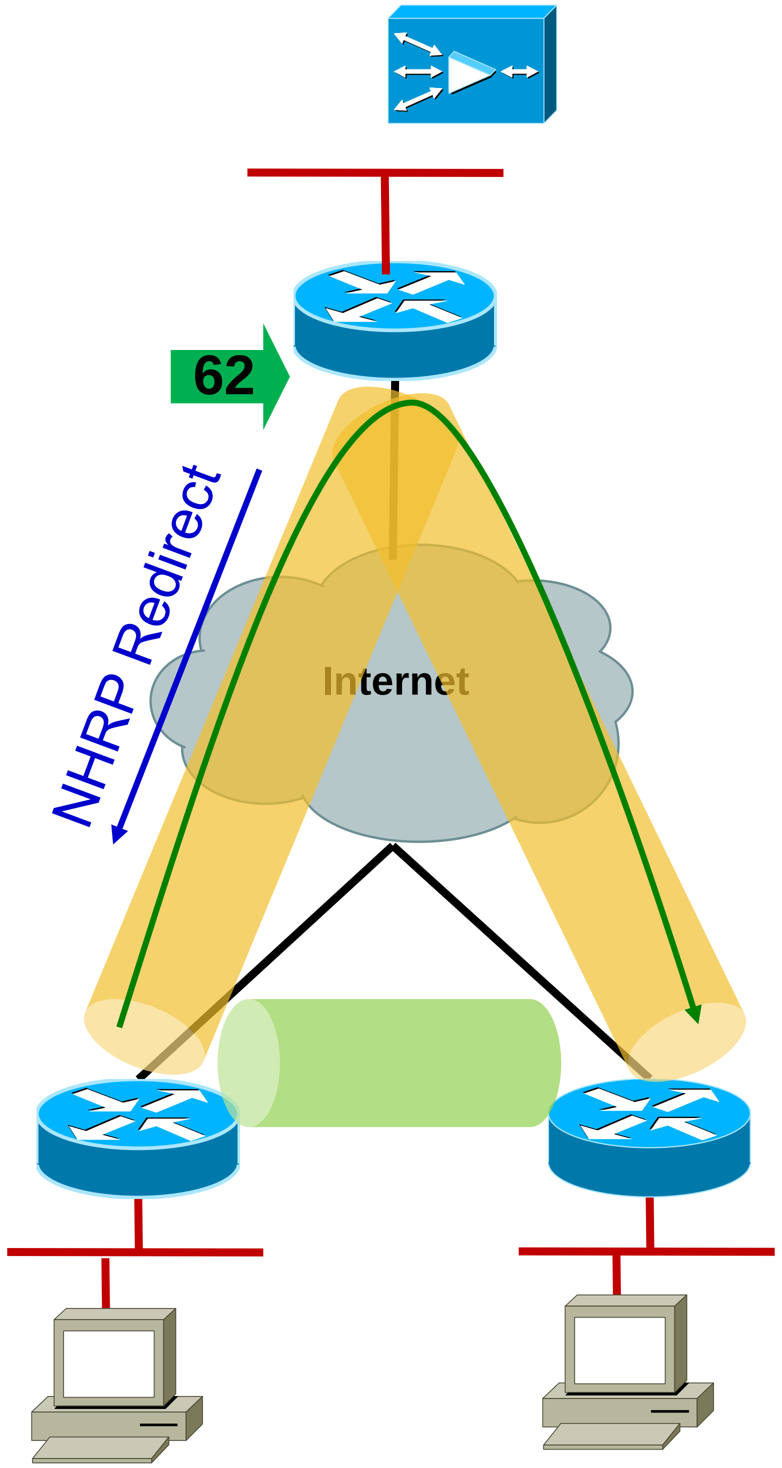

Agenda

- WAN Technologies & Solutions
  - WAN Transport Technologies
  - WAN Overlay Technologies
  - WAN Optimisation
  - Wide Area Network Quality of Service
- WAN Architecture Design Considerations
  - WAN Design and Best Practices
  - Secure WAN Communication with GETVPN
  - DMVPN Over Internet Deployment
- Summary

Key Takeaways

- Understand how WAN characteristics can affect your applications
  - Bandwidth, latency, loss
- Dual carrier designs can provide resiliency but have unique design considerations
- A QoS-enabled, highly-available network infrastructure is the foundation layer of the WAN architecture
- Encryption is a foundation component of all WAN designs and can be deployed transparently
- Understand the how to apply WCCPv2 in the branch network to enable WAN optimisation appliances.

Final Thoughts

- Get hands-on experience with the Walk-in Labs located in World of Solutions, booth 1042
- Come see demos of many key solutions and products in the main Cisco booth 2924
- Visit www.ciscoLive365.com after the event for updated PDFs, on-demand session videos, networking, and more!
- Follow Cisco Live! using social media:
  - Facebook: https://www.facebook.com/ciscoliveus
  - Twitter: https://twitter.com/#!/CiscoLive
  - LinkedIn Group: http://linkd.in/CiscoLI

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