Techniques for Enterprise Network Virtualization

BRKCAM-3002

Victor Moreno

Cisco Networkers 2007

HOUSEKEEPING

- We value your feedback, don’t forget to complete your online session evaluations after each session and complete the Overall Conference Evaluation which will be available online from Friday.
- Visit the World of Solutions on Level -01!
- Please remember this is a ‘No Smoking’ venue!
- Please switch off your mobile phones!
- Please remember to wear your badge at all times including the Party!
- Do you have a question? Feel free to ask them during the Q&A section or write your question on the Question form given to you and hand it to the Room Monitor when you see them holding up the Q&A sign.

Agenda

- Problem Definition
- Campus Virtualization Alternatives
- WAN Extensibility
- Shared Services and Inter-VPN Communication
- Data Center Integration

Network Virtualization
Simplicity and Agility in Managing Resources

- What is virtualization?
  - A *logical* rather than physical view of data, storage, network, and other resources presented independently of location, packaging, or capacity
  - One Network Supports many physical resources: simplifies operations, reduces cost
  - One Network Consolidates all types of resources for increased flexibility (data, voice, video, storage)
- Benefit: flexible configuration and management of all infrastructure resources to reduce costs and increase agility

Virtualization Required Across All Industries

Manufacturing	Healthcare	Government
		
Automation of Production Plants Integration of Sales Sites, Suppliers and Partners Video Surveillance	Individual "Hotel" Services for Patients Isolated medical Networks for Records, Services	Shared Buildings and Facilities across different Agencies: Police Fire Department Tax Administration

What is Network Virtualization?

- Virtualization: 1 to Many or Many to 1
- One network supports many virtual networks

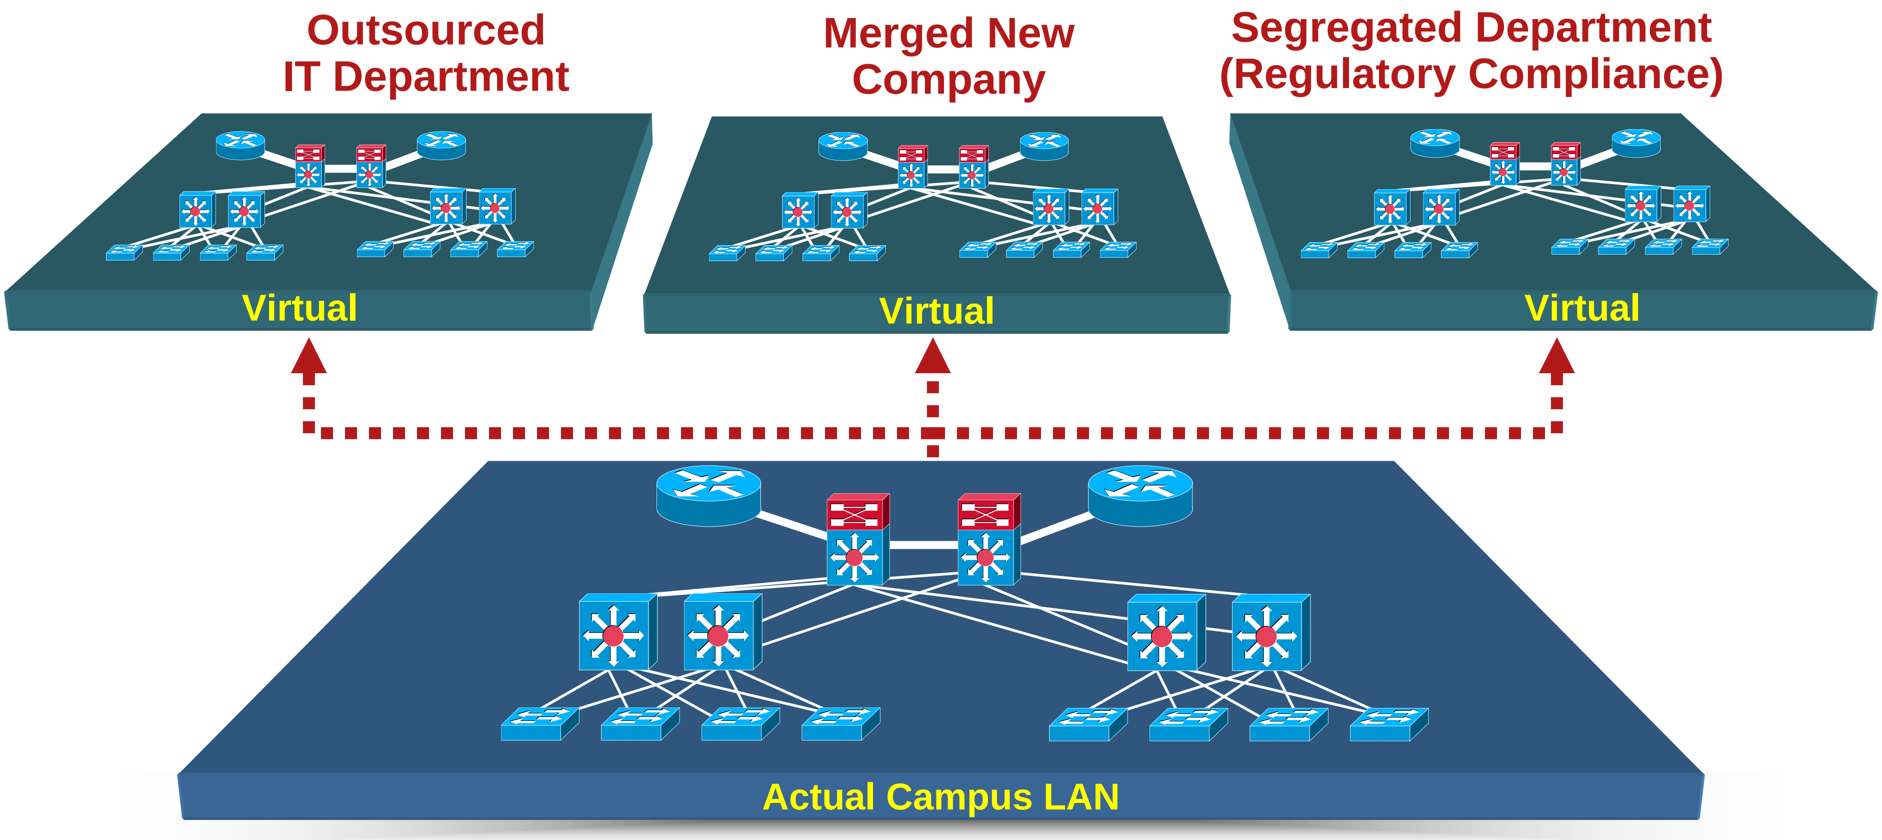

Problem Definition: Requirements

Broad range: From simple to elaborate

- Simple: Create segments for guest access and/or NAC quarantine
- Elaborate: IT department as a “network service provider”
  - Provide a private network for each ‘customer’
  - Leverage a shared infrastructure
  - Customized Routing per virtual network
  - Scalability and simplicity
  - Minimize operational overhead → Agility, manageability
  - Centralize security policies and access to shared services
  - Virtual Networks extensible over the WAN
- IT departments: From cost centers to revenue centers?
  - Potential to enhance enterprise business processes

Recommended Network Design

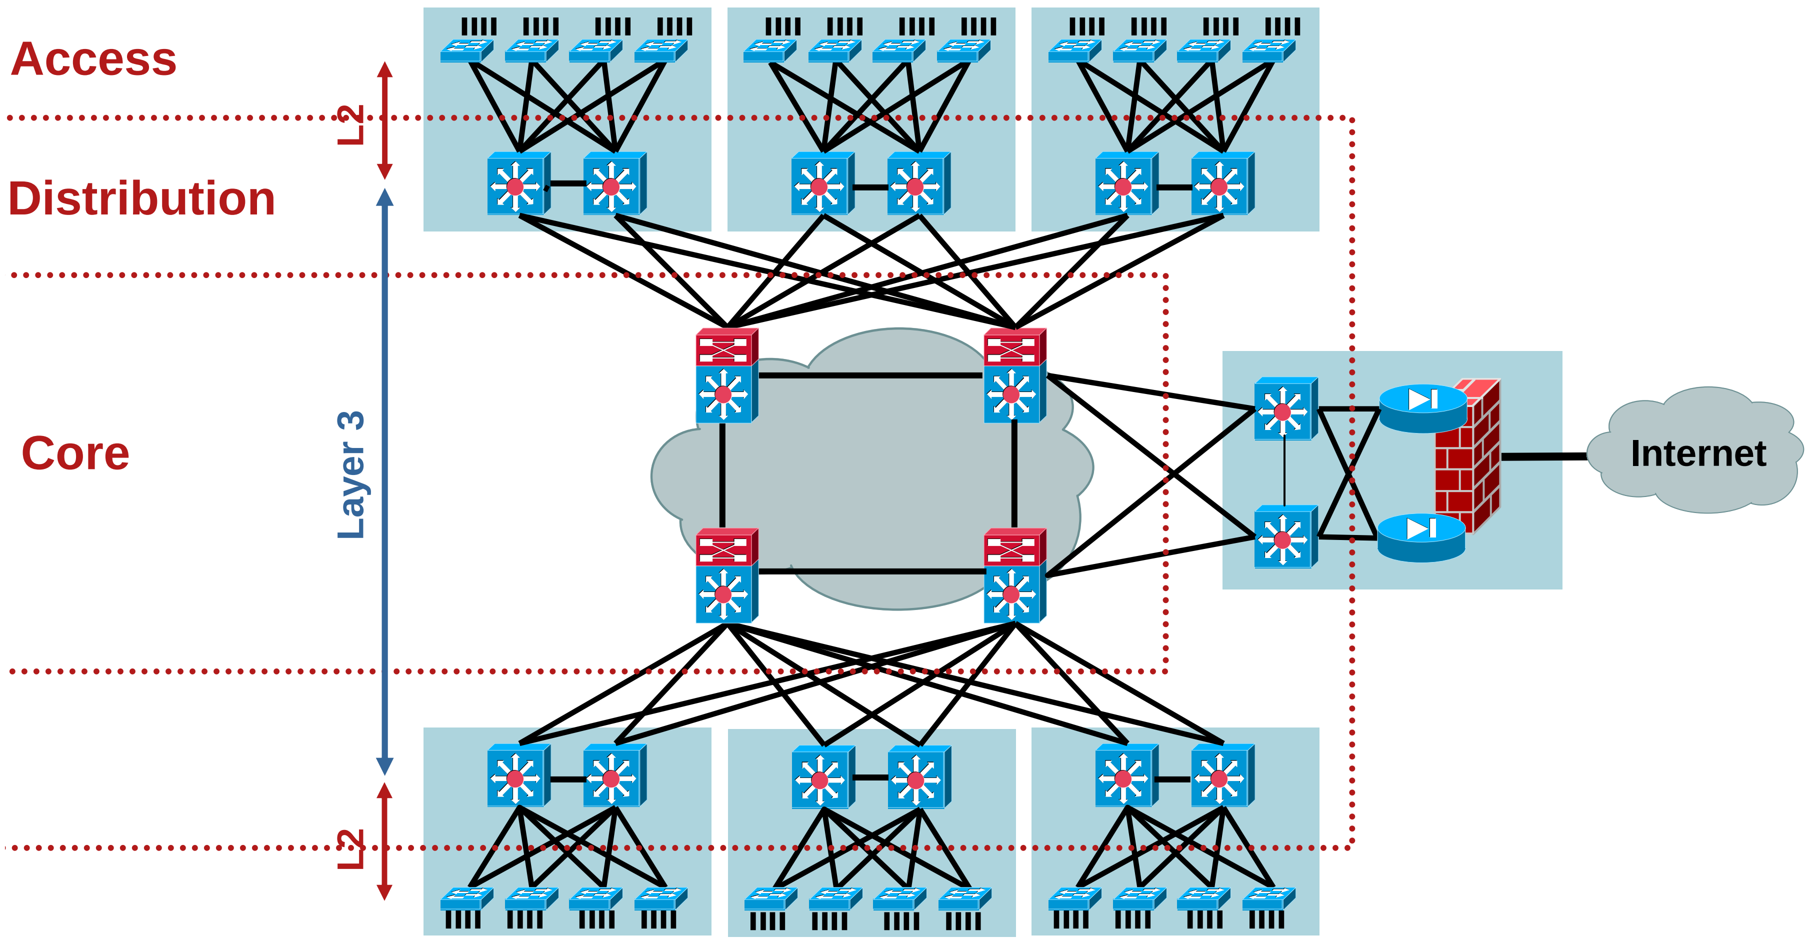

Modular, Hierarchical → Scalable yet Not Virtualized

Anatomy of a Virtualized Network

Preserve Hierarchy and Scalability

- Virtualized devices
- Virtualized services
- Virtualized data paths

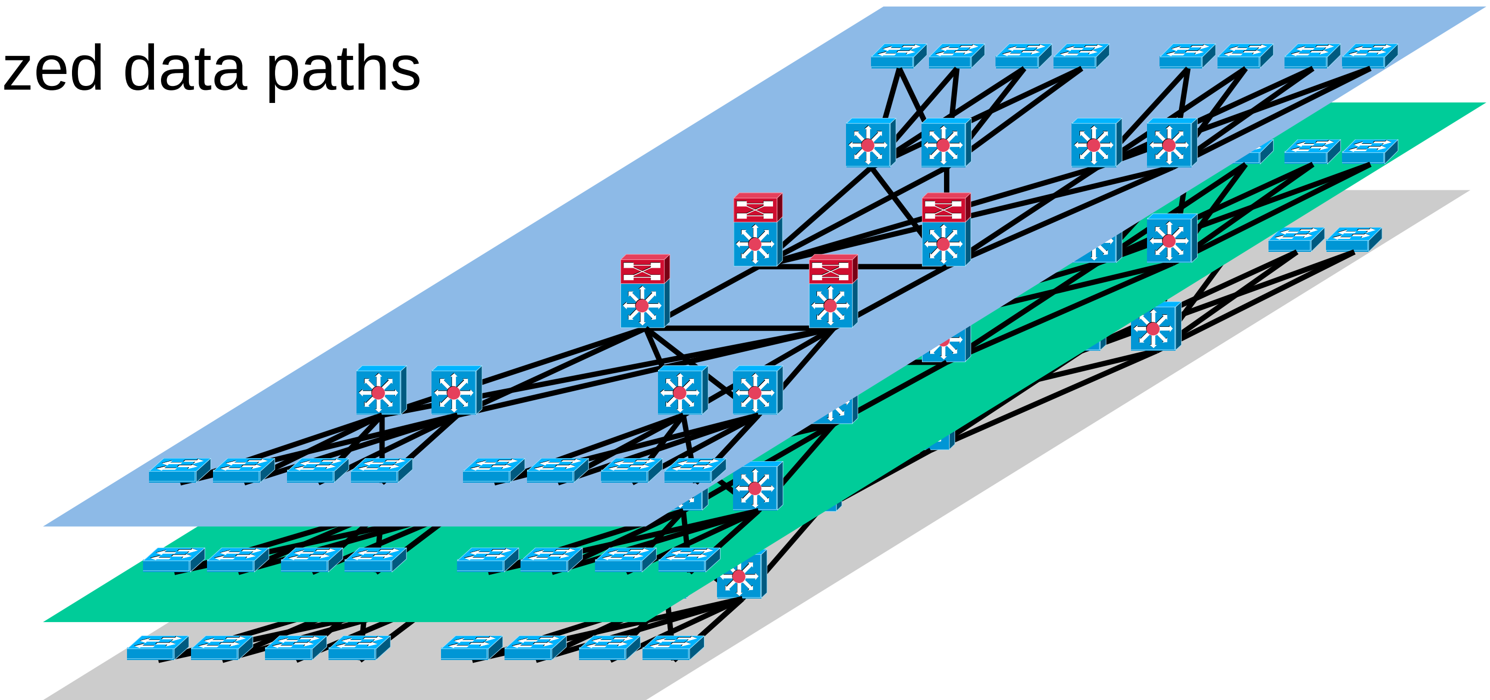

Virtualized Network Devices

- Switch virtualization—VLANs
- Router virtualization—Virtual Routing/ Forwarding (VRFs)

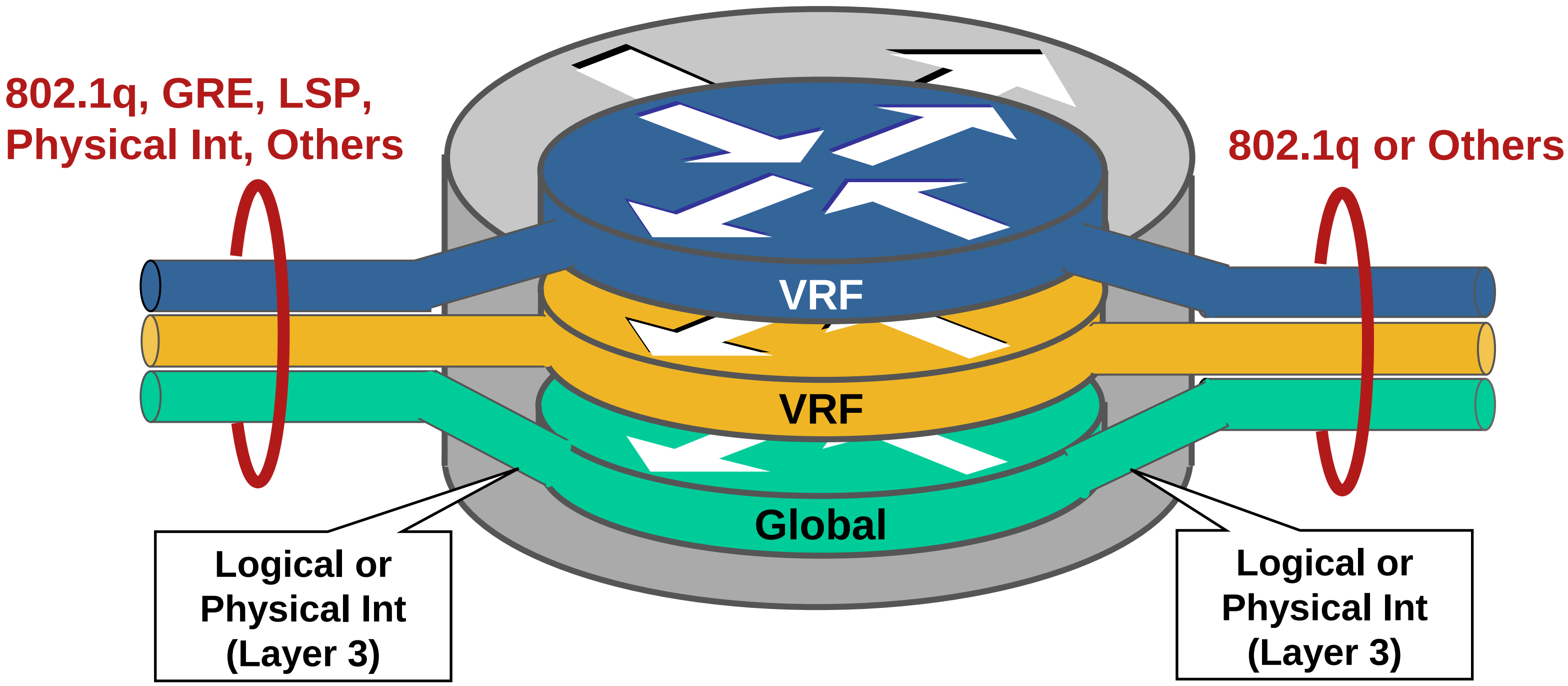

Virtualized Network Services

- Network services include:
  - Firewalls
  - Intrusion detection systems
  - VPN service modules (IPsec and SSL)
  - Load balancers
  - DHCP servers
  - DNS servers
- Levels of virtualization
  - VLAN awareness
  - L2-7 virtualization
    - E.g. firewall contexts, DHCP VPN awareness (CNR—Cisco Network Registrar)
  - Multi-service integration—Application Control Engine (ACE)

Data Path Virtualization

- Tags
  - 802.1q
  - Others (DSCP, CTS)
- Virtual circuits
  - ATM
  - Frame Relay
  - AToM L2 Circuits

Single Hop Data Path Virtualization

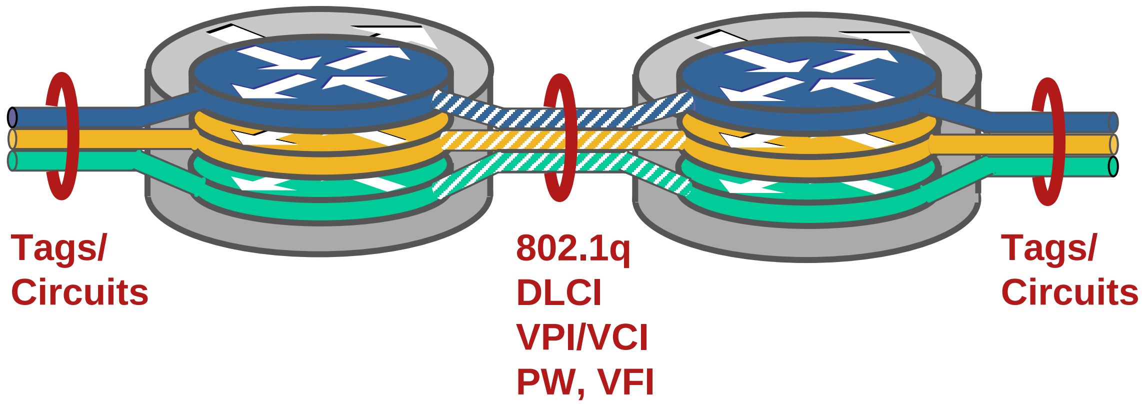

Multi-Hop Data Path Virtualization

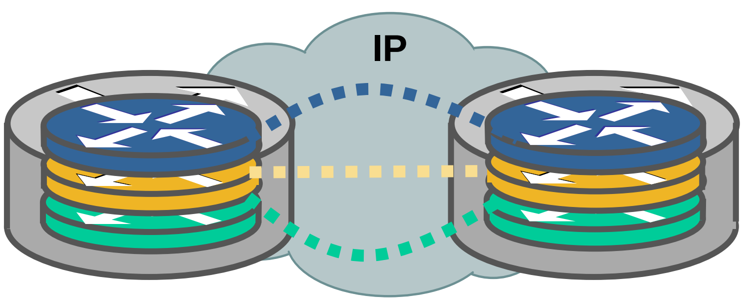

- Tunnels (connection oriented)
  - GRE/mGRE
  - L2TPv3
  - Label Switched Paths—LSP (MPLS)

Centralized Policies with Virtual Services

- Users assigned to VLANs
- VLANs mapped to VRFs
- VRFs interconnected
- Services centralized at VPN perimeter

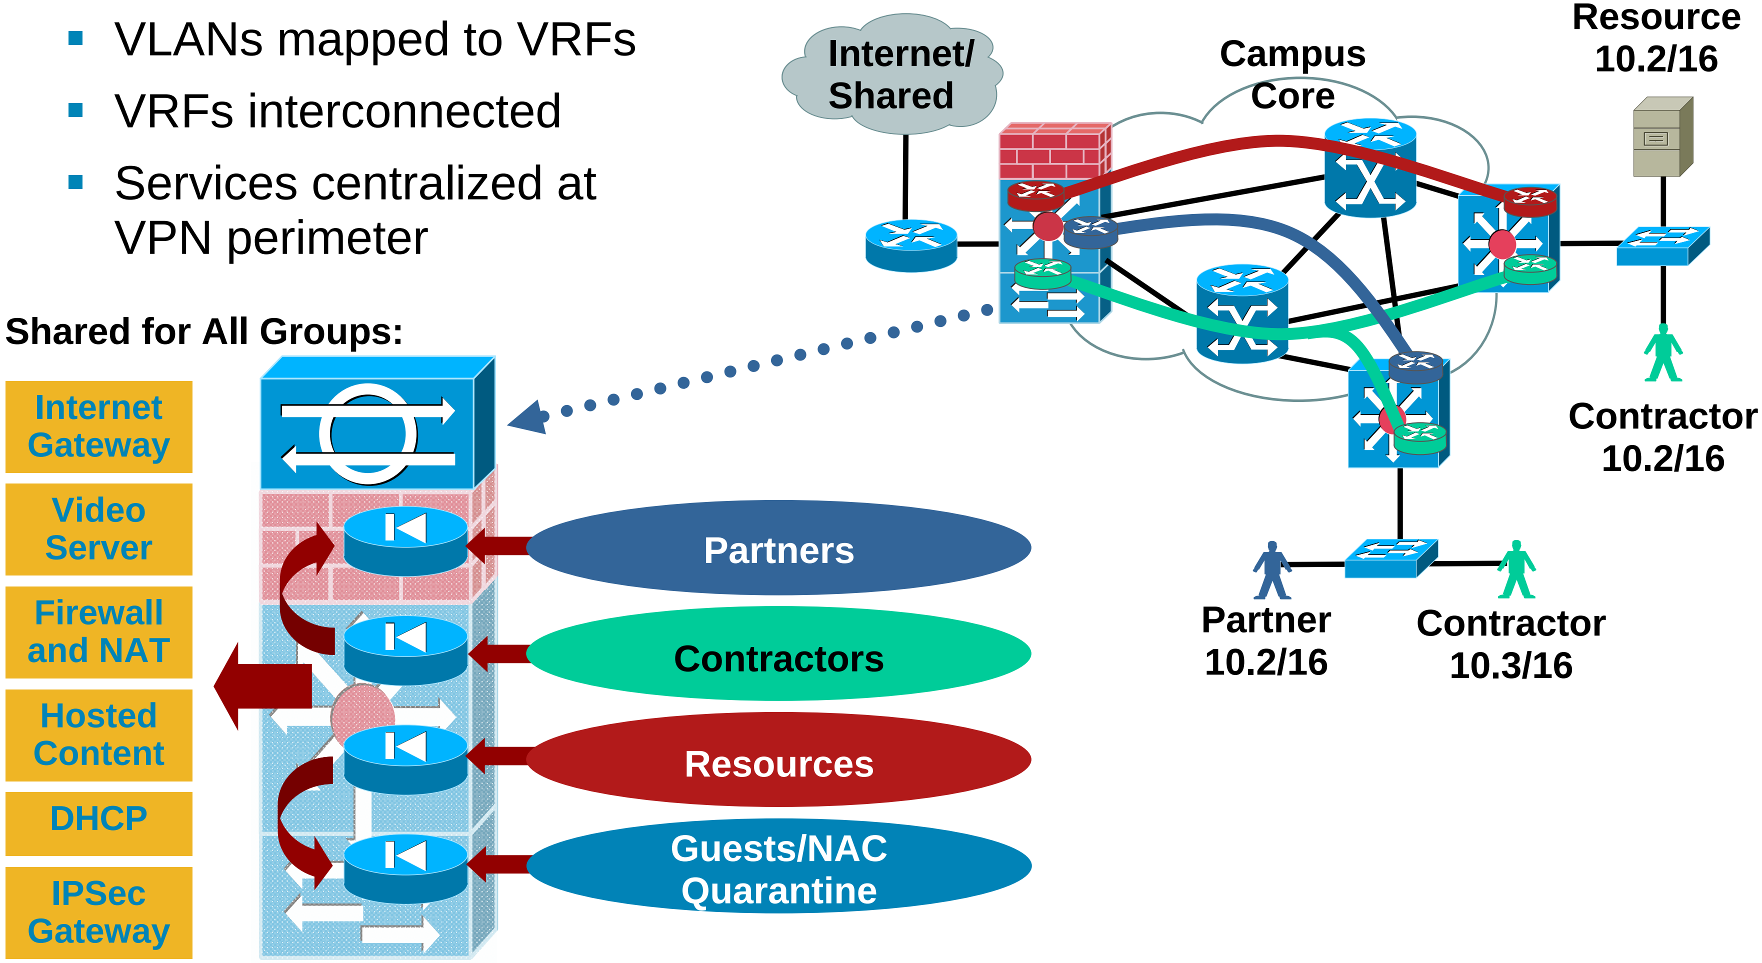

Authentication, Authorization, Posture
Access and Policy Control

- Identity-Based Networking Services (IBNS)
  - Authenticate the user/device based on:
    - 802.1x credentials
    - MAC address
    - Web authentication credentials
  - Authorize appropriate network access → VLAN assignment
- Network Access Control (NAC)
  - Identifies posture compliance of the device
  - Ensures device is not a hazard
  - Quarantines non-compliant devices for remediation → VLAN or ACL assignment

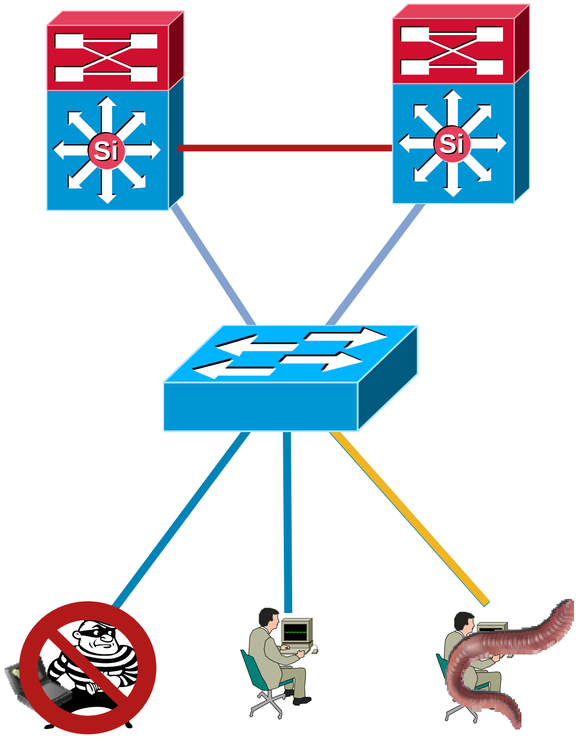

Edge Access Control

For In-Depth Information See Session BRKCAM-2007, Understanding Identity-Based Networking Services, Authentication, and Policy Enforcement

Agenda

- Problem Definition
- Campus Virtualization Alternatives
    1. VLANs
    2. ACLs
    3. VRFs + GRE Tunnel Mesh
    4. RFC2547 VPNs
    5. Hop-by-Hop Multi-VRF
- WAN Extensibility
- Shared Services and Inter-VPN Communication
- Data Center Integration

Enterprise Closed User Groups
End-to-End Virtualized Enterprise

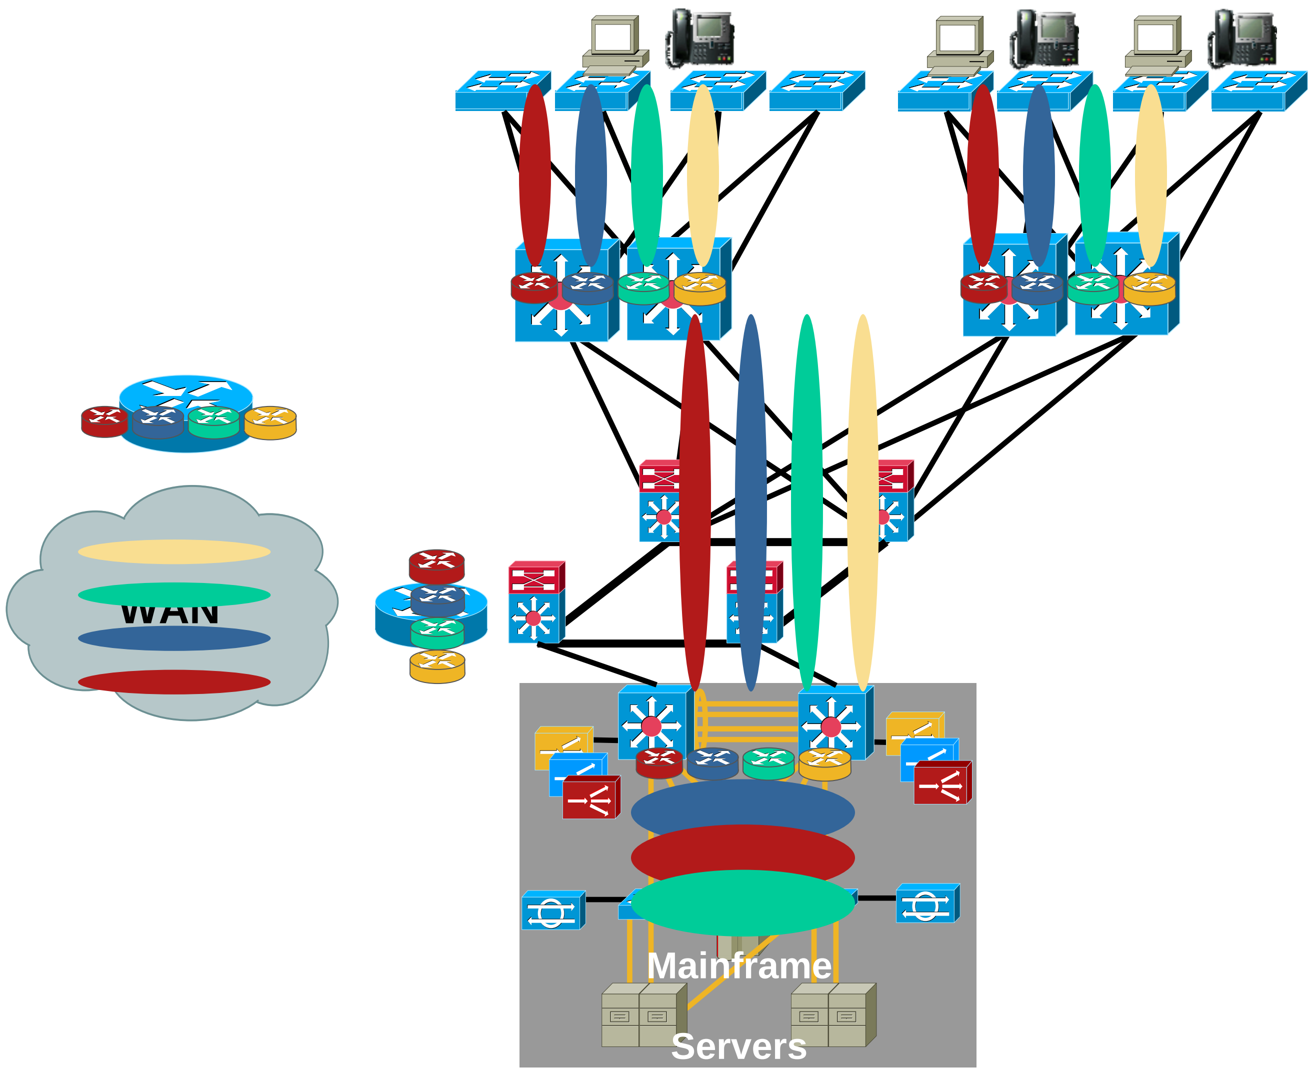

User Identification (per Port or 802.1x)

Per User Role L2 VLANs

L3 VRFs

802.1Q + VRFs
MPLS,
GRE (PBR/VRF),
L2TPv3

Virtualized Services:
FW, CSM

VLANs Partition Server Farms

VLAN: Logical Separation

1

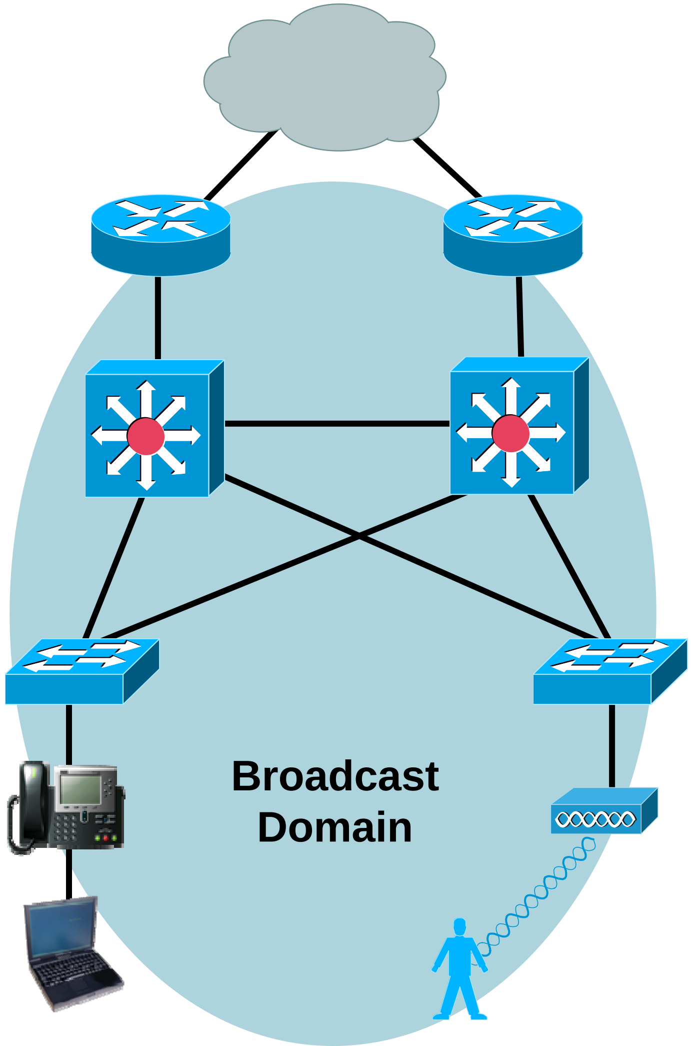

- Propagate broadcasts/failures
- Large Layer 2 domains hard to scale
- Complex spanning tree topology → hard to manage and troubleshoot
- Viable alternative for small networks

Access Control Lists

Distributed Versus Centralized Deployment

2

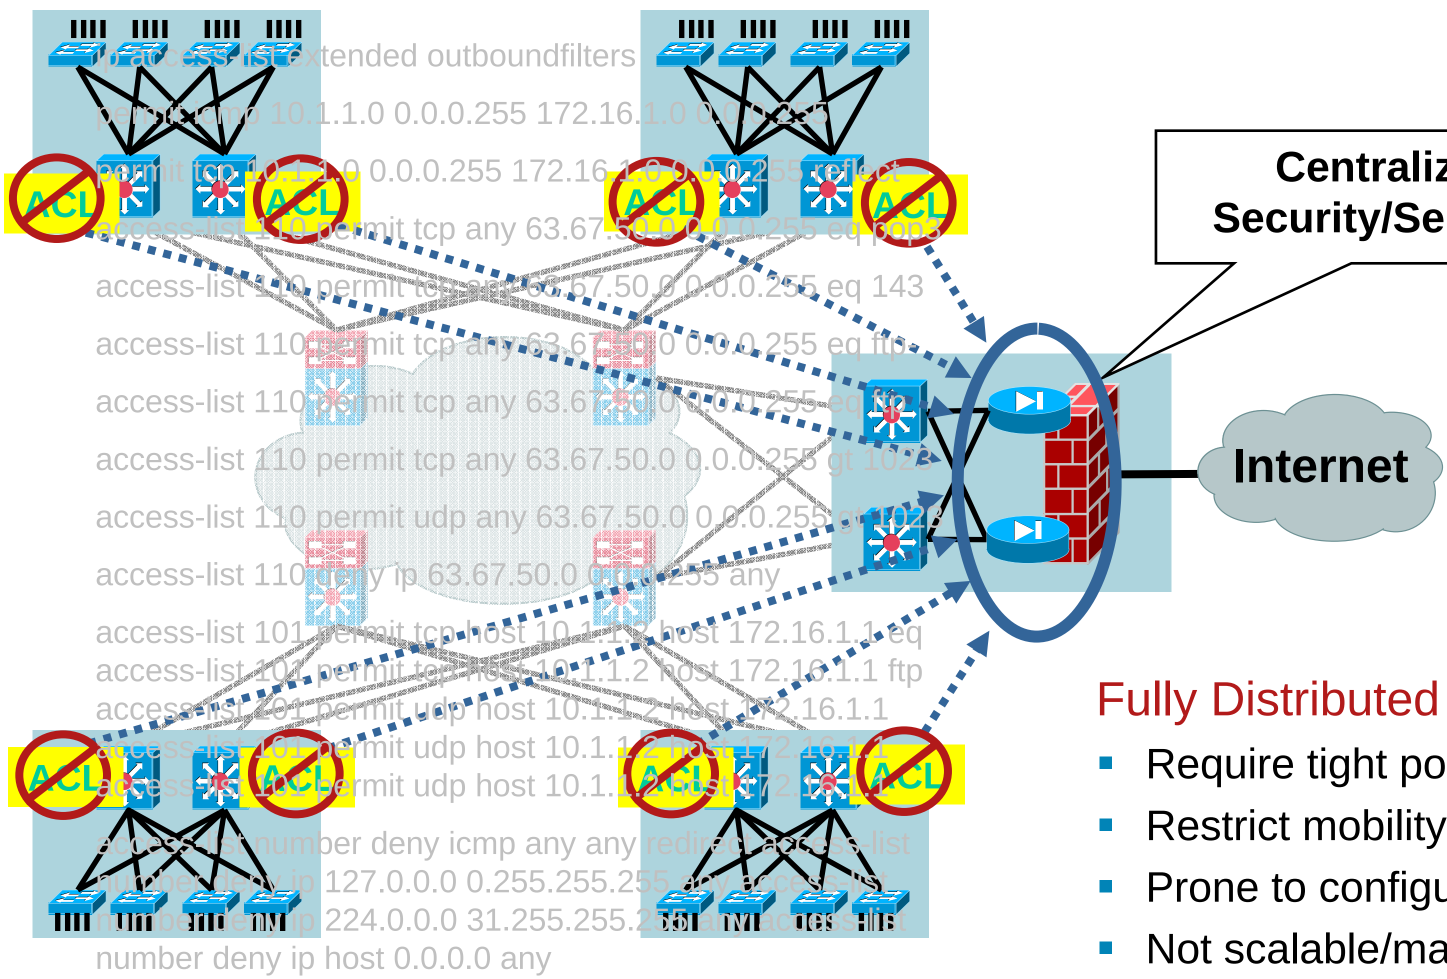

Fully Distributed ACLs

- Require tight policy control
- Restrict mobility
- Prone to configuration error
- Not scalable/manageable
- May suit specific requirements

GRE Protocol Tunneling + VRFs

Spoke to Hub Guest/Remediation Access

3

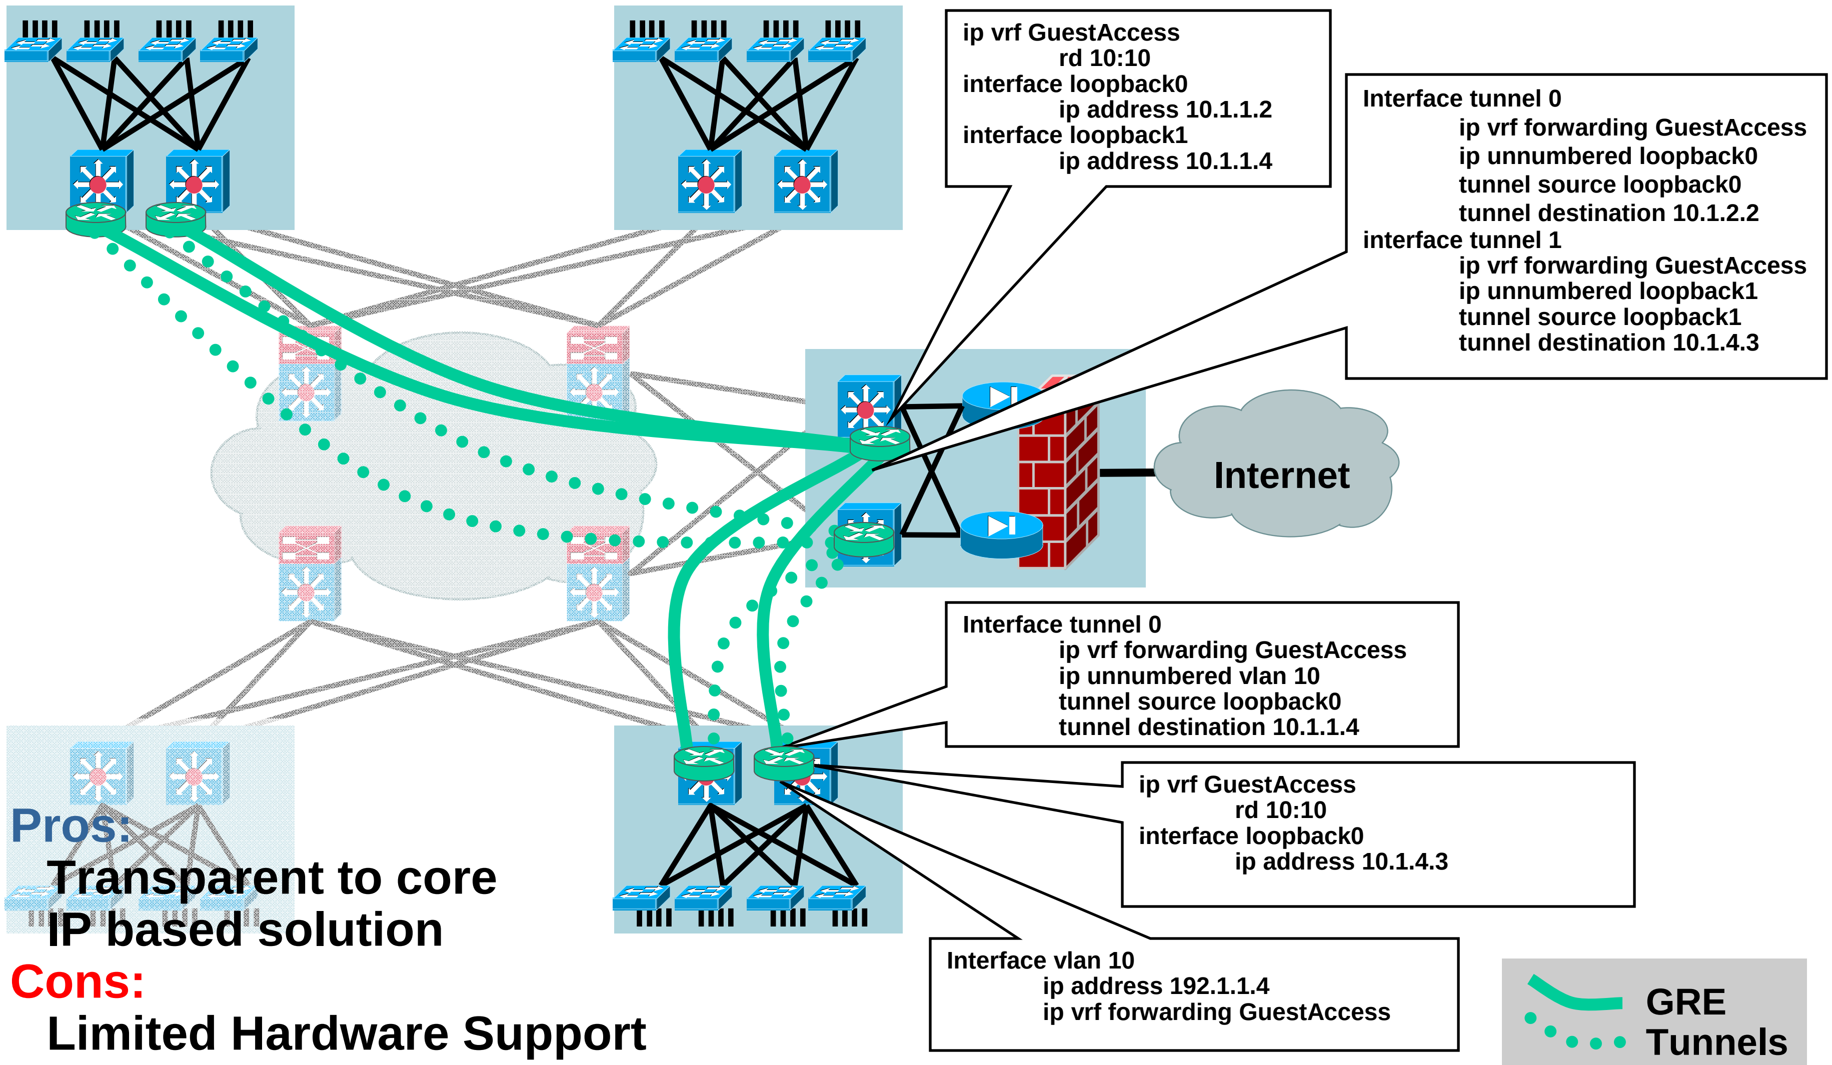

Pros:

- Transparent to core
- IP based solution

Cons:

- Limited Hardware Support
- Limited Scalability

MPLS-VPN—RFC2547 VPNs

Any-to-Any Connectivity

4

- Any-to-any connectivity per user group
- Highly scalable
- Each VPN is a separate IP cloud
- User-to-cloud connectivity
- Pervasive VPNs allow user mobility

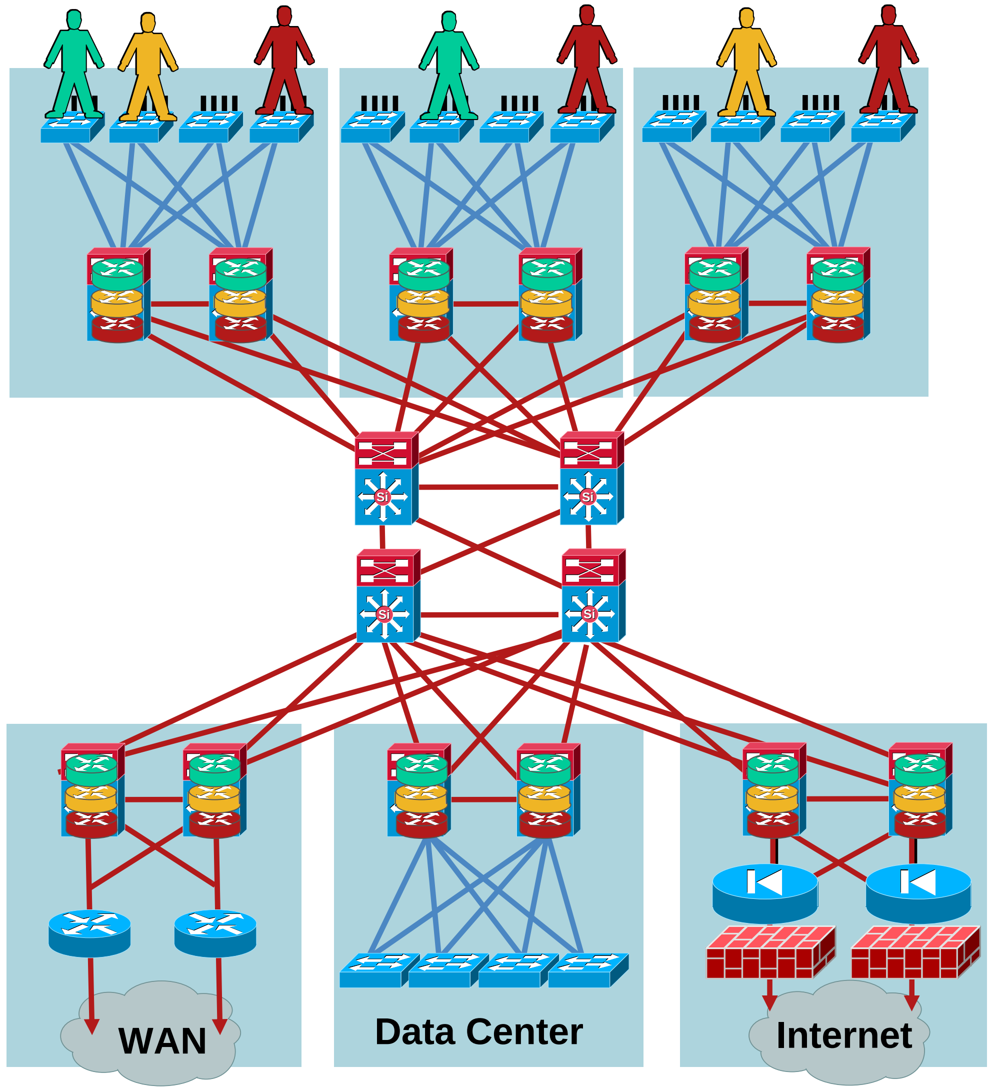

RFC2547 VPNs—Router Roles

Data Path/Forwarding Plane

- **PE** — **Provider Edge → Distribution Switch**
- **CE** — **Customer Edge**
- **P** — **Provider Equipment → Core Switch**
- **Multi-VRF**
- **802.1q trunks**
- **Label Switched Paths**

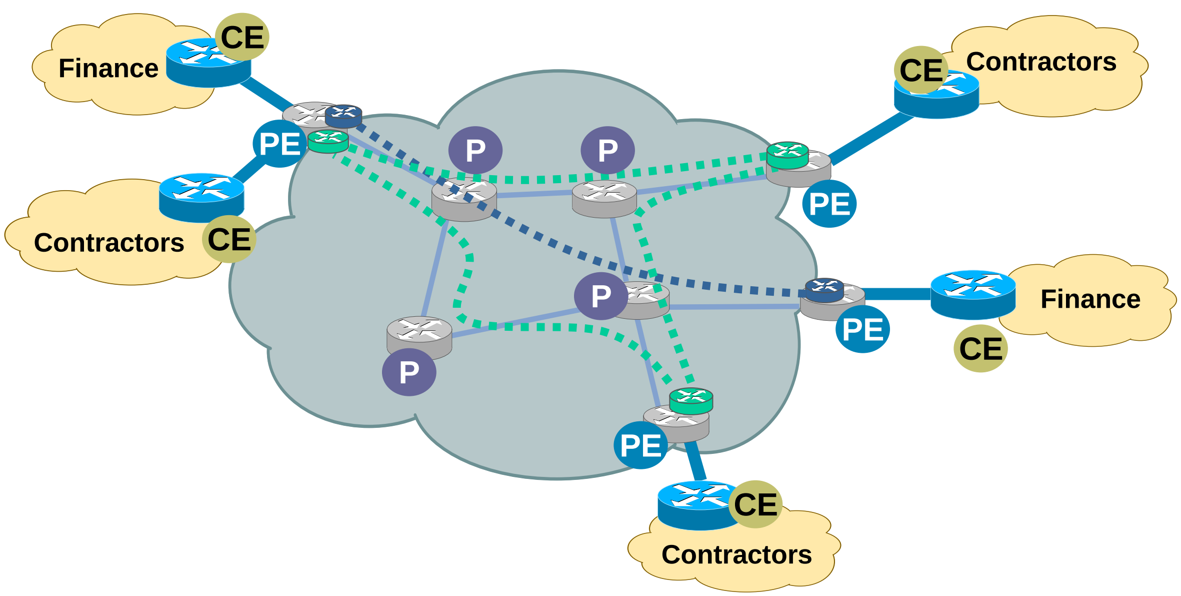

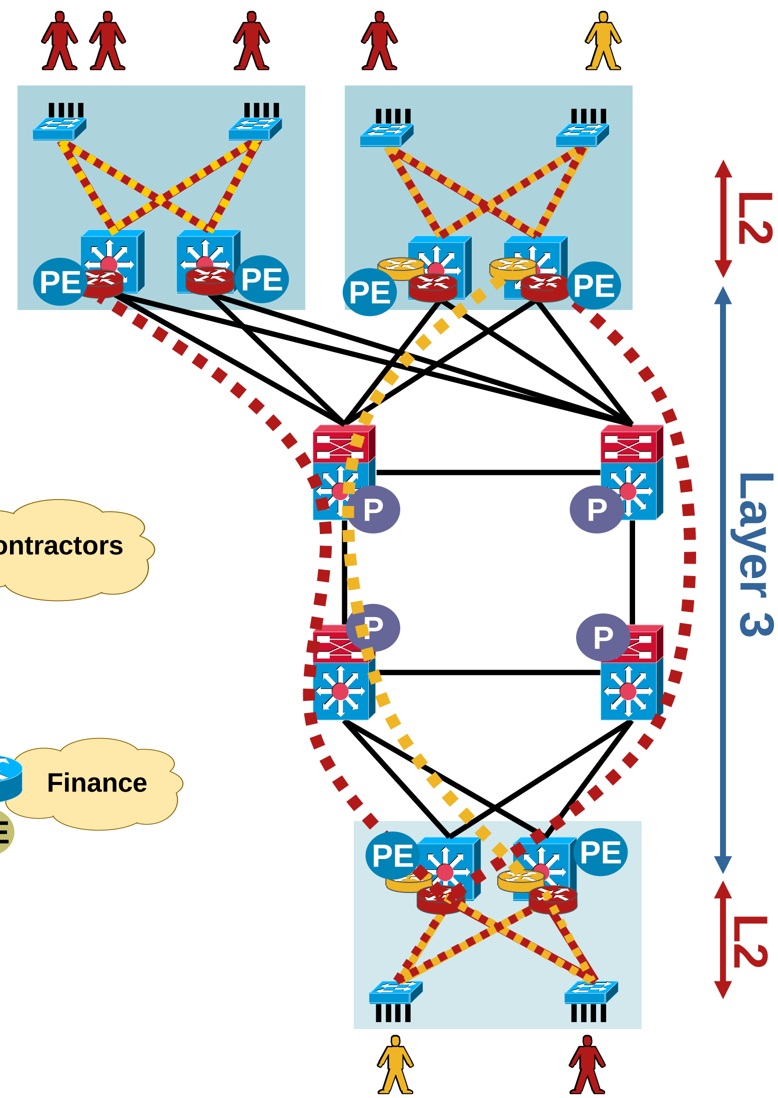

RFC2547 VPNs: Routing Peers
Control Plane

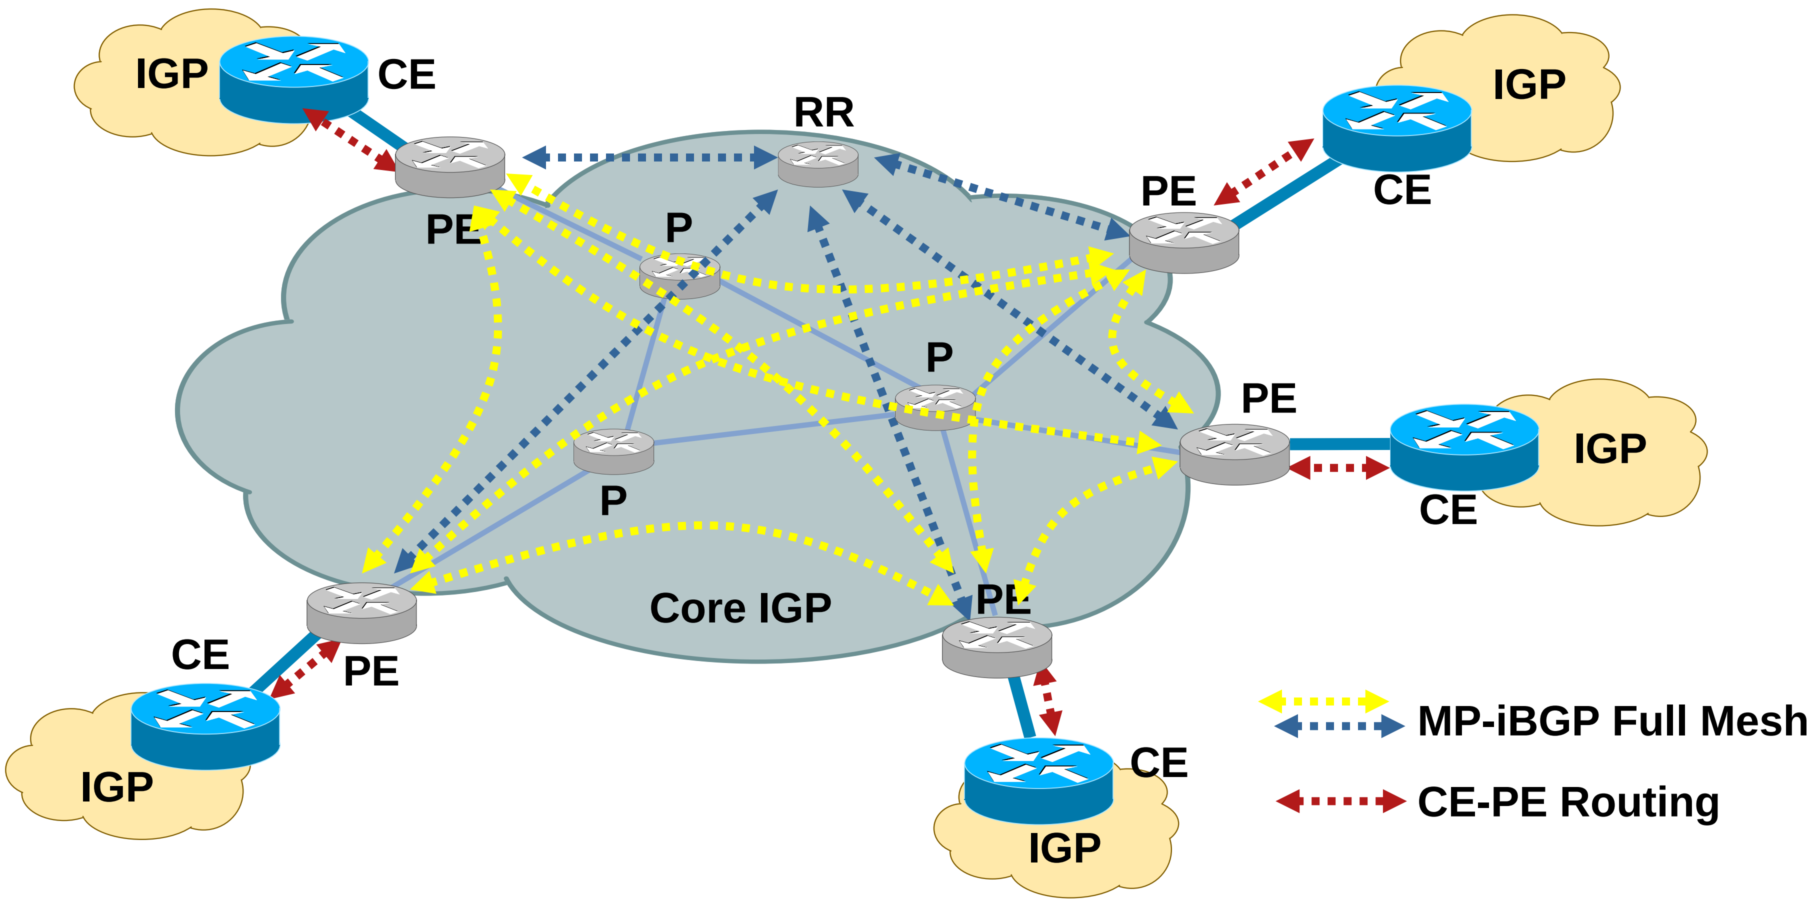

- PE routers handle all subscriber state (VRFs/VPNs)
    - Customer routes and VPNs are transparent to PE routers
- Core IGP provides connectivity between PE routers

PE: VRFs
Provider Edge VRFs

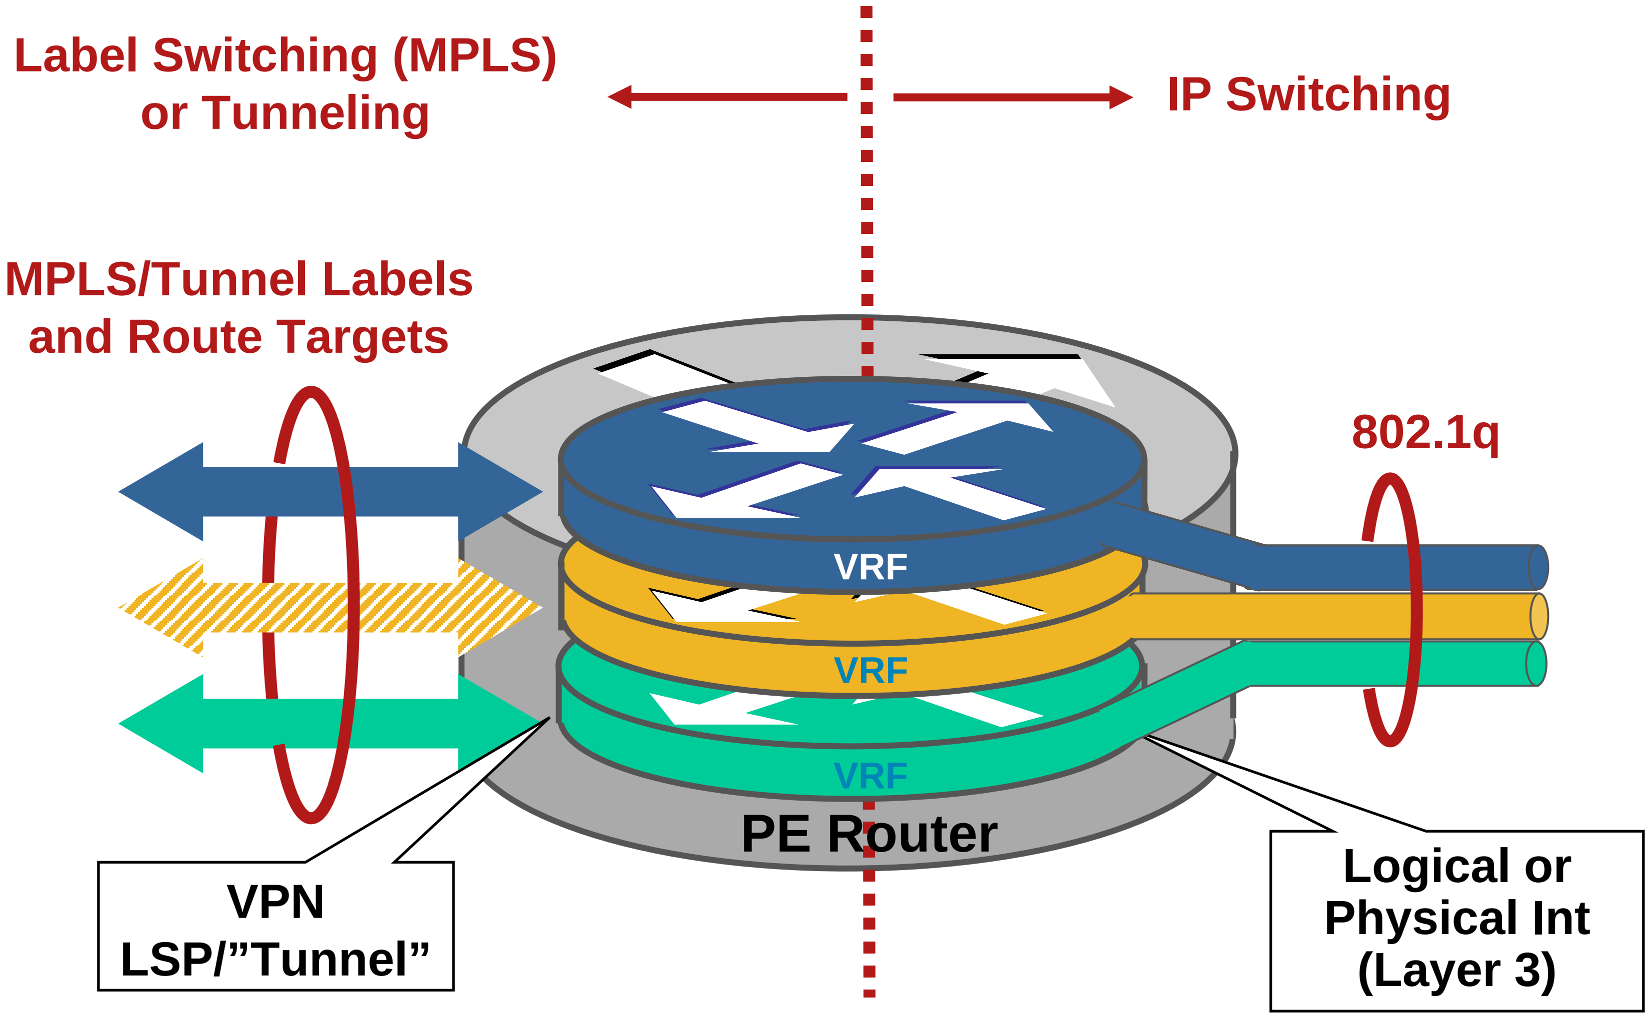

RFC 2547 VPNs over MPLS

- L2 access (no CE)
- VPN at the first L3 hop (distribution = PE)
- MP-iBGP at the distribution only (PE)
- MPLS in core and distribution (P and PE)
- Overlaid onto existing IGP

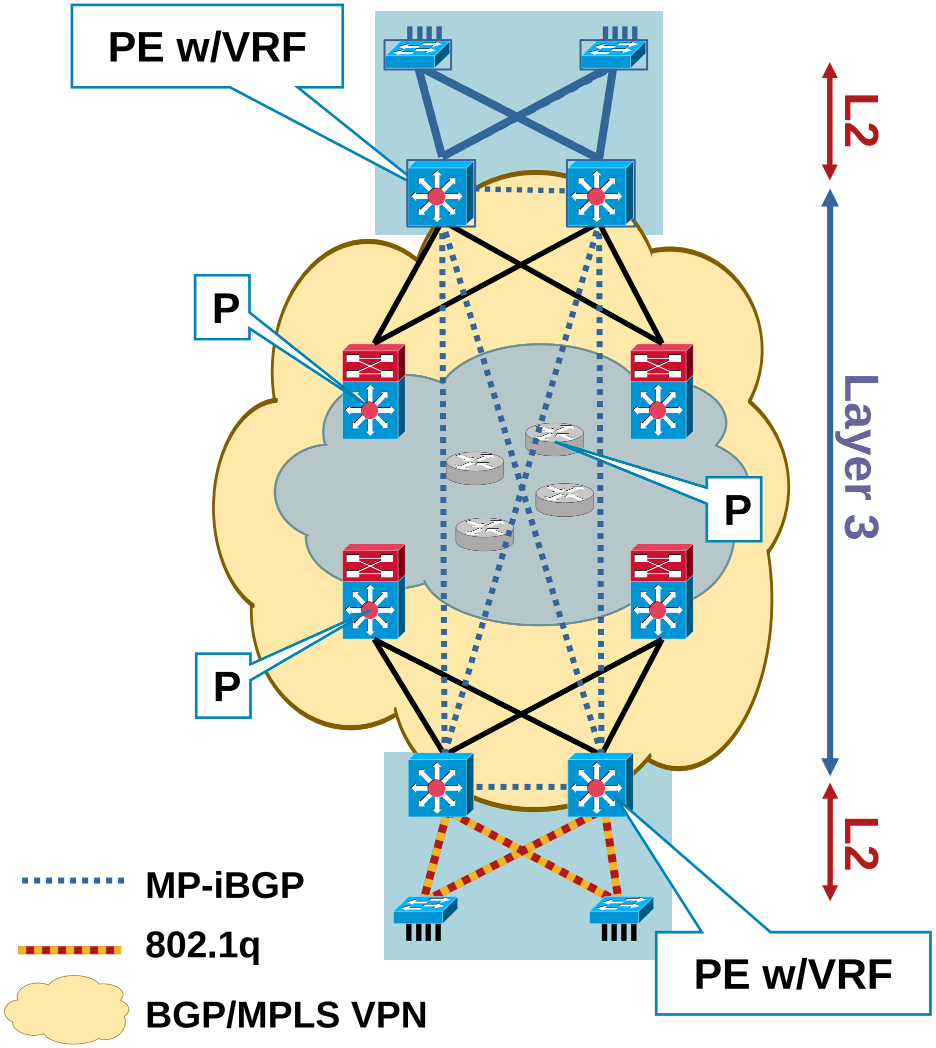

Configuration Summary (MPLS-Based RFC2547)

1. Configure the core (P and PE routers)
   - Configure an IGP
   - Enable MPLS switching
2. Configure PE routers
   - Configure MP-iBGP (route reflectors recommended)
   - Configure VRFs
     - Create VRFs
     - Configure route target imports/exports
   - Add interfaces to the VRFs
3. Configure CE routers (if in use)
   - Configure CE (lite) VRFs
     - Create CE VRFs
     - Add interfaces to CE VRFs
   - Configure PE-CE routing

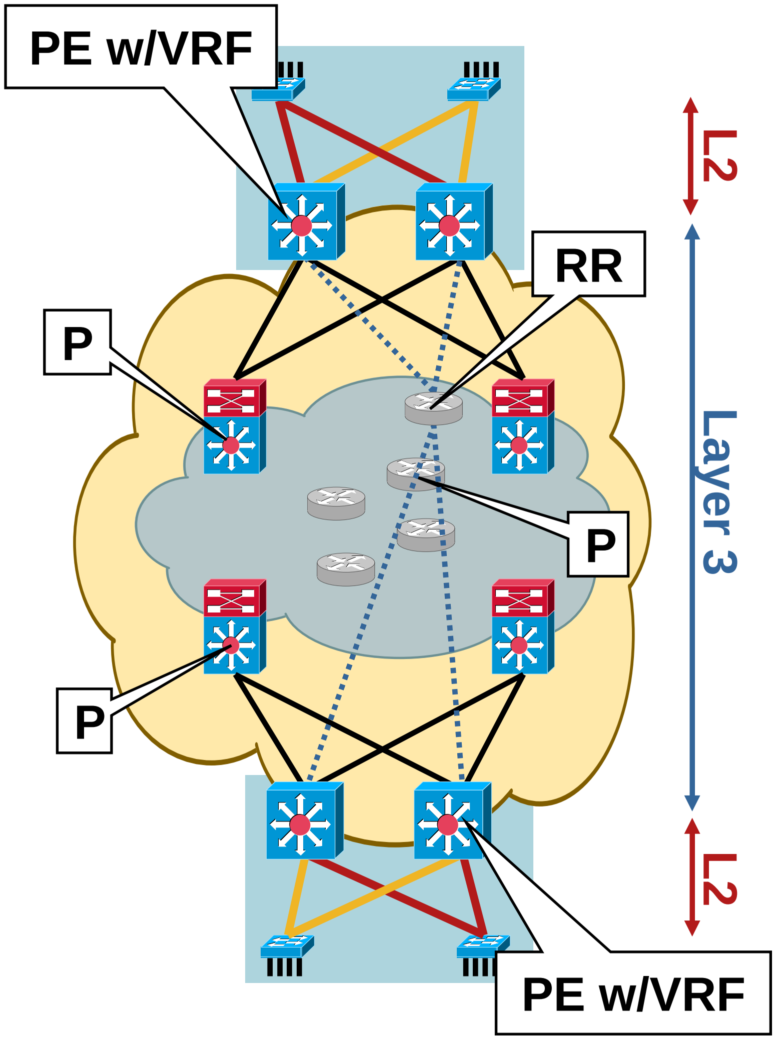

Configuration (Tag-Switching and P Routers)

```
interface Loopback0
 ip address 1.1.1.6 255.255.255.255
!
interface FastEthernet0/0
description Facing P router
 ip address 10.0.0.8 255.255.255.252
 mpls ip
```

```
interface Loopback0
 ip address 1.1.1.7 255.255.255.255
!
interface FastEthernet0/0
description Facing P router
 ip address 10.0.0.11 255.255.255.252
 mpls ip
```

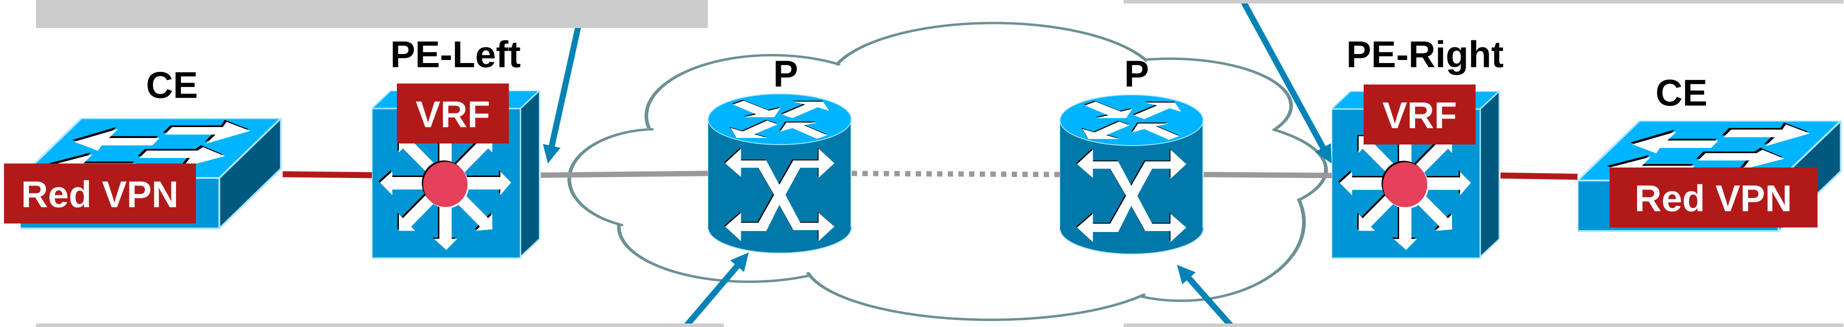


```
interface Loopback0
 ip address 1.1.1.3 255.255.255.255
!
interface FastEthernet0/0
description facing P router
 ip address 10.0.0.6 255.255.255.252
 mpls ip
!
interface FastEthernet1/0
description facing PE-Left
 ip address 10.0.0.9 255.255.255.252
 mpls ip
```

```
interface Loopback0
 ip address 1.1.1.4 255.255.255.255
!
interface FastEthernet0/0
description Facing P router
 ip address 10.0.0.14 255.255.255.252
 mpls ip
!
interface FastEthernet3/0
description Facing PE-Right
 ip address 10.0.0.10 255.255.255.252
 mpls ip
```

Configuration (VPN and PE-PE Routing)

```
ip vrf Red
 rd 100:33
 route-target import 100:33
 route-target export 100:33

interface Vlan11
 ip vrf forwarding Red
 ip address 10.20.1.1/24
```

```
ip vrf Red
 rd 100:33
 route-target both 100:33

interface Vlan11
 ip vrf forwarding Red
 ip address 10.20.4.1/24
```

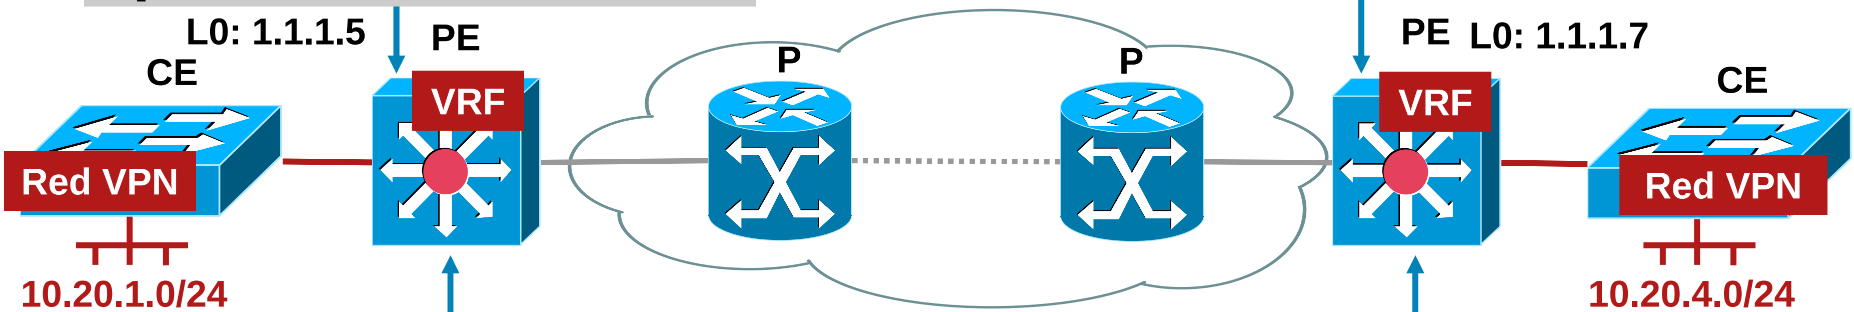


```
router bgp 100
 no bgp default ipv4-unicast

 neighbor 1.1.1.7 remote-as 100
 neighbor 1.1.1.7 update-source Loopback0

 address-family vpnv4
 neighbor 1.1.1.7 activate
 neighbor 1.1.1.7 send-community extended

address-family ipv4 vrf Red
network 10.20.1.0 mask 255.255.255.0
```

```
router bgp 100
 no bgp default ipv4-unicast

 neighbor 1.1.1.5 remote-as 100
 neighbor 1.1.1.5 update-source Loopback0

 address-family vpnv4
 neighbor 1.1.1.5 activate
 neighbor 1.1.1.5 send-community extended

address-family ipv4 vrf Red
network 10.20.4.0 mask 255.255.255.0
```

Operational Verification: VPN Routes (Ingress)

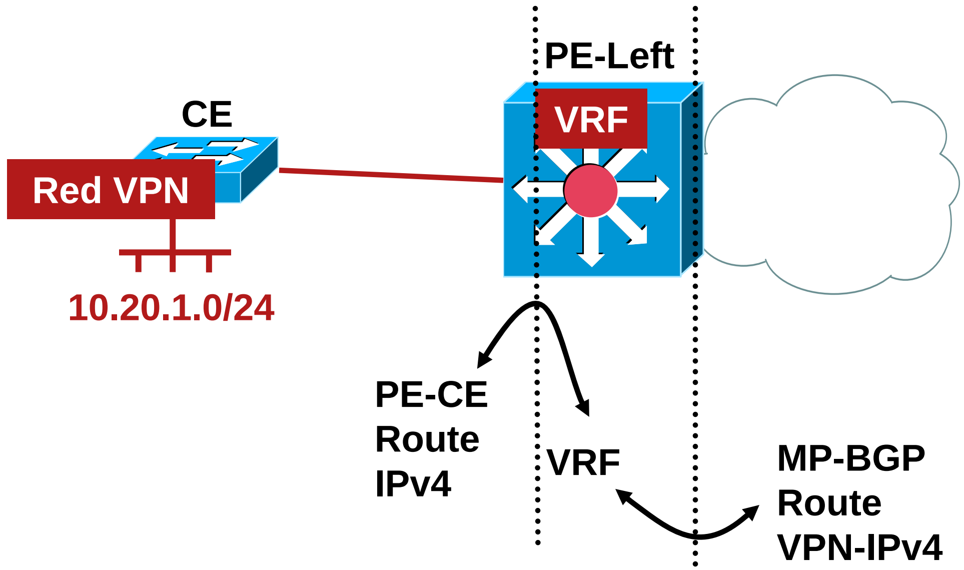


```
PE-Left#sh ip vrf
  Name              Default RD        Interfaces
  Red                 100:33            Vlan 12
  Blue                100:22            Vlan 11

PE-Left#sh ip vrf interface
Interface         IP-Address       VRF       Protocol
VLan12            10.20.1.1        Red          up
Vlan11            10.20.1.1        Blue         up
```

```
PE-Left#sh ip bgp vpnv4 vrf Red
BGP table version is 9, local router ID is 1.1.1.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:33 (default for vrf Red)
*> 10.20.1.0/24     0.0.0.0                  0         32768 i
*>i10.20.4.0/24     1.1.1.7                  0    100      0 i
```


RFC 2547 with Multi-VRF CE (VRF-Lite) at Distribution

4b

- L2 access
- Multi-VRF-CE at distribution
- BGP/MPLS VPNs in core only
- VRF-lite between core and distribution
- Labels substituted by 802.1q tags between distribution and core
- Multi-VRF CE could be used to deploy on a routed access model

Access = Multi-VRF CE

Distribution = PE

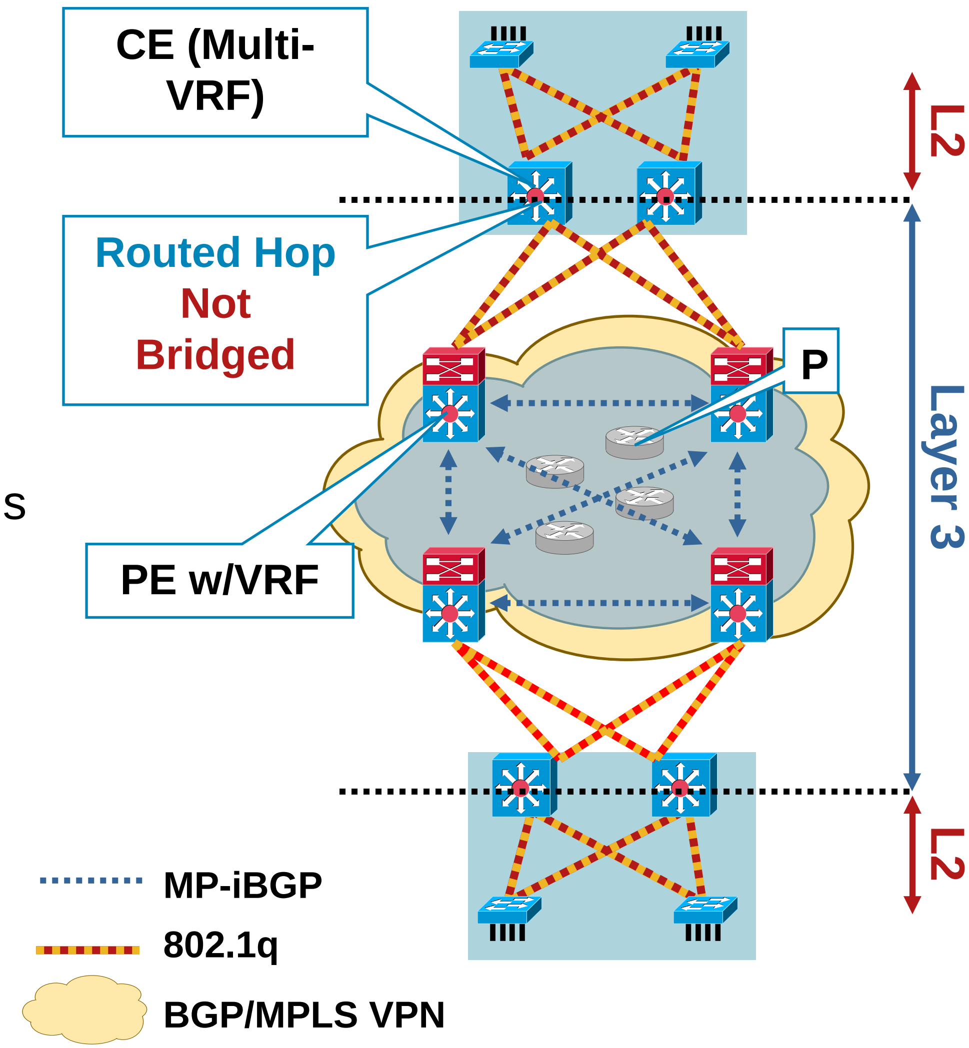

Multi-VRF CE (VRF-Lite)

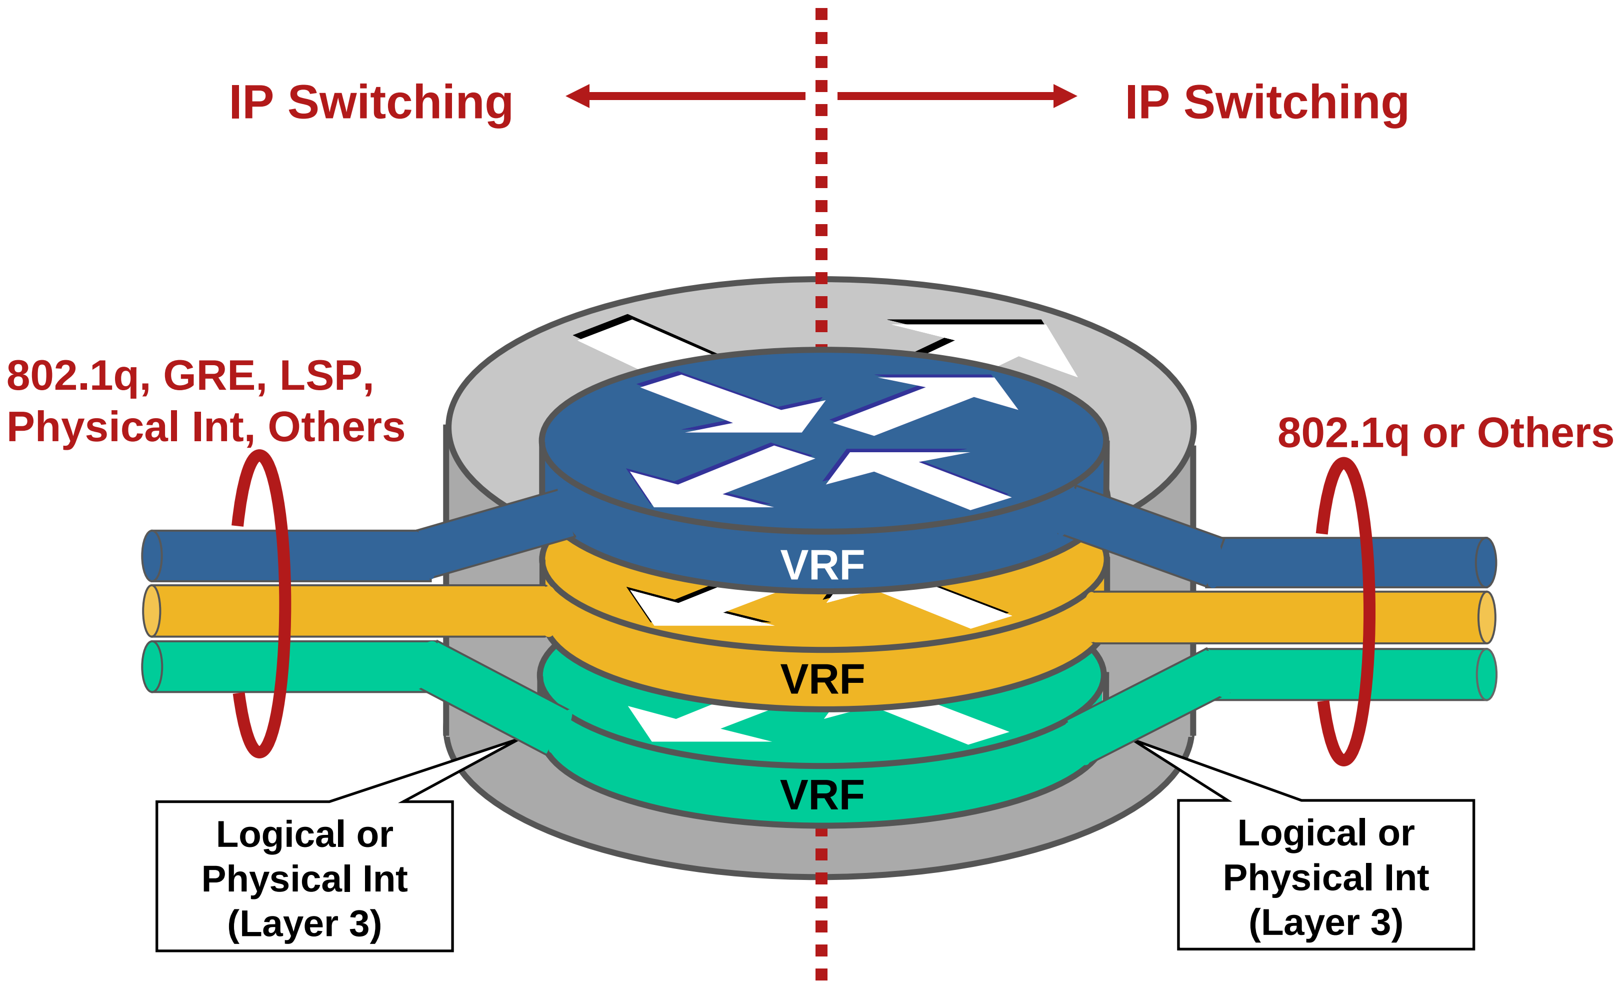

VRF-Lite—PE—CE Interaction

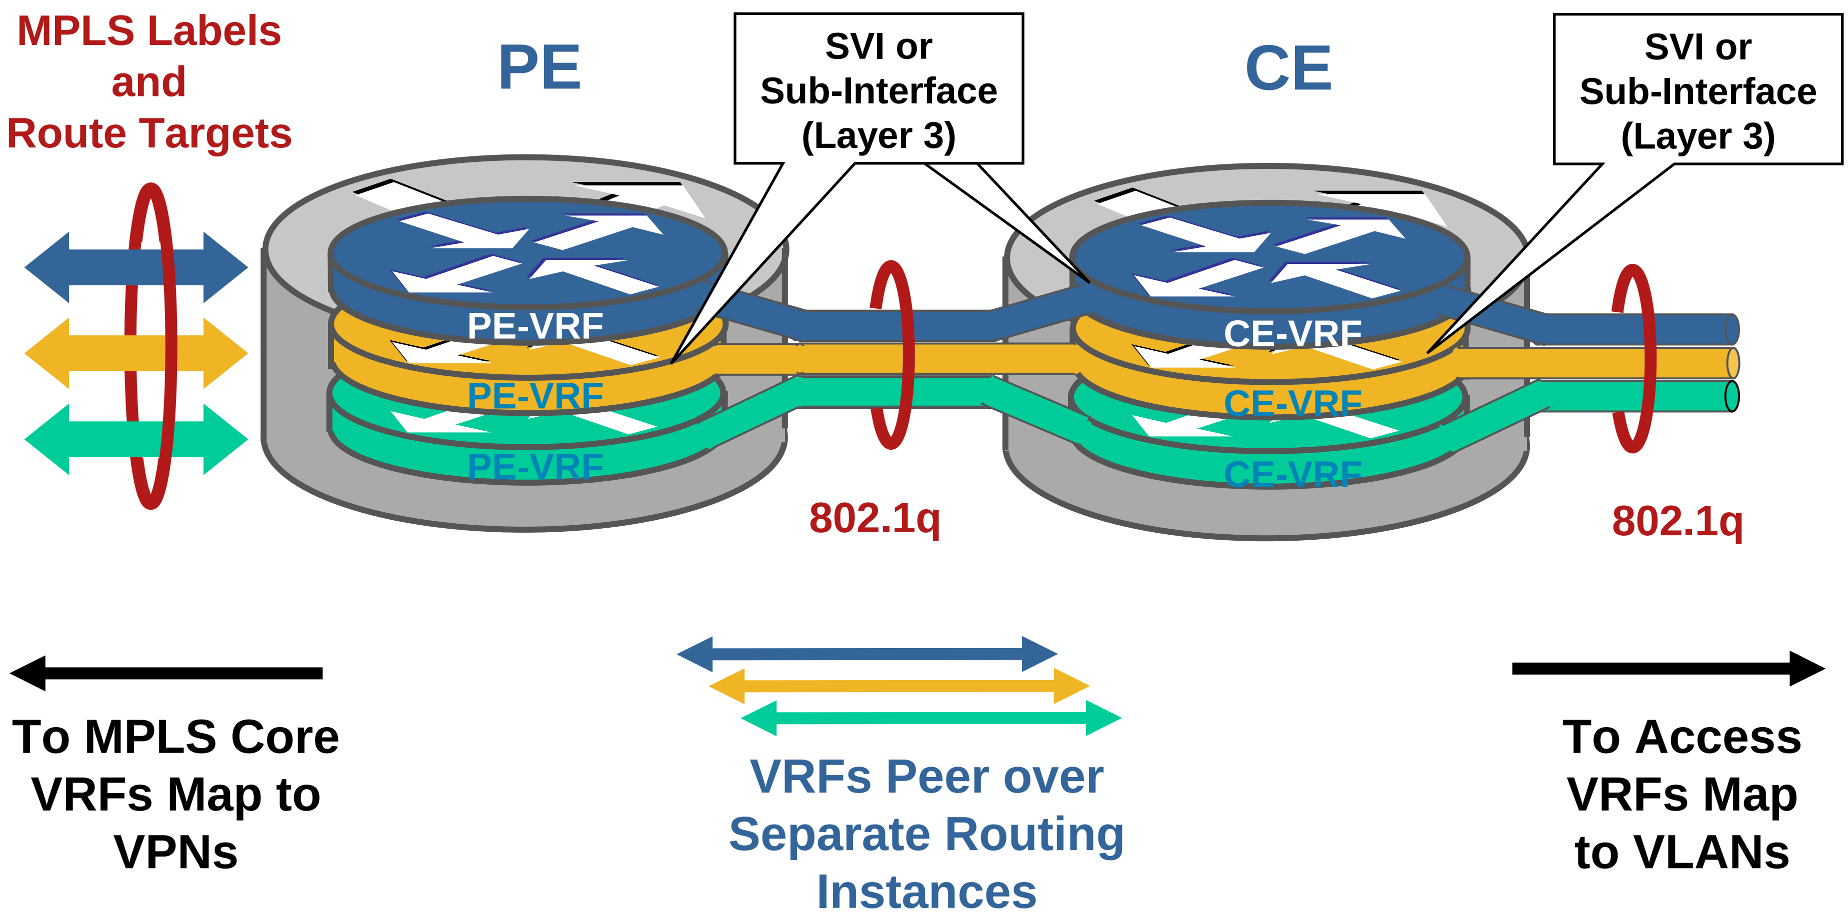

RFC 2547 with Multi-VRF CE at the Access

Routed Access

4c

- Routed access
- Multi-VRF-CE at the Access
- MP-iBGP at the distribution only (PE)
- MPLS in core and distribution (P and PE)
- 2547 VPNs overlaid onto existing core IGP
- Access is IP switched with multi-VRF
- PE-CE routing per VRF

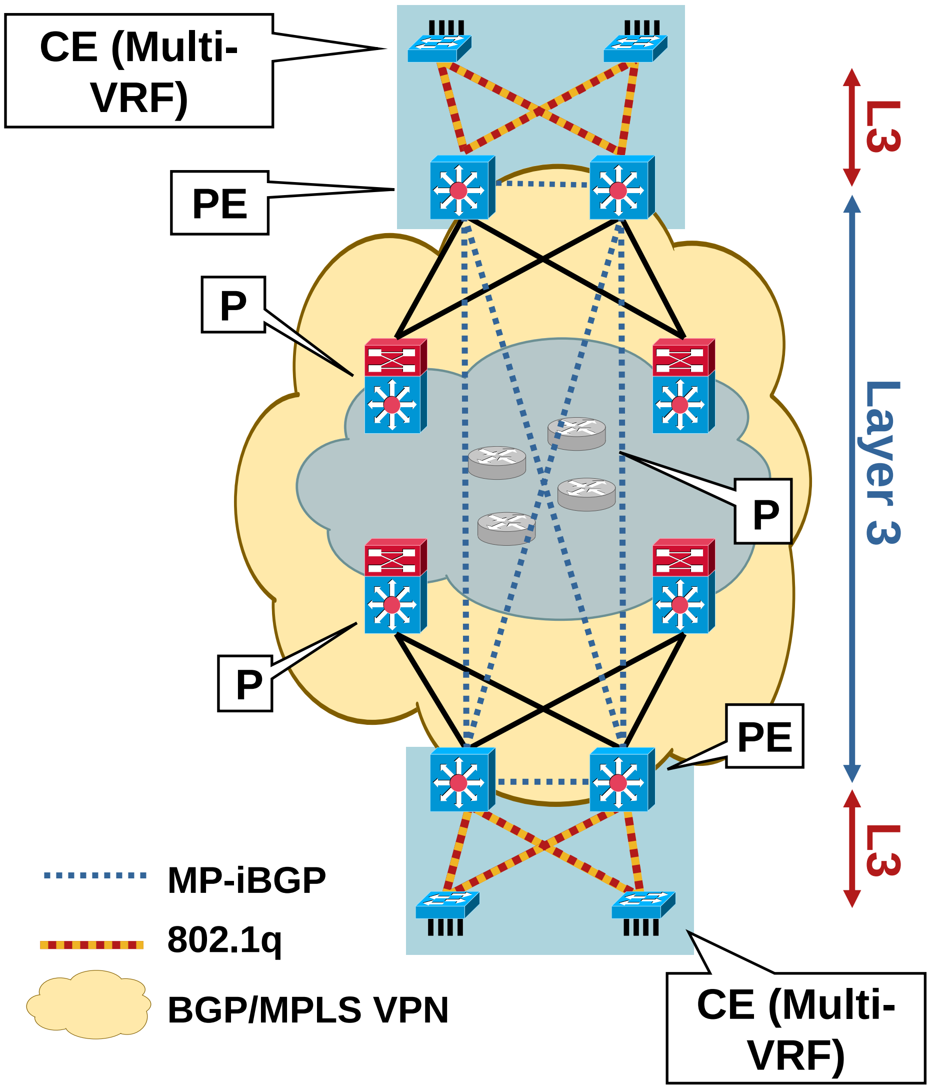

Multi-VRF CE (VRF-Lite)—End to End

- No BGP or MPLS
- VRF-lite on all routed hops: Core and distribution
- 802.1q tags provide single hop data path virtualization
- Every link is a 802.1q trunk
- These trunks do not extend VLANs throughout the campus
- Trunks used to virtualized data path between multiple virtual routers
- Every physical link carries multiple logical routed links
- Provisioning challenges:
  - Four links and three groups = 12 VLAN IDs
  - Four links and five groups = 20 VLAN IDs
  - VLAN IDs must match on both ends

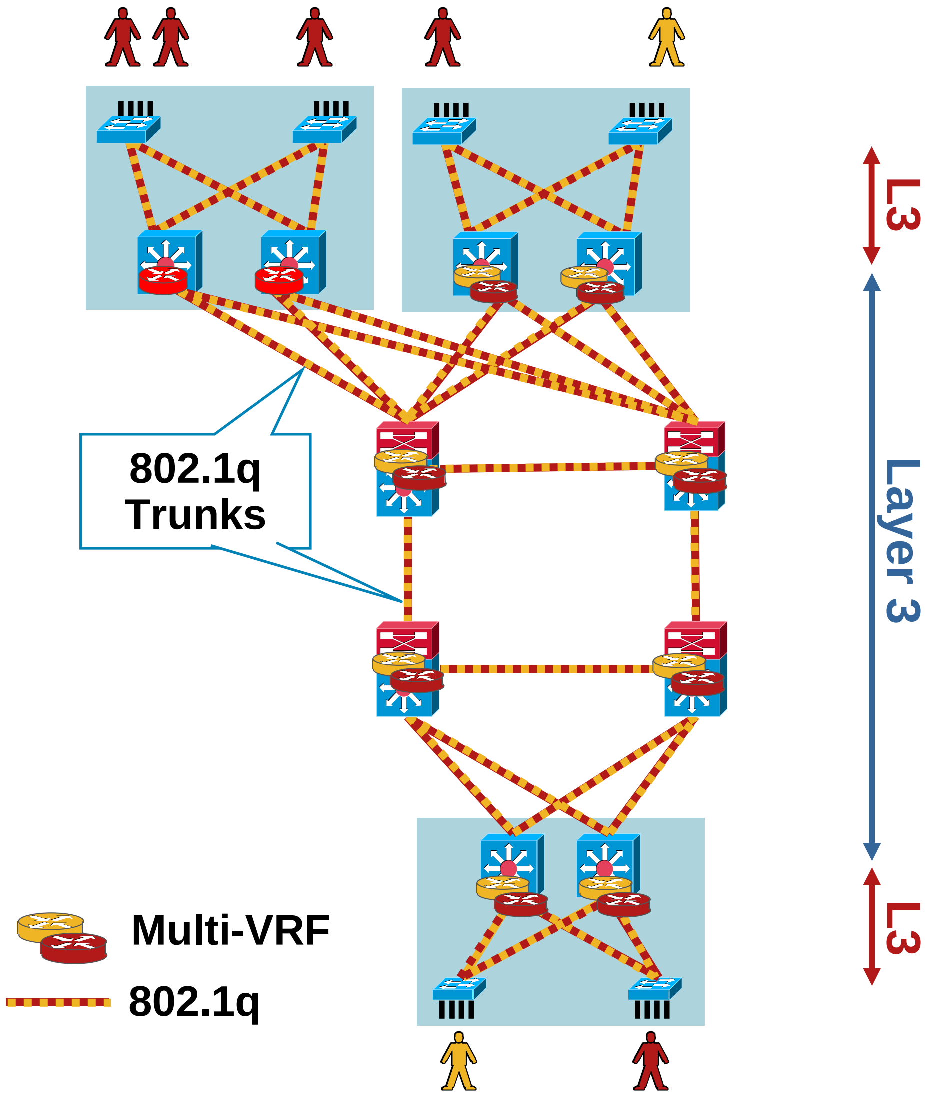

Multi-VRF (VRF-Lite) End to End

End-to-End VRF-Lite (802.1q Virtual Links)

- VRF-lite utilizes hop by hop 802.1q to VRF mapping to build a closed user group
- Association of VRF to VLAN is manually configured
- Each VRF Instance needs a separate IGP process (OSPF) or address family (EIGRP, RIPv2, MP-BGP)
- In this configuration Traffic is routed from each 802.1q VLAN to the associated 802.1q VLAN

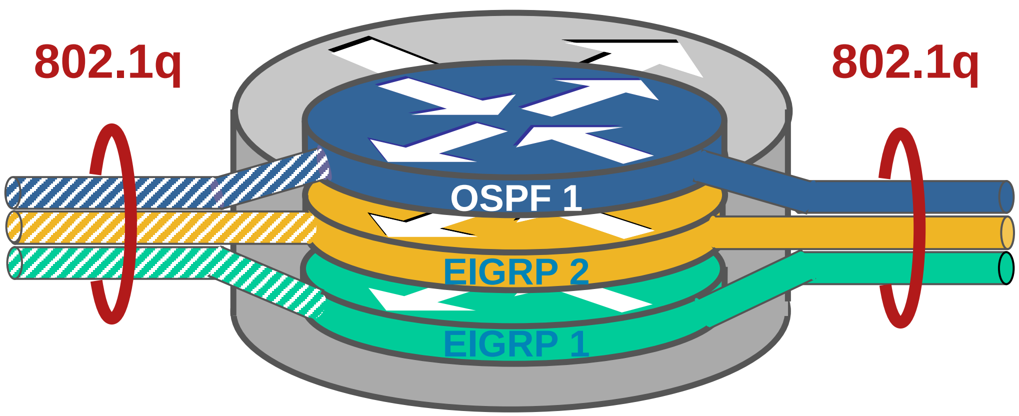

VRF-Lite Supported on 6500, 4500, 3560, and 3750

VRF-Lite End to End (802.1q Virtual Links)
Trunk with Switchport

Links Between Routers Defined as L2 Trunk with Switchports

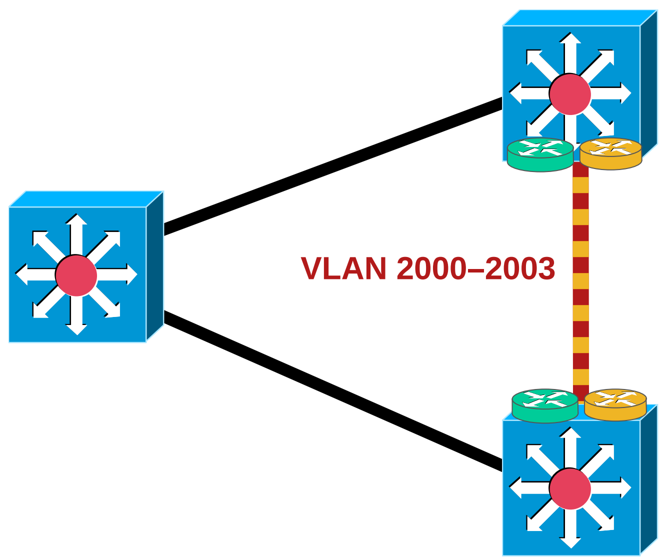


```
Catalyst-1
interface GigabitEthernet1/1
 description --- To Cat6500-1 ---
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 2000-2003
 switchport mode trunk
 spanning-tree portfast trunk
!
interface Vlan2000
 description --- Link to Cat6500-1
 ip address 10.149.12.2 255.255.255.0
 ip ospf network point-to-point
!
interface Vlan2001
 ip vrf forwarding VPN1
 ip address 1.1.12.2 255.255.255.0
 ip ospf network point-to-point
!
interface Vlan2002
 ip vrf forwarding VPN2
 ip address 2.2.12.2 255.255.255.0
 ip ospf network point-to-point
!
interface Vlan2003
 ip vrf forwarding VPN-SERVERS
 ip address 3.3.12.2 255.255.255.0
 ip ospf network point-to-point
!
```


VRF-Lite End to End (802.1q Virtual Links)
Trunk with Routed Ports

Links Between Routers or Defined as L3 Trunk with Sub-Interface

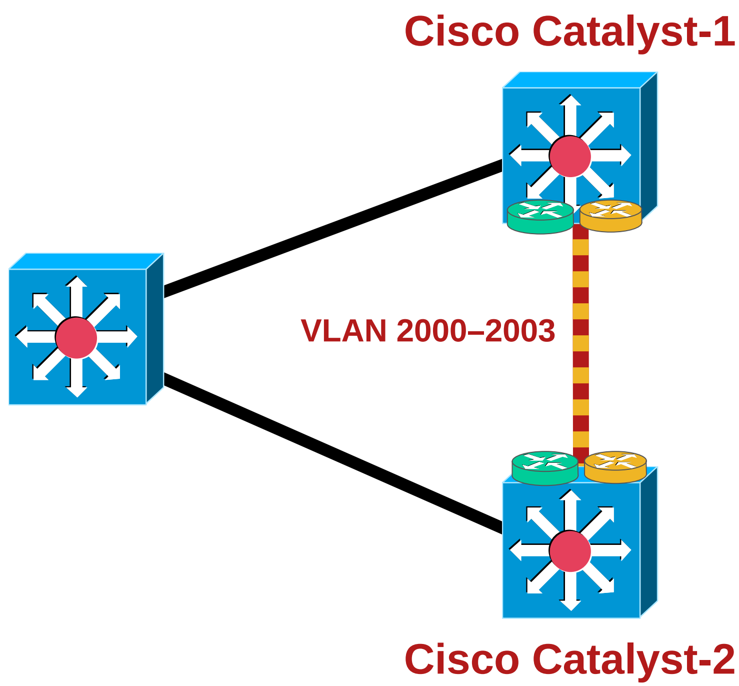

Currently Supported on Cisco Catalyst 6500 Only

```
Catalyst-2
interface GigabitEthernet6/1
 no ip address
!
interface GigabitEthernet6/1.2000
 encapsulation dot1Q 2000
 ip address 10.149.12.1 255.255.255.0
 ip ospf network point-to-point
!
interface GigabitEthernet6/1.2001
 encapsulation dot1Q 2001
 ip vrf forwarding VPN1
 ip address 1.1.12.1 255.255.255.0
 ip ospf network point-to-point
!
interface GigabitEthernet6/1.2002
 encapsulation dot1Q 2002
 ip vrf forwarding VPN2
 ip address 2.2.12.1 255.255.255.0
 ip ospf network point-to-point
!
interface GigabitEthernet6/1.2003
 encapsulation dot1Q 2003
 ip vrf forwarding VPN-SERVERS
 ip address 3.3.12.1 255.255.255.0
 ip ospf network point-to-point
!
```


VRF-Lite End to End (802.1q Virtual Links)
Routing Processes

Separate OSPF Processes per VRF
or
Separate EIGRP Address-families per VRF

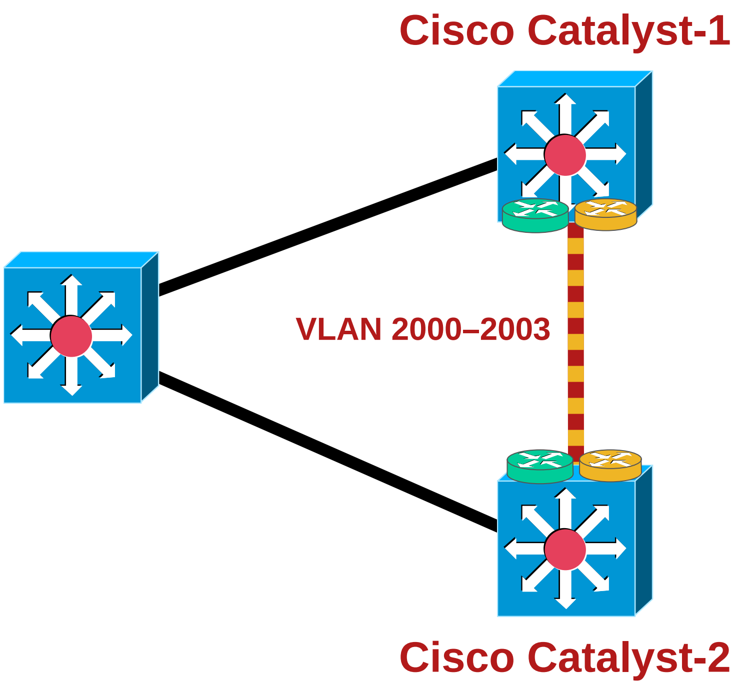


```
router ospf 1 vrf VPN1
 network 1.0.0.0 0.255.255.255 area 0
 network 10.0.0.0 0.255.255.255 area 0
 no passive-interface vlan 2001
!
router ospf 2 vrf VPN2
 network 2.0.0.0 0.255.255.255 area 0
 network 20.0.0.0 0.255.255.255 area 0
 no passive-interface vlan 2002
!
```

```
router eigrp 200
 address-family ipv4 vrf VPN1
  network 1.0.0.0
  network 10.0.0.0
  no auto-summary
  exit-address-family
 address-family ipv4 vrf VPN2
  network 2.0.0.0
  network 20.0.0.0
  no auto-summary
  exit-address-family
```


VRF-Lite End to End

Collapsed Core and Distribution

5b

- Collapsed core/distribution
- L2 access
- No BGP or MPLS
- VRF-lite at core
- Limited scalability

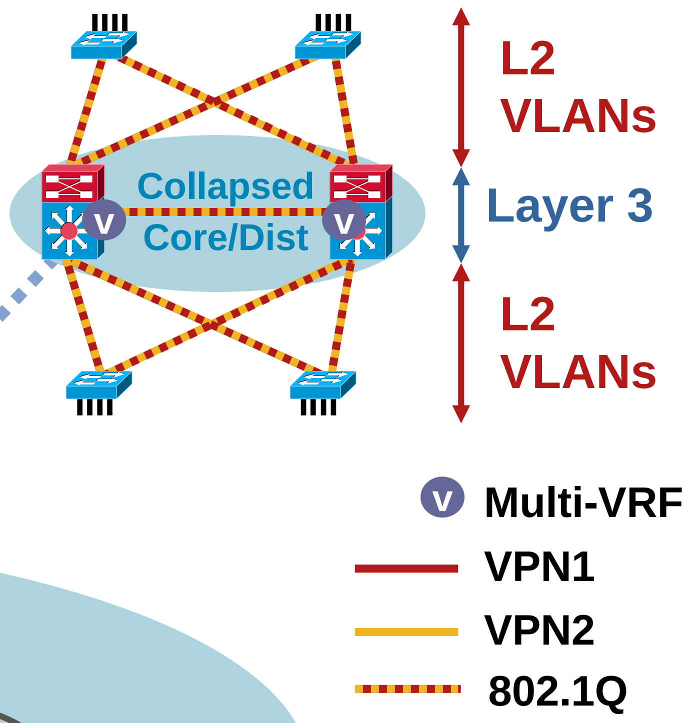

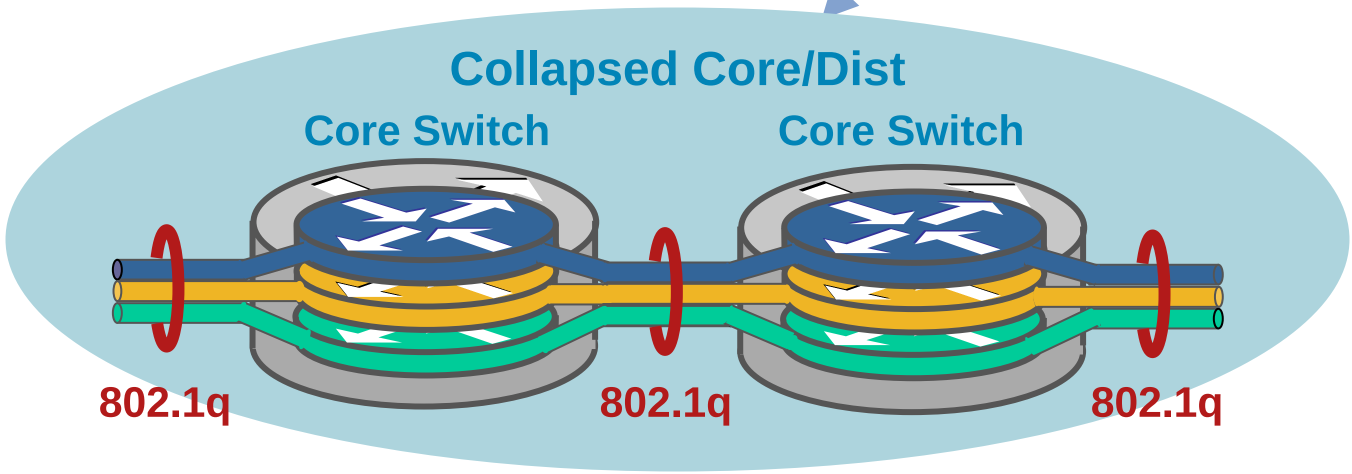

Campus Virtualization Alternatives
Summary

- End-to-end campus virtualization =

  Layer 2 virtualization (VLANs) +

  Layer 3 virtualization (Interconnected VRFs)

- Different alternatives include:
  1. VLANs
  2. ACLs
  3. VRFs + GRE Tunnel Mesh
  4. RFC2547 VPNs
  5. Hop-by-hop multi-VRF

- RFC2547 is the most scalable solution today

Agenda

- Problem Definition
- Campus Virtualization Alternatives
- WAN Extensibility
- Shared Services and Inter-VPN Communication
- Data Center Integration

Extensibility over the WAN

Groups Must Be Extensible over:

- The "private" WAN/MAN
  - L2 Services
  - IP Services
- The internet

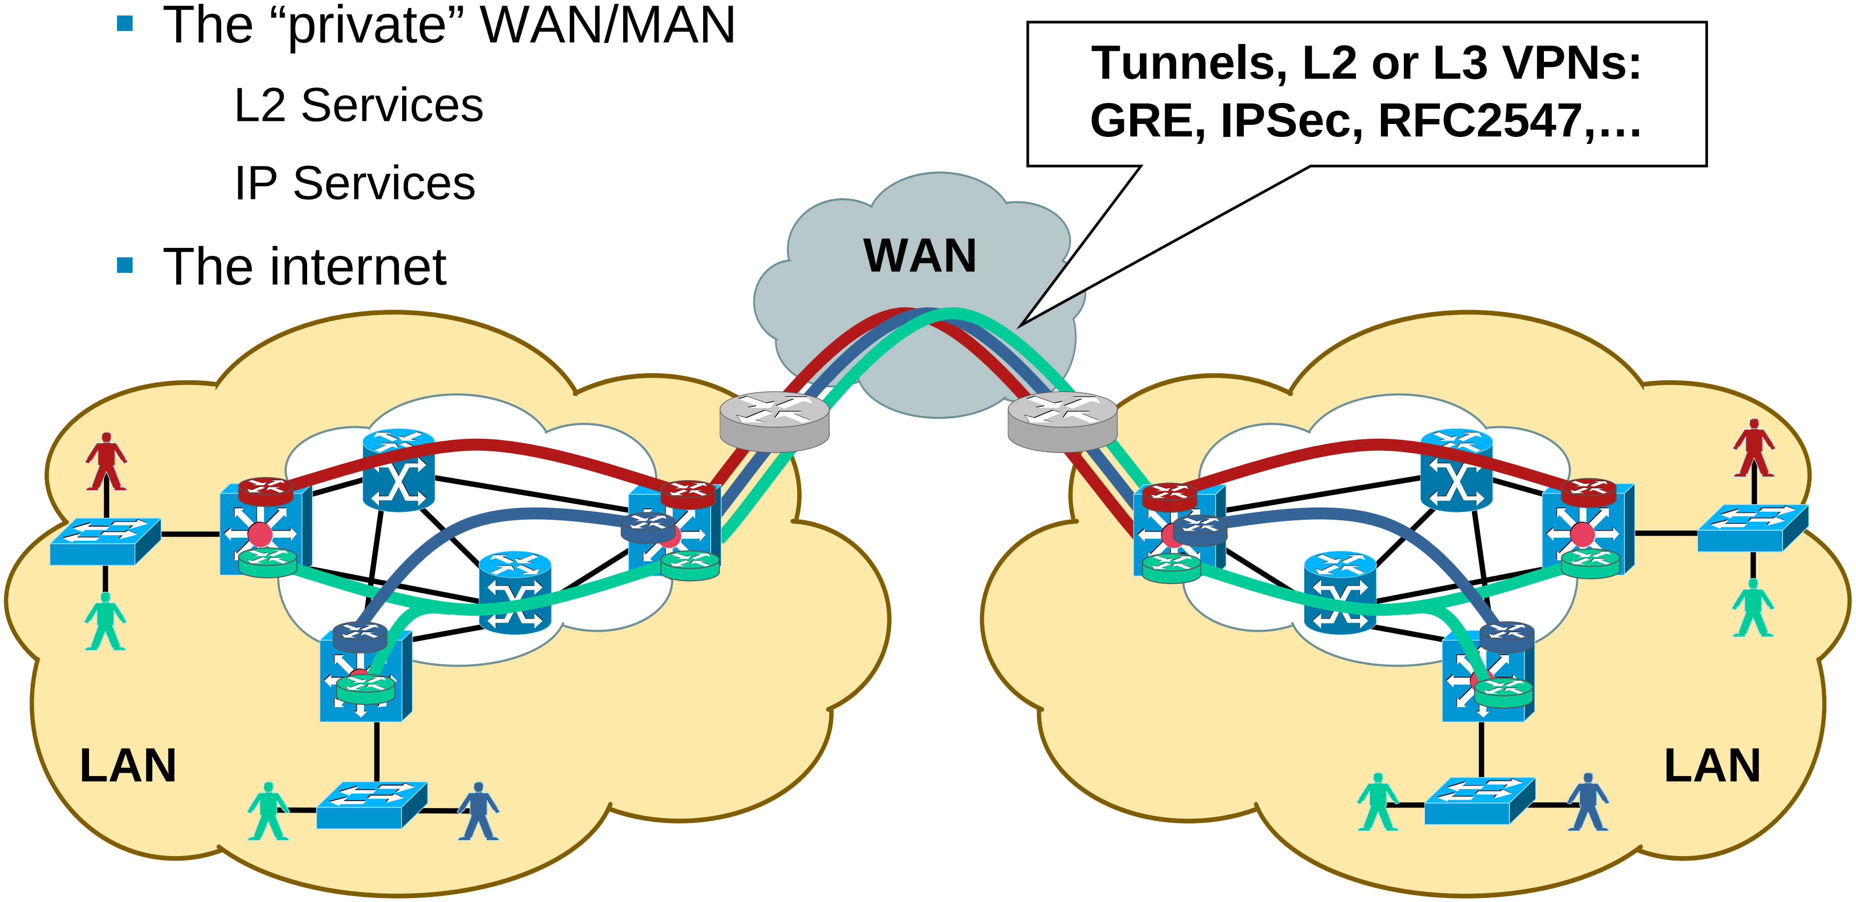

WAN Extensibility

PE-CE Routing over IP Cloud: L2 (FR/ATM) or L3 (MPLS)

- L2: Single SP VC

  GRE tunnels segment data path

- IP Service:

  GRE tunnels create 1:1 vrf connect

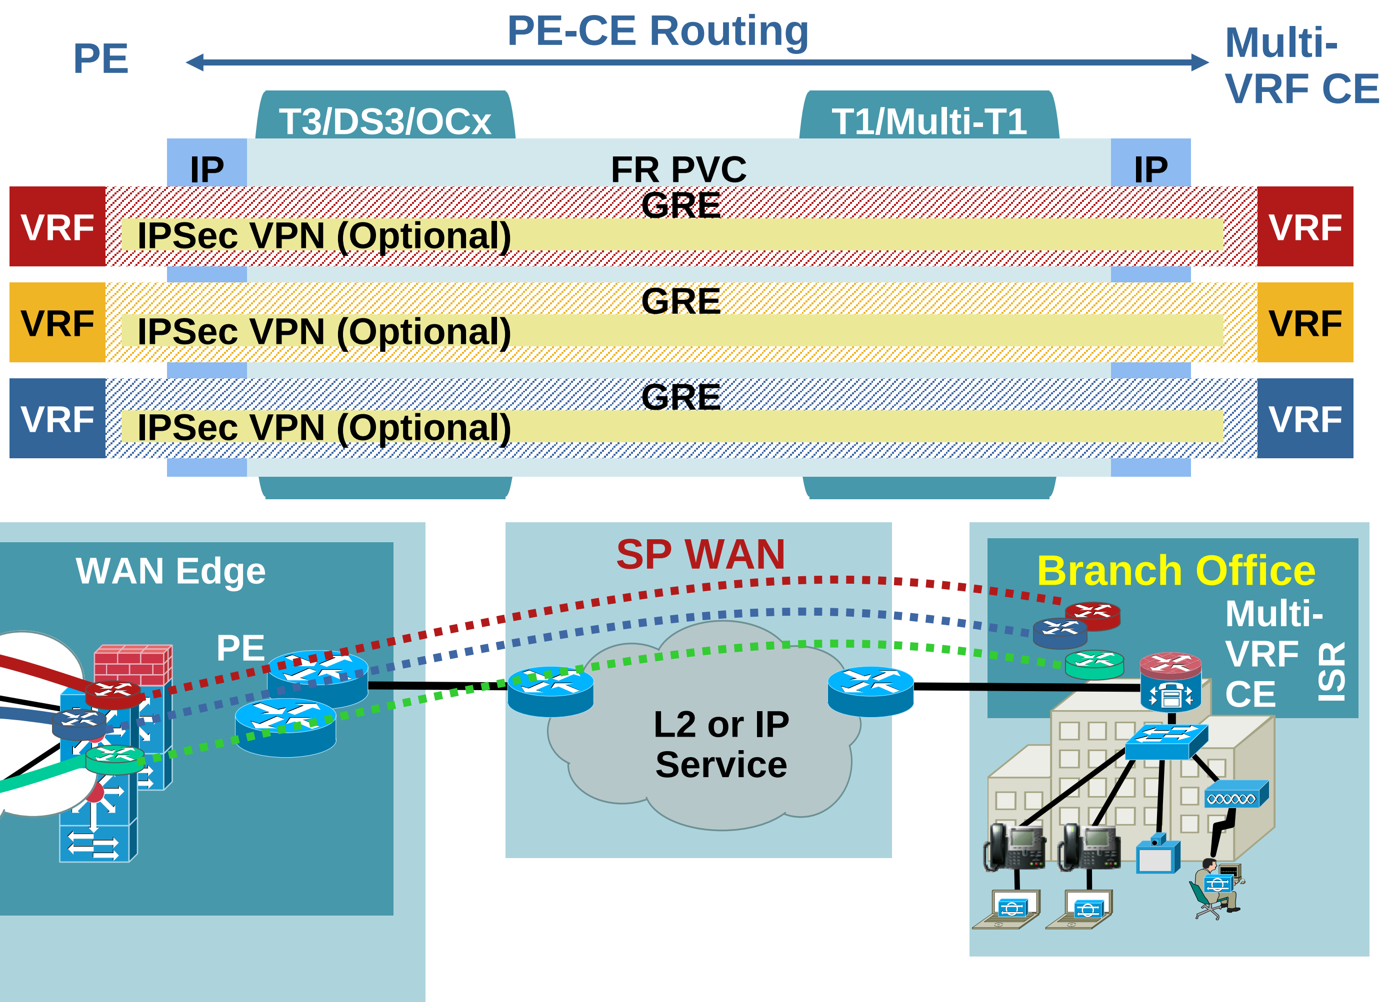

WAN Extensibility

MPLS over L2 (PPP, Frame, ATM, Leased Line)

- Single SP VC
- MPLS cloud tunneled over the WAN
- GRE encapsulation optional (required for IPsec)
- LSPs segment data path

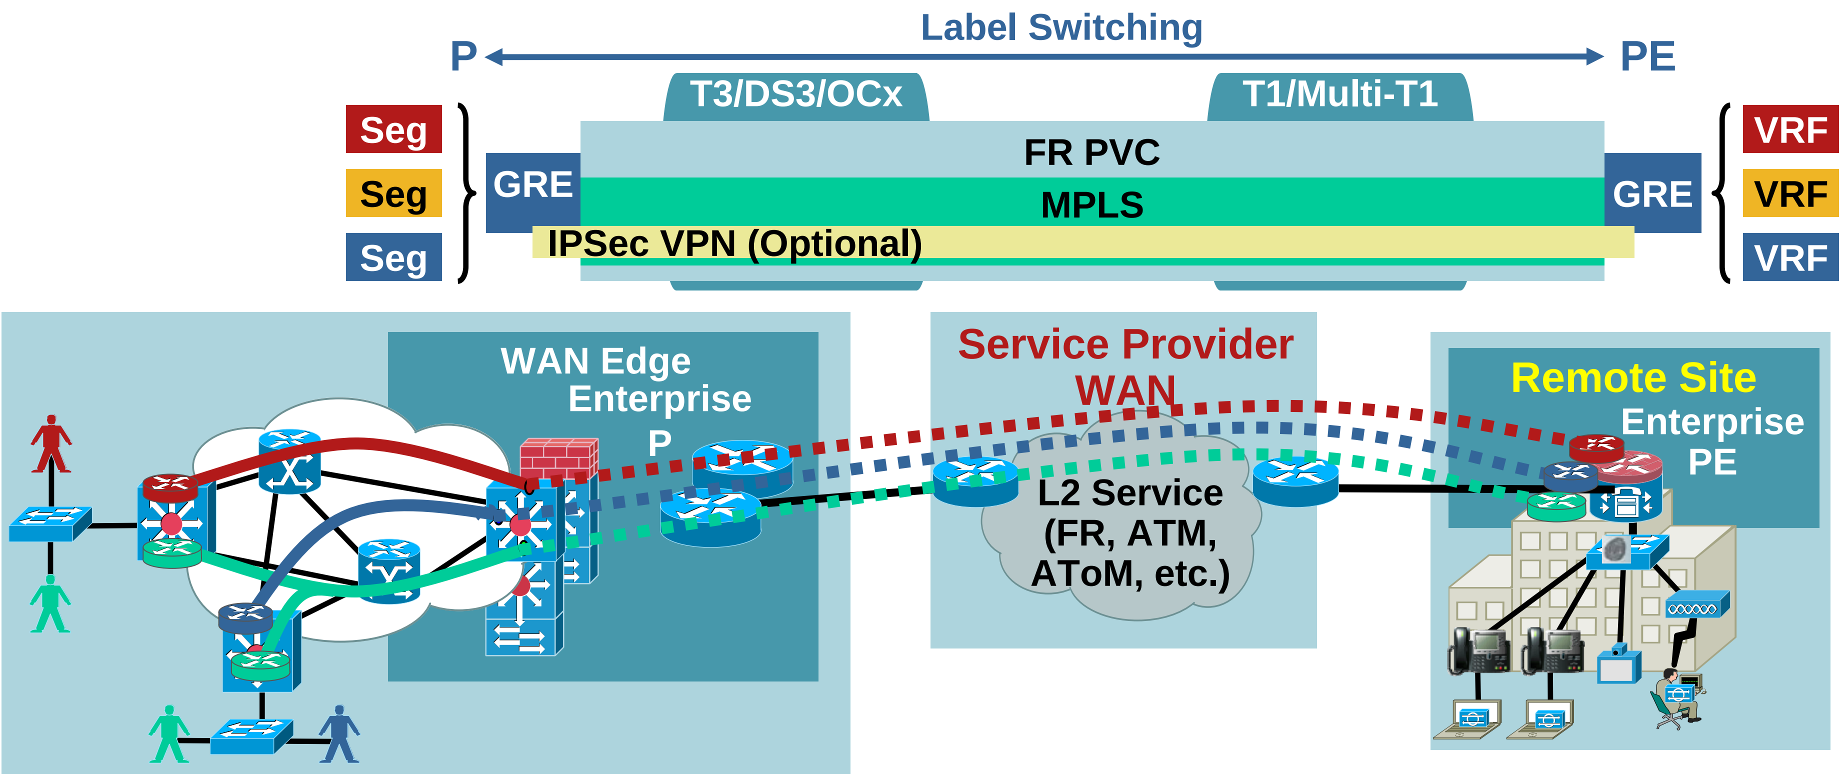

WAN Extensibility

MPLS over Tunnel Overlay over IP Cloud

- Single SP IP VPN
- Enterprise deploys a tunnel overlay
- Tunnel interfaces are label switched
- LSPs segment data path

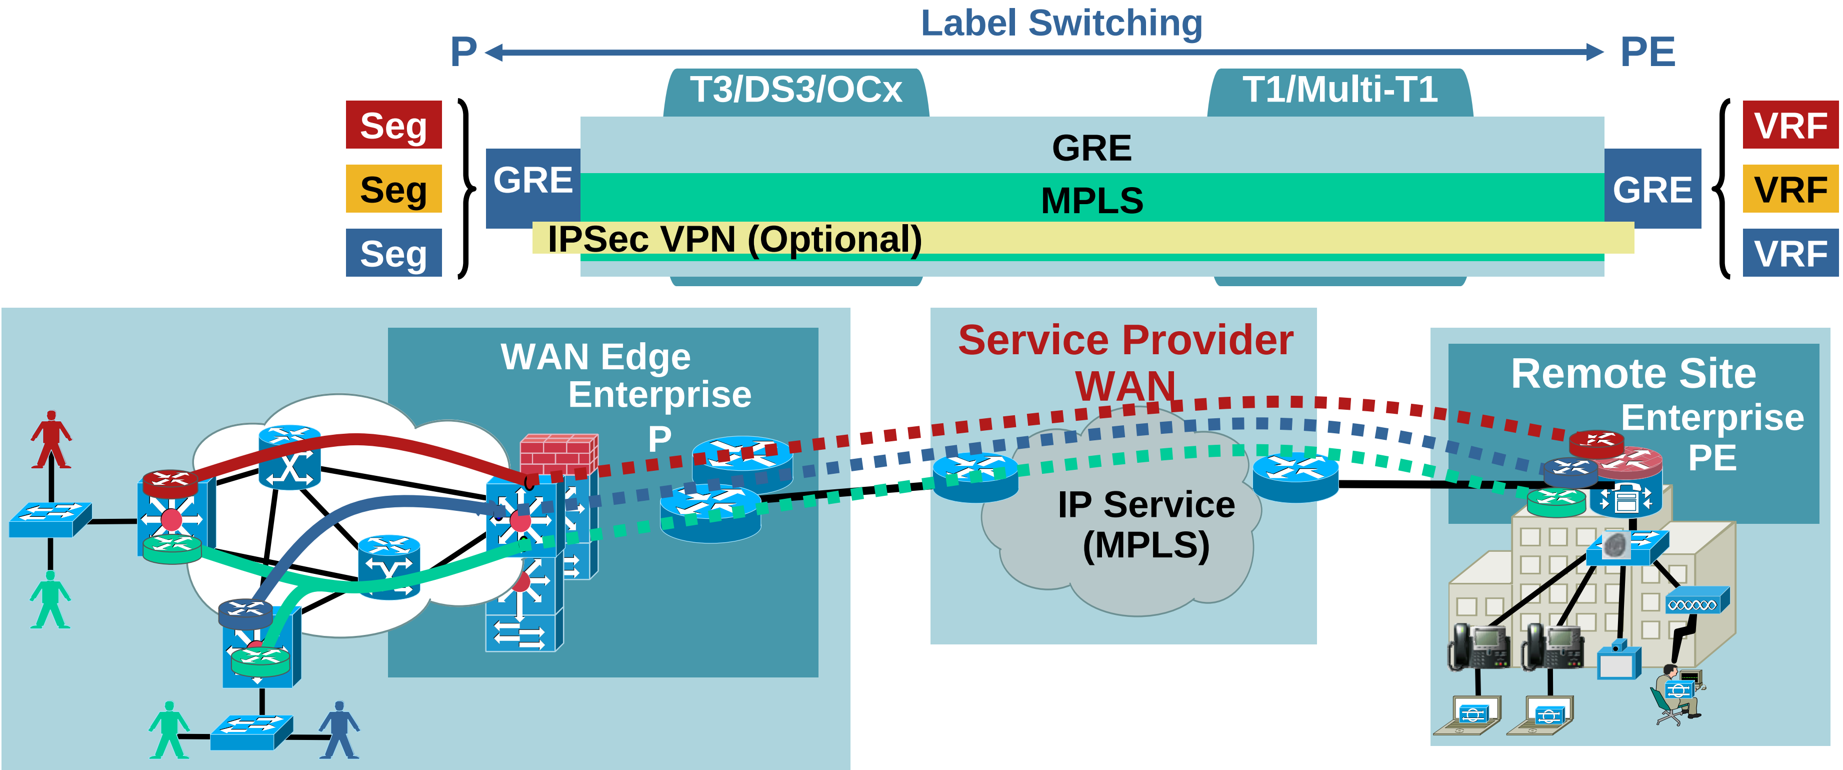

WAN Extensibility
Summary

- The virtual networks must be extended over the WAN
- We discussed several alternatives:
  - Per VPN GRE tunnels for PE-CE routing on private circuits
  - MPLS over private L2 circuits
  - MPLS over tunnel overlay over IP service
- Other alternatives could include:
  - Carrier-supporting-carrier (if the service was available)
  - RFC2547 over DMVPN
- The choice depends largely on the Enterprise’s WAN contracts and existing circuits
- Next Generation MPLS VPN MAN Design Guide:
  - http://www.cisco.com/go/srnd
  - http://www.cisco.com/univercd/cc/td/doc/solution/esm/ngmane.pdf

Agenda

- Problem Definition
- Campus Virtualization Alternatives
- WAN Extensibility
- Shared Services and Inter-VPN Communication
  1. FW/ACL Controlled Services
  2. Inter-VRF Route Leaking → Extranets
- Data Center Integration

Shared Services and Inter VPN Communication

Two Basic Models

1. Controlled by firewalls/ACLs

    Provides protected access to shared services

    Provides protected communication between VRFs

    Is equivalent to interconnecting separate IP networks

    → Routing between networks occurs at specific GWY points

2. Route leaking between VRFs using a BGP process

    Provides un-protected communication between VRFs

    Allows extranet creation for shared services access

    Populates routing tables to enable reachability between VPNs

    → Routing between networks is optimal

    → No inter VPN policy enforcement possible

Shared Services and Inter-VPN Communication
FW + Fusion Router (Internet)

1

- Fusion router:
  - Inter-VPN connectivity
  - Shared resource connectivity
    - → Internet, servers, etc.
- FW contexts:
  - VPN isolation/protection
  - Per VPN policies: ACL, NAT …
  - 256 contexts per FW
  - Map to VLANs
- Shared services available:
  - On their own VPN (distributed)
  - Off the transit router or DMZ (centralized)
  - Access is always centralized

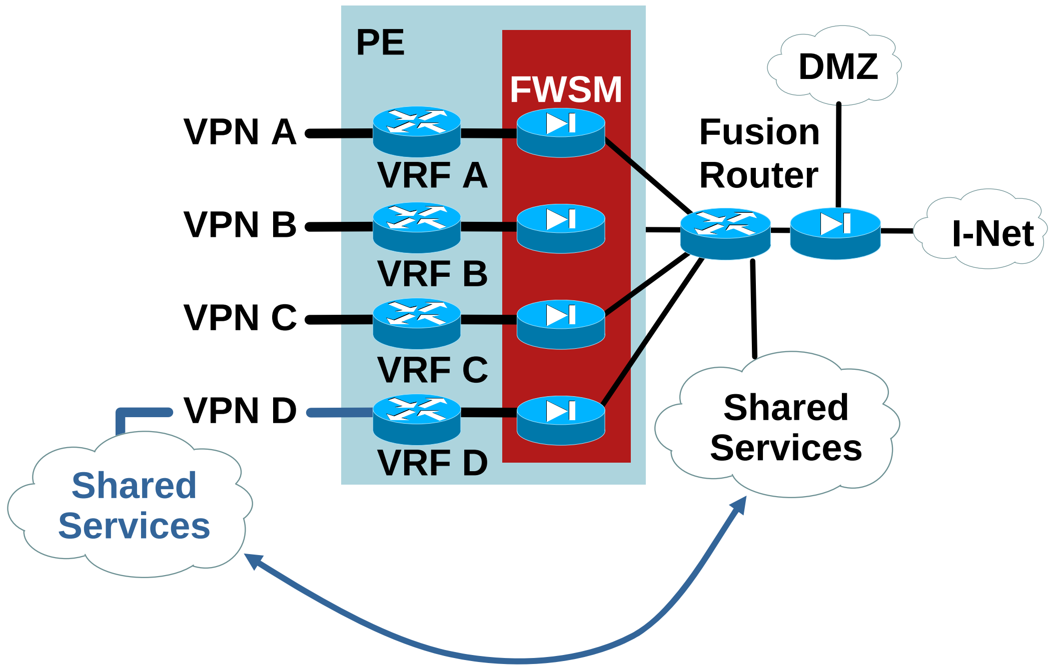

Context Functionality Also Available on PIX & ASA

Inter-VPN Communication
FW in Single Routed Mode

1a

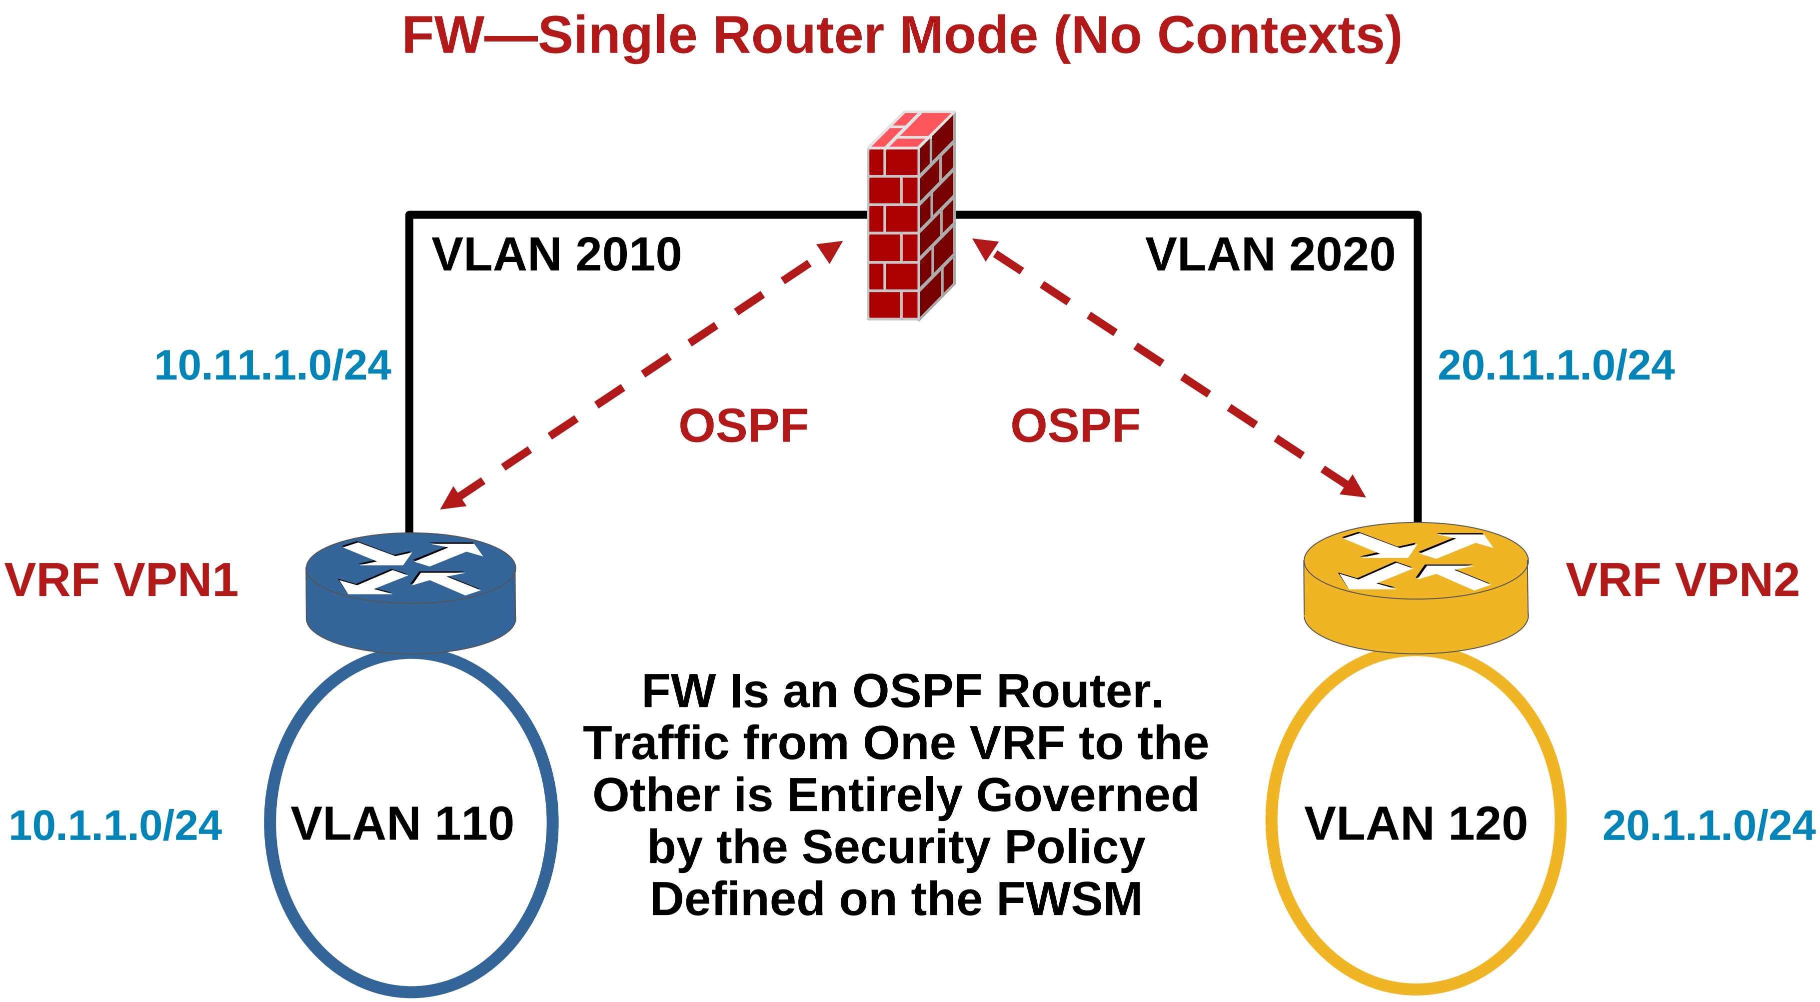

Inter-VPN Communication

FW in Transparent Mode

1b

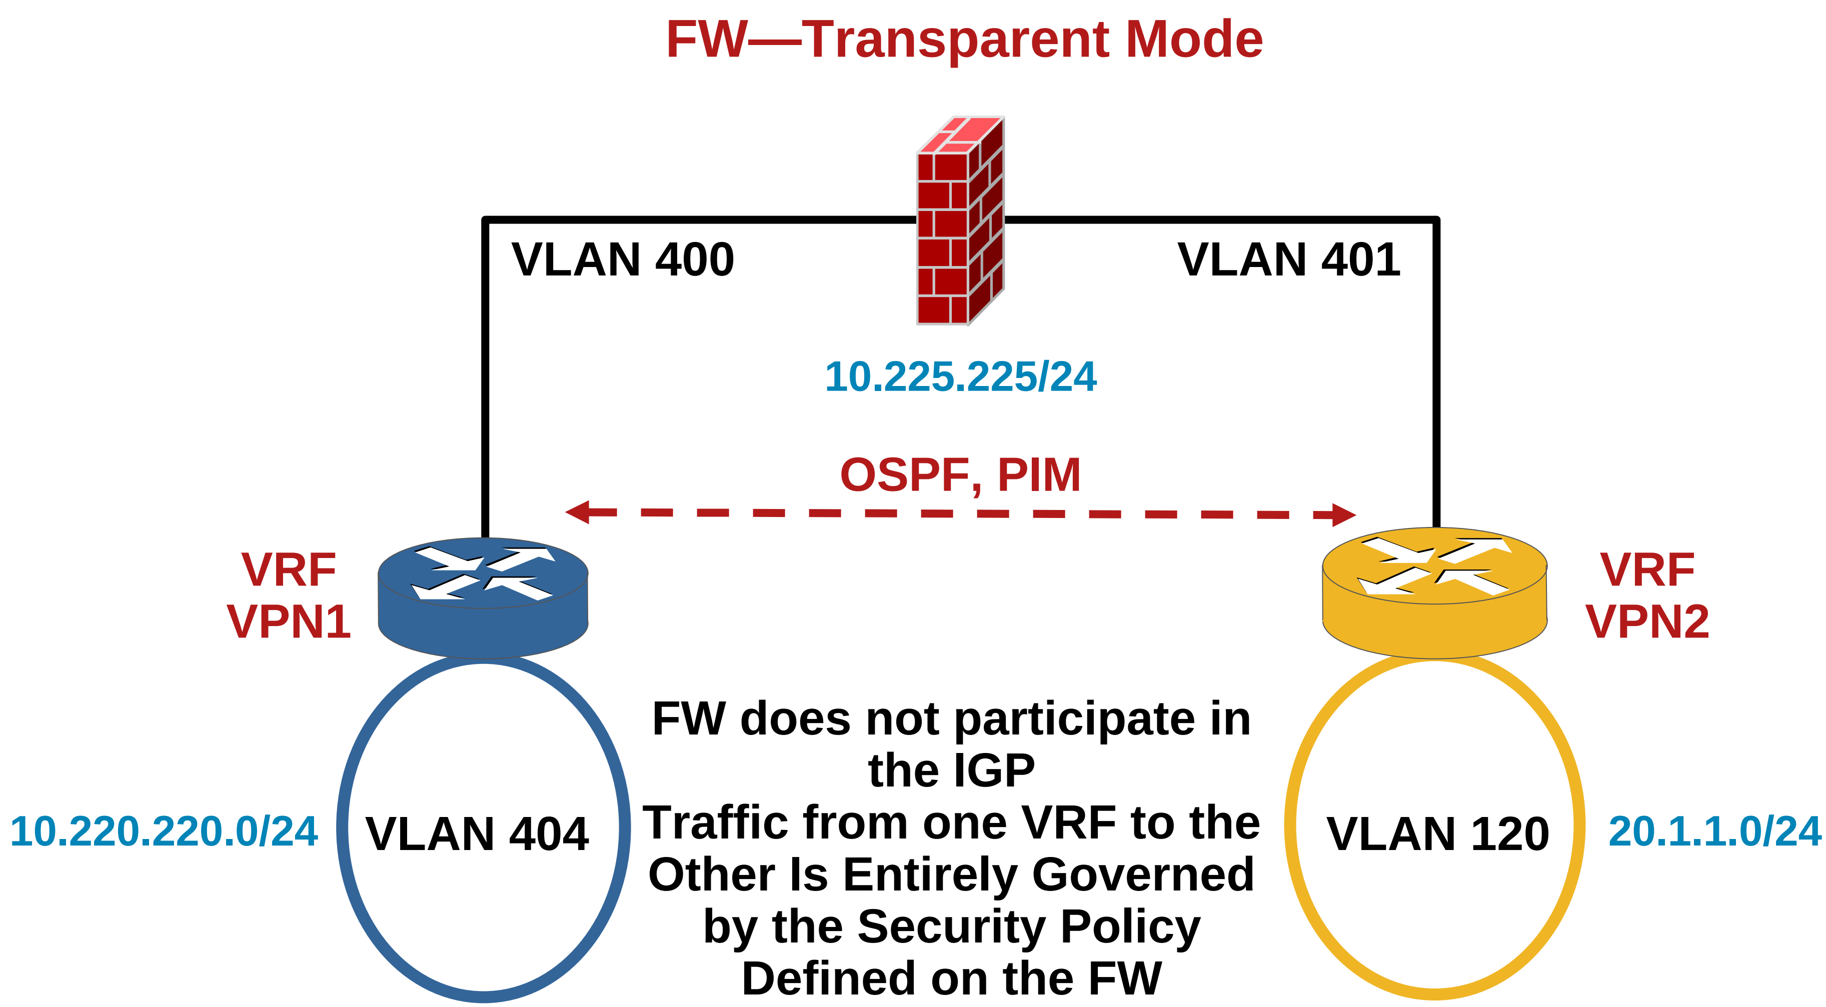

Inter-VPN Communication
Multi-Context Transparent Mode—Pairs

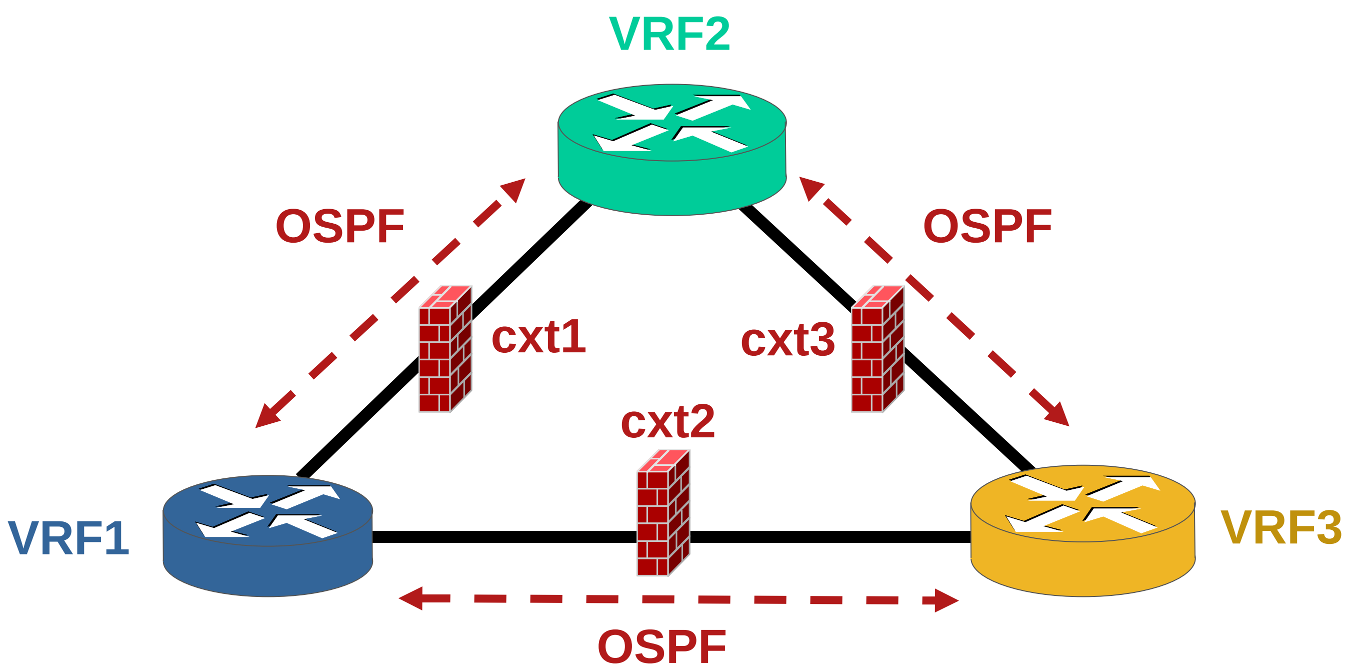

- One context per VRF pair, Transparent mode
- Filtering rules have to be done multiple times for each VRF pair

Inter-VPN Communication
Multi-Context Transparent Mode—Pairs

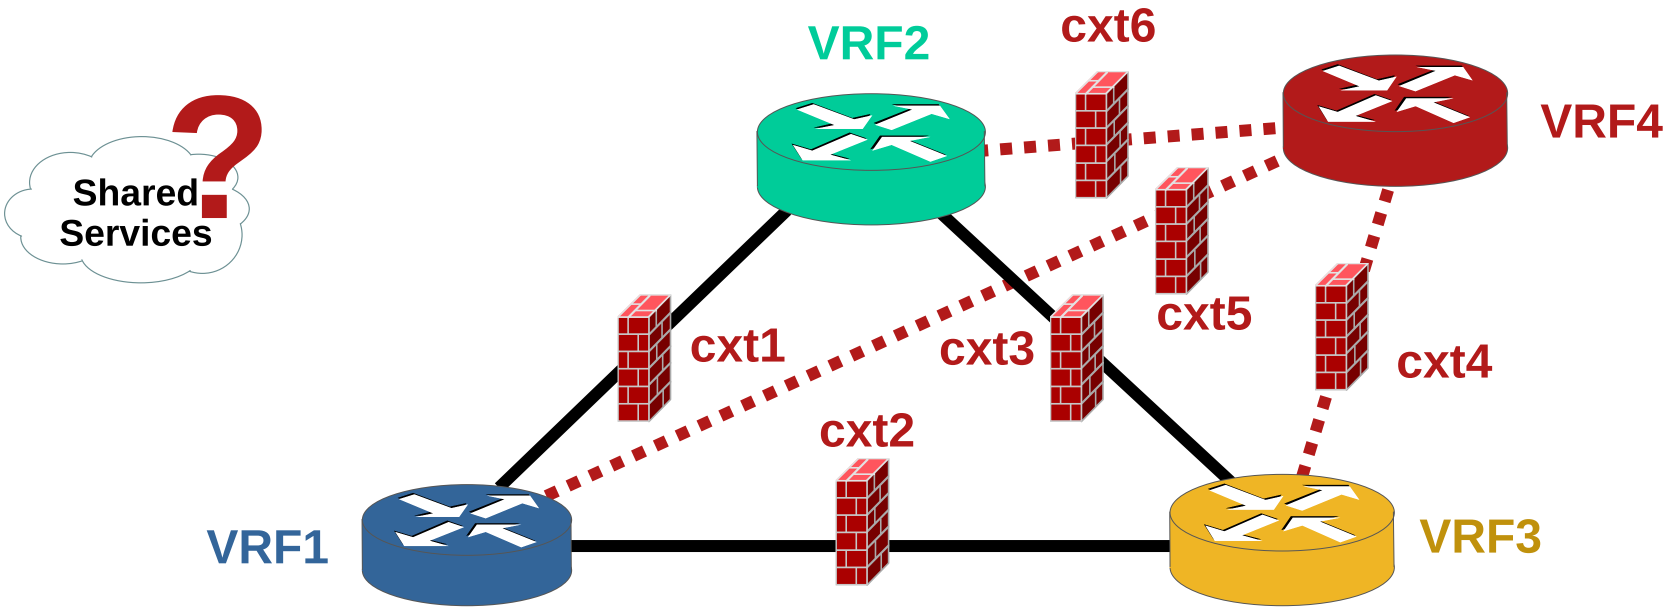

- One context per VRF pair, Transparent mode
- Filtering rules have to be done multiple times for each VRF pair
- Very limited scalability → an alternative is required
- How should shared services be reached?

Inter-VPN Communication

Transparent Mode—Fusion Router/VRF

1c

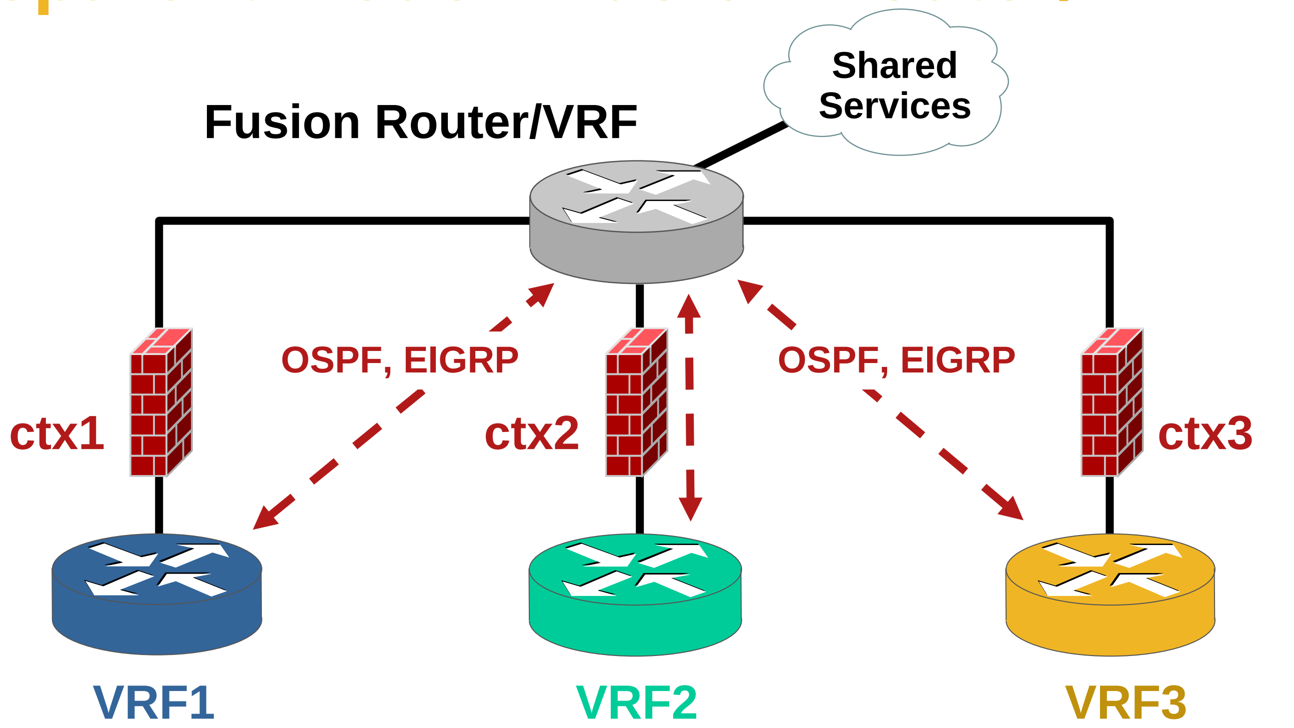

- Fusion Router/VRF (hub and spoke): All interVPN traffic must go through this Router/VRF
- FW Contexts could be managed per VPN
- Routing protocol between VRFs could be EIGRP to allow route filtering capabilities

Routed mode, NAT and Fusion VRF

Single Internet Access

1d

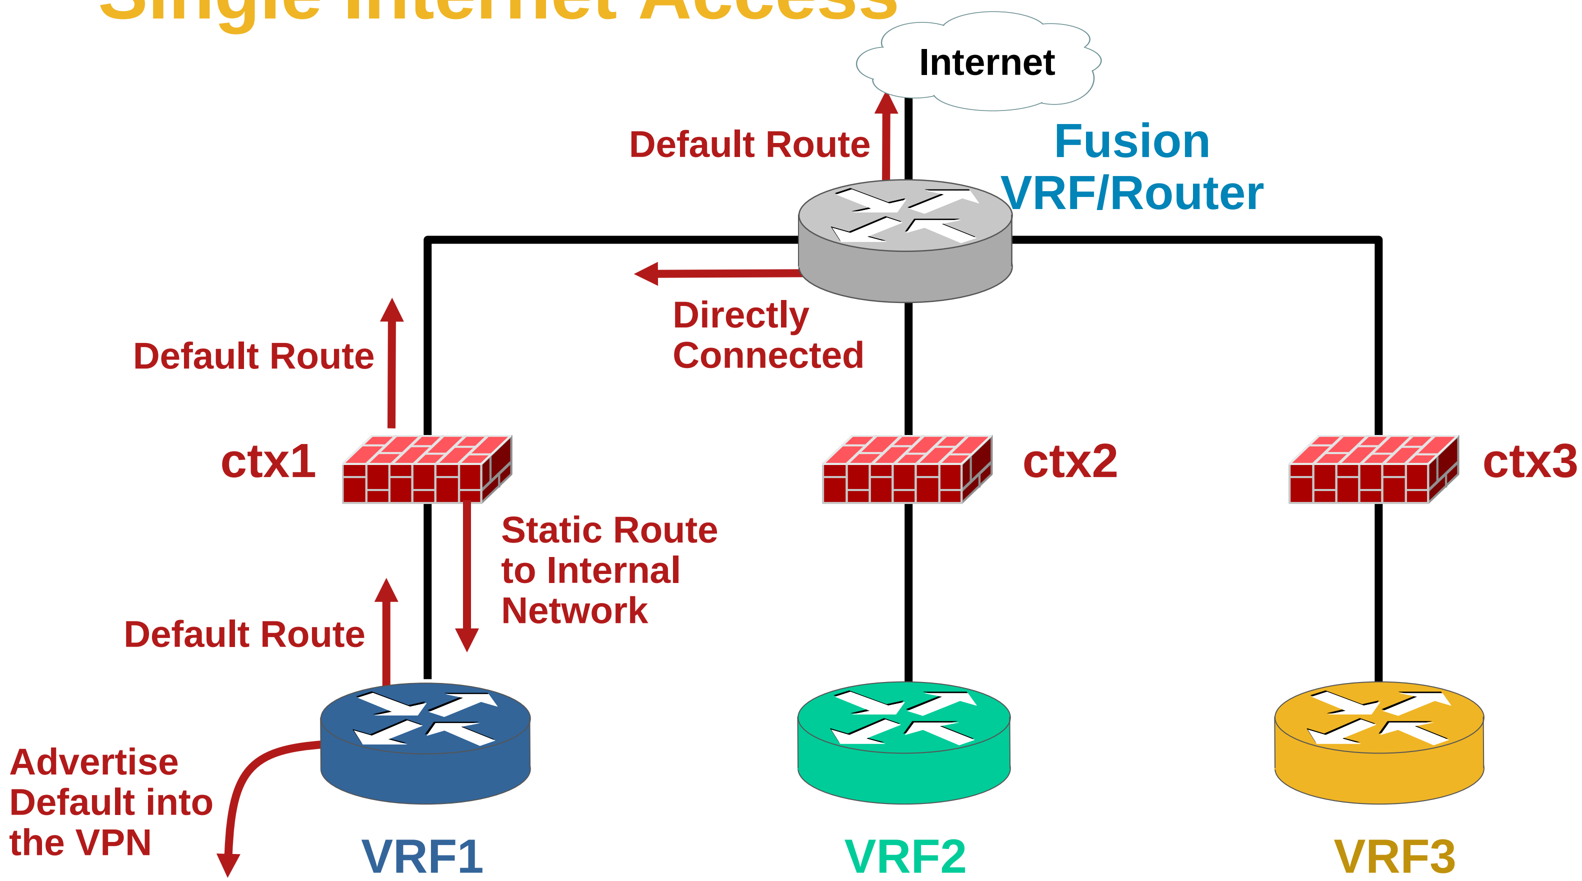

- Fusion VRF (hub and spoke): All inter-VPN traffic must go through this VRF
- Fusion VRF provides connectivity to shared services (internet)

Routed mode, NAT and Fusion Router

Redundant Internet Access

1e

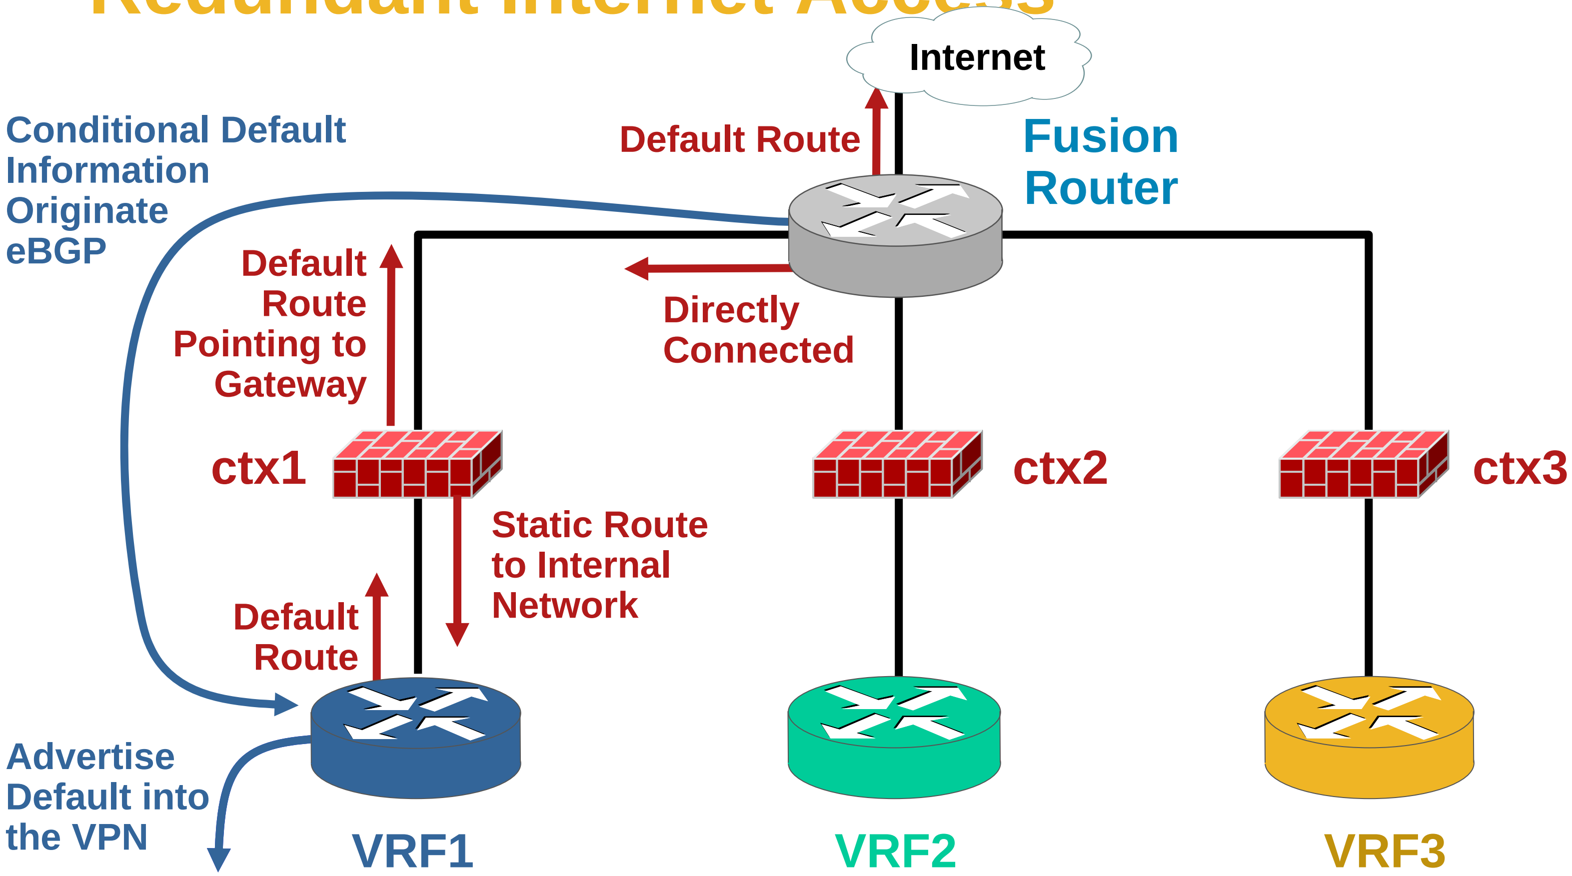

- Fusion Router (hub and spoke): All inter-VPN traffic must traverse Fusion Router
- When internet becomes unreachable, the Default route is not included in eBGP → Other default routes can then become active

Fusion VRF Single Device Implementation
eBGP peering between VRFs on a single router

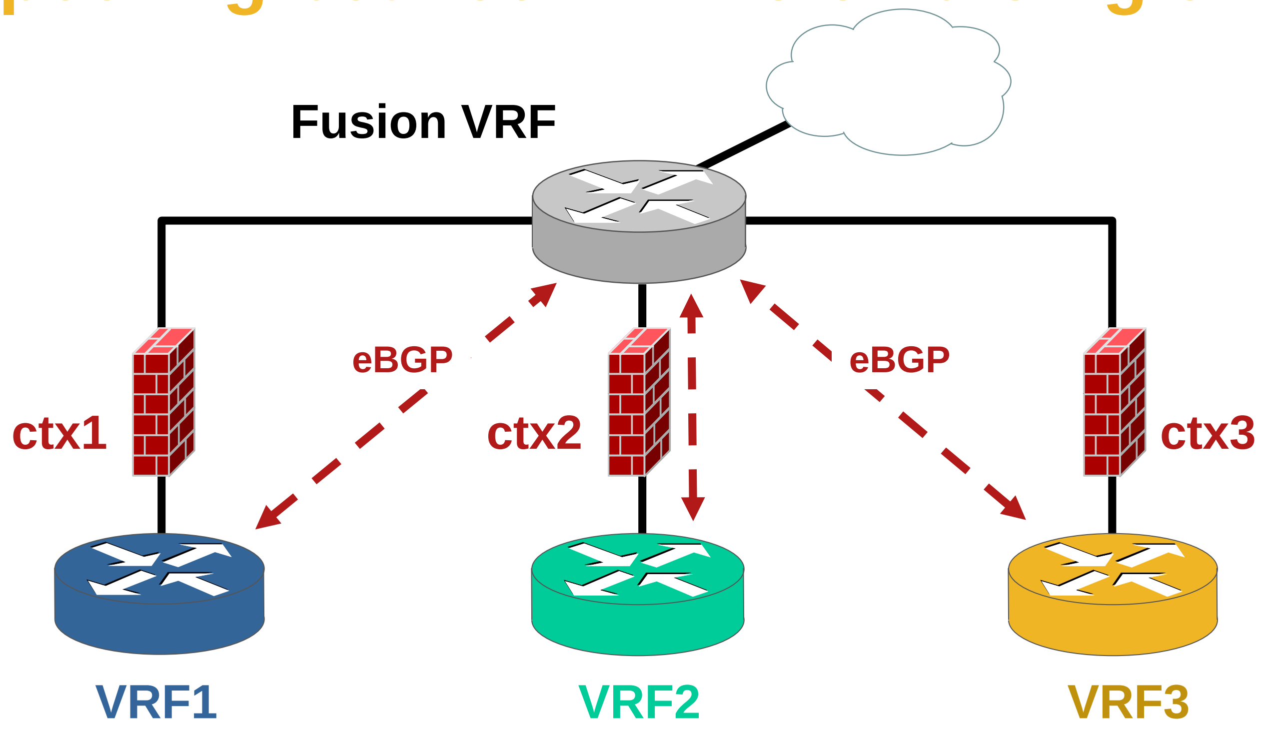

- All VRFs (including Fusion) reside on the same physical device
- eBGP peering within the same device requires:
  - BGP router-id per VRF
  - Multi-AS support for BGP

Fusion VRF – eBGP same box

Configuration Sample

```
router bgp 1
 no synchronization
 bgp router-id 10.149.149.1
 no auto-summary
 !
 address-family ipv4 vrf VRF2
 neighbor 1.1.1.1 remote-as 10
 neighbor 1.1.1.1 local-as 20 no-prepend replace-as
 neighbor 1.1.1.1 ebgp-multihop 2
 neighbor 1.1.1.1 update-source Loopback20
 neighbor 1.1.1.1 activate
 no synchronization
 bgp router-id 2.2.2.2
 exit-address-family
 !
 address-family ipv4 vrf VRF1
 neighbor 2.2.2.2 remote-as 20
 neighbor 2.2.2.2 local-as 10 no-prepend replace-as
 neighbor 2.2.2.2 ebgp-multihop 2
 neighbor 2.2.2.2 update-source Loopback10
 neighbor 2.2.2.2 activate
 no synchronization
 bgp router-id 1.1.1.1
 exit-address-family
```

Dual AS Support
Available on
Cisco 7600 12.2(33)SRA
Catalyst 6500 12.2(33)SXH (Future)

BGP router-id per VRF
Available on
Cisco 7600 12.2(33)SRA
Catalyst 6500 12.2(33)SXH (Future)

Agenda

- Problem Definition
- Campus Virtualization Alternatives
- WAN Extensibility
- Shared Services and Inter-VPN Communication
    1. FW/ACL Controlled Services
    2. Inter-VRF Route Leaking → Extranets
- Data Center Integration

Some Background—Understanding VRFs

2

Route Distinguishers

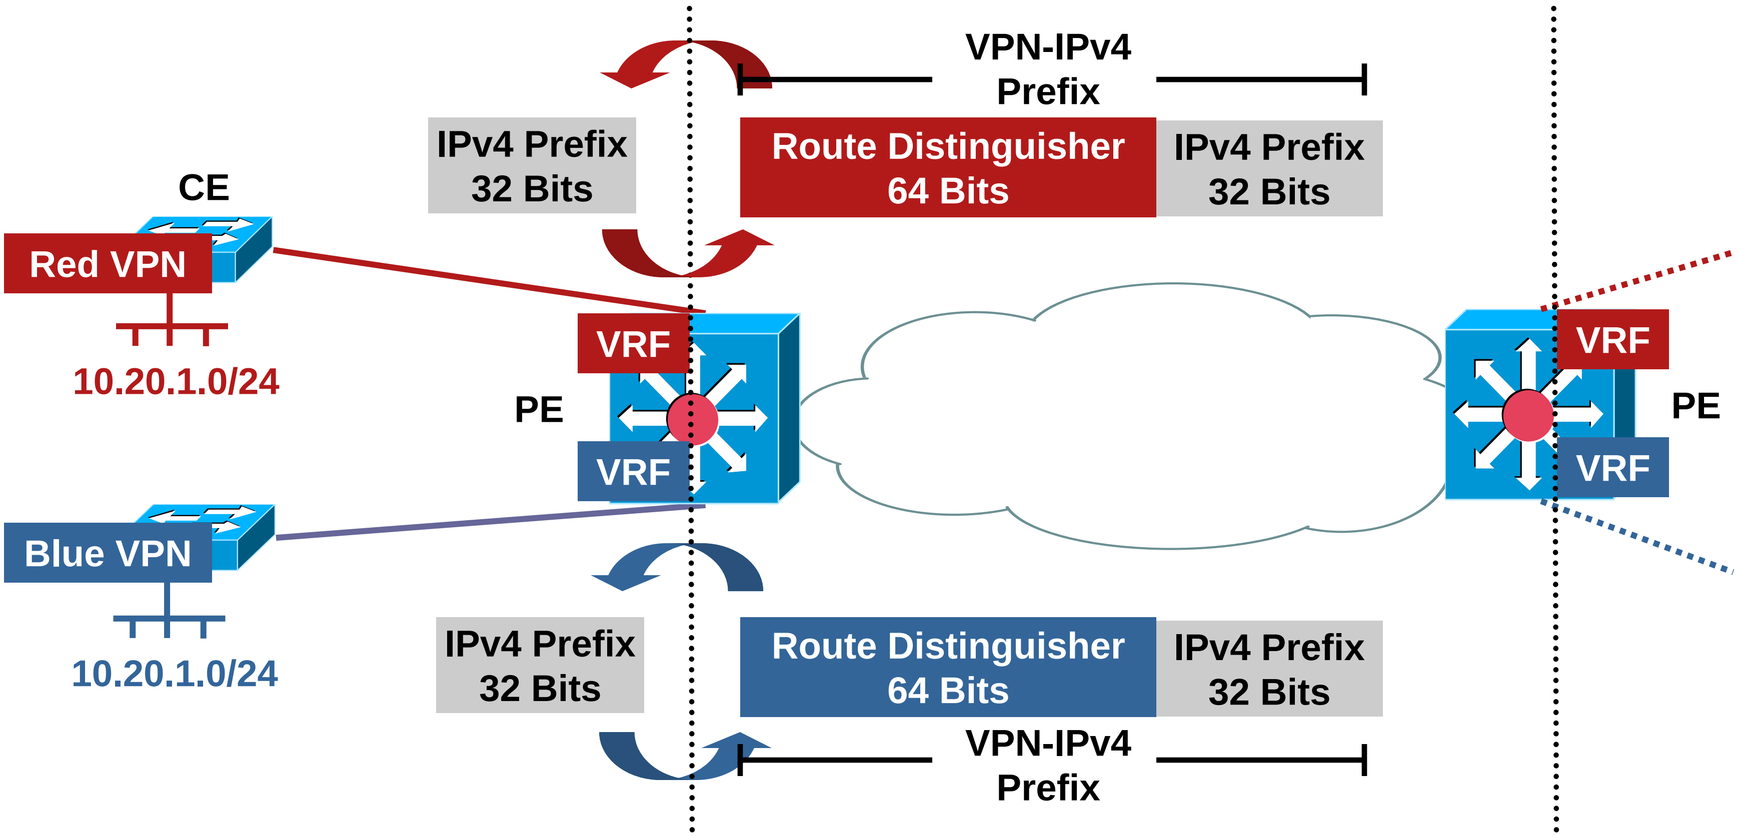

Route Distinguisher (RD)

- VPN-IPv4 prefix = RD + IPv4 prefix
- Locally significant

Understanding VRFs
Route Targets

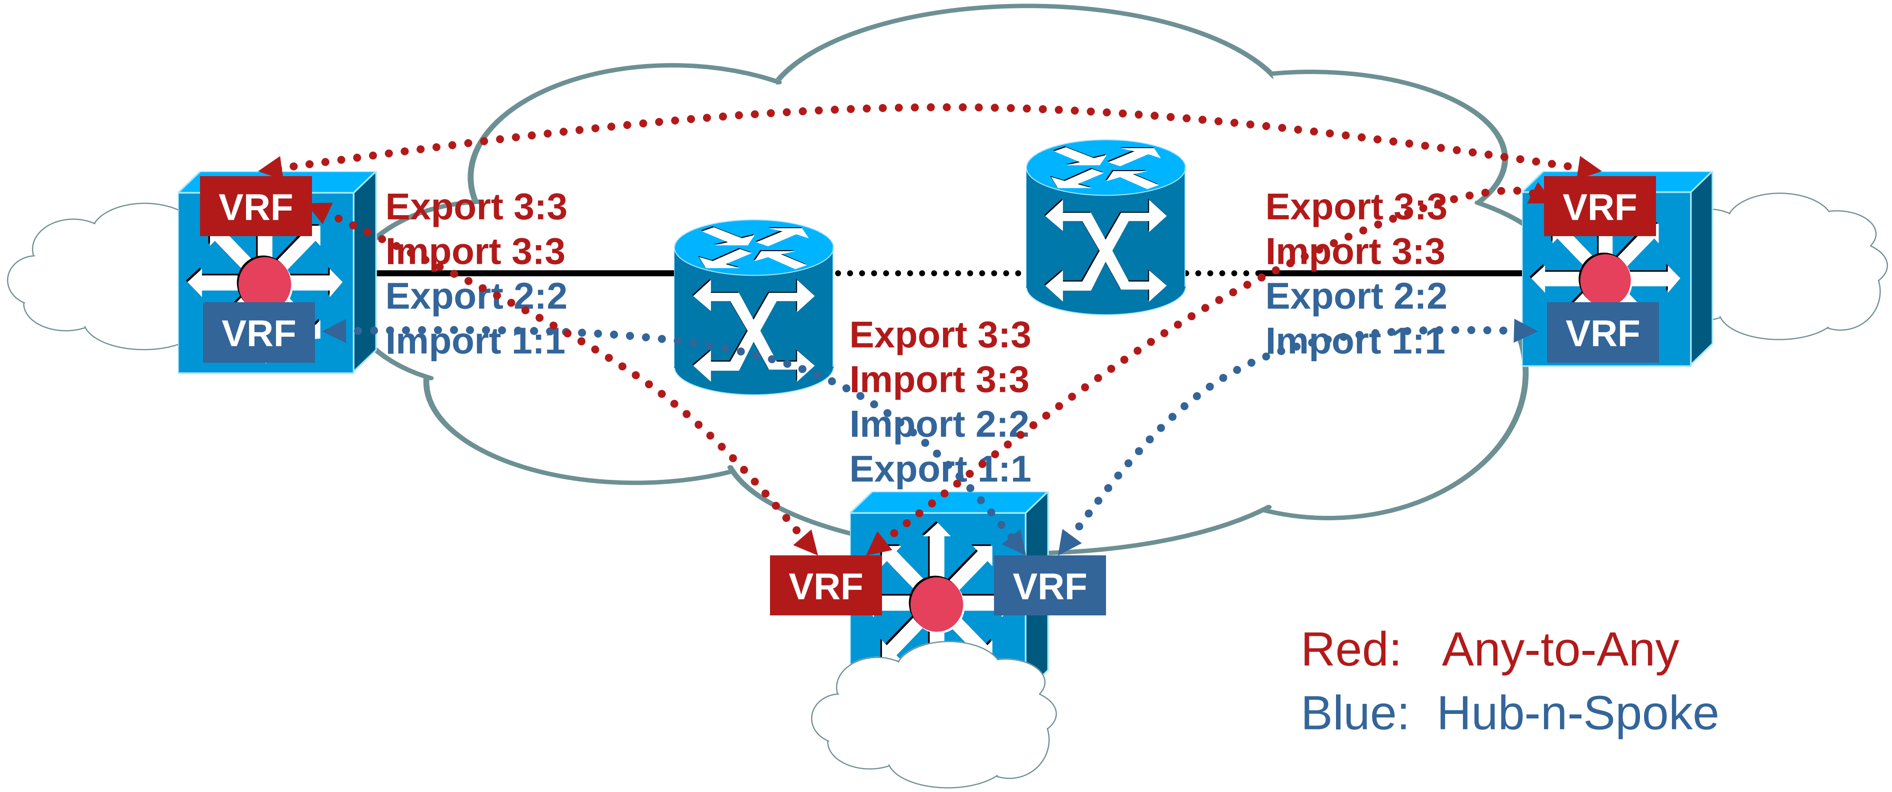

- Import/export routes to/from MP-BGP updates
- Globally significant—creates the VPN
- Allows hub and spoke connectivity (central services)

Shared Services Extranet VPN

Multiple-Box Extranet Implementation

2a

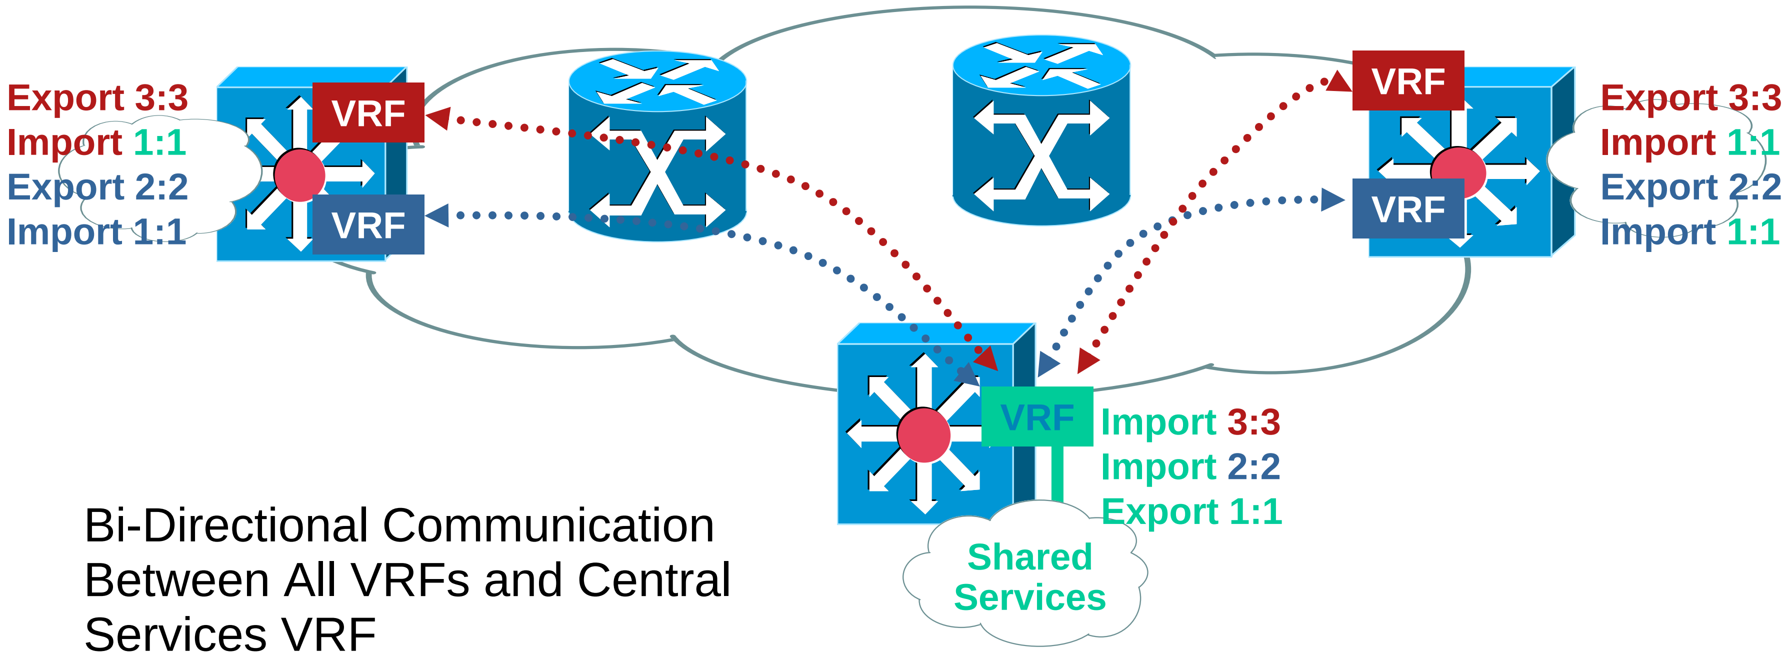

Bi-Directional Communication Between All VRFs and Central Services VRF

- Central services routes imported into both VRF red and blue (1:1)
- Central VRF imports routes for blue and red subnets (3:3, 2:2)
- No routes exchanged between blue/red
- No transitivity: imported routes are not “re-exported”

  → Blue and red remain isolated

Route Leaking Between VRFs

Single Box Extranet—Using a BGP Process

2b

a BGP Process to Leak the Routes Between VRFs

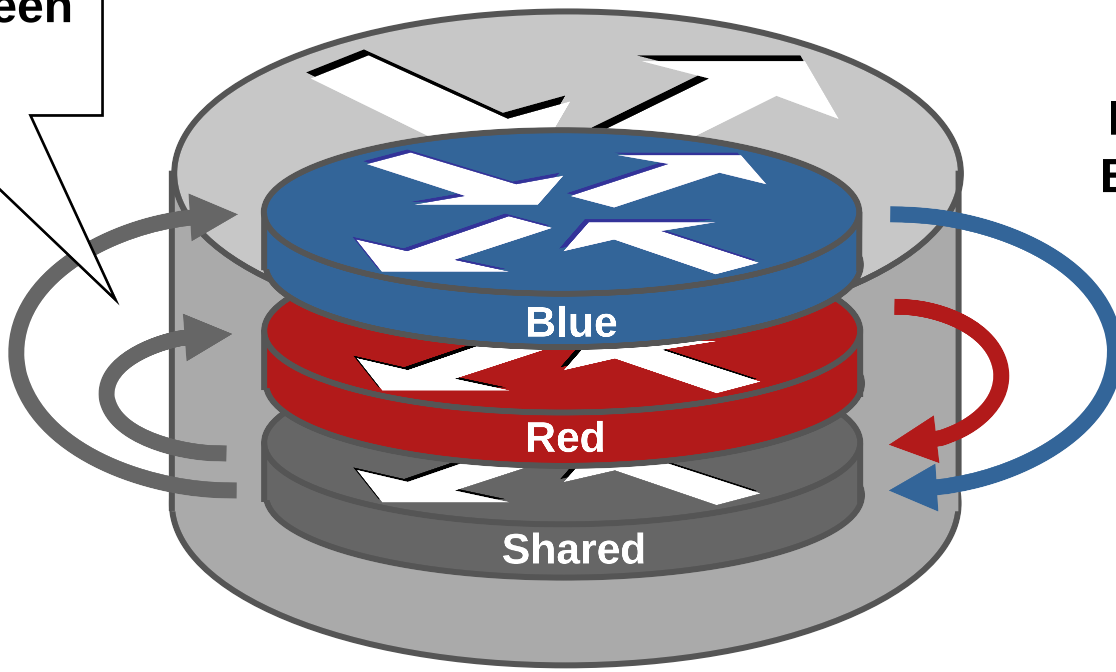

Import-Export Between VRFs Using RT

Single Box Extranet Implementation

VRF Configuration—Services Extranet VPN

```
ip vrf SERVICES
 rd 10:10
 route-target export 1:1
 route-target import 1:1
 route-target import 3:3
 route-target import 2:2
!
ip vrf RED
 rd 30:30
 route-target export 3:3
 route-target import 3:3
 route-target import 1:1
!
ip vrf BLUE
 rd 20:20
 route-target export 2:2
 route-target import 2:2
 route-target import 1:1
```


Single Box Extranet Implementation
BGP Process

```
router bgp 65001
 bgp log-neighbor-changes
 !
 address-family ipv4 vrf BLUE
 redistribute ospf 2
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf RED
 redistribute ospf 1
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf SERVICES
 redistribute ospf 3
 no auto-summary
 no synchronization
 exit-address-family
!
```

- Need a BGP process to leak the routes between VRFs
- Don't need any BGP neighbors/ sessions

Shared Services and Inter-VPN Communication
Summary

Two Basic Models to Share Services:

1. FW/ACL controlled
   - Allows address overlap
   - Allows per VPN policies
   - Secure inter-VRF communication
   - Higher complexity and higher flexibility
2. Inter-VRF route leaking → Extranet
   - No address overlap
   - Single shared policy
   - Open inter-VRF communication
   - Less complexity and less flexibility

Agenda

- Problem Definition
- Campus Virtualization Alternatives
- WAN Extensibility
- Shared Services and Inter-VPN Communication
- Data Center Integration

Virtualization and the Data Center

Many Aspects to a Single Buzz Word…

- Virtual connectivity services
  - IP/MPLS and VPN
  - Virtual Route Forwarding VRFs
- Virtualized front-end
  - VLANs and private VLANs
  - Virtual intelligent services (firewall, L4–7, etc.)
- Compute virtualization
  - Clustering, GRID, virtualization software
- Virtualized storage
  - Virtual SANs (VSANs)
  - Network-hosted storage virtualization software

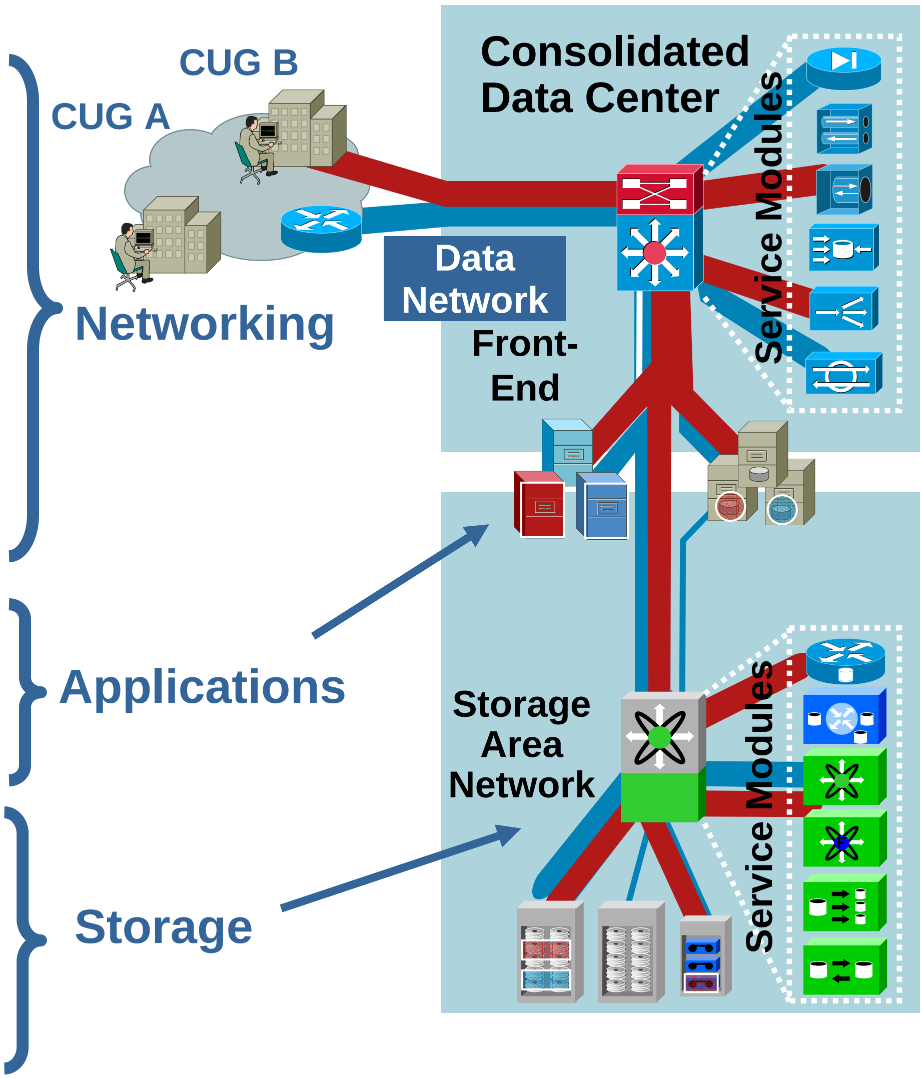

Data Center Aggregation
Interface Between Campus and Data Center

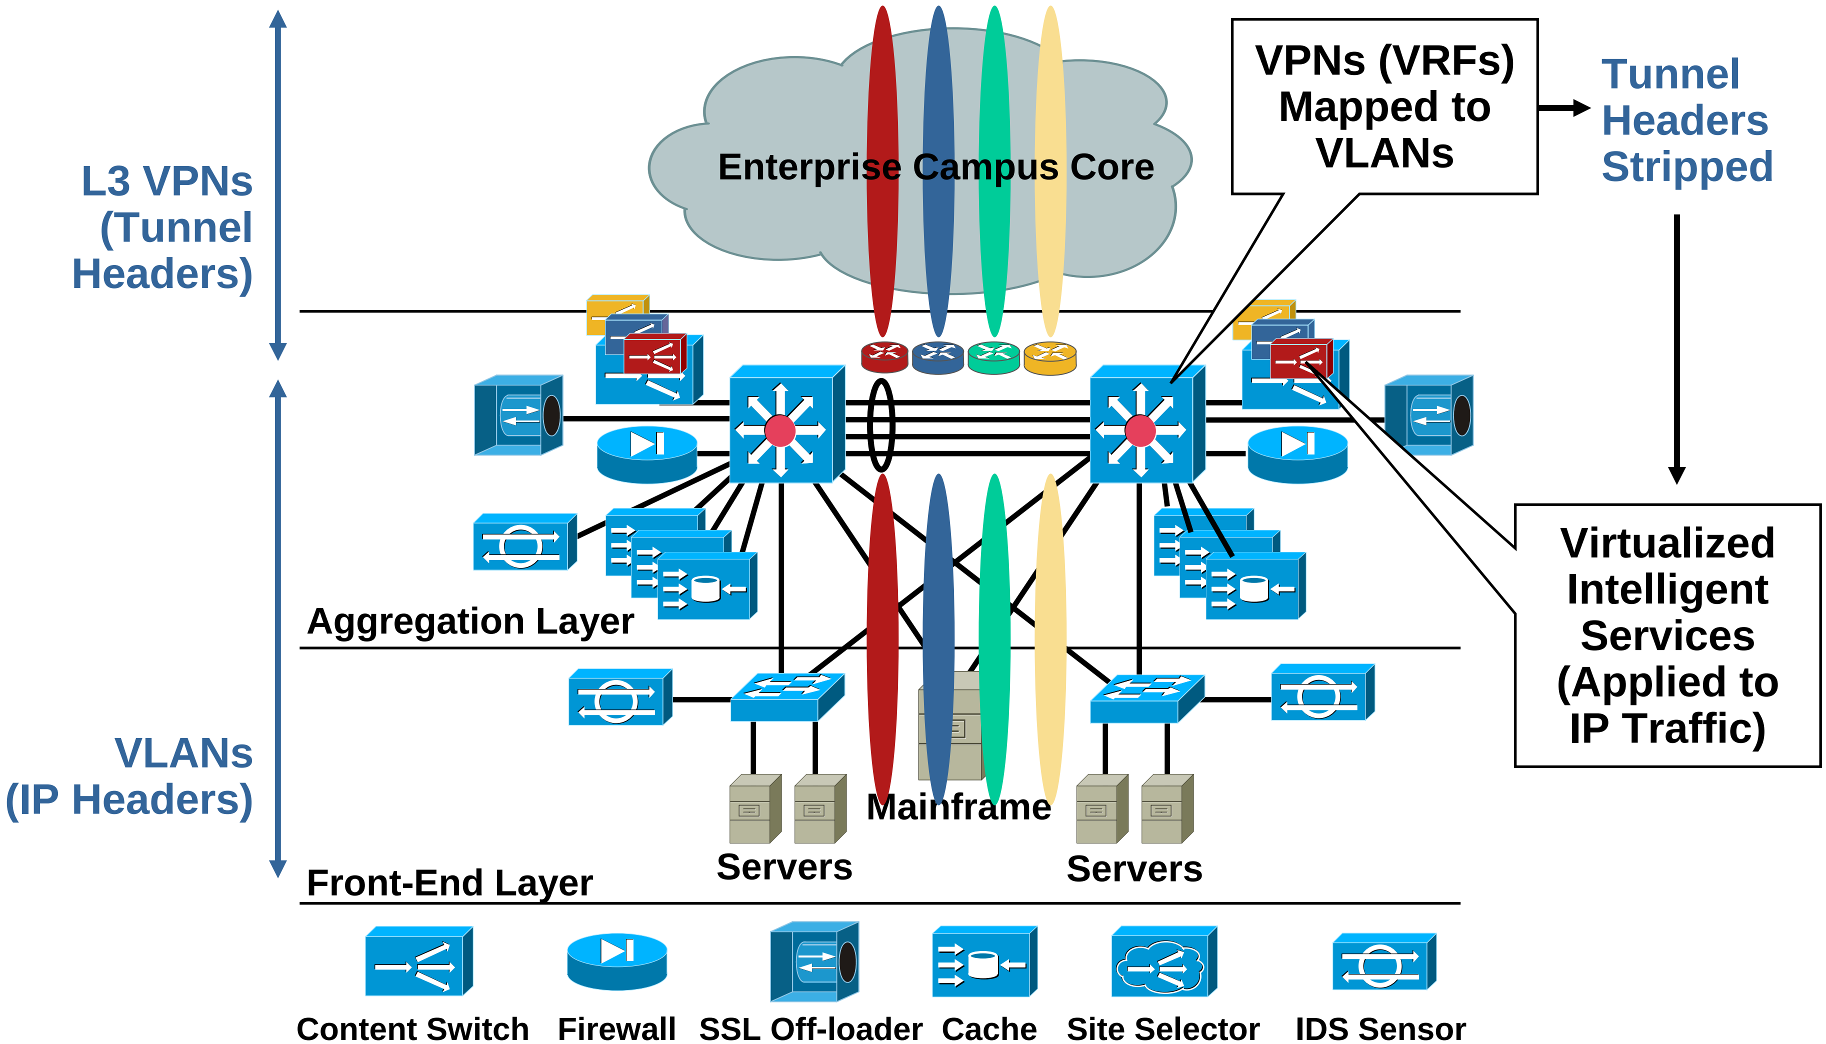

Data Center as a Shared Service on an Extranet VRF

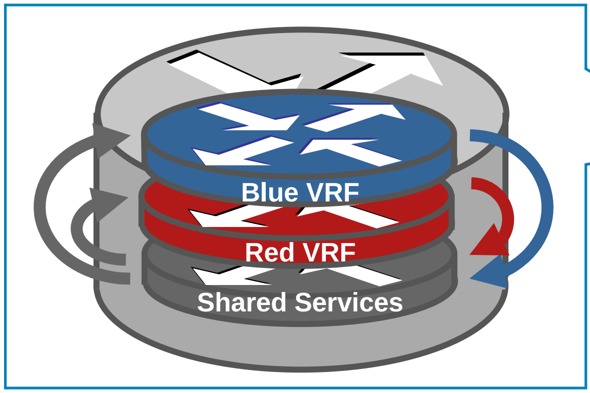

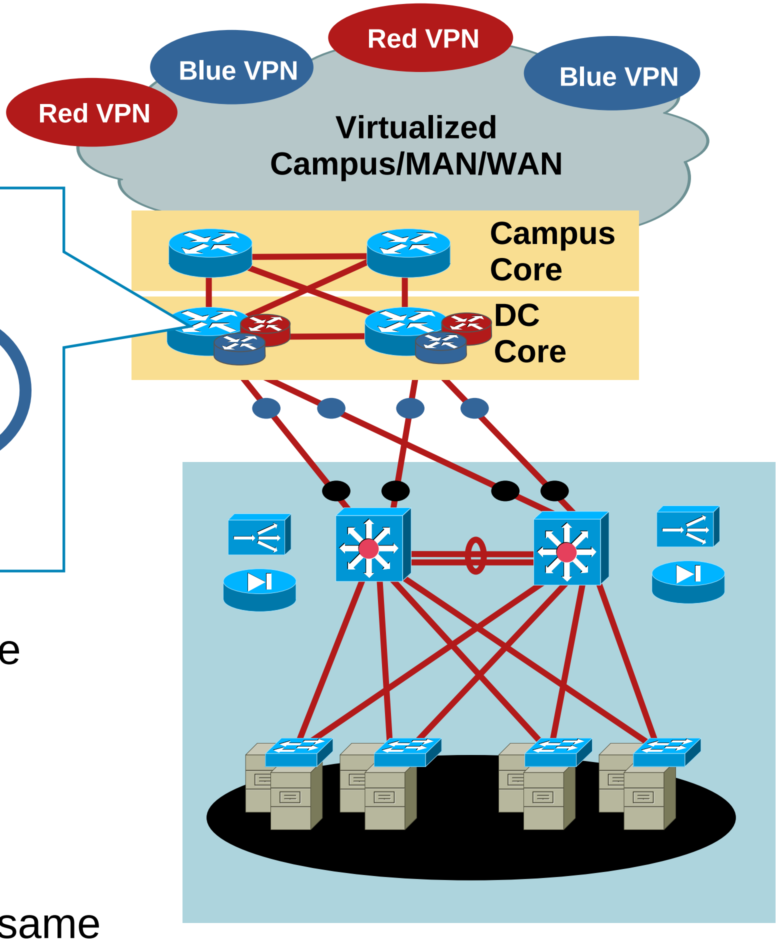

- VRFs terminate at the DC core
- Virtualization has no impact on DC front-end design
- All groups share the same front-end policies

- L3 Interface w/o VRF Enabled
- L3 Interface with Extranet-VRF Enabled

Shared Server Farm—Extranet VRF

- VRFs terminated at the DC core
- Extranet VRF
- Single shared server farm
- Only load balanced traffic to CSM
- Shared policies
- DoS load sharing between FW and Load Balancer
- PBR for return traffic
- FW transparent recommended

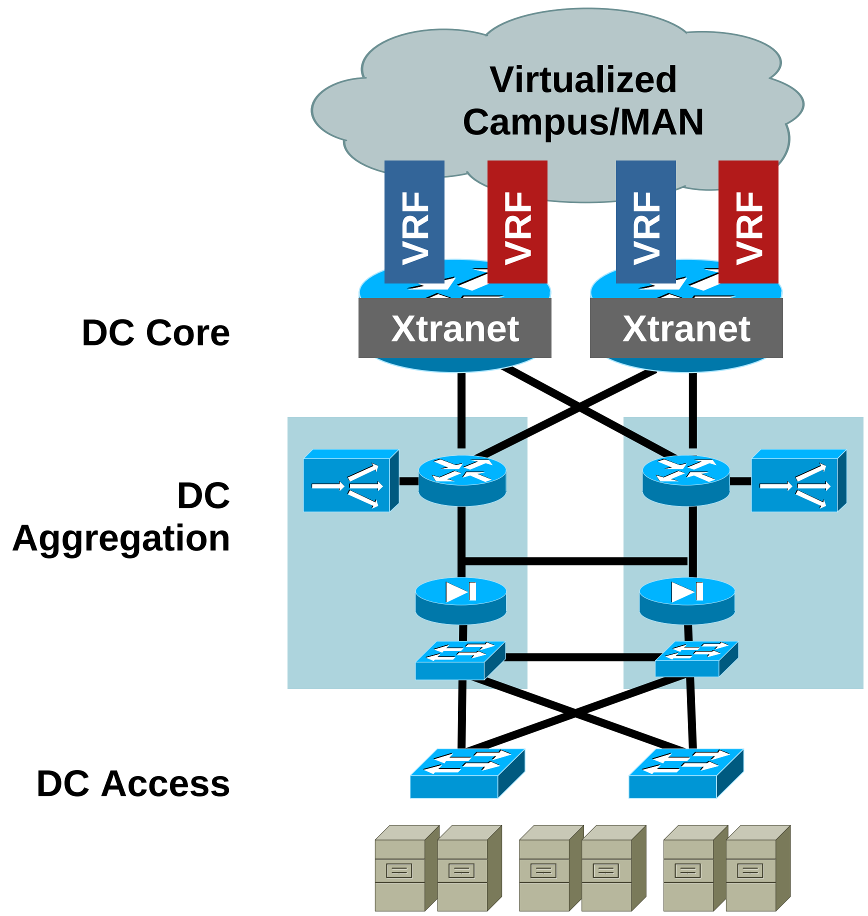

DC Front-End Segmentation Using ACE and FWSM 3.1

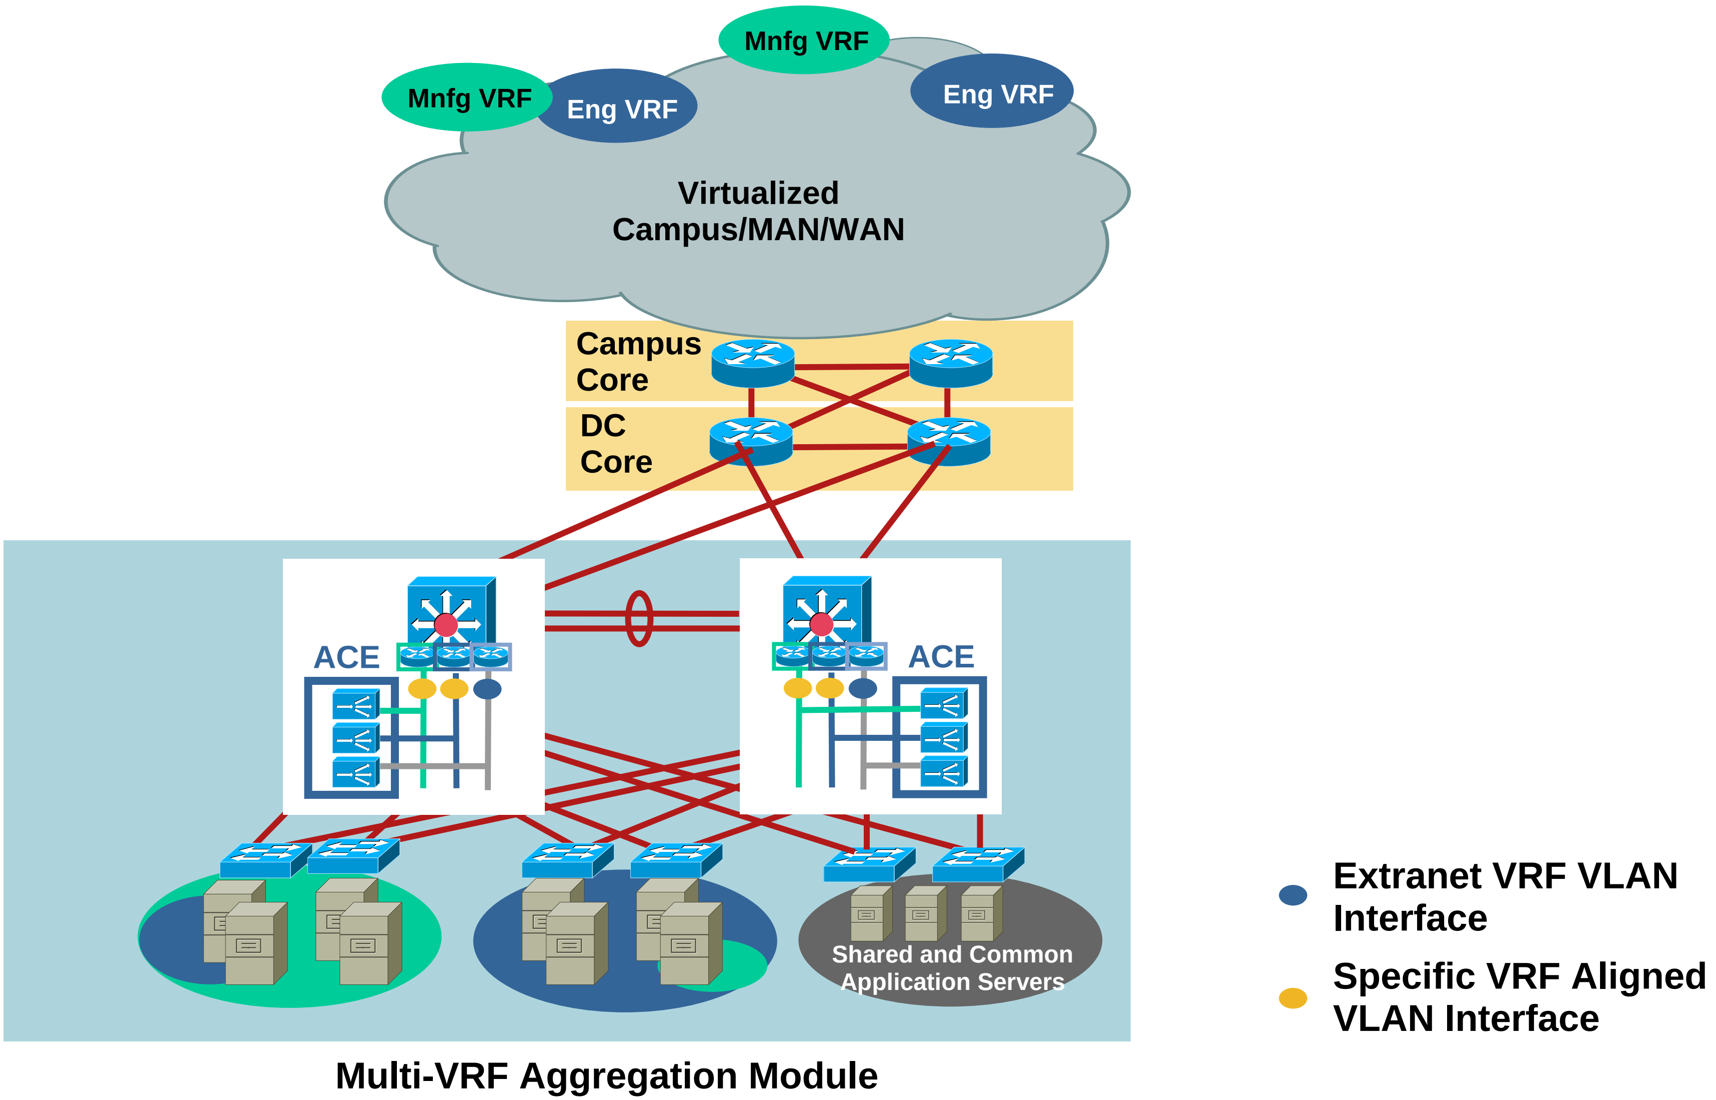

One-Arm Design with Dedicated Server Farms

- VRFs terminated at MSFC
- Dedicated context per outside VLAN
- Dedicated server farms in a partitioned data center
- Some DoS load sharing
- Address Re-use possible
- Per context management
- SourceNAT for return traffic
- ACE performance 15.5 Gbps

  → Only select traffic to ACE

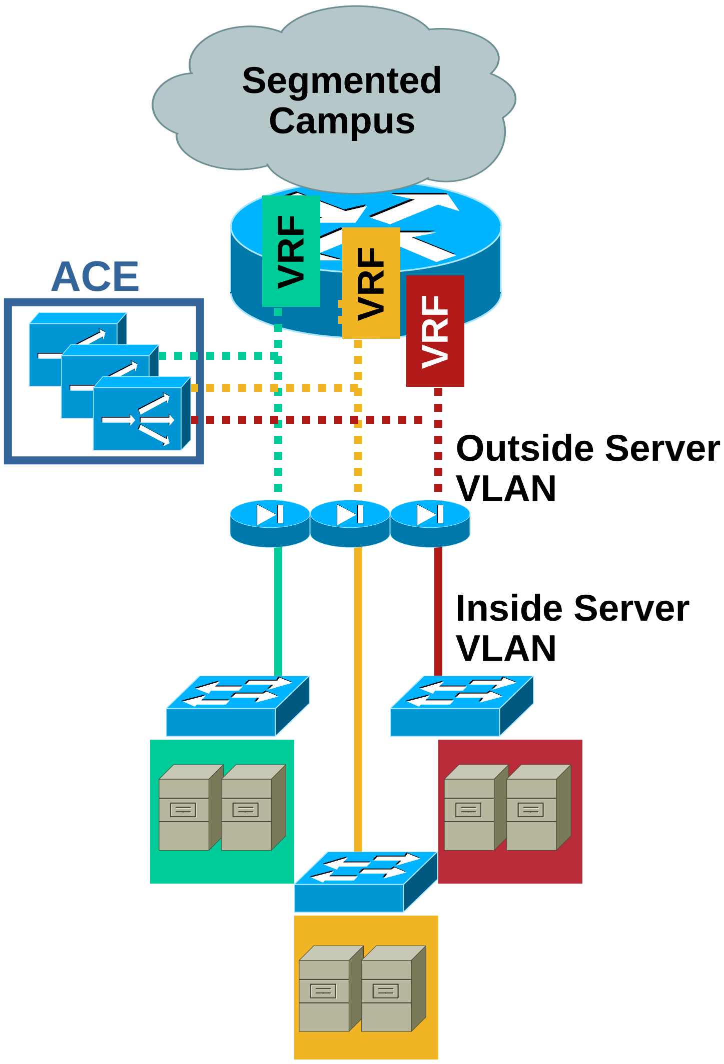

In-line Design with Dedicated Server Farms

- VRFs terminated at MSFC
- Dedicated context per inside VLAN
- Dedicated server farms in a partitioned data center
- No DoS load sharing
- Address re-use possible
- Per Context management
- ACE performance 15.5 Gbps
  - → System performance limit

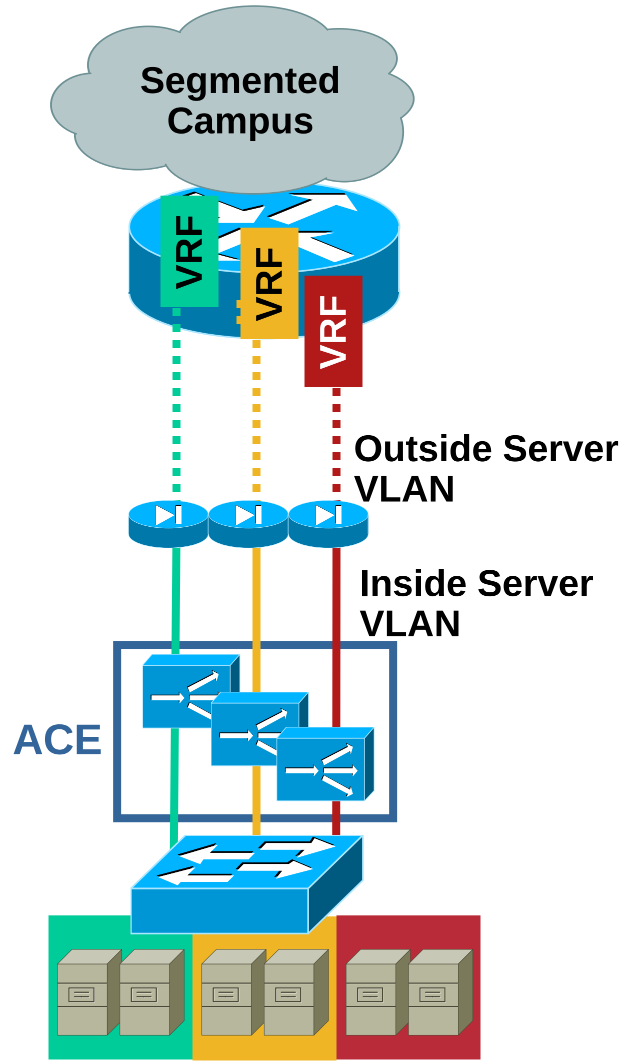

One-Arm Design with Shared Server Farm

- VRFs terminated at the MSFC
- FW transparent contexts
- Single shared server farm
- Single instance CSM
- Per context policies
- No DoS load sharing
- PBR for return traffic
- CSM performance four Gbps
  - → Only select traffic to CSM

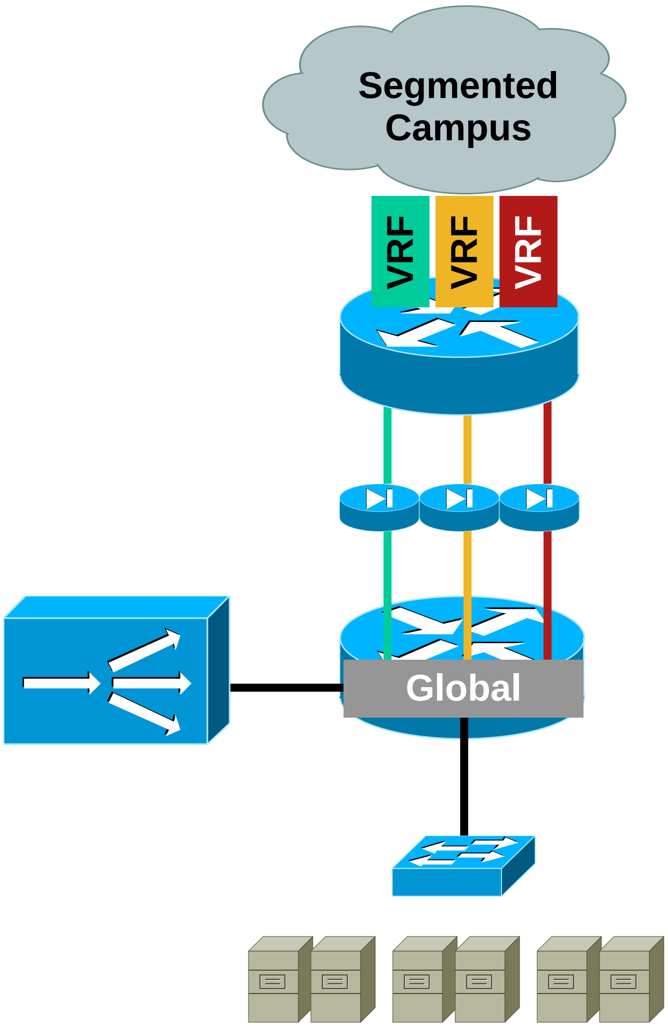

DC Front-End Segmentation

Routed Server Farm

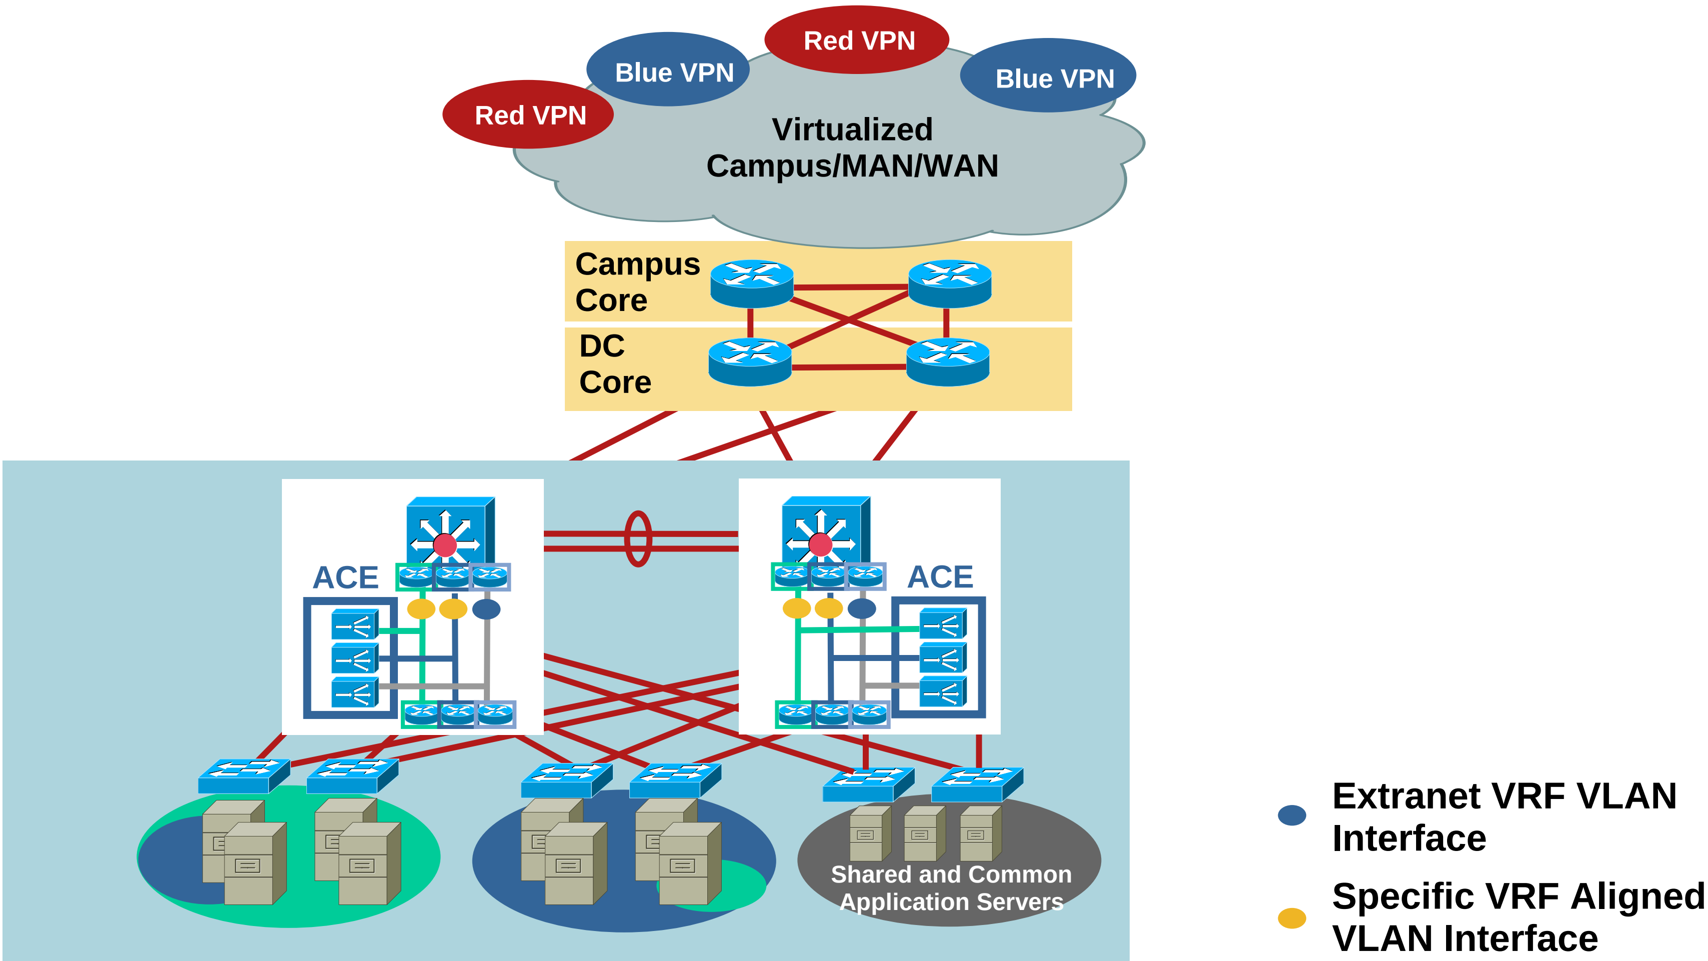

Multi-VRF Aggregation Module

Segmented Server Farms
Routed Server Farm—Detail

- Northside
  - Campus VRFs terminated at MSFC
  - Dedicated context per inside VLAN
  - Dedicated server farms in a partitioned data center
  - No DoS load sharing
  - Address re-use possible
  - Per context management
  - ACE performance 15.5 Gbps
    - → System performance limit
- Southside
  - 802.1q trunks
  - Topologies:
    - Layer 2 loop free
    - Routed access
  - Server-to-server traffic within a VRF does not impact the services

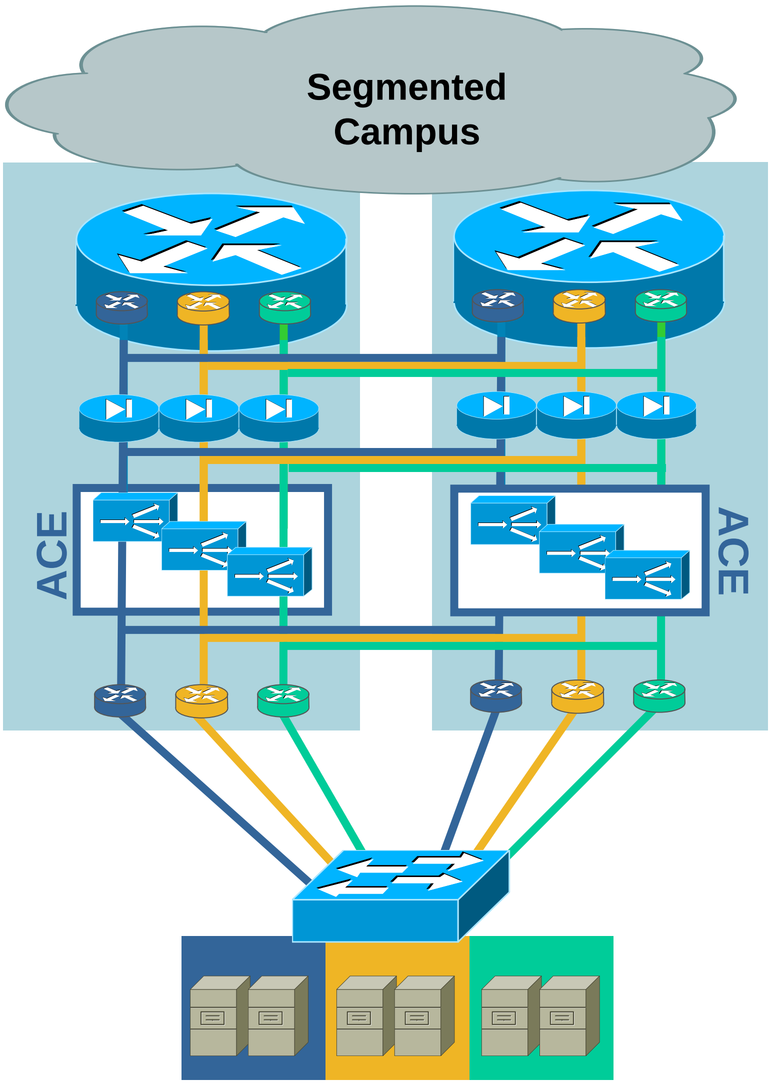

Data Center Integration
Summary

- Many groups will use the same datacenter front-end
- Serverfarms can be:
  - Shared
  - Dedicated
- Server module placement must be done carefully:
  - VRF-aware PBR, RHI limitations
  - Virtualization support → ACE and FW 3.1 recommended
- A single virtualized front-end can provide both dedicated and shared services
- Data Center Solution Reference Network Designs:
  - http://www.cisco.com/go/SRND

Deployment Considerations

…Some Things to Look out For

IP Communications Considerations	Recommendations
• Cross-VN (Virtual Net) traffic patterns Peer-to-peer requirements Firewall fix-ups	• Do you really need separate virtual networks? No → keep ipc in the global table Yes → use inter-CUG proxies
• IPC services not VRF-aware/compatible • Can they co-exist with VRFs on the same device?	• Multi-box solution today • Single-box (VRF-aware) solution road-mapped
• IPC services not virtualized • Can I create multiple service instances?	• Services can be offered as a common shared service if necessary
• End-point integration • Will phones work in a virtual network?	• Different end-points have different access control capabilities
Management Considerations	**Recommendations**
▪ Management	• Design according to Best Practices • IP Solutions Center for MPLS VPNs • CiscoWorks being evaluated

Deployment Considerations
…Some Things to Look out For

Multicast Considerations	Recommendations
• Understand your traffic patterns Point-to-multipoint Multipoint-to-multipoint Inter-CUG or intra-CUG	• Maintain point-to-multipoint patterns when possible • Use Extranet MVPN for efficient inter-CUG replication
• Some VRF-lite platforms may not support multicast	• Choose appropriate platform • Check platform roadmap

Cisco IOS® Feature Considerations	Recommendations
• Not all features are VRF-aware on all platforms • e.g., PBR, dDNS, RHI, multicast	• Choose platforms accordingly e.g., ACE has VRF-aware RHI • Design around the limitations e.g., deploy ACE in-line to avoid PBR

Enterprise Virtual Networks Summary

End-to-End Virtualized Enterprise

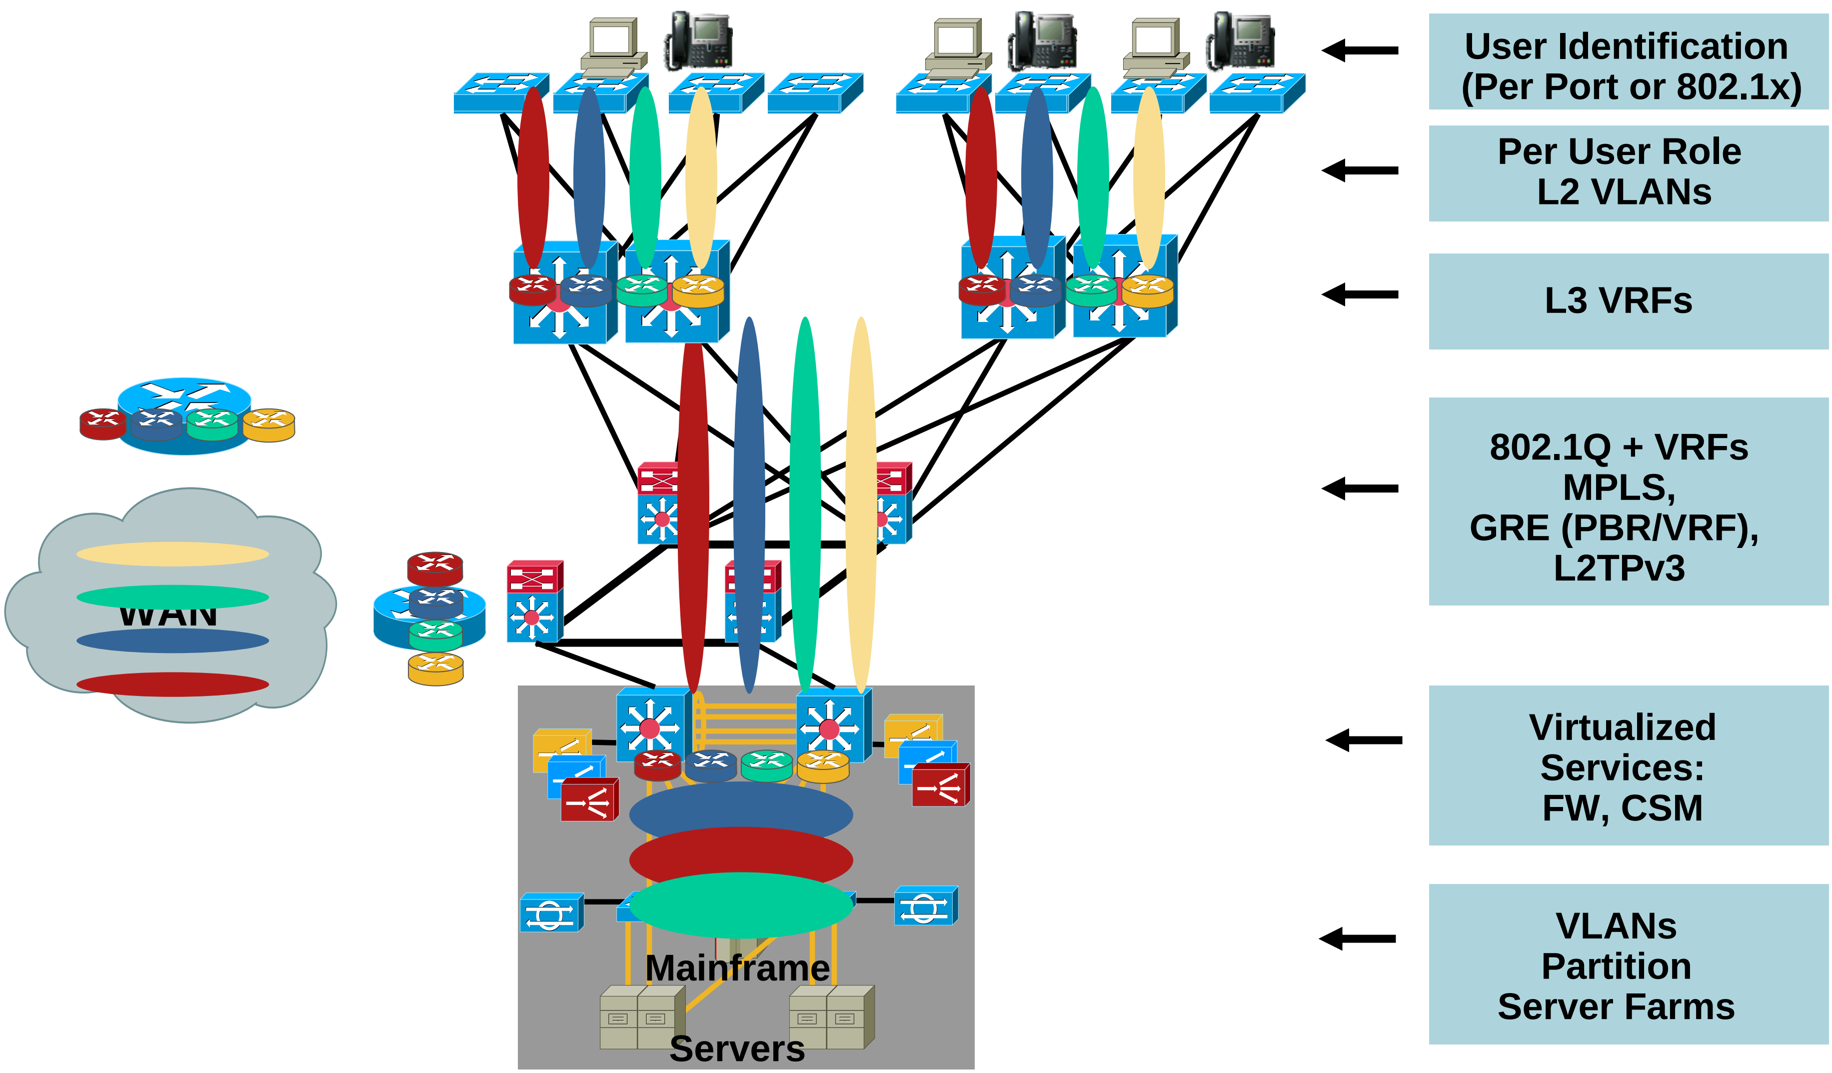

Meet the Experts
Campus and Wireless Evolution

- Mark Montanez
  Corporate Dev Consulting Engineer

- Tim Szigeti
  Technical Leader

- Sujit Ghosh
  Technical Mktg Eng

- Victor Moreno
  Technical Leader

- Mike Herbert
  Technical Leader

Recommended Reading

BRKCAM -3002

- Definitive MPLS Network Designs
- MPLS and VPN Architectures
- Network Virtualization

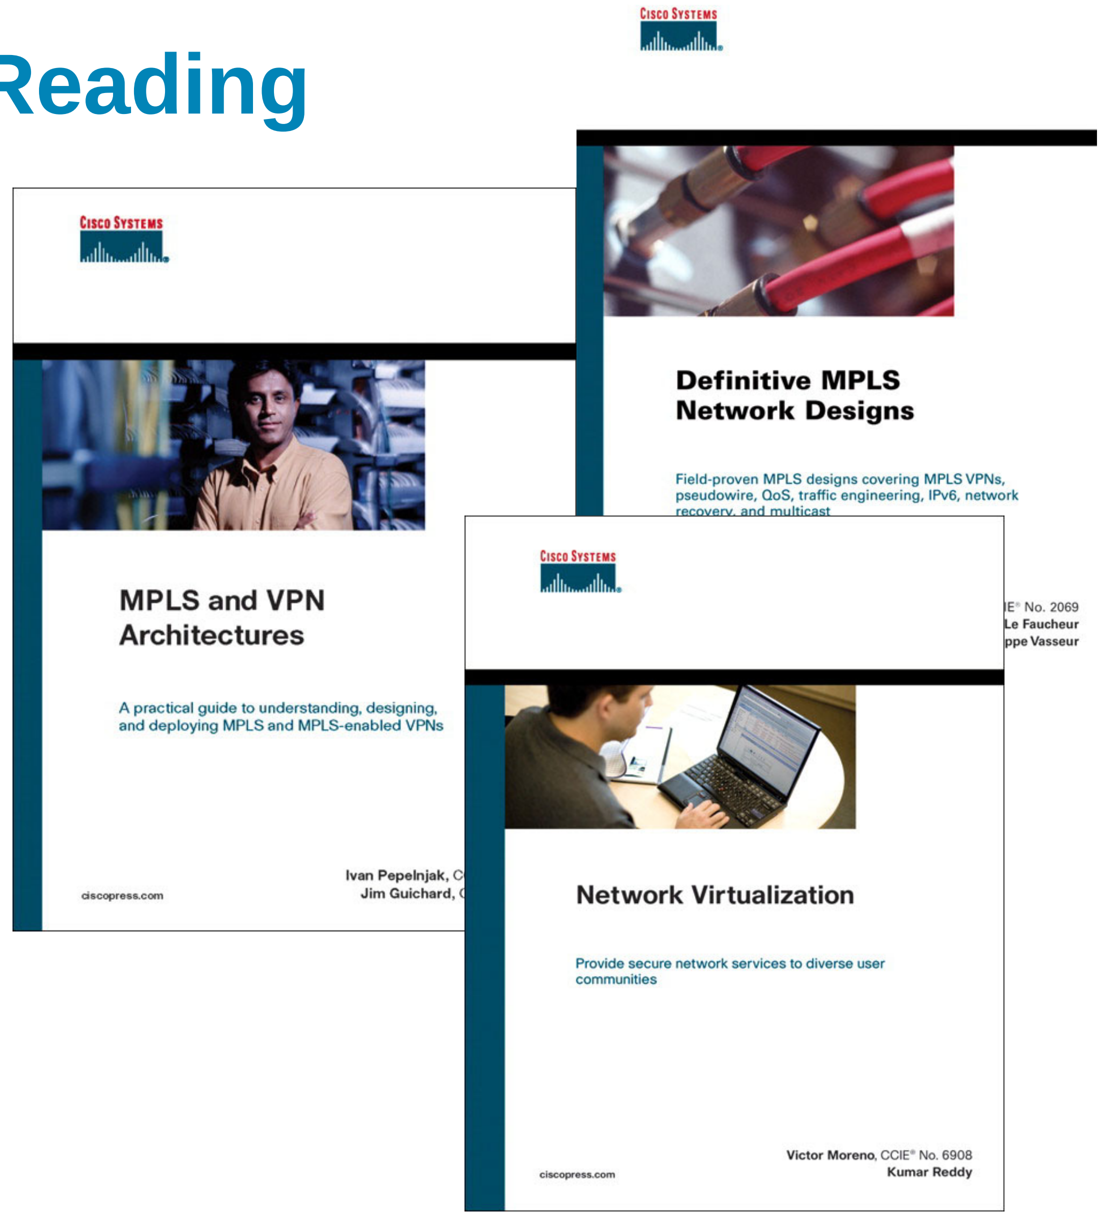

Available in the Cisco Company Store

Q and A

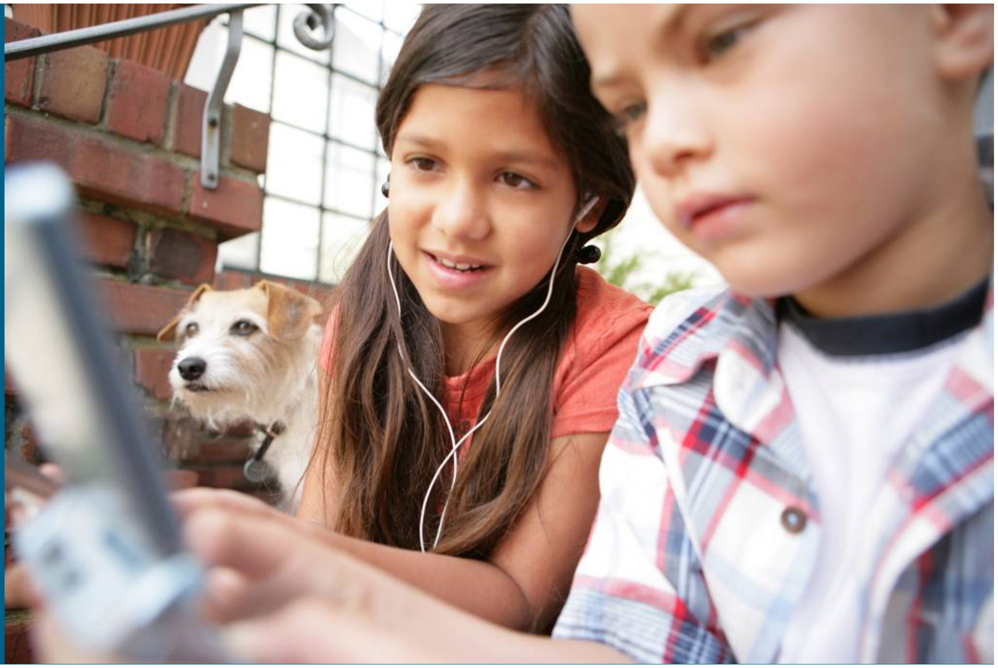

Related Sessions

- **BRKCAM-2003** Guest Access Services for Wired and Wireless Architectures
- **BRKCAM-2007** Understanding Identity-Based Networking Services, Authentication, and Policy Enforcement
- **BRKCAM-2001** Multilayer Campus Architectures and Design Principles
- **BRKCAM-3004** Deploying a Fully Routed Enterprise Campus Network (Routing in the Access Layer)
- **BRKDCT-2001** Data Center Networking - Architecture and Design Guidelines
- **LABNMS-2002** IP Solution Center—MPLS Management Lab
- **LABIPM-2002** Enabling MPLS in Enterprise Networks
- **BRKIPM3014** Advanced MPLS Deployment in Enterprise Networks

Acronyms

- VRF—Virtual Routing/ Forwarding
- CNR—Cisco Network Registrar
- Rd—Route Distinguisher
- RT—Route Target
- BGP—Border Gateway Protocol
- iBGP—Internal BGP
- MP-BGP—Multipoint BGP
- GRE—Generic Routing Encapsulation
- mGRE—Multipoint GRE
- OSPF—Open Shortest Path First
- NAC—Network Access Control
- ACL—Access Control List
- IGP—Internal Gateway Protocol
- CUG—Closed User Group
- VPN—Virtual Private Network
- FW—Firewall
- DMZ—De-Militarized Zone
- PE—Provider Edge Router
- P—Provider Router
- CE—Customer Edge Router
- IT—Information Technology
- DHCP—Dynamic Host Configuration Protocol
- DNS—Dynamic Name Services