
TOMORROW starts here.
cisco
Cisco live!

Deploying a Virtualised Campus Network Infrastructure

BRKCRS-2033

Geoff Yates

Systems Engineer

Cisco live!

Clear Message for Virtualisation

Qld to spend $7.4 billion fixing nearly all IT systems

By Allie Coyne on Jun 11, 2013 9:53 AM
Filed under Software

Like 39 Tweet 29 g+1 13 Share 36 3 Comments

IT audit report finds "systemic business risk".

The Queensland Government will need to replace ninety percent of its IT systems within five years, with the overall project to cost $7.4 billion, more than $2 billion over the initial forecast.

The state's new IT minister Ian Walker tabled the long-awaited IT Audit and the government's response to Parliament on Friday last week. The audit had been due for release last year but was held back multiple times.

The five-month audit covered 900 projects and 10,000 systems. It cost $5.2 million and required 32 public servants.

The report also made the following recommendations, which the government has agreed to:

- Cancel unused mobile and fixed telephone services, optimise data plans, consolidate telco accounts and increase printer efficiencies
- Decommission unused systems and exit its Travel Management System
- Initiate and maintain a program of rigorous application of business continuity planning for all business critical systems
- Never modify commercially-provided commodity applications to meet unique business requirements
- Conduct basic technical upgrades for high-risk payroll, finance, systems
- Further analyse the Health finance system replacement
- Establish an externally-managed desktop arrangement, and
- Study the options for a single-government data network for all agencies.

Clear Message for Virtualisation

Study the options for a single-government data network for all agencies.

Agenda

- Virtualisation solves these Challenges
- Virtualisation Architectures
- Case Study
- Industry Trends
- Putting it all Together

Legend

Informational Icons:

"For Your Reference" – these slides are used to help you configure a particular feature or technology solution

"Emerging Technology" – self explanatory

BRKCRS-2033

"Where to Learn More" – for additional details, please see the indicated presentation

Network Connections:

Routed Connections in "Red" – L3

Switched Connections in "Black" – L2

Agenda

- Virtualisation solves these Challenges
- Virtualisation Architectures
- Case Study
- Industry Trends
- Putting it all Together

Why Virtualise?

- Unique security policies per logical domain
- Traffic isolation per application, group, service etc…
- Logically separates traffic using one physical infrastructure

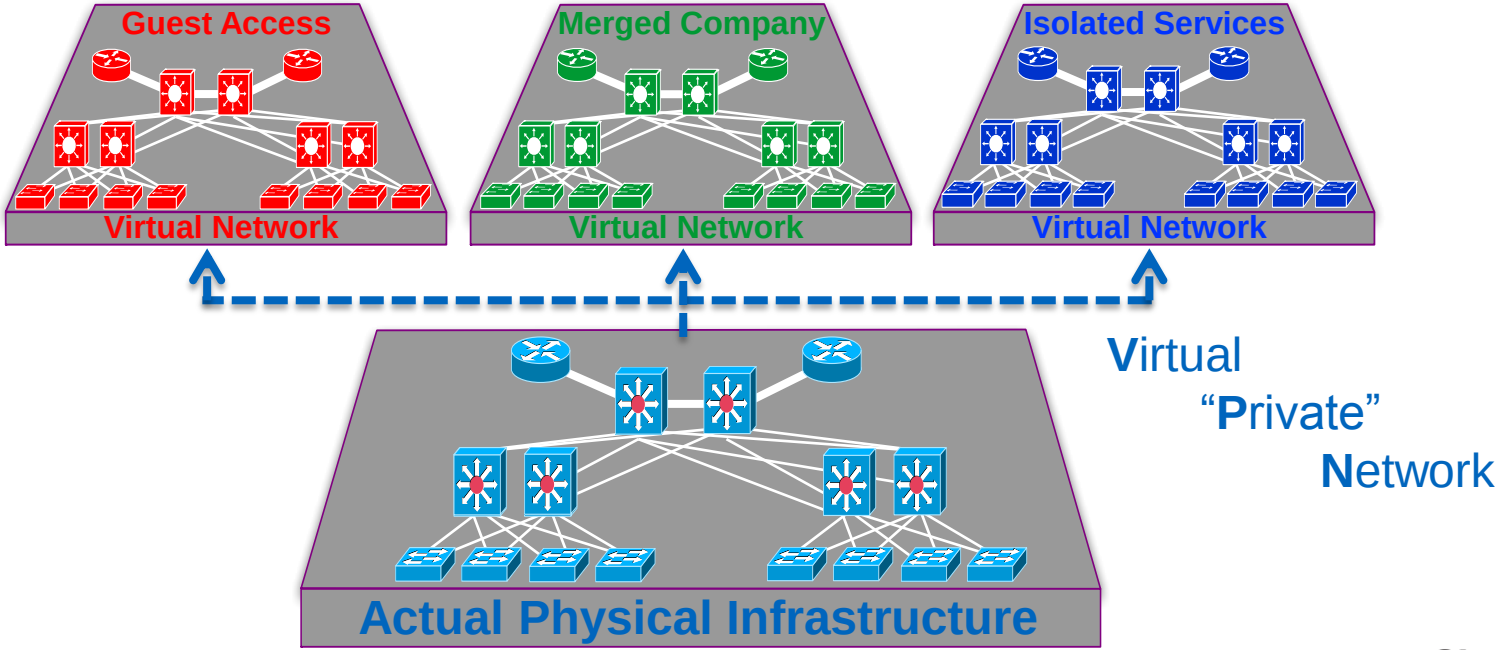

Virtualisation Benefits

- Groups and services are logically separated
  - Telephony systems, building control, surveillance
  - Security Policies are unique to each virtual group/service
- Regulatory compliance
  - HIPAA
  - PCI
  - SOX
  - etc…

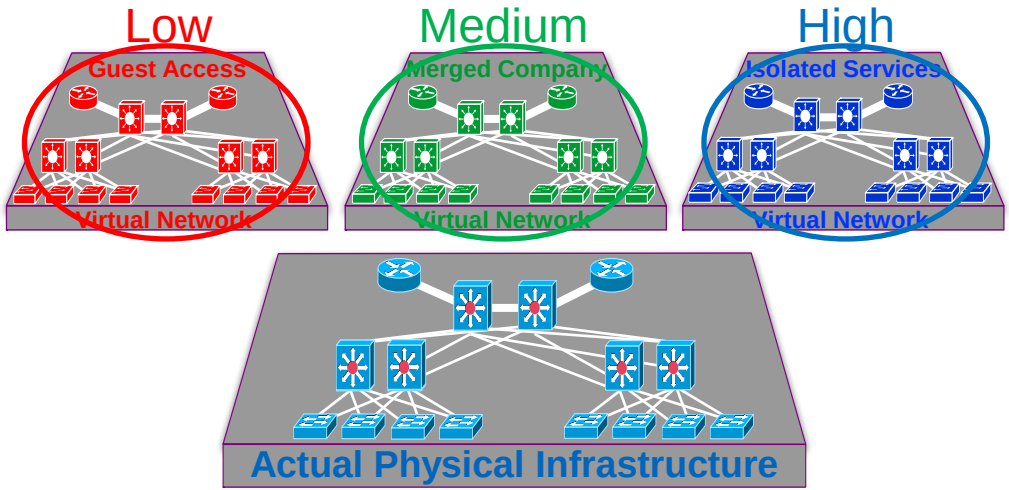

Agenda

- Virtualisation solves these Challenges
- Virtualisation Architectures
- Case Study
- Industry Trends
- Putting it all Together

Network Virtualisation

Components

Service	Access Control	Path Isolation	Services Edge
	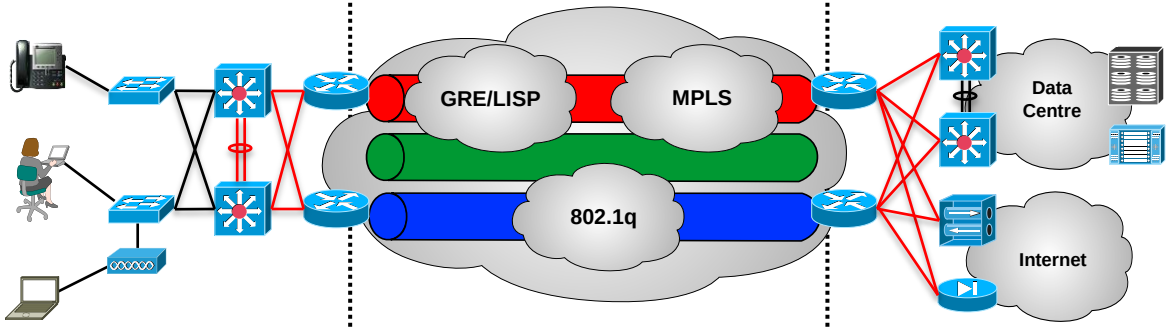		
Functions	▪ Authenticate device attempting to gain network access ▪ Authorise device into a partition (VLAN) ▪ Deny access to unauthenticated devices	▪ Maintain traffic partitioned over Layer 3 infrastructure ▪ Transport traffic over isolated Layer 3 partitions ▪ Map Layer 3 isolated path to VLANs / VRFs in access and services edge	▪ Provide access to services Shared Dedicated ▪ Apply policy per partition ▪ Isolate application environments if necessary

Access Control

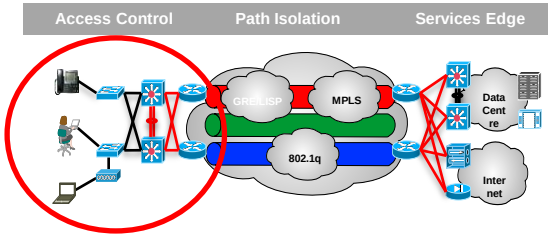

Authentication - Who are you?

- Client-based
  - 802.1X – assigned to VLAN
  - Identity Services Engine (ISE)
- Clientless
  - Web authentication
  - MAC-addressed based
  - Identity Services Engine (ISE)
- Static control
  - Port security (static VLAN, ACL, MAC, etc…)

Authorisation - Where can you go?

- VLAN / VRF
- ACL, Security Group Tags (SGT), Security Group ACLs (SGACL)
- Policy enforcement via Identity Services Engine (ISE)

Identity Services Engine

Primary Features and Benefits

Comprehensive Secure Access

Operational Efficiency

Device Profiling and Posture

Contextual Identity (Intelligent Identity)

Policy Management

Network Enforcement and Control Point

Device Virtualisation

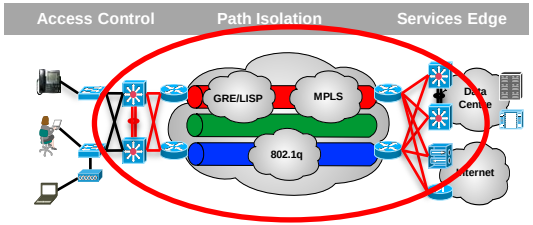

Virtually multiple devices

- Control plane virtualisation
- Data plane virtualisation
- Services virtualisation

Device virtualisation

- One physical device
  - Switch
  - Router
  - Firewall
  - Etc…

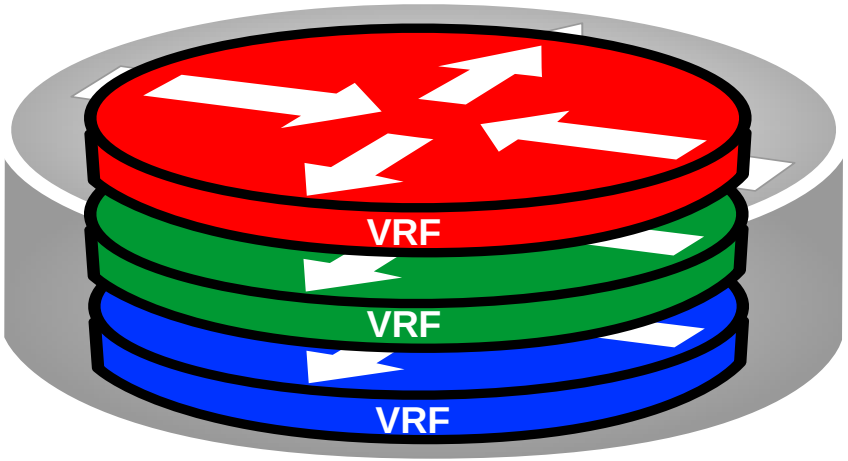

VRF: Virtual Routing and Forwarding

Device Virtualisation

Connecting to a VRF – Client Side

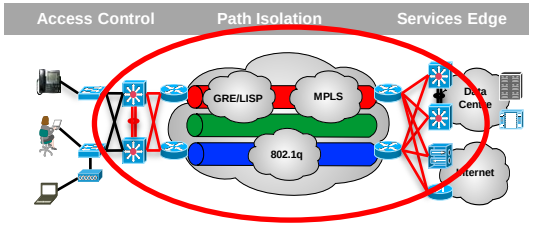

- Physical interface
  - Ethernet

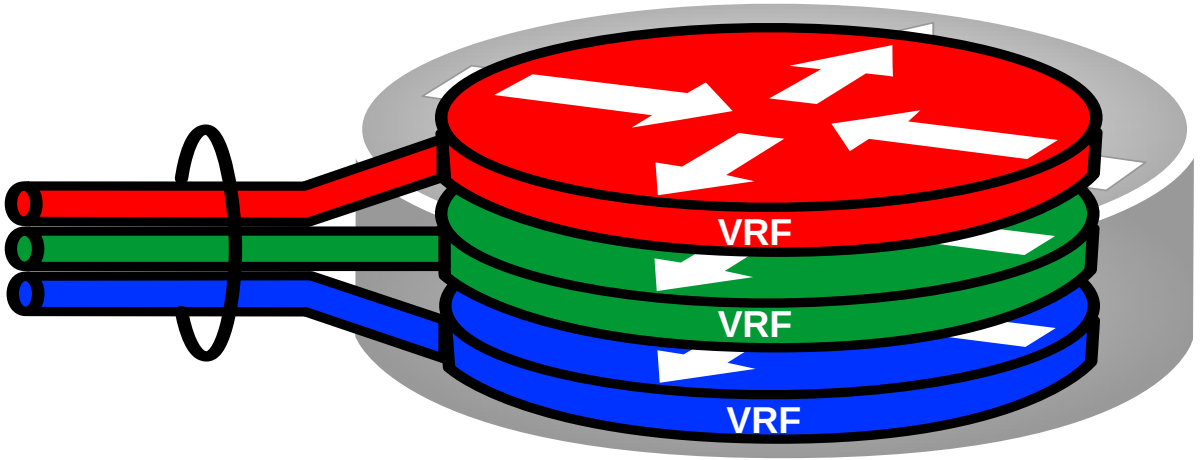

- Logical interface
  - VLAN - 802.1q trunk

Path Isolation

Data Path Virtualisation – Network Side

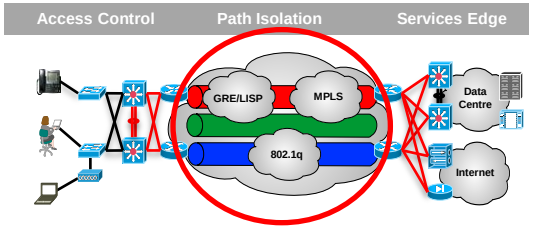

- Hop-by-Hop
  - VRF-Lite End-to-End
  - EVN (Easy Virtual Network)
  - 802.1q for Separation
- Multi-Hop
  - VRF-Lite + GRE
  - VRF-Lite + LISP
  - GRE/LISP for Separation
- Multi-Hop
  - MPLS-VPN
  - MPLS Labels for Separation

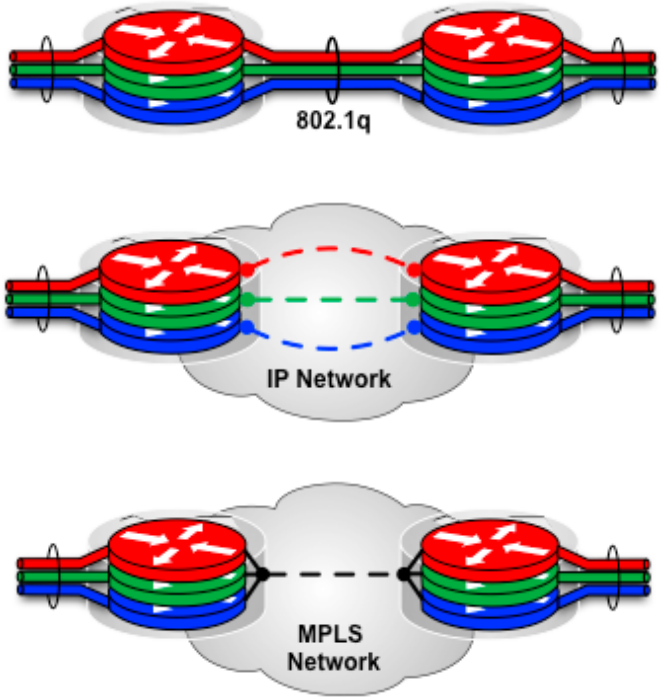

Services Edge

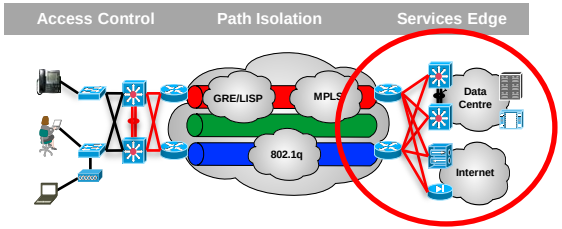

Sharing Services Between VPNs

- Unnecessary to duplicate services per group
  - E-mail, DNS, LDAP, Storage, etc…
- Economical
- Efficient and manageable
- Policies centrally deployed

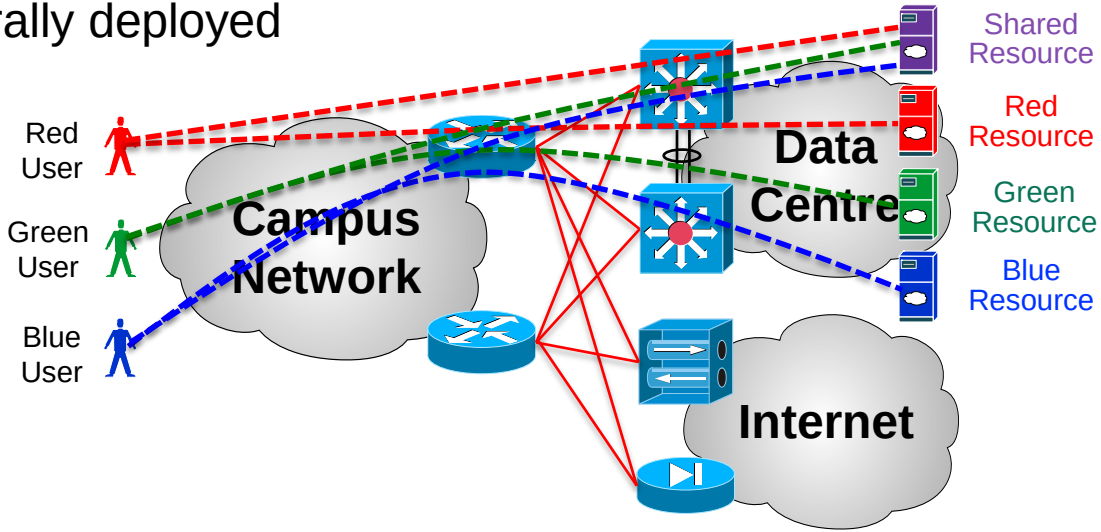

Services Edge

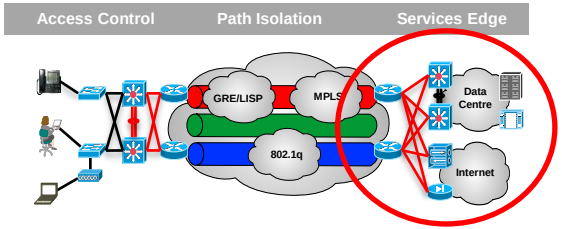

Sharing Resources

- Firewall (multi-context) - FWSM / ASA / ASA Module
- Server Load Balancing (multi-context) - ACE
- IPSec / SSL VPN - Router (F-VRF) / ASA VLAN mapping

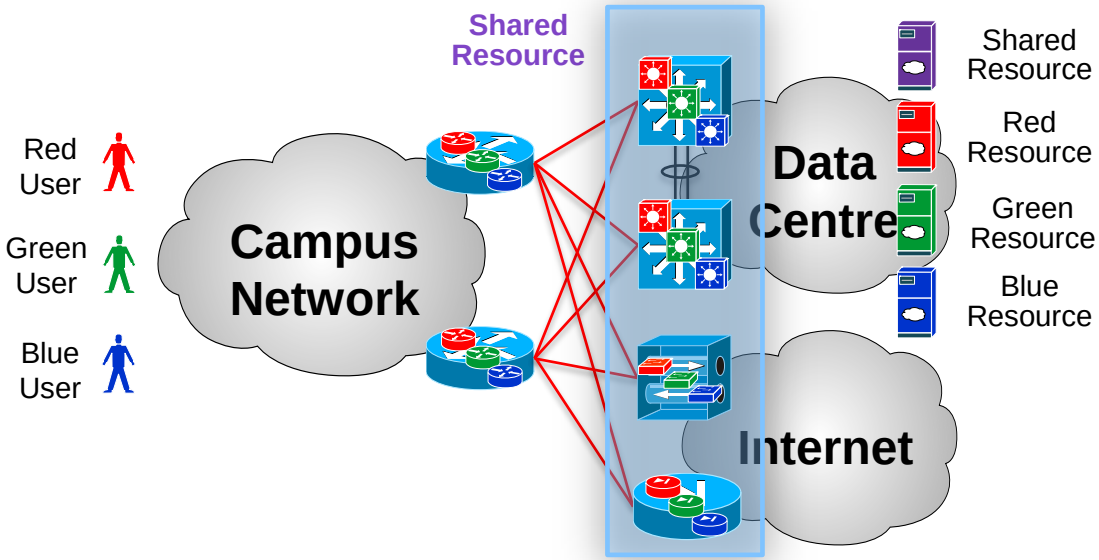

VRF-Lite and GRE tunnels

VRF-Lite and GRE Tunnels

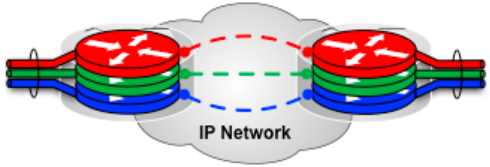

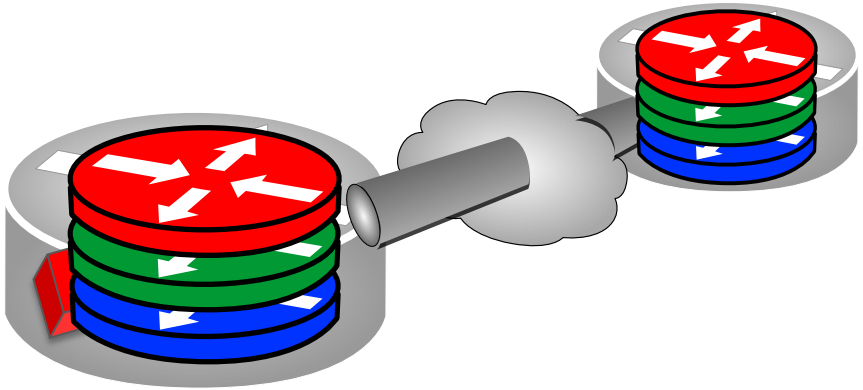

GRE encapsulation represent 24 extra bytes or 28 if a key is present

VRF-Lite and GRE Tunnels

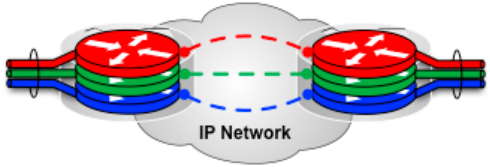

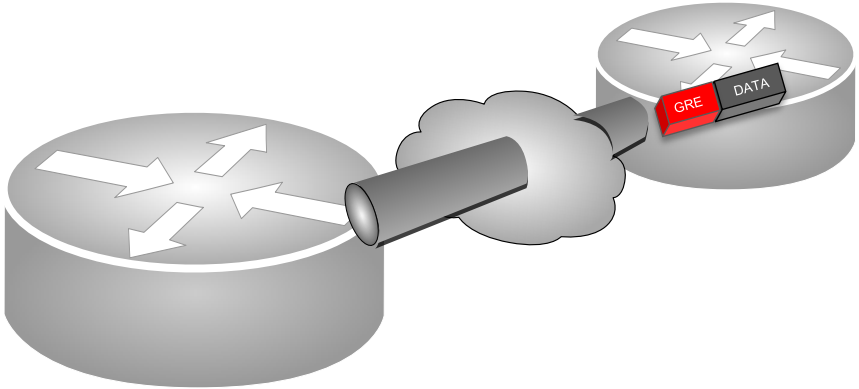

GRE encapsulation represents 24 extra bytes or 28 if a key is present

VRF-Lite and GRE Tunnels

Deployment Summary

- Infrastructure
  - Recommended for hub-and-spoke requirements
  - Limited scale for single or few VPN applications (guest access, NAC remediation)
  - GRE supported in HW on Catalyst 6500 and Nexus 7K
- Application and Services
  - Multiple VRF-aware services available
- Learning Curve
  - Familiar routing protocols can be used
  - IP Based solution

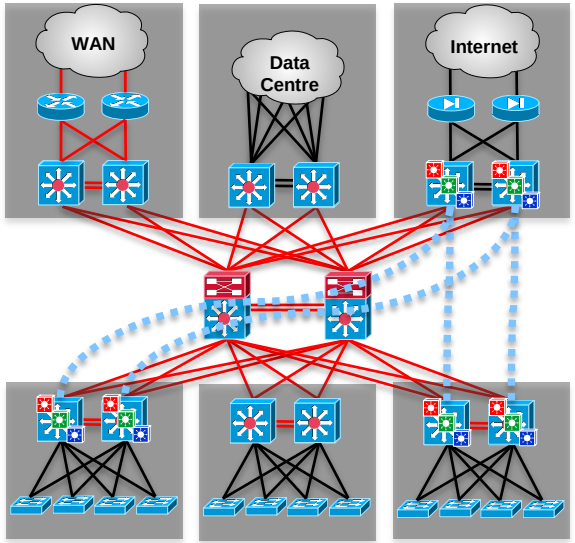

VRF-Lite and Easy Virtual Network (EVN)

VRF-Lite/EVN End-to-End

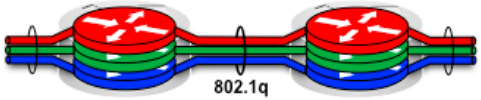

- Packets processed per VRF
- Unique Control Plane and Data Plane

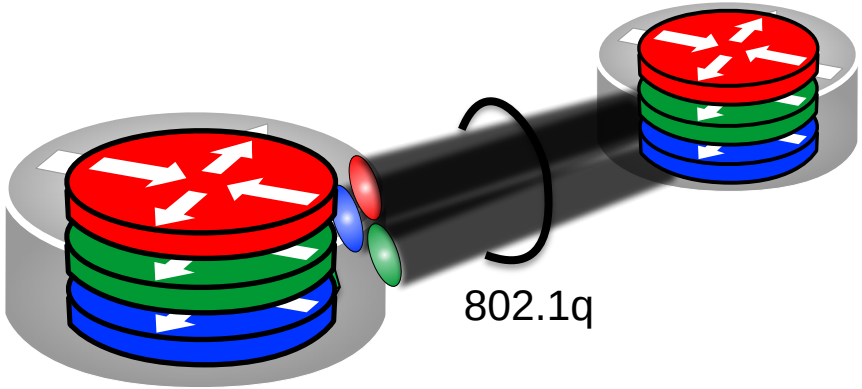

VRF-Lite/EVN End-to-End

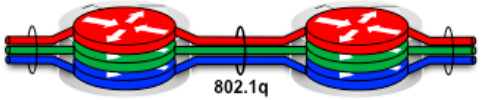

- Packets processed per VRF
- Unique Control Plane and Data Plane

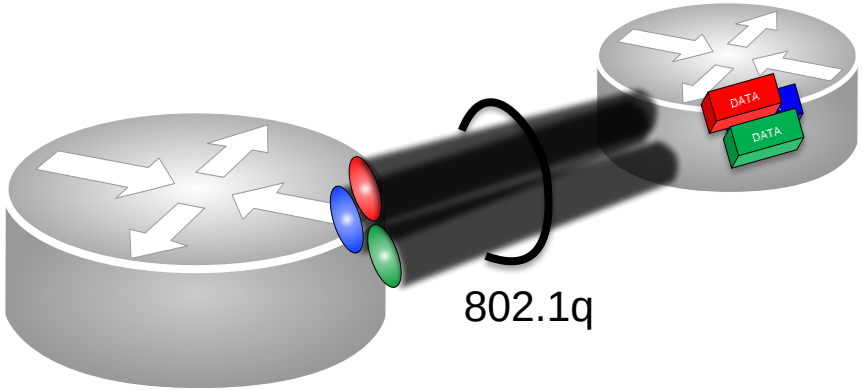

VRF-Lite/EVN

Client-Side Configuration

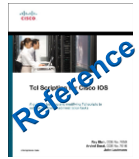


```
vrf definition GRN
 !
 address-family ipv4
!
 address-family ipv6
!
vrf definition RED
 !
 address-family ipv4
!
 address-family ipv6


interface Vlan17
 vrf forwarding GRN
 ip address 172.17.8.8 255.255.255.0
 ipv6 address 2001:17:8::8/64
!
interface Vlan16
 vrf forwarding RED
 ip address 172.16.8.8 255.255.255.0
 ipv6 address 2001:16:8::8/64
```

Defining the
VRFs
IPv4 and IPv6

Client-side
Interface

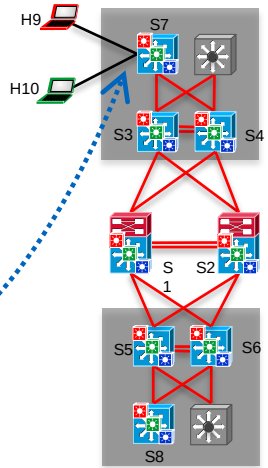

Currently no IPv6 support for EVN

VRF-Lite

Network-Side Configuration

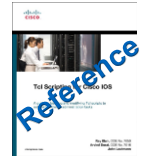


```
interface Ethernet0/0.16
 vrf forwarding RED
 encapsulation dot1Q 16
 ip address 172.16.85.8 255.255.255.0
 ipv6 address 2001:16:85::8/64
!
interface Ethernet0/0.17
 vrf forwarding GRN
 encapsulation dot1Q 17
 ip address 172.17.85.8 255.255.255.0
 ipv6 address 2001:17:85::8/64
!
!
interface Ethernet0/1.16
 vrf forwarding RED
 encapsulation dot1Q 16
 ip address 172.16.86.8 255.255.255.0
 ipv6 address 2001:16:86::8/64
!
interface Ethernet0/1.17
 vrf forwarding GRN
 encapsulation dot1Q 17
 ip address 172.17.86.8 255.255.255.0
 ipv6 address 2001:17:86::8/64
```

Assign IPv4 and v6 addresses

Network side interface

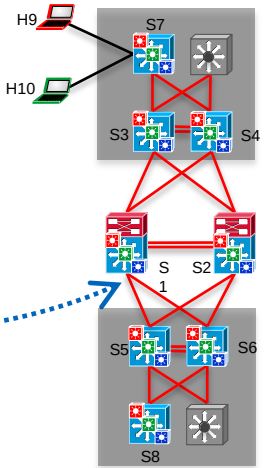

EVN

Network-Side Configuration

```
vrf definition GRN
 vnet tag 102
 !
 address-family ipv4
!
vrf definition RED
 vnet tag 101
 !
 address-family ipv4



interface Ethernet0/0
 vnet trunk
 ip address 192.168.74.7 255.255.255.0
!
interface Ethernet0/1
 vnet trunk
 ip address 192.168.73.7 255.255.255.0
!
```

VRF Definition and VNET tag association

Network-side interfaces

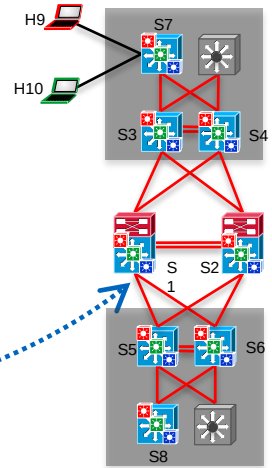

EVN

Derived Configuration

```
#show derived-config

interface Ethernet0/0
 vnet trunk
 ip address 192.168.74.7 255.255.255.0
!
interface Ethernet0/0.101
 description Subinterface for VNET RED
 vrf forwarding RED
 encapsulation dot1Q 101
 ip address 192.168.74.7 255.255.255.0
!
interface Ethernet0/0.102
 description Subinterface for VNET GRN
 vrf forwarding GRN
 encapsulation dot1Q 102
 ip address 192.168.74.7 255.255.255.0
```

Physical interface

Network Side

Sub-interfaces created automatically

Descriptions added

Reuse of IP address – logically separated on trunk

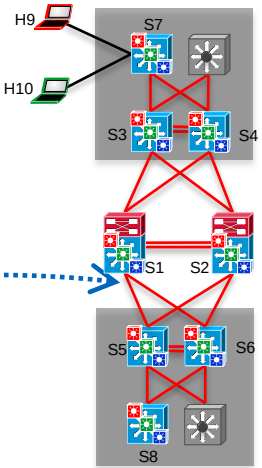

EVN

Traffic Example

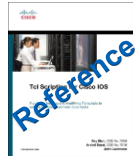


```
H9#traceroute 172.16.8.11
Type escape sequence to abort.
Tracing the route to 172.16.8.8
VRF info: (vrf in name/id, vrf out name/id)
  1 172.16.7.7 (RED,RED/101) 0 msec 1 msec 1 msec
  2 192.168.74.4 (RED/101,RED/101) 1 msec 0 msec 1 msec
  3 192.168.42.2 (RED/101,RED/101) 1 msec 0 msec 0 msec
  4 192.168.52.5 (RED/101,RED/101) 1 msec 1 msec 0 msec
  5 192.168.85.8 (RED/101,RED) 2 msec 5 msec 4 msec
  6 172.16.8.11 5 msec *  5 msec

H10#traceroute 172.17.8.12
Type escape sequence to abort.
Tracing the route to 172.17.8.12
VRF info: (vrf in name/id, vrf out name/id)
  1 172.17.7.7 (GRN,GRN/102) 0 msec 0 msec 1 msec
  2 192.168.73.3 (GRN/102,GRN/102) 1 msec 0 msec 1 msec
  3 192.168.32.2 (GRN/102,GRN/102) 5 msec 5 msec 5 msec
  4 192.168.52.5 (GRN/102,GRN/102) 6 msec 5 msec 5 msec
  5 192.168.85.8 (GRN/102,GRN) 5 msec 5 msec 4 msec
  6 172.17.8.12 5 msec *  5 msec
```

Traceroute indicates every L3 hop and provides VRF name and VLAN

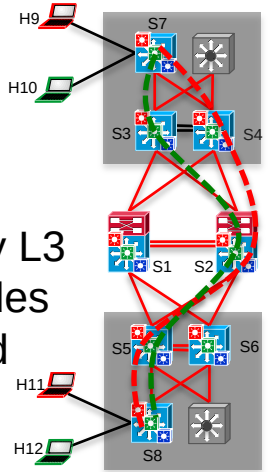

VRF-Lite End-to-End

Summary

- Deployment
  - End-to-End IP based Solution
  - Easy migration from existing campus architecture
  - Any to any connectivity within VPNs
  - 8 or less VRFs recommended
  - Supported on Catalyst 6500, 4500E/X, 3000 families, and Nexus 7000
- Application and Services
  - Multiple VRF-aware Services available
- Learning Curve
  - Familiar routing protocols
  - IP Alternative to MPLS

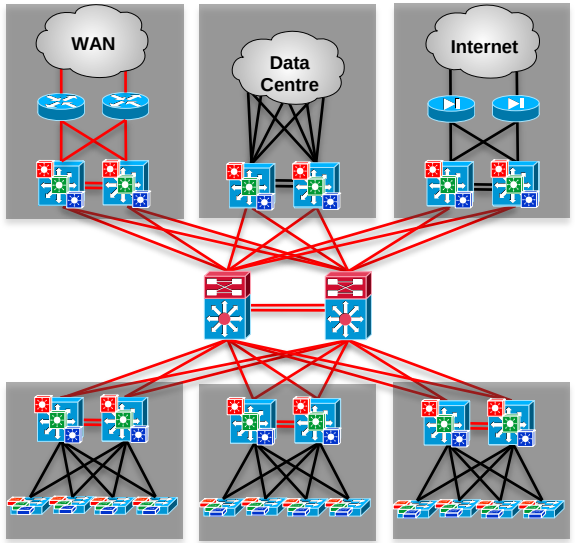

EVN
Summary

- Deployment
  - End-to-End IP based Solution
  - Easy integration with VRF-Lite
  - Any to any connectivity within VPNs
  - Route replication
  - Supported on ASR1K, Sup2T, Cat4K, ISR-G2
  - 32 or less VRFs supported
- Applications and Services
  - Multiple VRF-aware services available
- Learning Curve
  - Familiar routing protocols can be used
  - IP Alternative to MPLS

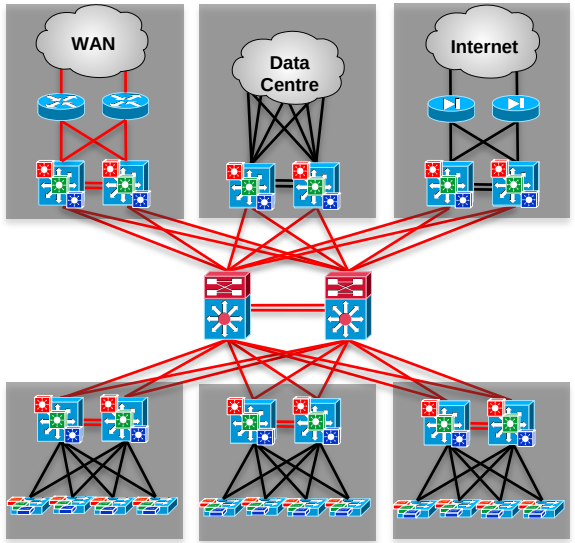




MPLS-VPN

Test Diagram

MPLS-VPN

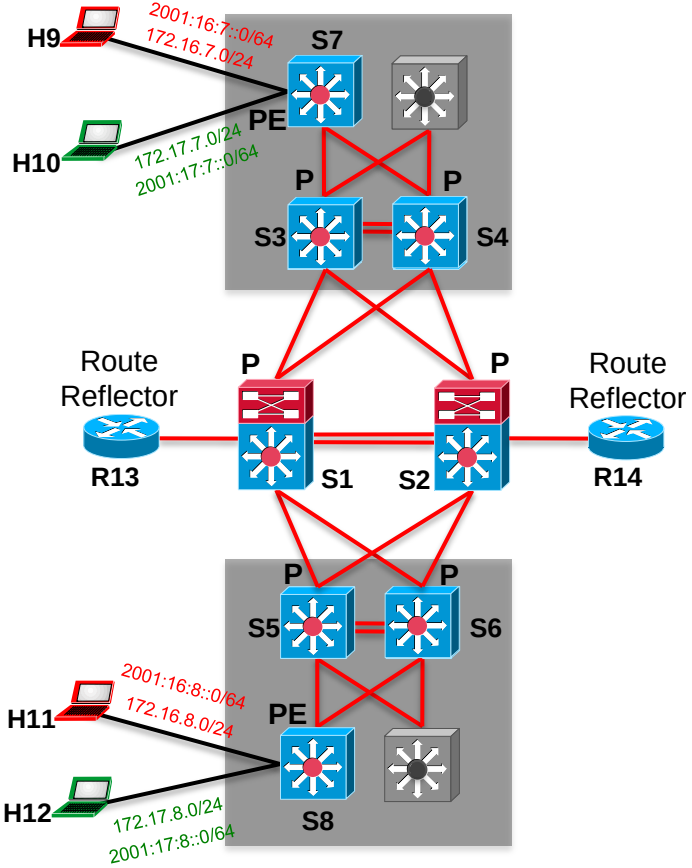

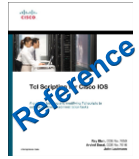

MPLS-VPN

Overview

- P (Provider) router = Label Switching Router (LSR) = core router
  - Runs an IGP and LDP
- PE (Provider Edge) router = edge router (LSR)
  - Runs an IGP, LDP and MP-BGP
- CE (Customer Edge) router
  - Connects customer network to MPLS network

MPLS-VPN

BGP Scalability – iBGP Neighbour Relationships

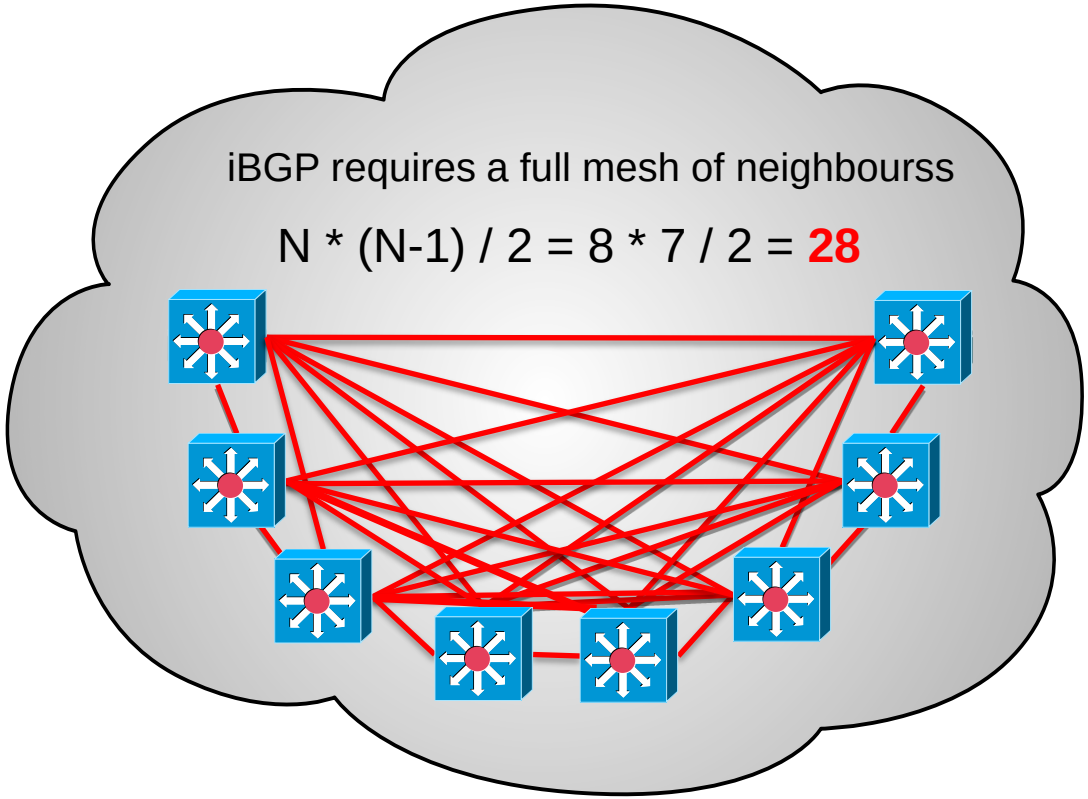

MPLS-VPN

BGP Scalability – Route Reflectors

- Use “purpose-built” RRs
- Don’t place RRs in data path
- Geographically diverse

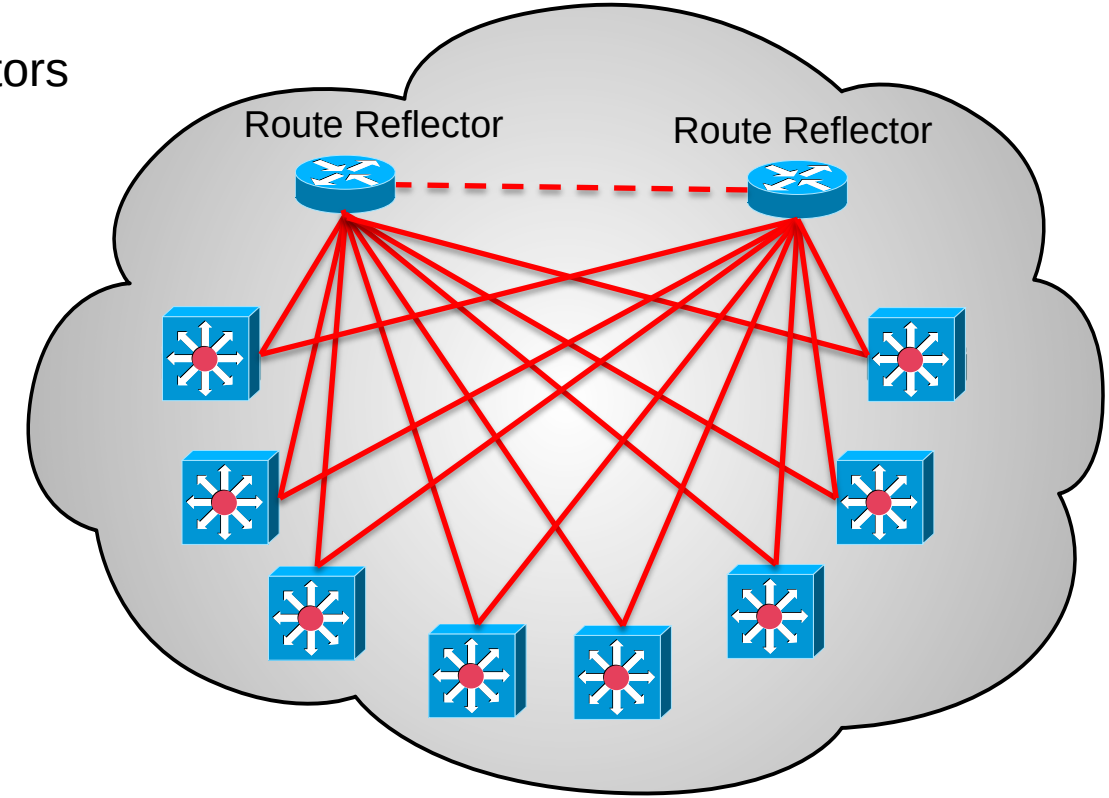

MPLS-VPN

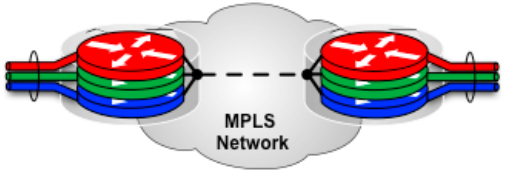

Label Stack

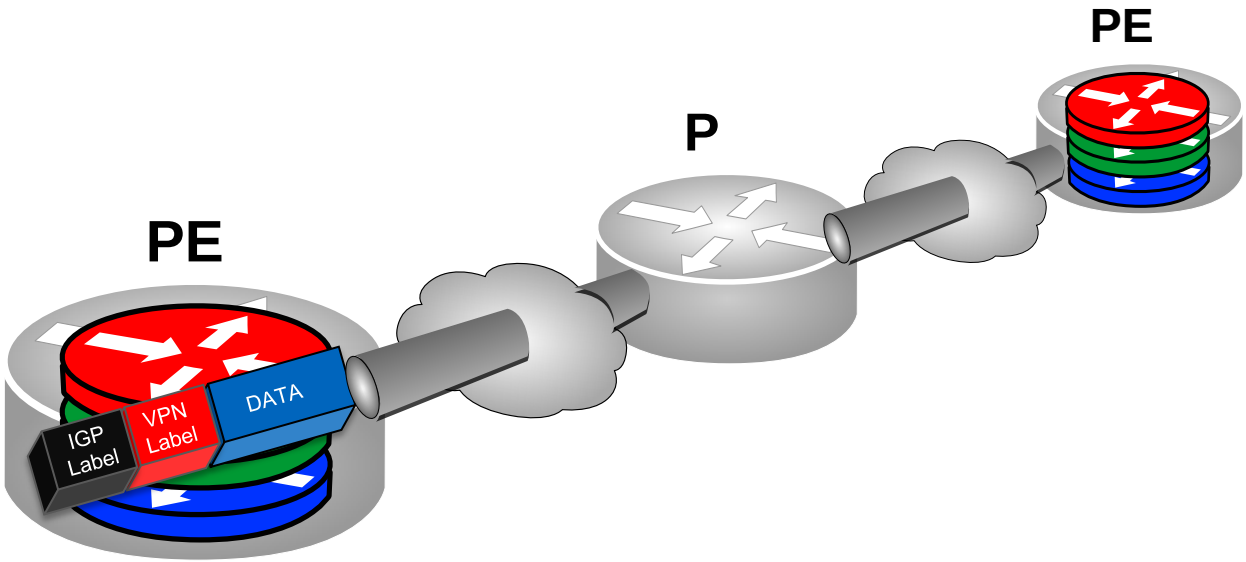

MPLS VPN packet format

MPLS-VPN

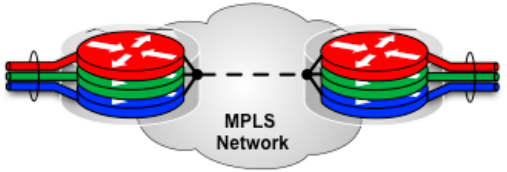

Label Stack

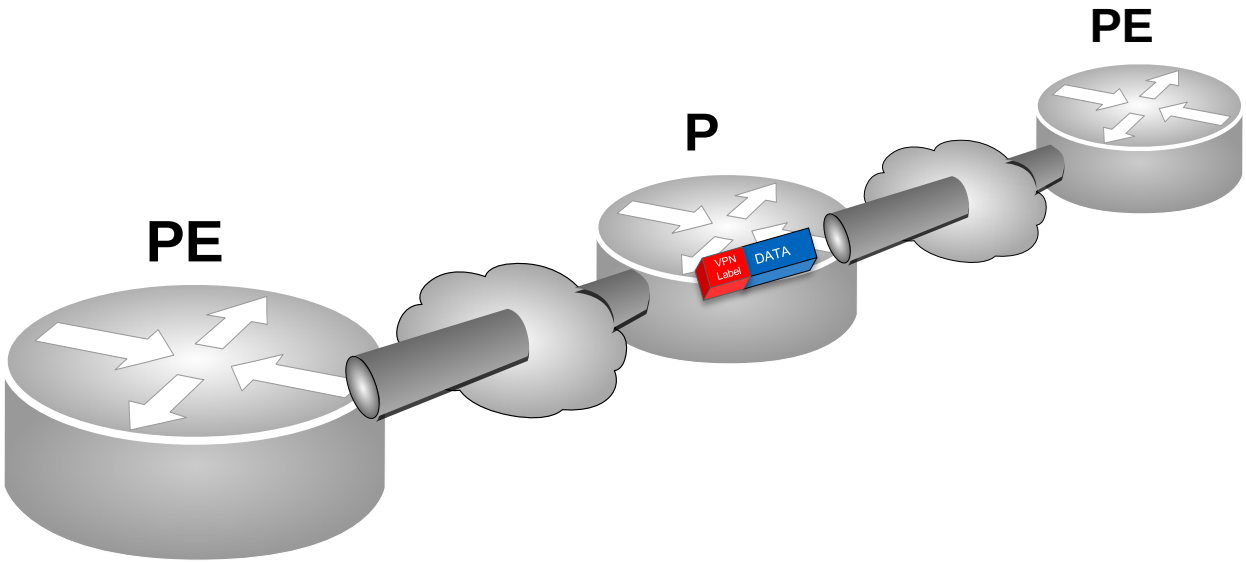

4 Byte IGP Label	4 Byte VPN Label	Original Packet

MPLS VPN packet format

MPLS-VPN

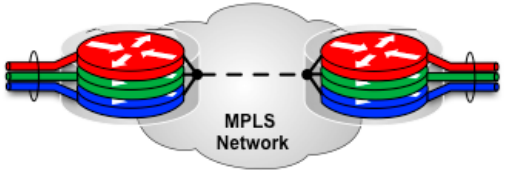

Label Stack

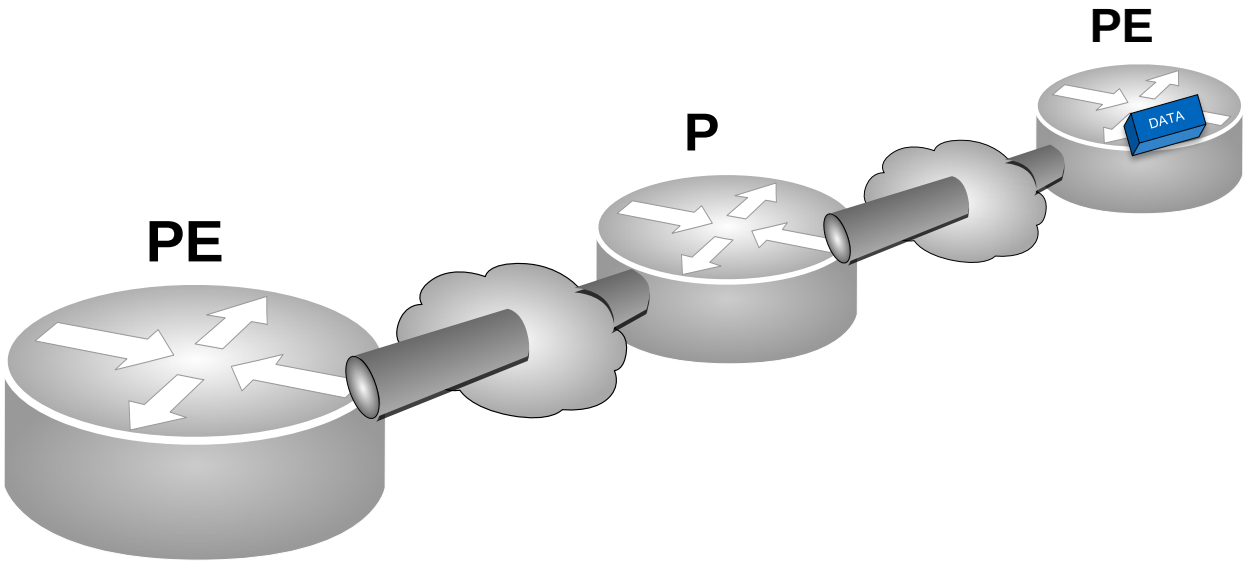

4 Byte IGP Label	4 Byte VPN Label	Original Packet

MPLS VPN packet format

MPLS-VPN – Label Exchange

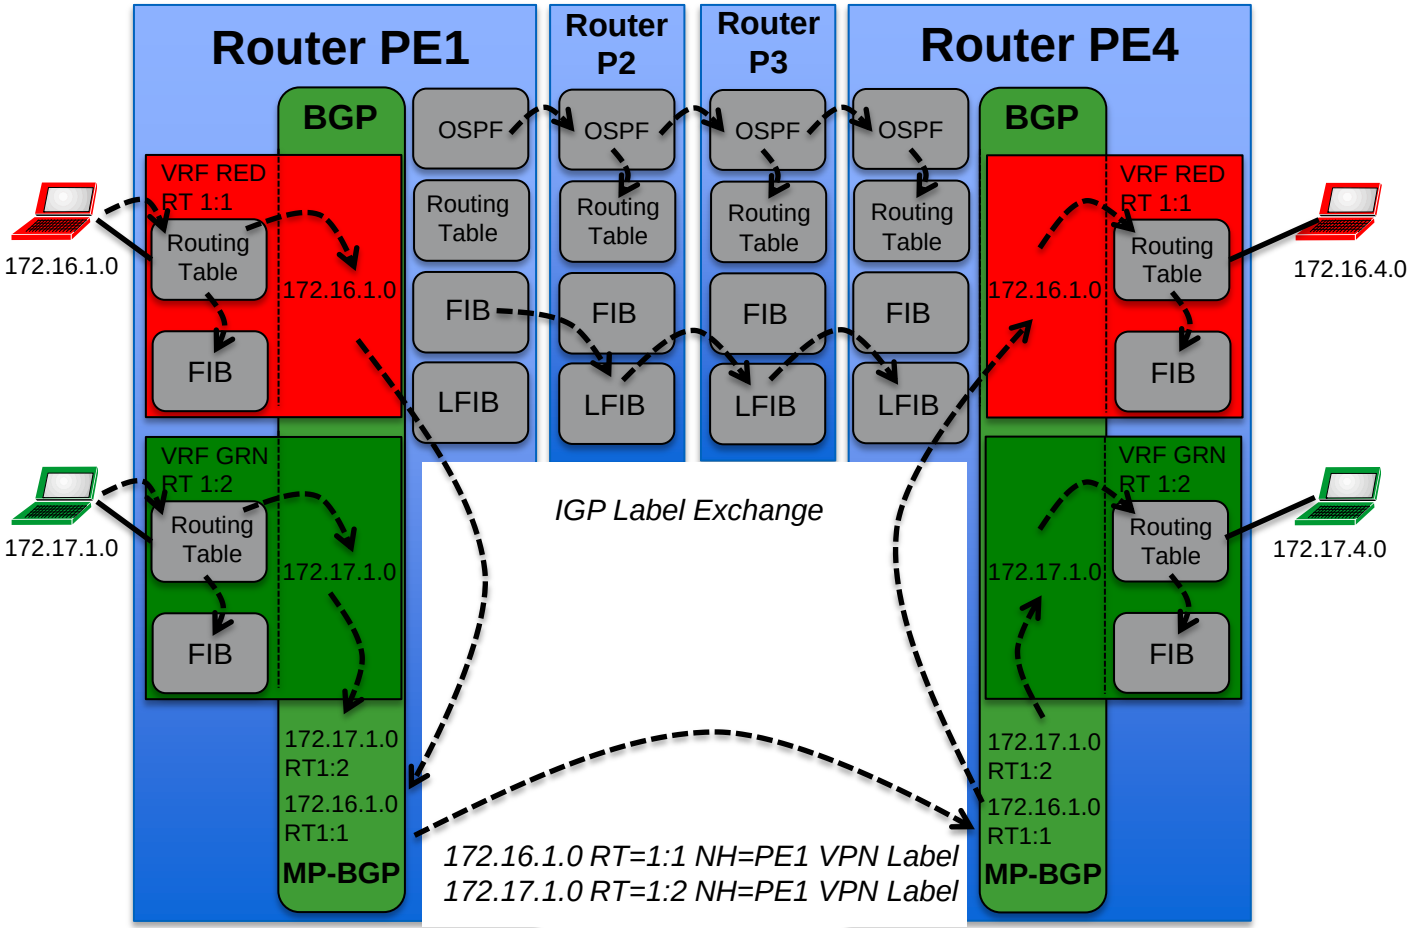

MPLS-VPN – Packet Flow

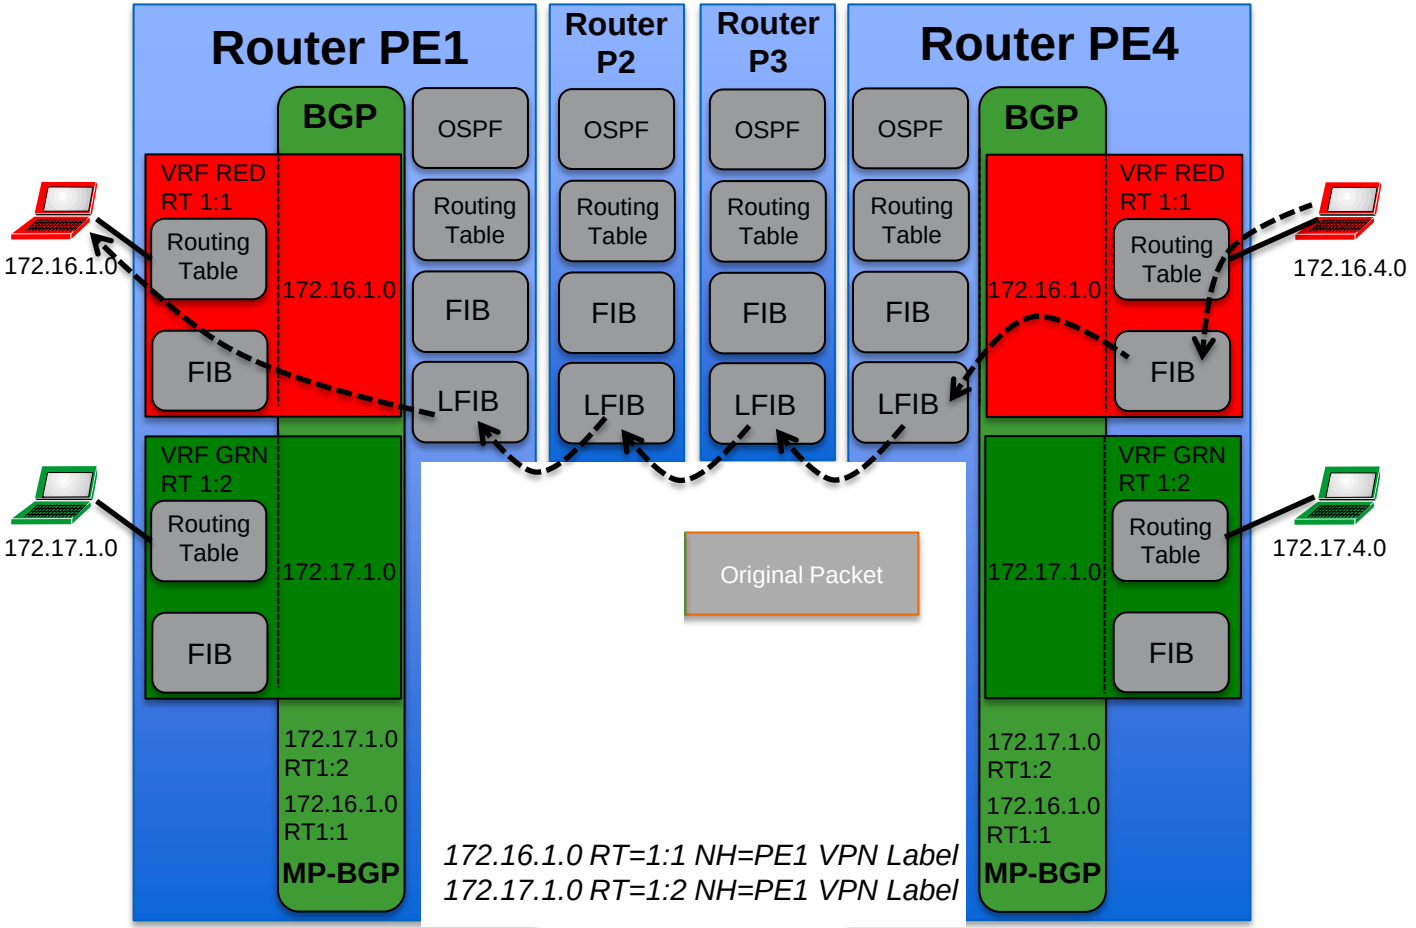

MPLS-VPN

Configuration (PE)

```
vrf definition GRN
 rd 1:2
 !
 address-family ipv4
  route-target export 1:2
  route-target import 1:2
 exit-address-family
 !
 address-family ipv6
  route-target export 1:2
  route-target import 1:2
 exit-address-family
!
vrf definition RED
 rd 1:1
 !
 address-family ipv4
  route-target export 1:1
  route-target import 1:1
 exit-address-family
 !
 address-family ipv6
  route-target export 1:1
  route-target import 1:1
 exit-address-family
```

Defining the VRFs IPv4 and IPv6

RD is required for BGP

Import and Export to populate VRF routing table

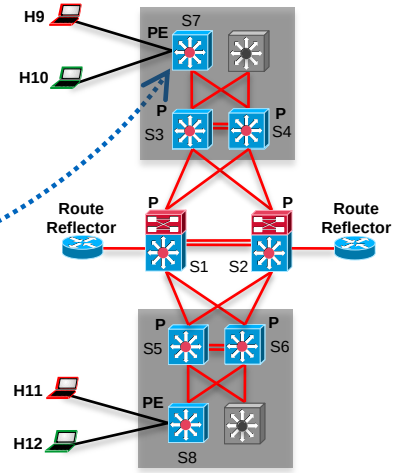

MPLS-VPN

Configuration (PE)

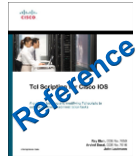


```
interface Loopback0
 ip address 192.168.0.8 255.255.255.255

interface Ethernet0/0
 ip address 192.168.85.8 255.255.255.0
 mpls ip
!
interface Ethernet0/1
 ip address 192.168.86.8 255.255.255.0
 mpls ip
!
router eigrp 1
 network 192.168.0.0 0.0.255.255

interface Ethernet0/2
 vrf forwarding GRN
 ip address 172.17.8.8 255.255.255.0
 ipv6 address 2001:17:8::8/64
!
interface Ethernet0/3
 vrf forwarding RED
 ip address 172.16.8.8 255.255.255.0
 ipv6 address 2001:16:8::8/64
```

Host-route on loopback for directed LDP session

Network Side Interfaces

IGP for propagation of loopbacks

Client Side Interface

IPv4 and IPv6 address assignment

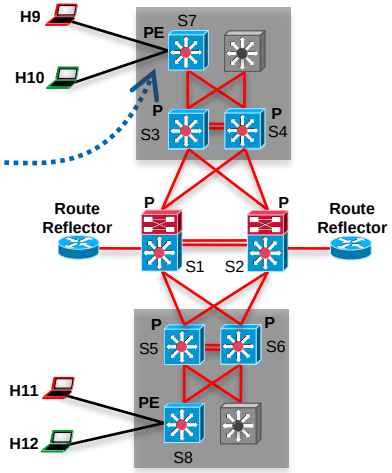

MPLS-VPN

Configuration (PE)

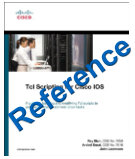


```
router bgp 65000
 neighbor 192.168.0.13 remote-as 65000
 neighbor 192.168.0.13 update-source Loopback0
 neighbor 192.168.0.14 remote-as 65000
 neighbor 192.168.0.14 update-source Loopback0
 !
 address-family vpnv4
  neighbor 192.168.0.13 activate
  neighbor 192.168.0.13 send-community extended
  neighbor 192.168.0.14 activate
  neighbor 192.168.0.14 send-community extended
!
 address-family vpnv6
  neighbor 192.168.0.13 activate
  neighbor 192.168.0.13 send-community extended
  neighbor 192.168.0.14 activate
  neighbor 192.168.0.14 send-community extended
```

BGP base configuration

VPNv4 configuration

VPNv6 configuration

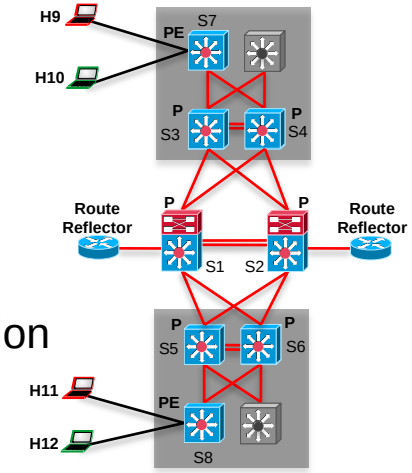

MPLS-VPN

Configuration (PE)

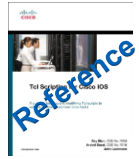


```
router bgp 65000
!
 address-family ipv4 vrf GRN
  redistribute connected
!
 address-family ipv6 vrf GRN
  redistribute connected
!
 address-family ipv4 vrf RED
  redistribute connected
!
 address-family ipv6 vrf RED
  redistribute connected
```

VRF address-family

Redistribute locally connected routes

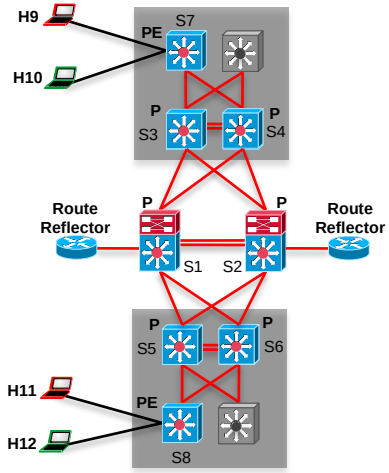

MPLS-VPN

Configuration – Route Reflector (RR)

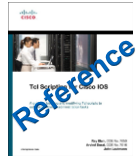


```
router bgp 65000
 no bgp default route-target filter
 neighbor AS65000 peer-group
 neighbor AS65000 remote-as 65000
 neighbor AS65000 update-source Loopback0
 neighbor AS65000 route-reflector-client
 neighbor 192.168.0.7 peer-group AS65000
 neighbor 192.168.0.8 peer-group AS65000
 !
 address-family vpnv4
  neighbor AS65000 send-community extended
  neighbor AS65000 route-reflector-client
  neighbor 192.168.0.7 activate
  neighbor 192.168.0.8 activate
!
 address-family vpnv6
  neighbor AS65000 send-community extended
  neighbor AS65000 route-reflector-client
  neighbor 192.168.0.7 activate
  neighbor 192.168.0.8 activate
```

BGP base configuration

Route-target filter to allow all VPN routes in

VPNv4 configuration

VPNv6 configuration

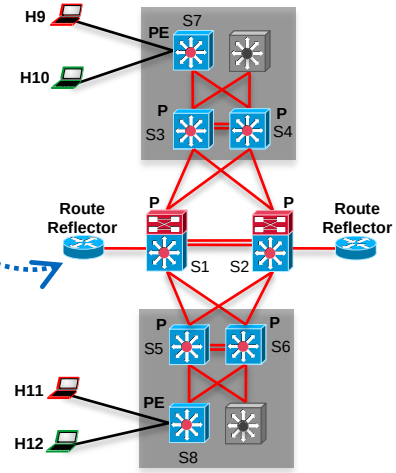

MPLS-VPN

Traffic Example

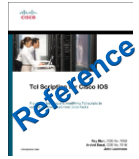


```
H9#trace 172.16.8.11
Tracing the route to 172.16.8.11
VRF info: (vrf in name/id, vrf out name/id)
  1 172.16.7.7 0 msec 4 msec 4 msec
  2 192.168.74.4 [MPLS: Labels 22/22 Exp 0] 0 msec 4 msec 2 msec
  3 192.168.41.1 [MPLS: Labels 22/22 Exp 0] 0 msec 1 msec 0 msec
  4 192.168.61.6 [MPLS: Labels 22/22 Exp 0] 1 msec 1 msec 1 msec
  5 172.16.8.8 1 msec 1 msec 5 msec
  6 172.16.8.11 1 msec *  0 msec

H10#trace 172.17.8.12
Tracing the route to 172.17.8.12
VRF info: (vrf in name/id, vrf out name/id)
  1 172.17.7.7 2 msec 0 msec 0 msec
  2 192.168.73.3 [MPLS: Labels 22/20 Exp 0] 1 msec 0 msec 0 msec
  3 192.168.32.2 [MPLS: Labels 22/20 Exp 0] 1 msec 1 msec 1 msec
  4 192.168.62.6 [MPLS: Labels 22/20 Exp 0] 1 msec 1 msec 0 msec
  5 172.17.8.8 1 msec 1 msec 1 msec
  6 172.17.8.12 0 msec *  1 msec
```

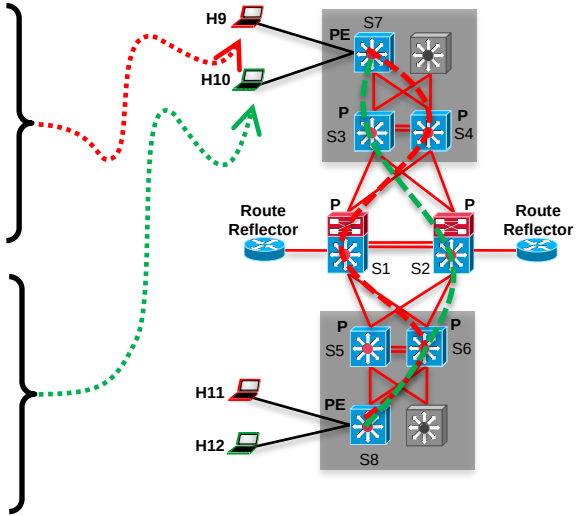

Traceroute indicates labels

The hosts in this example (H9/H10) are IOS routers

MPLS-VPN

Traffic Example

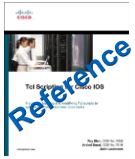


```
H9#trace 2001:16:8::11
Tracing the route to 2001:16:8::11
  1 2001:16:7::7 1 msec 0 msec 4 msec
  2 ::FFFF:192.168.73.3 [MPLS: Labels 22/23 Exp 0] 0 msec 0 msec 0 msec
  3 ::FFFF:192.168.32.2 [MPLS: Labels 22/23 Exp 0] 1 msec 1 msec 2 msec
  4 ::FFFF:192.168.62.6 [MPLS: Labels 22/23 Exp 0] 1 msec 1 msec 1 msec
  5 2001:16:8::8 0 msec 0 msec 0 msec
  6 2001:16:8::11 1 msec 5 msec 1 msec


H10#trace 2001:17:8::12
Tracing the route to 2001:17:8::12
  1 2001:17:7::7 4 msec 5 msec 4 msec
  2 ::FFFF:192.168.74.4 [MPLS: Labels 22/21 Exp 0] 2 msec 1 msec 0 msec
  3 ::FFFF:192.168.42.2 [MPLS: Labels 22/21 Exp 0] 1 msec 1 msec 0 msec
  4 ::FFFF:192.168.62.6 [MPLS: Labels 22/21 Exp 0] 0 msec 0 msec 1 msec
  5 2001:17:8::8 0 msec 1 msec 1 msec
  6 2001:17:8::12 1 msec 1 msec 1 msec
```

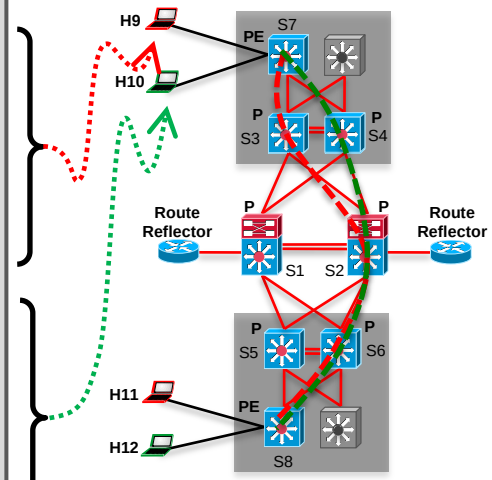

Traceroute indicates labels

IPv4 core only

The hosts in this example (H9/H10) are IOS routers

MPLS-VPN

Considerations

- Deployment
  - Highly scalable
  - Purpose-built route-reflectors recommended
  - Any-to-any connectivity within VPNs
  - Pseudo-wire support (DCI/Legacy applications)
  - Supported on Catalyst 6500 (Sup720 and Sup32 – no DFC3A/PFC3A), Sup2T, Nexus 7000, ME3750, ME3600/3800 and ASR9K
- Application and Services
  - Multiple VRF-aware Services available
- Learning Curve
  - MPLS
  - Multi-Protocol BGP

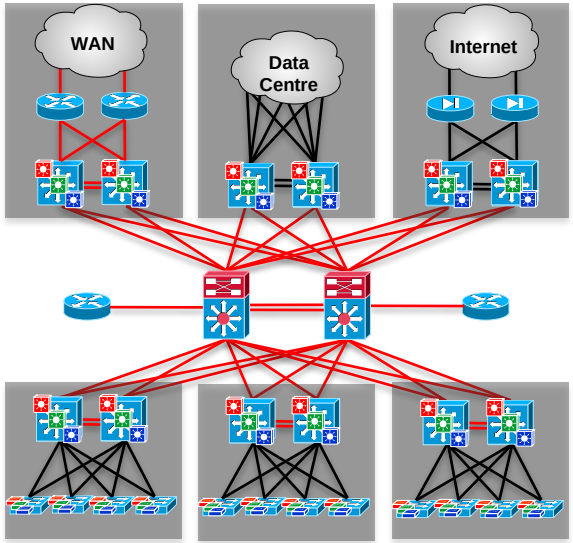




Solid Design

Solid Design

What's Required?

- Hierarchical Network Design
  - Core, Distribution, Access
- Redundancy, Load balancing
  - FHRP – HSRP, VRRP, GLBP
  - Redundant paths
  - CEF L3/L4 Load Balancing
- Minimise Protocol Exchanges
  - Summarise routes to core
  - Passive interfaces on Access
  - Hard-set Trunks and Channels
- L2 **Convergence and Security**
  - Use RSTP+, Set STP Roles (Root, Backup)
  - STP Toolkit (RootGuard, STP priorities, BPDU Guard)
  - Control Plane Policing (CPP)
  - Catalyst Integrated Security Features (CISF)

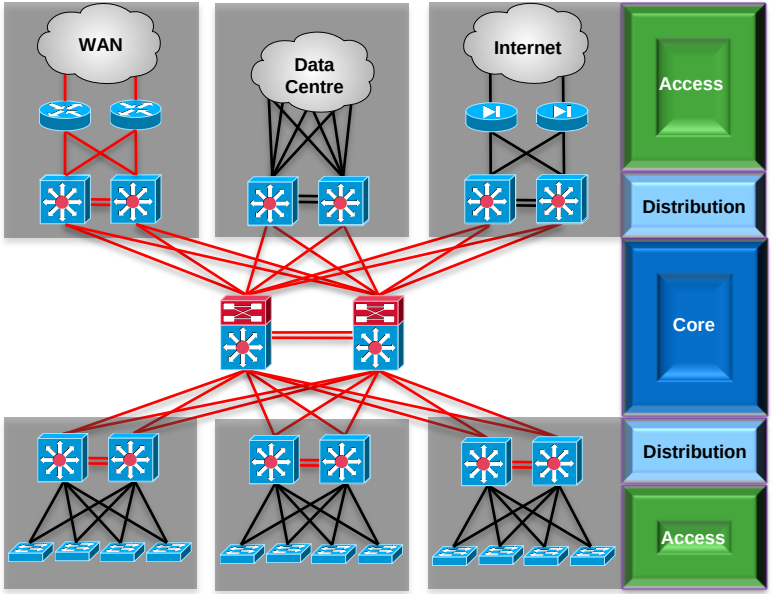




Additional Virtualised Services

Authentication

802.1X with Dynamic VLAN Assignment

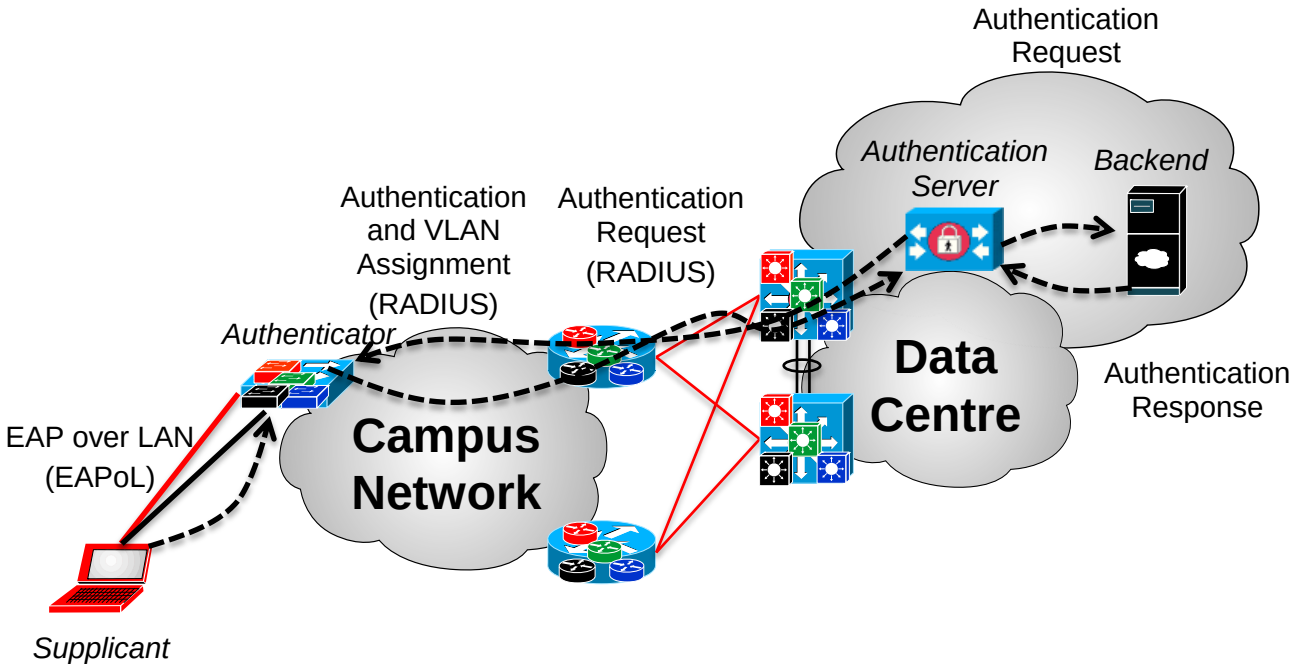

Unicast Shared Services

Using Route-Leaking

- Service sharing (DHCP, DNS, etc…)

- Leverage the BGP route-target mechanism for route leaking
  - No support for overlapping IP addresses across VPNs

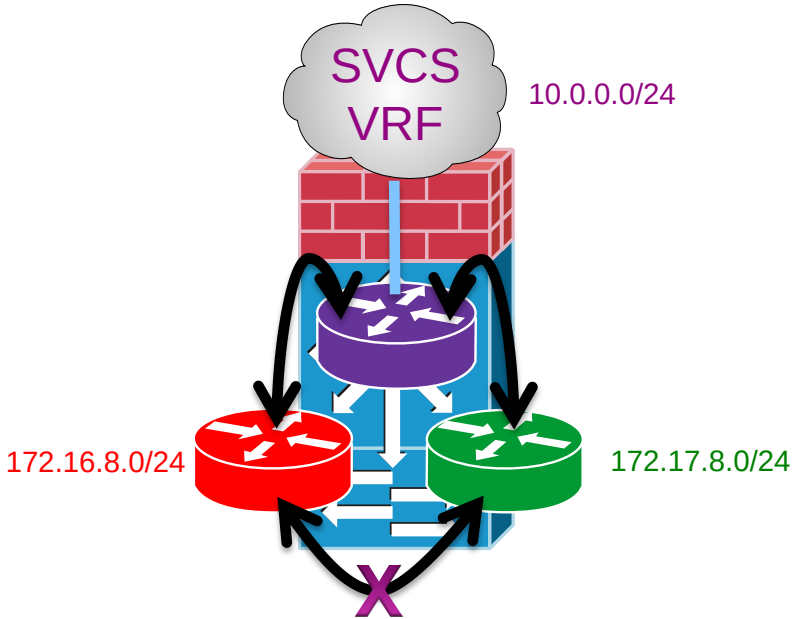

Unicast Shared Services

MPLS-VPN Configuration

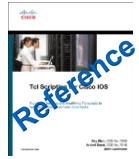


```
vrf definition SVCS
 rd 1:100
 !
 address-family ipv4
  route-target export 1:100
  route-target export 1:1
  route-target export 1:2
  route-target import 1:100
  route-target import 1:1
  route-target import 1:2
!
 address-family ipv6
  route-target export 1:100
  route-target export 1:1
  route-target export 1:2
  route-target import 1:100
  route-target import 1:1
  route-target import 1:2
```

Defining the VRFs IPv4 and IPv6

RD is required for BGP

Import and Export to populate VRF routing table

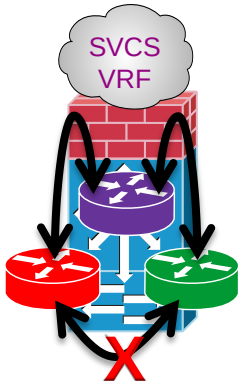

Unicast Shared Services

MPLS-VPN Verification

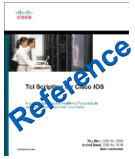


```
S8#show ip route vrf RED
      10.0.0.0/24 is subnetted, 1 subnets
B        10.0.0.0 [200/0] via 192.168.0.7, 00:16:35
      172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        172.16.8.0/24 is directly connected, Ethernet0/3
L        172.16.8.8/32 is directly connected, Ethernet0/3

S8#show ip route vrf GRN
      10.0.0.0/24 is subnetted, 1 subnets
B        10.0.0.0 [200/0] via 192.168.0.7, 00:16:42
      172.17.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        172.17.8.0/24 is directly connected, Ethernet0/2
L        172.17.8.8/32 is directly connected, Ethernet0/2

S8#show ipv6 route vrf RED
B   2001:10::/64 [200/0]
     via 192.168.0.7%default, indirectly connected
C   2001:16:8::/64 [0/0]
     via Ethernet0/3, directly connected
L   2001:16:8::8/128 [0/0]
     via Ethernet0/3, receive
L   FF00::/8 [0/0]
     via Null0, receive

S8#show ipv6 route vrf GRN
B   2001:10::/64 [200/0]
     via 192.168.0.7%default, indirectly connected
C   2001:17:8::/64 [0/0]
     via Ethernet0/2, directly connected
L   2001:17:8::8/128 [0/0]
     via Ethernet0/2, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

Each VRF contains local and shared routing information

Unicast Shared Services

EVN

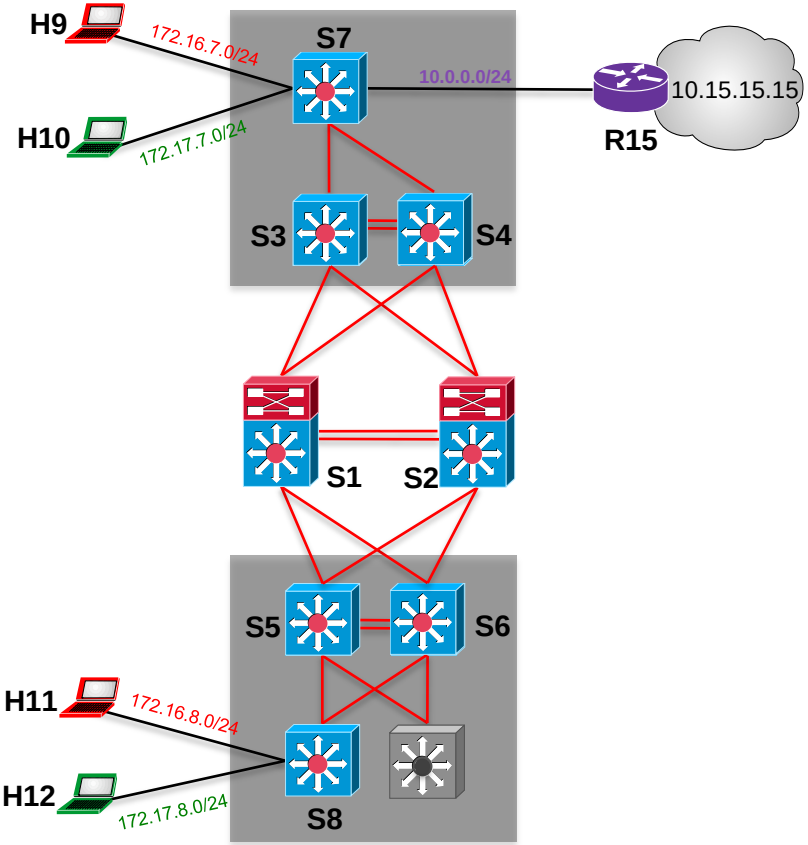

Unicast Shared Services

EVN Configuration

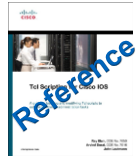


```
vrf definition GRN
 vnet tag 102
 !
 address-family ipv4
  route-replicate from vrf SVCS unicast all
!
vrf definition RED
 vnet tag 101
 !
 address-family ipv4
  route-replicate from vrf SVCS unicast all
!
vrf definition SVCS
 vnet tag 100
 !
 address-family ipv4
  route-replicate from vrf RED unicast all route-map RED-IMPORT
  route-replicate from vrf GRN unicast all route-map GRN-IMPORT

route-map RED-IMPORT permit 10
 match ip address RED-ACL
!
route-map GRN-IMPORT permit 10
 match ip address GRN-ACL

ip access-list standard GRN-ACL
 permit 172.17.0.0 0.0.255.255
ip access-list standard RED-ACL
 permit 172.16.0.0 0.0.255.255
```

Defining the IPv4 VRFs, assign a tag and configure route replication

Create route-maps and access-lists

Unicast Shared Services

EVN Configuration

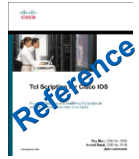


```
router eigrp LAB
 !
 address-family ipv4 unicast vrf RED autonomous-system 16
  !
  topology base
   redistribute vrf SVCS eigrp 100
  exit-af-topology
  network 172.16.0.0
  network 192.168.0.0 0.0.255.255
!
 address-family ipv4 unicast vrf GRN autonomous-system 17
  !
  topology base
   redistribute vrf SVCS eigrp 100
  exit-af-topology
  network 172.17.0.0
  network 192.168.0.0 0.0.255.255
!
address-family ipv4 unicast vrf SVCS autonomous-system 100
  !
  topology base
   redistribute vrf RED eigrp 16
   redistribute vrf GRN eigrp 16
  exit-af-topology
  network 10.0.0.0
```

Redistribute routing information

Unicast Shared Services

EVN Verification

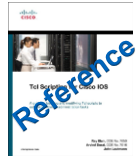


```
S7#routing-context vrf SVCS
S7%SVCS#sh ip route

Routing Table: SVCS
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C        10.0.0.0/24 is directly connected, Ethernet1/0
L        10.0.0.7/32 is directly connected, Ethernet1/0
D        10.15.15.0/24 [90/409600] via 10.0.0.15, 01:19:53, Ethernet1/0
      172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C   +    172.16.7.0/24 is directly connected (RED), Ethernet0/3
L   +    172.16.7.7/32 is directly connected (RED), Ethernet0/3
D   +    172.16.8.0/24
           [90/384000] via 192.168.74.4 (RED), 02:00:56, Ethernet0/0.101
           [90/384000] via 192.168.73.3 (RED), 02:00:56, Ethernet0/1.101
      172.17.0.0/16 is variably subnetted, 3 subnets, 2 masks
C   +    172.17.7.0/24 is directly connected (GRN), Ethernet0/2
L   +    172.17.7.7/32 is directly connected (GRN), Ethernet0/2
D   +    172.17.8.0/24
           [90/384000] via 192.168.74.4 (GRN), 02:00:55, Ethernet0/0.102
           [90/384000] via 192.168.73.3 (GRN), 02:00:55, Ethernet0/1.102
```

Imported RED routes

Imported GRN routes

Shared Services Edge

Fusion Router

- A Fusion router provides:
  - Inter-VPN connectivity
  - Protected access to shared resources
- Use a Firewall for:
  - VPN isolation/protection
  - Application of per VPN policies
  - Leverage multi-context functionality
- Firewall modes of operation
  - FW in Transparent Mode
  - FW in Routed Mode

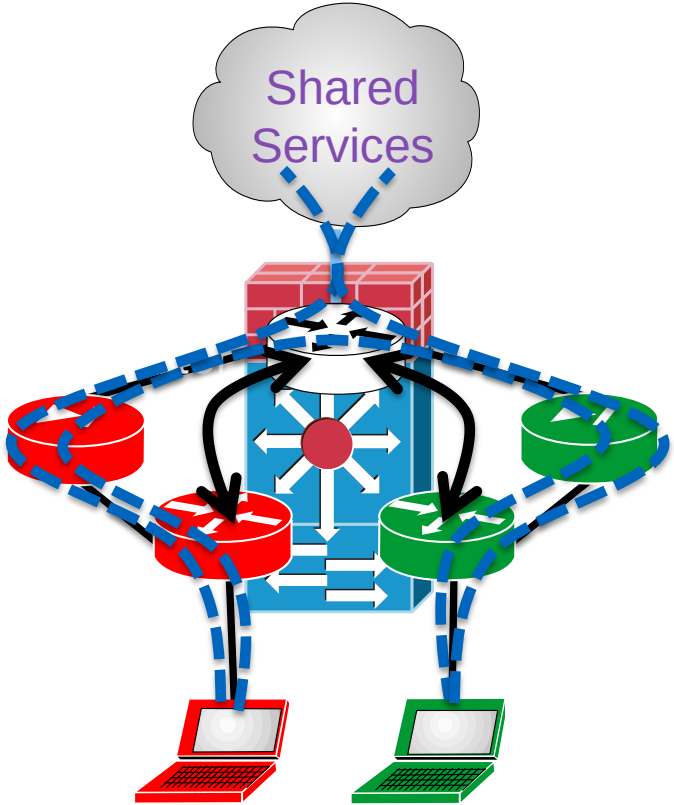

Protected Services

Deploying Firewall Contexts in Routed Mode

- Firewall acts as L3 hop
  - ASA 9.0 supports OSPFv2 and EIGRP
- Use BGP over-the-top of the firewall context
  - Static routes are still required!
- A “Fusion” VRF may be used

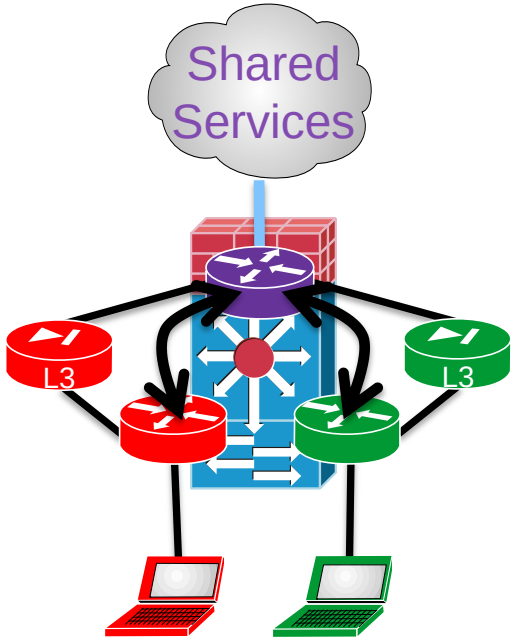

Protected Services

Deploying Firewall Contexts in Transparent Mode

- Firewall acts as L2 bridge
- Peering protocols:
  - Use IGP (EIGRP or OSPF) for VRF-lite deployments
  - Use BGP for MPLS-VPN scenarios
- A “Fusion” VRF may be used
  - Define MAC addresses on switch interfaces

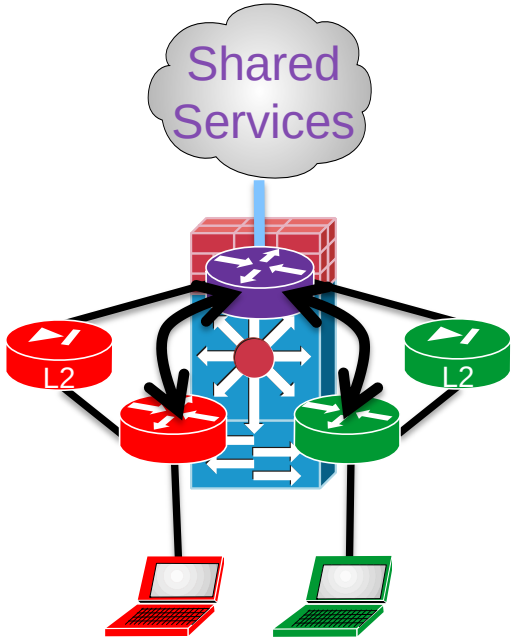

Multicast Shared Services

Multicast Overview

- Multicast crosses VRF boundaries
- Ensure RFP check is successful
  - Route-Leaking
  - VRF Fallback
  - VRF Select

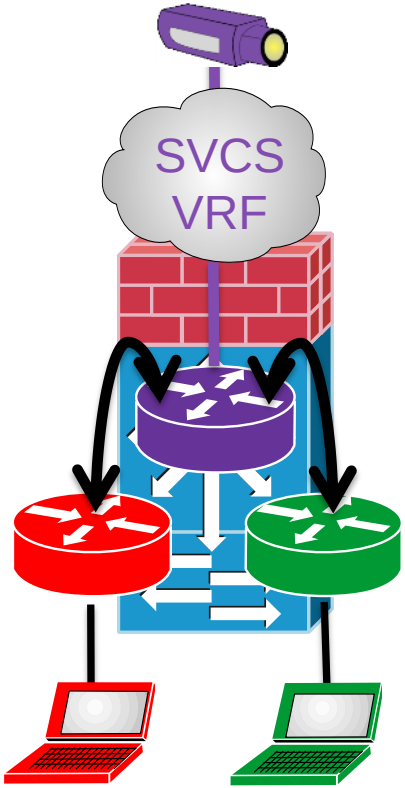

MPLS VPN and Multicast

Concept and Fundamentals

- Enable multicast in the core
- The MPLS Core forms a Default MDT for each given VRF defined on the PE
- A High-bandwidth source for that customer starts sending traffic
- Interested receivers 1 & 2 join that High Bandwidth source
- The Data-MDT is formed for this High-Bandwidth source

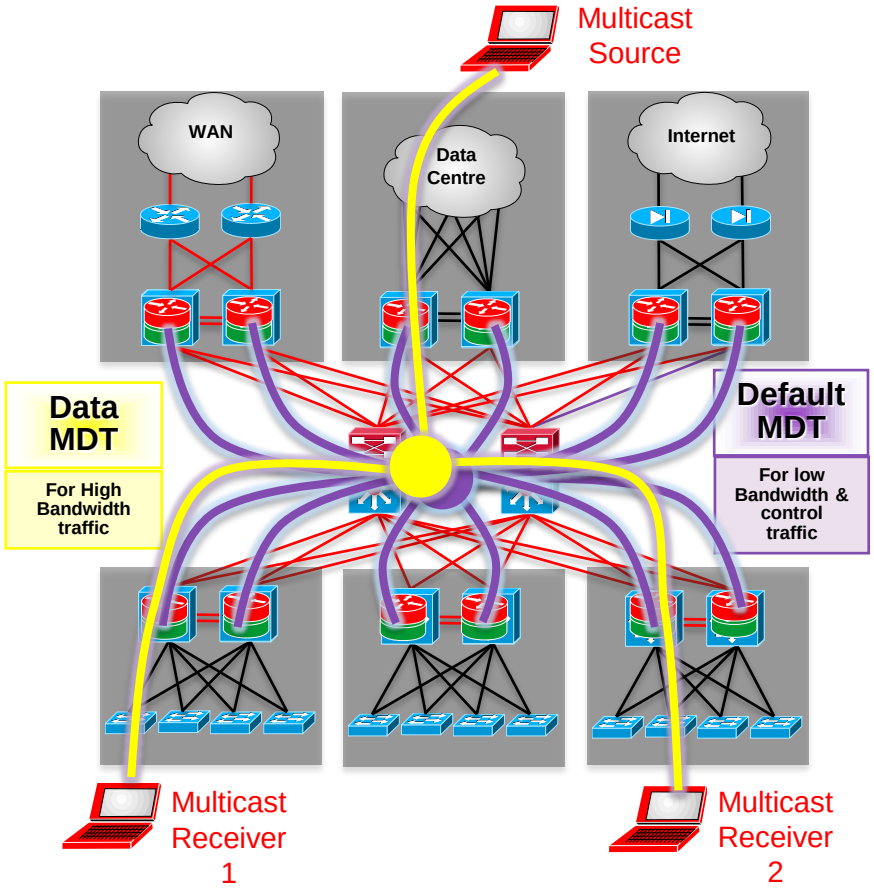

Multicast Shared Services

Three ways to perform Extranet with IP Multicast today

- BGP Route-Target Import
  - Uses BGP or EVN to exchange routes between VRFs
  - No overlapping IP addresses
- VRF Fallback
  - Used when the route doesn’t exist in receiver VRF
  - Con: VRF Fallback can’t be used with a default unicast route
  - Con: Can’t be used if source addresses overlap between VRFs
- VRF Select
  - Statically assigns a VRF to RPF for a multicast group range
  - Pro: Can be used with overlapping source addresses

QoS and Network Virtualisation

Overview

- Classify and mark traffic at the edge
- Traffic is queued/shaped according to DSCP values or MPLS EXP bits
- MPLS EXP only offer 8 classes
- Choose the appropriate class of service
  - Web – Best effort/scavenger
  - Voice – Priority
  - Other – you decide

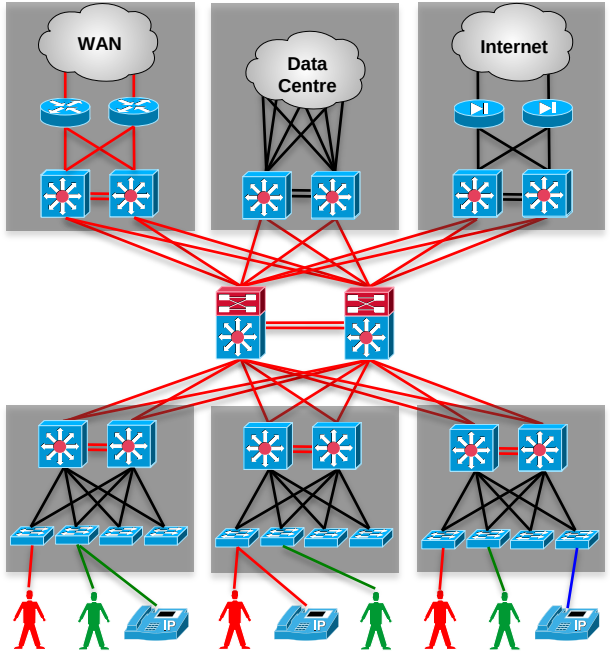

DC Integration

Data Centre

Integration

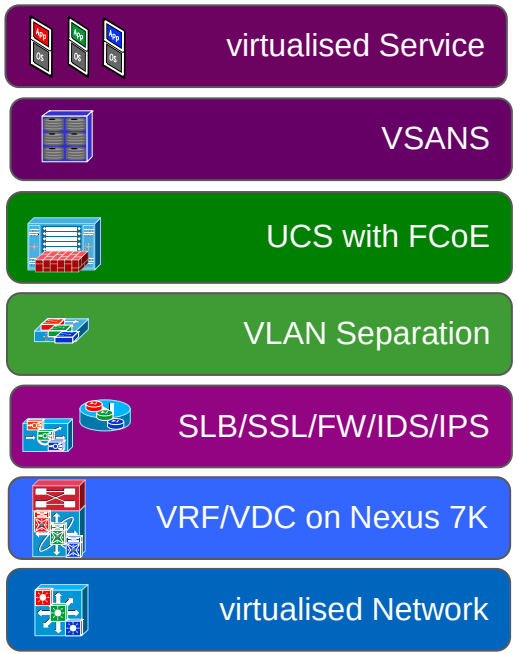

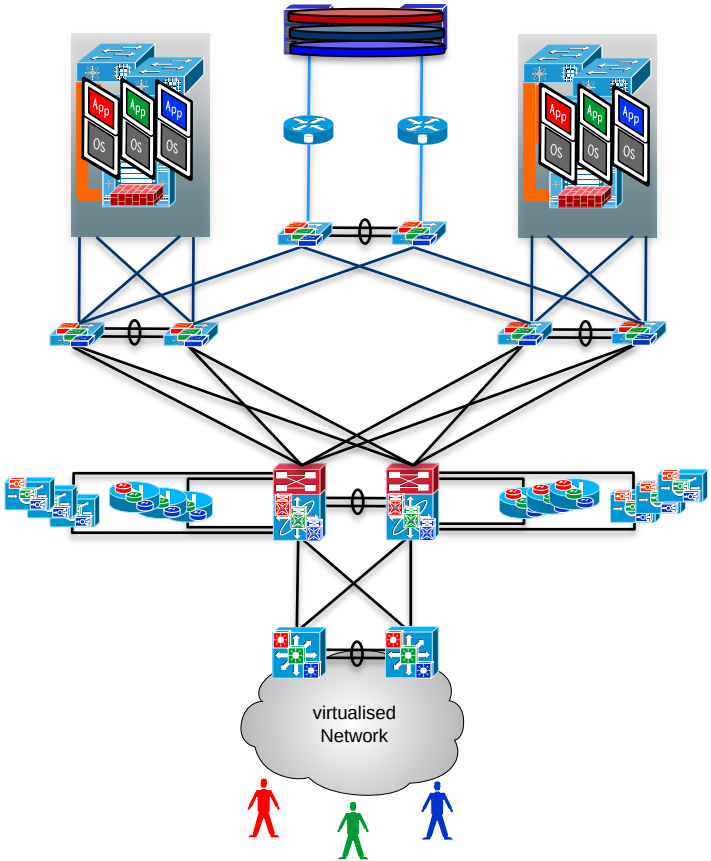

Agenda

- Virtualisation solves these Challenges
- Virtualisation Architectures
- Case Study
- Industry Trends
- Putting it all Together

VRF-lite End-to-End

Pros:

- No MP-BGP configuration
- L3 to the edge
- Minimise impact on distribution layer)
- Lower cost solution
- VSS

Cons:

- Adding VRFs is arduous
- Limited scalability
- Import/export of routes requires additional equipment

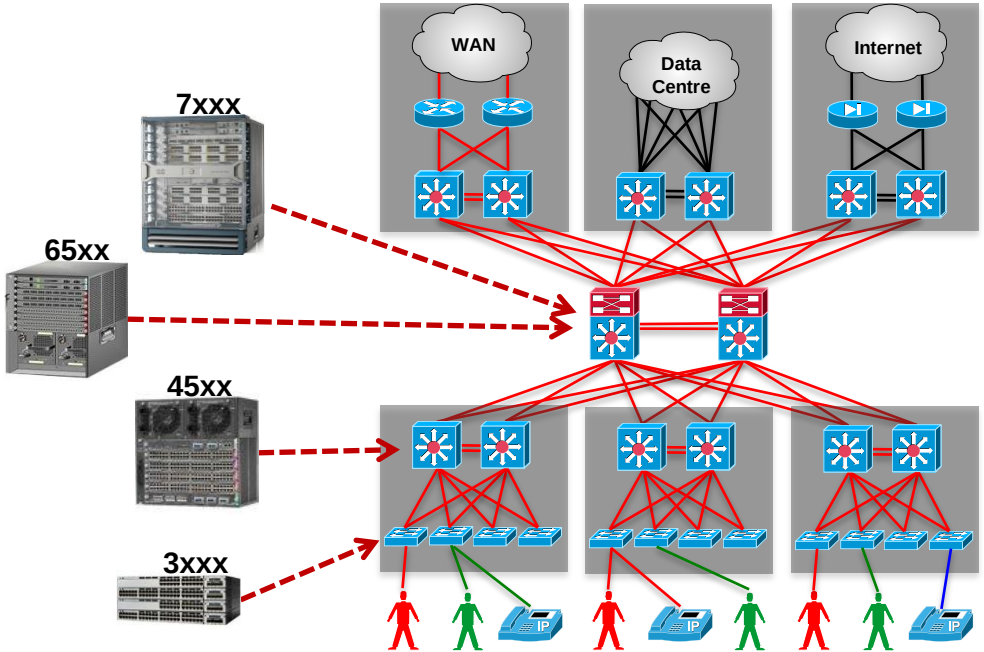

EVN w/ L2 Access

Pros:

- No MP-BGP configuration
- Lower cost solution
- VSS

Cons:

- Limited product support (today)
- No IPv6 support (today)
- FHRP on distribution devices

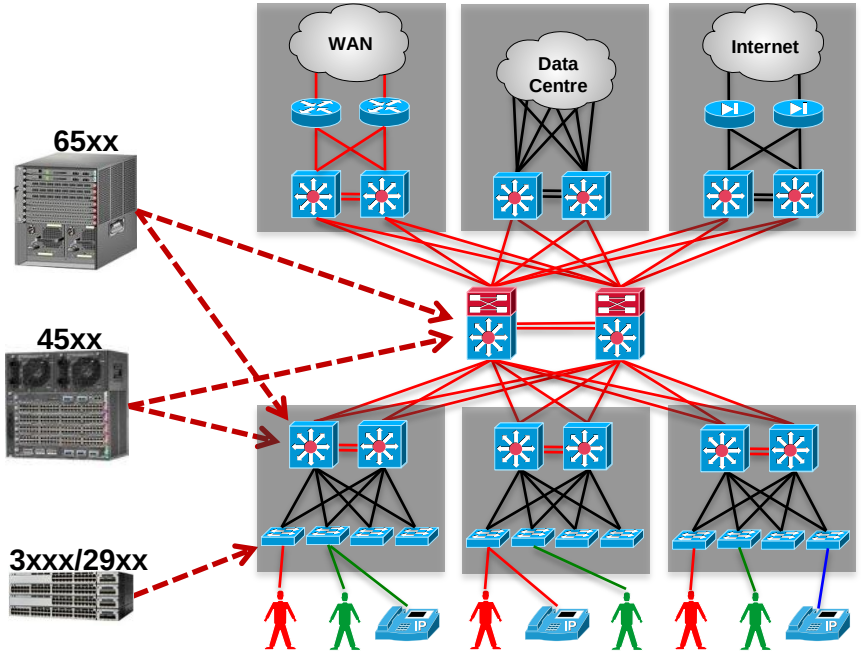

MPLS-VPN w/ L2 Access

Pros:

- Very scalable
- Pseudo-wire support
- IPv6 support (6VPE)
- VSS

Cons:

- MP-BGP configuration
- Multicast configuration is complex
- FHRP on distribution devices

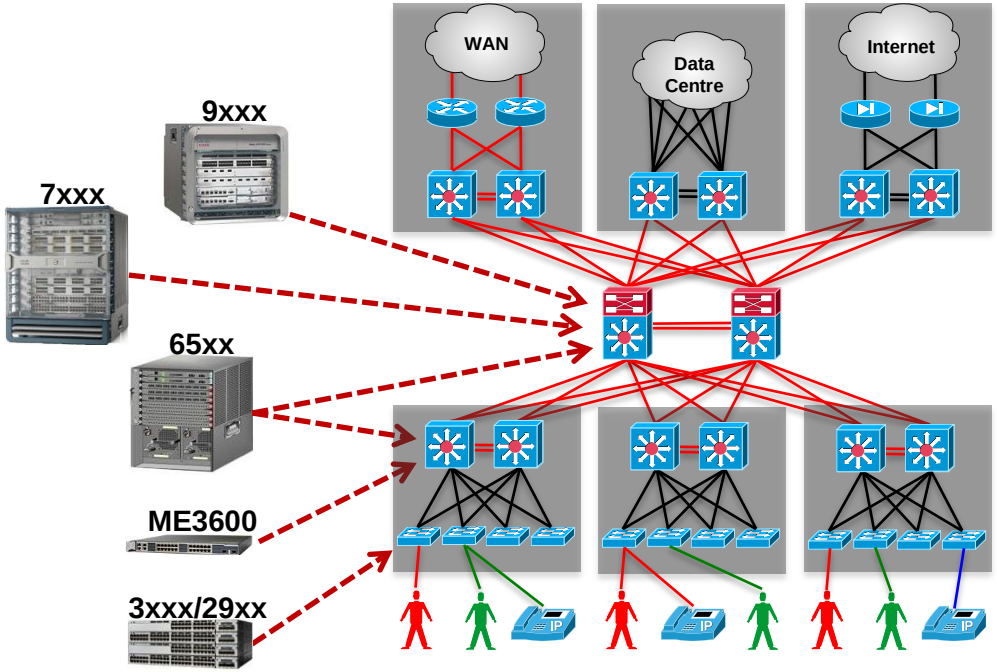

MPLS-VPN w/ L3 VRF-lite/EVN Access

Pros:

- L3 to the edge
- Minimise impact on distribution layer (FHRP)

Cons:

- Complex route redistribution

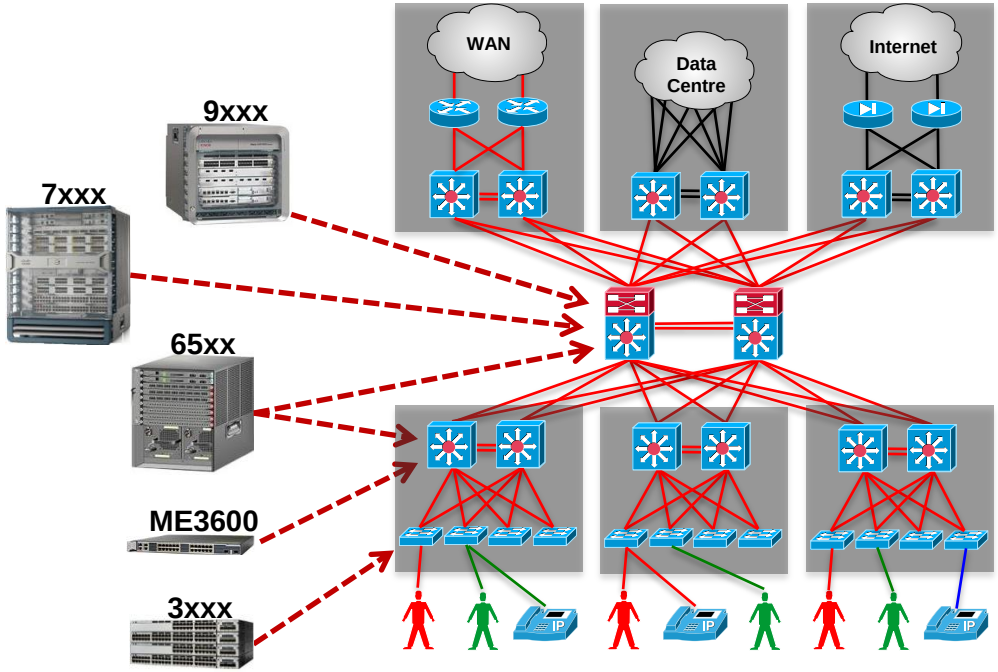

Agenda

- Virtualisation solves these Challenges
- Virtualisation Architectures
- Case Study
- Industry Trends
- Putting it all Together

Locator/ID Separation Protocol (LISP)

What is LISP?

Summary

- Originally conceived to address Internet scaling challenges
- Locator/Identity split creates a “level of indirection” by using two namespaces – hosts and locators
- Similar to DNS
- LISP involves an host-to-locator lookup…

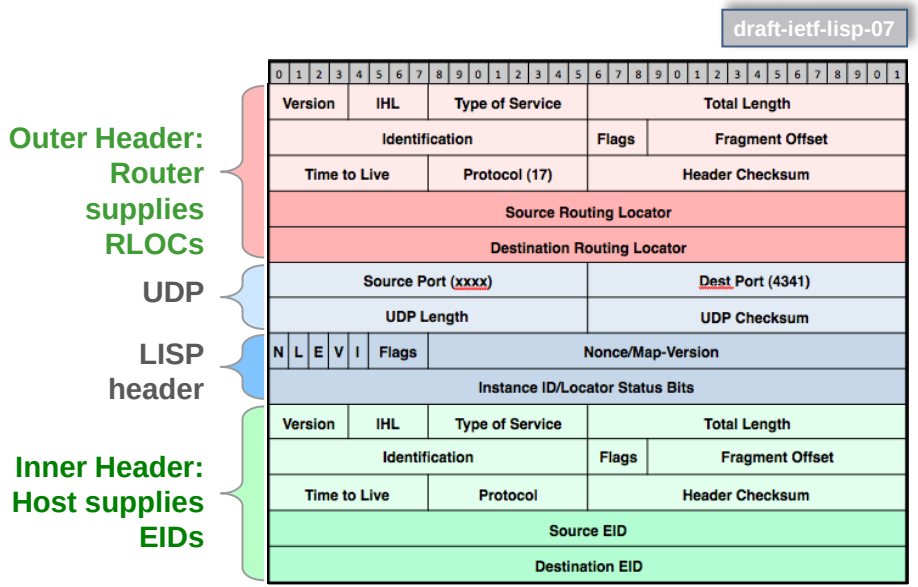

What are the Components of LISP?

- LISP Loc/ID Split namespaces
  - EID (Endpoint Identifier) is the IP address of a host
  - RLOC (Routing Locator) is the IP address of the LISP router
  - EID-to-RLOC is the mapping
- MS/MR
  - Map-Resolver and Map-Server (similar to DNS Resolver and DNS Server)
- ITR – Ingress Tunnel Router
  - Receives packets from site-facing interfaces
  - Encapsulation to remote LISP sites or native-forward to non-LISP sites
- ETR – Egress Tunnel Router
  - Receives packets from core-facing interfaces
  - De-capsulation and deliver packets to local EIDs at site

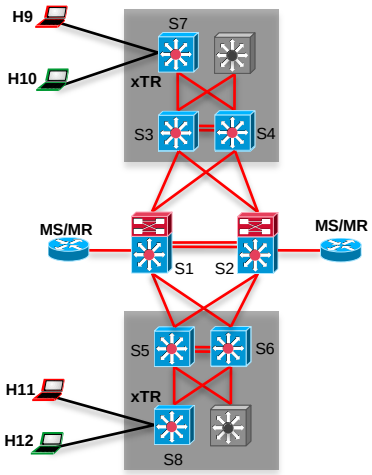

Test Diagram

LISP

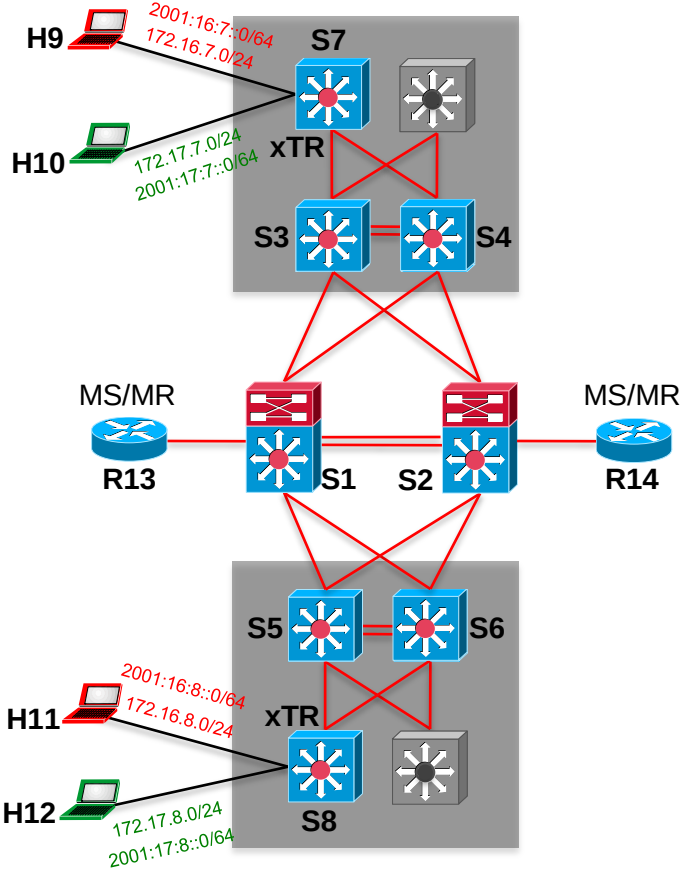

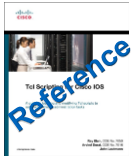

LISP
Configuration (xTR)

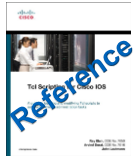


```
vrf definition GRN
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
vrf definition RED
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
```

GRN VRF

RED VRF

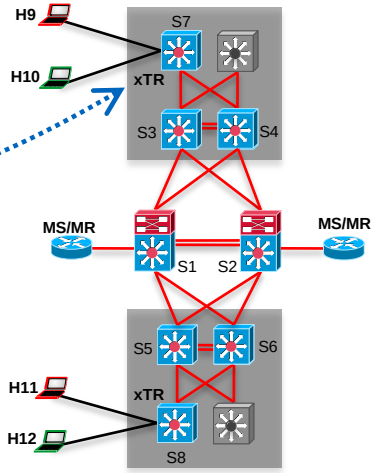

LISP
Configuration (xTR)

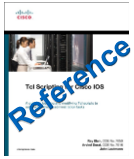


```
router lisp
 eid-table default instance-id 0
!
 eid-table vrf RED instance-id 101
  database-mapping 172.16.7.0/24 192.168.0.7 priority 1 weight 1
  database-mapping 2001:16:7::/64 192.168.0.7 priority 1 weight 1
!
 eid-table vrf GRN instance-id 102
  database-mapping 172.17.7.0/24 192.168.0.7 priority 1 weight 1
  database-mapping 2001:17:7::/64 192.168.0.7 priority 1 weight 1
!
 ipv4 itr map-resolver 192.168.0.13
 ipv4 itr map-resolver 192.168.0.14
 ipv4 itr
 ipv4 etr map-server 192.168.0.13 key R7
 ipv4 etr map-server 192.168.0.14 key R7
 ipv4 etr
 ipv6 itr map-resolver 192.168.0.13
 ipv6 itr map-resolver 192.168.0.14
 ipv6 itr
 ipv6 etr map-server 192.168.0.13 key R7
 ipv6 etr map-server 192.168.0.14 key R7
 ipv6 etr
```

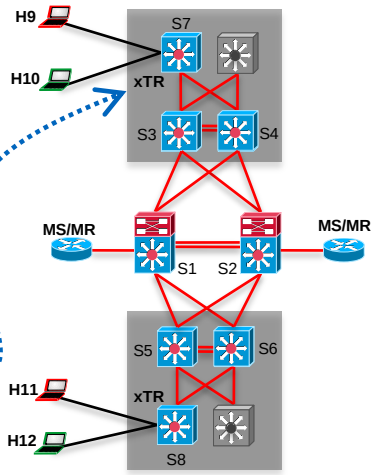

LISP
Configuration (MS/MR)

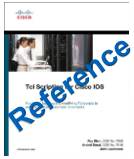


```
router lisp
 site R7
  authentication-key R7
  eid-prefix instance-id 101 172.16.7.0/24
  eid-prefix instance-id 101 2001:16:7::/64
  eid-prefix instance-id 102 172.17.7.0/24
  eid-prefix instance-id 102 2001:17:7::/64
!
 site R8
  authentication-key R8
  eid-prefix instance-id 101 172.16.8.0/24
  eid-prefix instance-id 101 2001:16:8::/64
  eid-prefix instance-id 102 172.17.8.0/24
  eid-prefix instance-id 102 2001:17:8::/64
!
 ipv4 map-server
 ipv4 map-resolver
 ipv6 map-server
 ipv6 map-resolver
```

R7 configuration

R8 configuration

MS/MR configuration

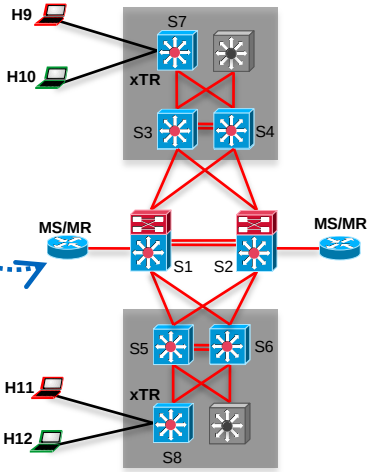

LISP

Traffic Example

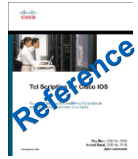


```
H9#trace ipv6 2001:16:8::11
Type escape sequence to abort.
Tracing the route to 2001:16:8::11

  1 2001:16:7::7 1 msec 15 msec 10 msec
  2 2001:16:8::8 1 msec 0 msec 1 msec
  3 2001:16:8::11 0 msec 0 msec 1 msec

H10#trace ipv6 2001:17:8::12
Type escape sequence to abort.
Tracing the route to 2001:17:8::12

  1 2001:17:7::7 1 msec 12 msec 9 msec
  2 2001:17:8::8 1 msec 0 msec 1 msec
  3 2001:17:8::12 0 msec 0 msec 1 msec
```

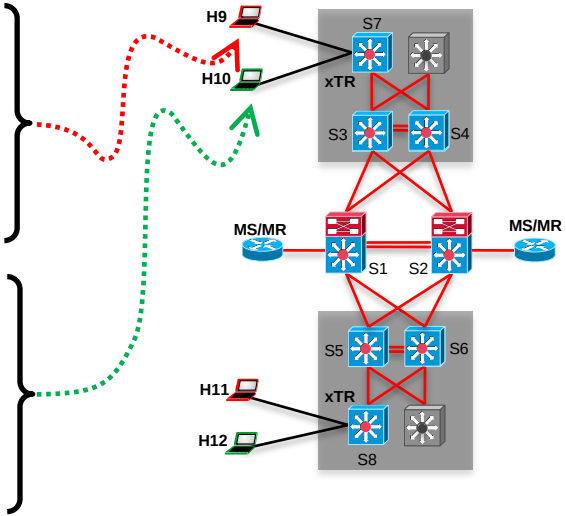

The hosts in this example (H9/H10) are IOS routers

LISP – Traffic Capture

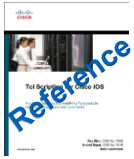


```
S1#
================================================================================
16:02:53.215 GMT Sun Oct 21 2012                  Relative Time: 23.871998
Packet 30 of 223                                  In: Ethernet0/2

Ethernet Packet:  626 bytes
      Dest Addr: AABB.CC00.0120,   Source Addr: AABB.CC00.0501
      Protocol: 0x0800

IP    Version: 0x4,  HdrLen: 0x5,  TOS: 0xC0 (Prec=Internet Contrl)
      Length: 612,   ID: 0x095B,   Flags-Offset: 0x4000 (don't fragment)
      TTL: 253,   Protocol: 17 (UDP),   Checksum: 0x9B0D (OK)
     Source: 192.168.85.8,     Dest: 192.168.0.7

UDP   Src Port: 3330,   Dest Port: 4341
      Length: 592,   Checksum: 0x0000 ERROR: CC99

Data:
    0 : C874 764E 0000 6501 45C0 0240 3DD7 0000 FE06 14EC  .tvN..e.E..@=.......
   20 : AC10 080B AC10 0709 0017 CCAE CF9A 2BDF 1358 647A  ..............+..Xdz
   40 : 5010 0FF8 A893 0000 6574 312F 320D 0A20 6E6F 2069  P.......et1/2.. no i
   60 : 7020 6164 6472 6573 730D 0A20 7368 7574 646F 776E  p address.. shutdown
   80 : 0D0A 210D 0A69 6E74 6572 6661 6365 2045 7468 6572  ..!..interface Ether
  100 : 6E65 7431 2F33 0D0A 206E 6F20 6970 2061 6464 7265  net1/3.. no ip addre
       … deleted for brevity
```

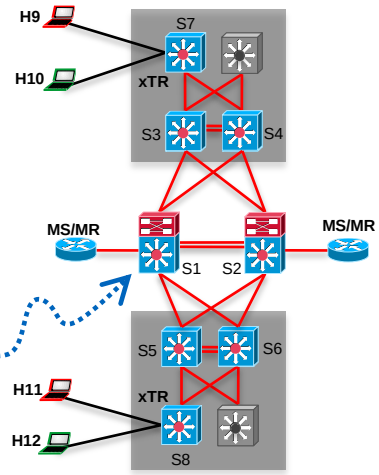

Telnet traffic from H9 to H11 captured at S1




FabricPath

What is FabricPath?

Overview

- Layer 2 routing technology
  - Eliminates spanning-tree
  - Uses IS-IS to route MAC addresses
  - Unicast – Broadcast – Multicast
  - Uses up to 16 equal-cost multipath links (ECMP)
- Fabric
  - Externally appears as a single switch
  - Internally the FabricPath protocol ties the elements together
  - Extend VLANs without limitation
- Virtualisation
  - L2 separation (VLANs)

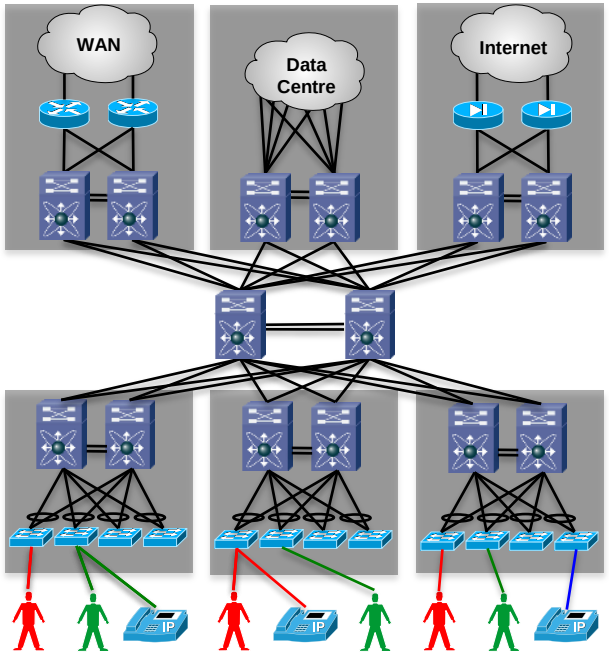

FabricPath Encapsulation

16-Byte MAC-in-MAC Header

- Switch ID – Unique number identifying each FabricPath switch
- Sub-Switch ID – Identifies devices/hosts connected via VPC+
- LID – Local ID, identifies the destination or source interface
- FTag (Forwarding tag) – Unique number identifying topology and/or distribution tree
- TTL – Decremented at each switch hop to prevent frames from looping infinitely

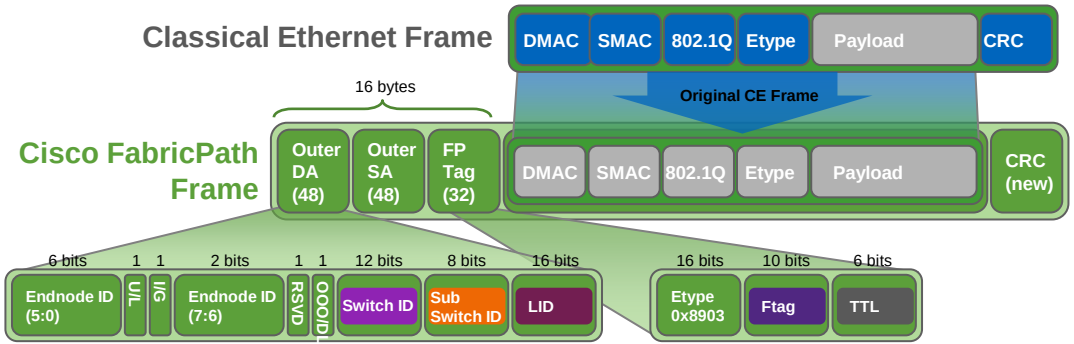

Software Defined Networking (SDN)

What is SDN?

A technology that decouples the control plane from the data plane

- Control plane (Controller)
  - Software that programs the data plane (hardware)
- Data plane
  - Hardware elements in the network (routers, switches, firewalls, etc…)

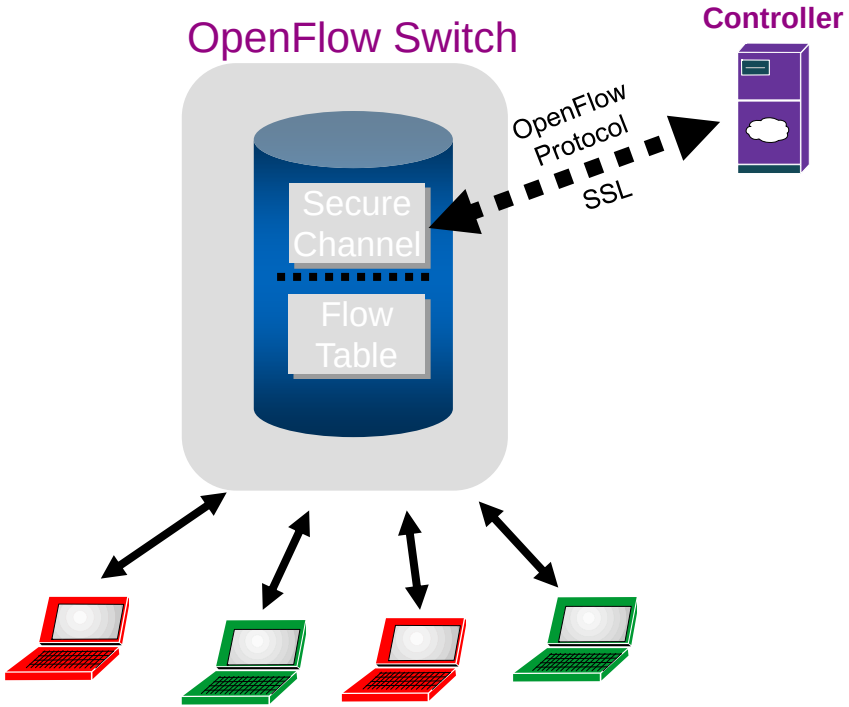

How Does SDN Provide Virtualisation?

- Slicing
  - A sandbox for a given Department/Group/Service
  - Virtual networks over a single common physical network
  - Intra-Slice management by Slice Owner (slice based management)
  - Per slice views available to Slice Owner
  - Isolation between slices
- Isolation
  - VLAN (OpenFlow v1.0)
  - MPLS (OpenFlow v1.3)

Multicast Label Distribution Protocol (mLDP)

Why is mLDP Better?

- Current MVPN Implementation
  - Supports P2MP only
  - Requires core to run PIM
  - Uses GRE to encapsulate traffic - limiting scale
  - Signalling is periodic
  - LSPs are built from head-end to tail-end

- mLDP
  - Supports P2MP and MP2MP
  - PIM is not required in the core
  - Native LDP mapping
  - No periodic signalling
  - Supports FRR through unicast P2P TE
  - LSPs are built from tail-end to head-end

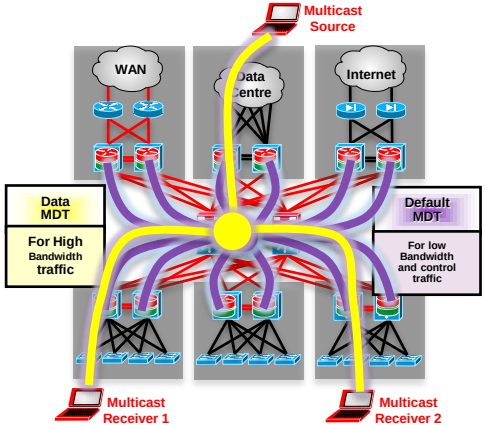

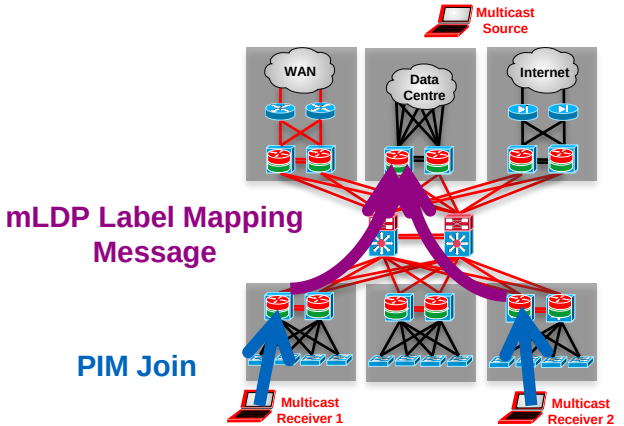

Agenda

- Virtualisation solves these Challenges
- Virtualisation Architectures
- Case Study
- Industry Trends
- Putting it all Together

Network Virtualisation

Putting It All Together

- Extending VPNs over MAN/WAN cloud
- VLANs Partition Server Farms
- virtualised Services: Firewall, ACE
- VRF-Lite + GRE, VRF-Lite End-to-End, MPLS VPN
- L3 VRFs
- Per User Role L2 VLANs
- User Identification (Static/NAC/Identity)

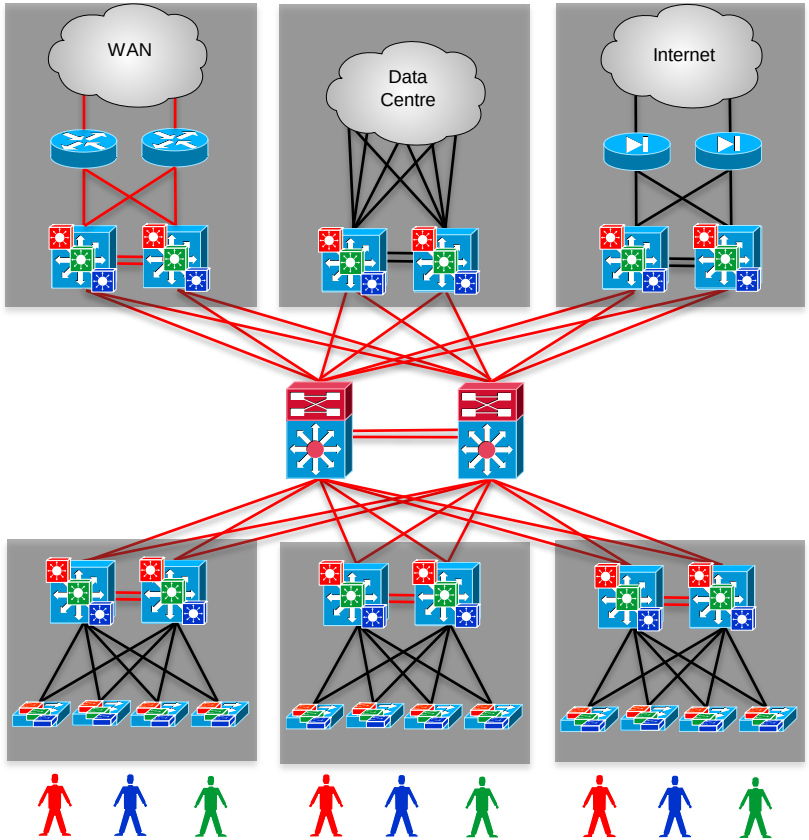

Network Virtualisation

Where to go for more information

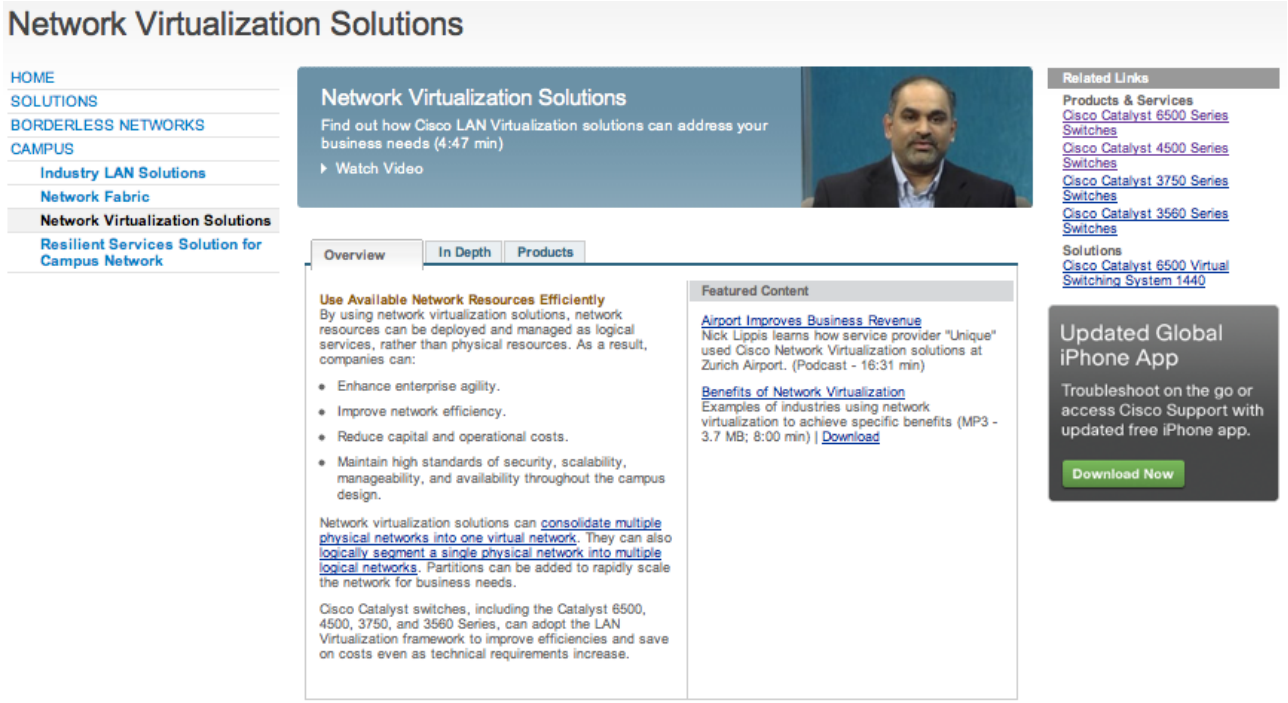

Clear Message for Virtualisation

Qld to spend $7.4 billion fixing nearly all IT systems

By Allie Coyne on Jun 11, 2013 9:53 AM
Filed under Software

IT audit report finds "systemic business risk".

The Queensland Government will need to replace ninety percent of its IT systems within five years, with the overall project to cost $7.4 billion, more than $2 billion over the initial forecast.

The state's new IT minister Ian Walker tabled the long-awaited IT Audit and the government's response to Parliament on Friday last week. The audit had been due for release last year but was held back multiple times.

The five-month audit covered 900 projects and 10,000 systems. It cost $5.2 million and required 32 public servants.

The report also made the following recommendations, which the government has agreed to:

- Cancel unused mobile and fixed telephone services, optimise data plans, consolidate telco accounts and increase printer efficiencies
- Decommission unused systems and exit its Travel Management System
- Initiate and maintain a program of rigorous application of business continuity planning for all business critical systems
- Never modify commercially-provided commodity applications to meet unique business requirements
- Conduct basic technical upgrades for high-risk payroll, finance, systems
- Further analyse the Health finance system replacement
- Establish an externally-managed desktop arrangement, and
- Study the options for a single-government data network for all agencies.




Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2014 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 March 12:00pm - 2:00pm

Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations.
www.CiscoLiveAPAC.com


CISCO
TM