

Cisco Day at the Movies

GLBP & VRF-lite



Tim Thomas
Customer Solutions Architect

Agenda

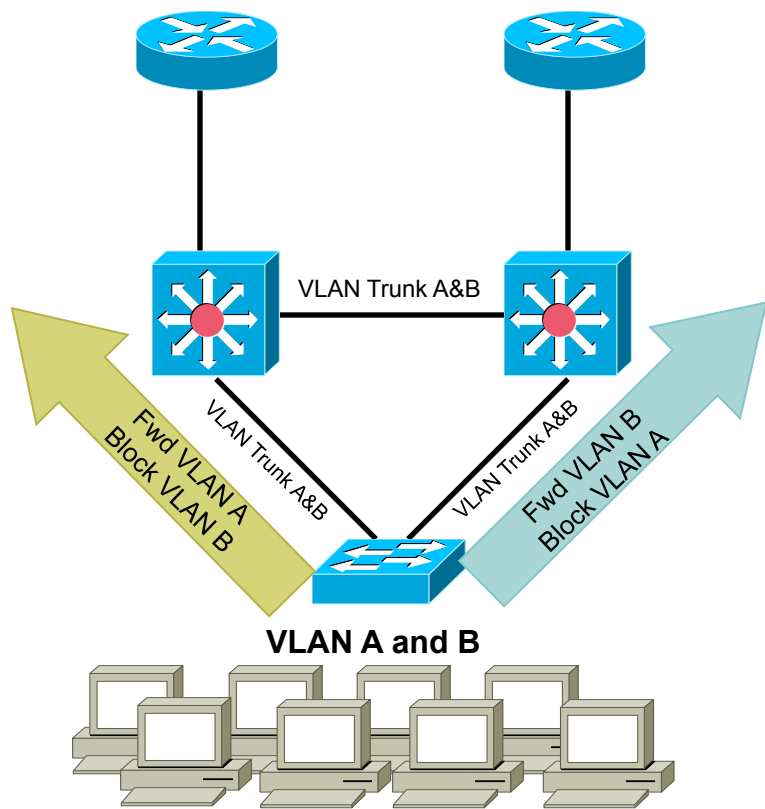
- Gateway Load-Balancing Protocol (GLBP)
- Network Virtualization with VRF-lite

First Hop Routing Protocols

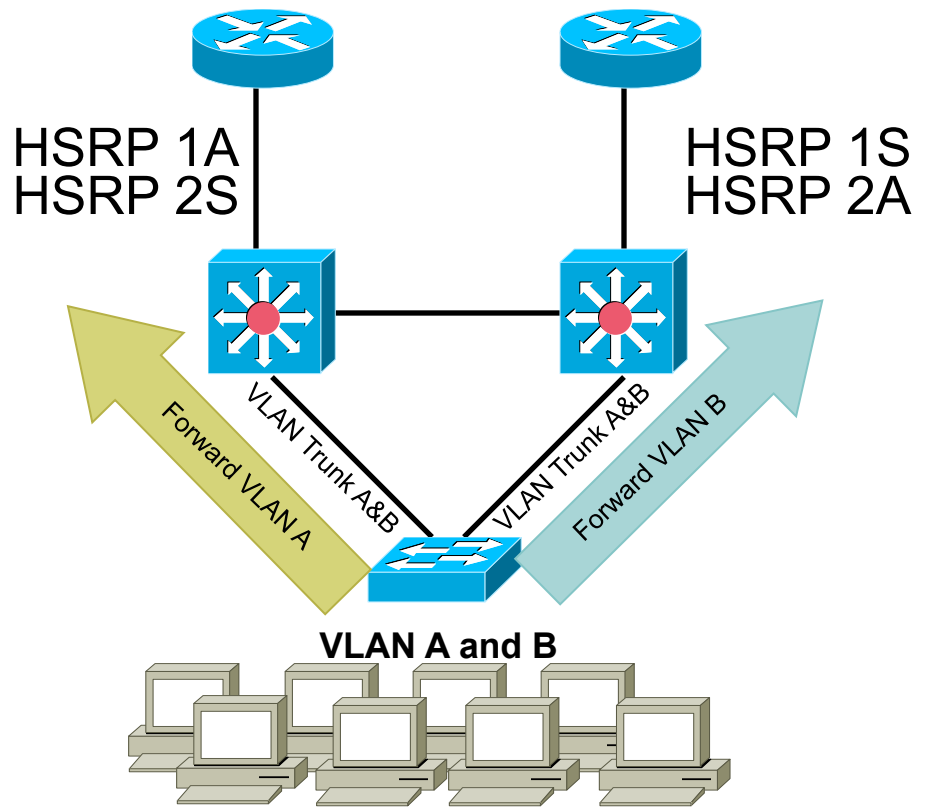
- Hot Standby Router Protocol (HSRP)
 - Cisco informational RFC 2281 (March 1998)
 - Patented: US Patent 5,473,599, December 5, 1995
- Virtual Router Redundancy Protocol (VRRP)
 - IETF Standard RFC 2338 (April 1998)
 - Now made obsolete by www.ietf.org/rfc/rfc3768.txt
- Gateway Load Balancing Protocol (GLBP)
 - Cisco innovation, load sharing, patent pending

Previous Multi-VLAN Load Balancing Methods

Layer-2 Mode Load Balancing



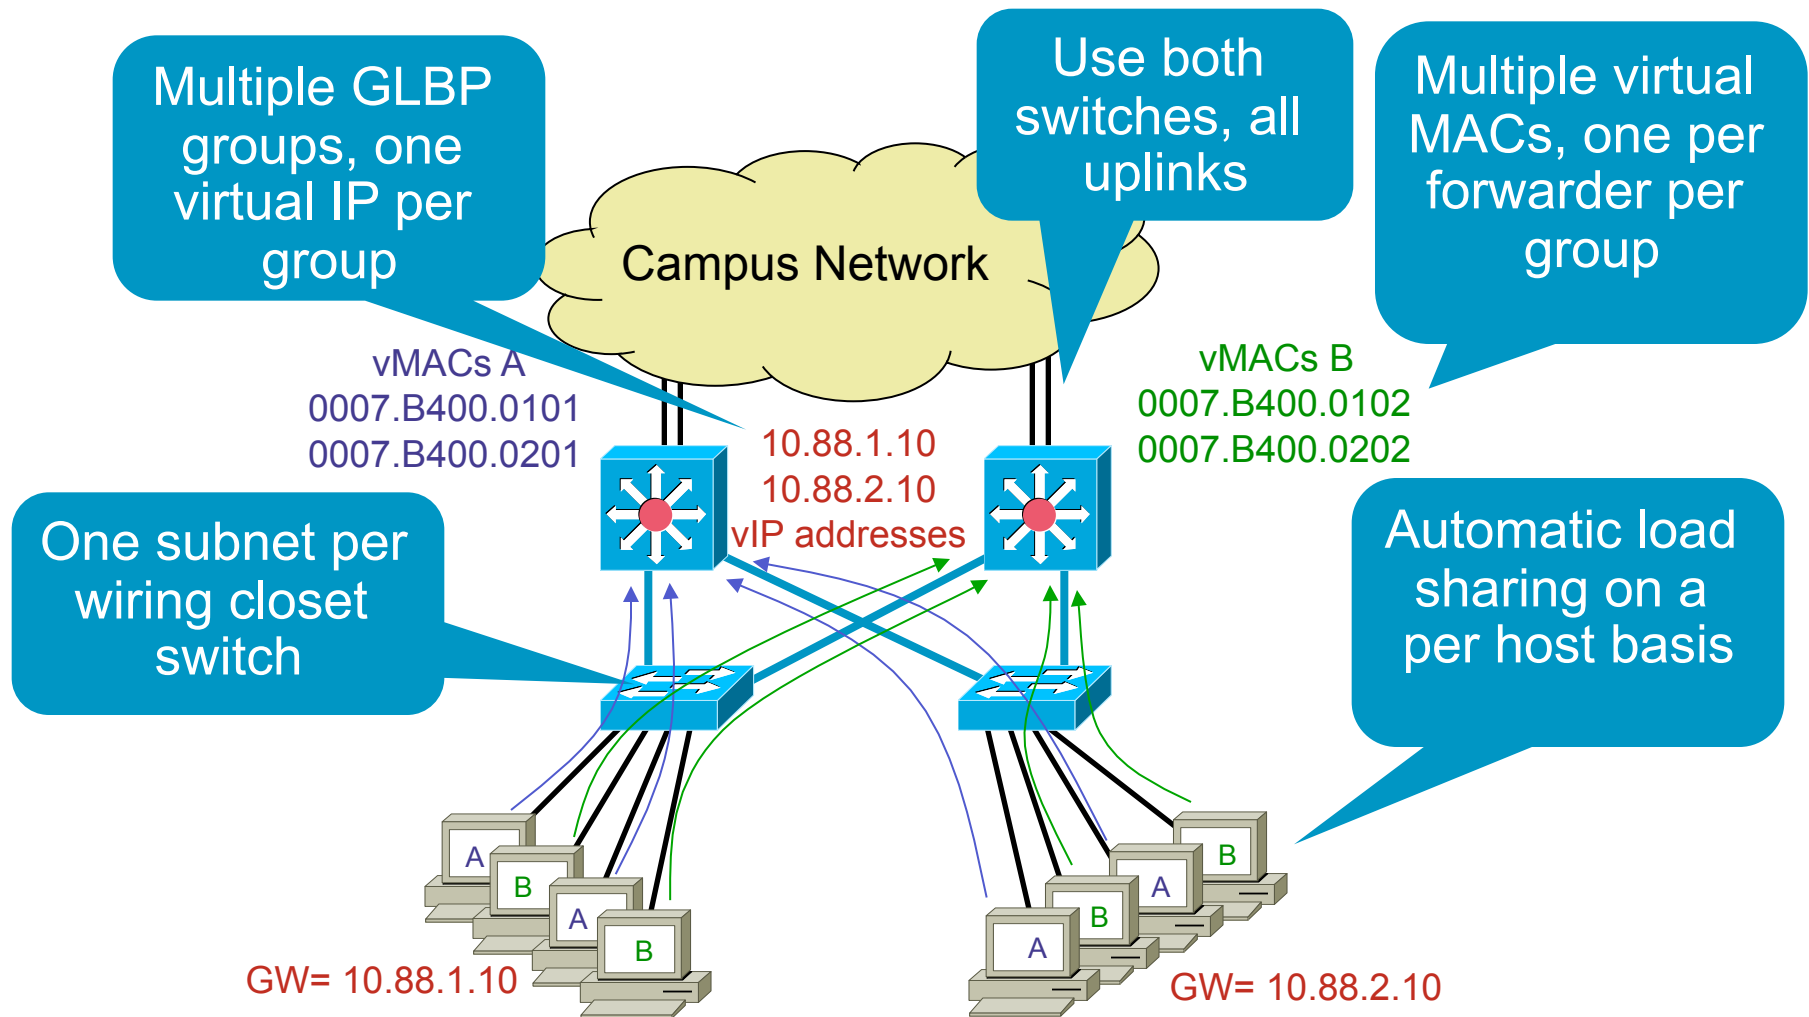
Layer-3 Mode Load Balancing



Gateway Load Balancing Protocol

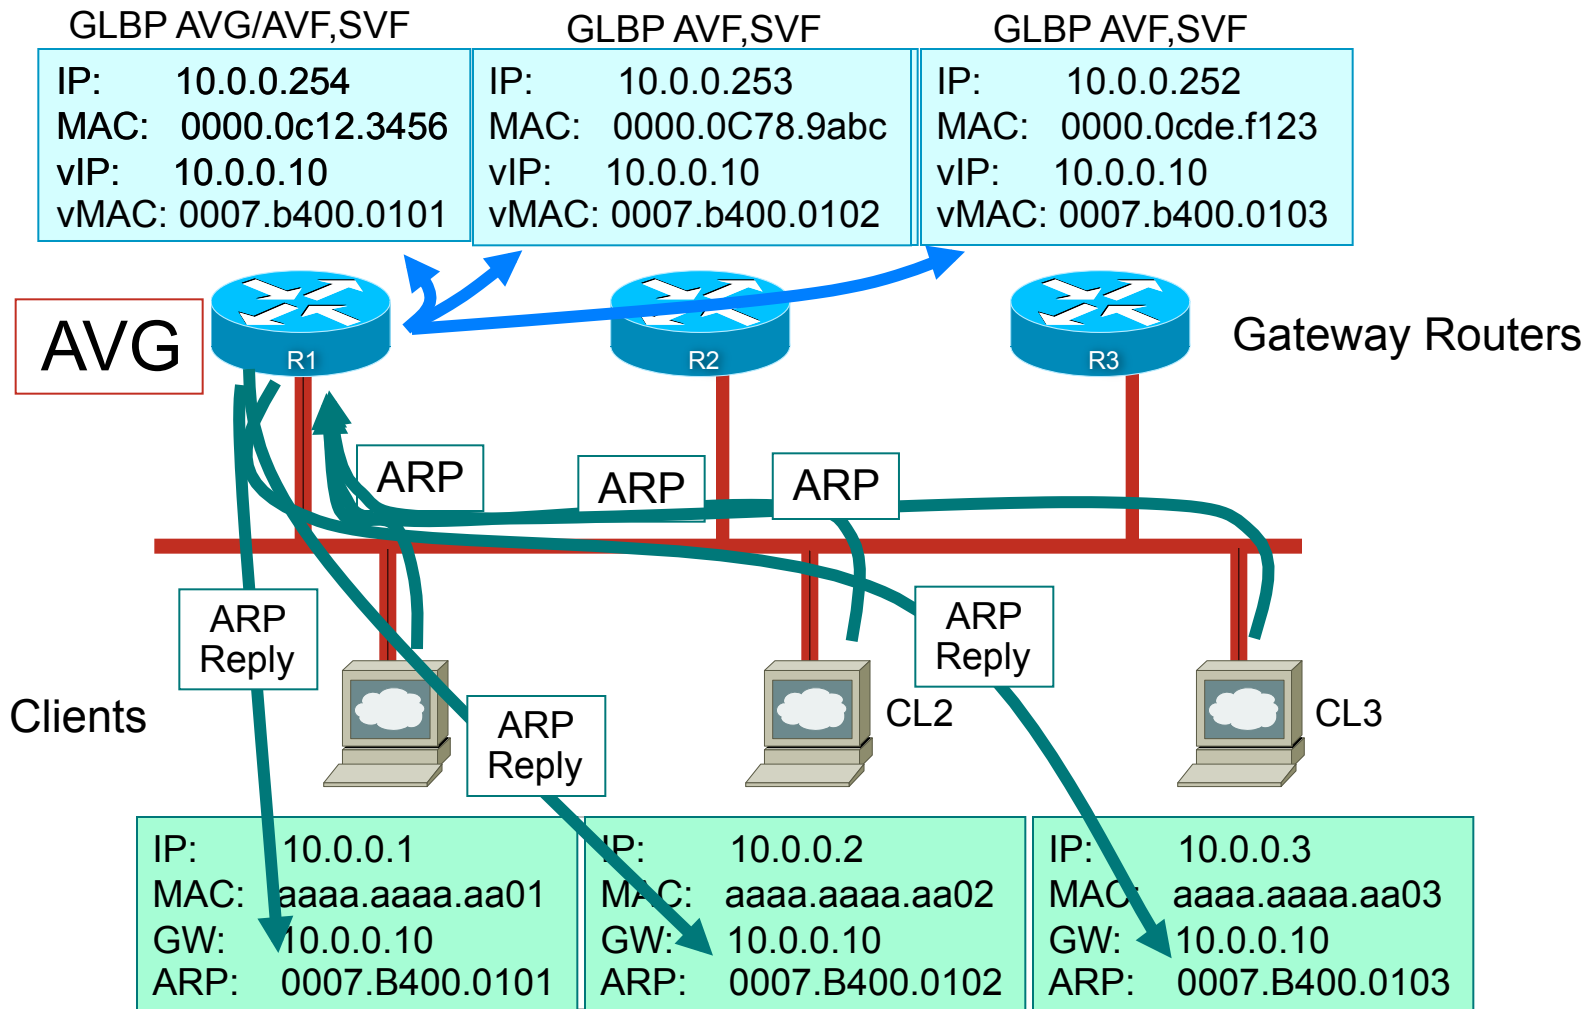
- Cisco innovation (patent pending)
- GLBP goes beyond both HSRP and VRRP
 - Previously, backup Layer-3 devices in the HSRP or VRRP group remained inactive, leaving underutilized capacity
- With GLBP, ALL L3 devices in the GLBP group actively participate in packet forwarding
 - Without allocating additional subnets
 - Without configuring multiple groups per subnet
 - Without pre-directing end stations to specific gateways (VIP addresses)
- The intelligence is in the network
 - No extra administrative burden
 - Better return on investment
 - Fully utilize resources, reduce potential for packet loss

GLBP Campus Deployment



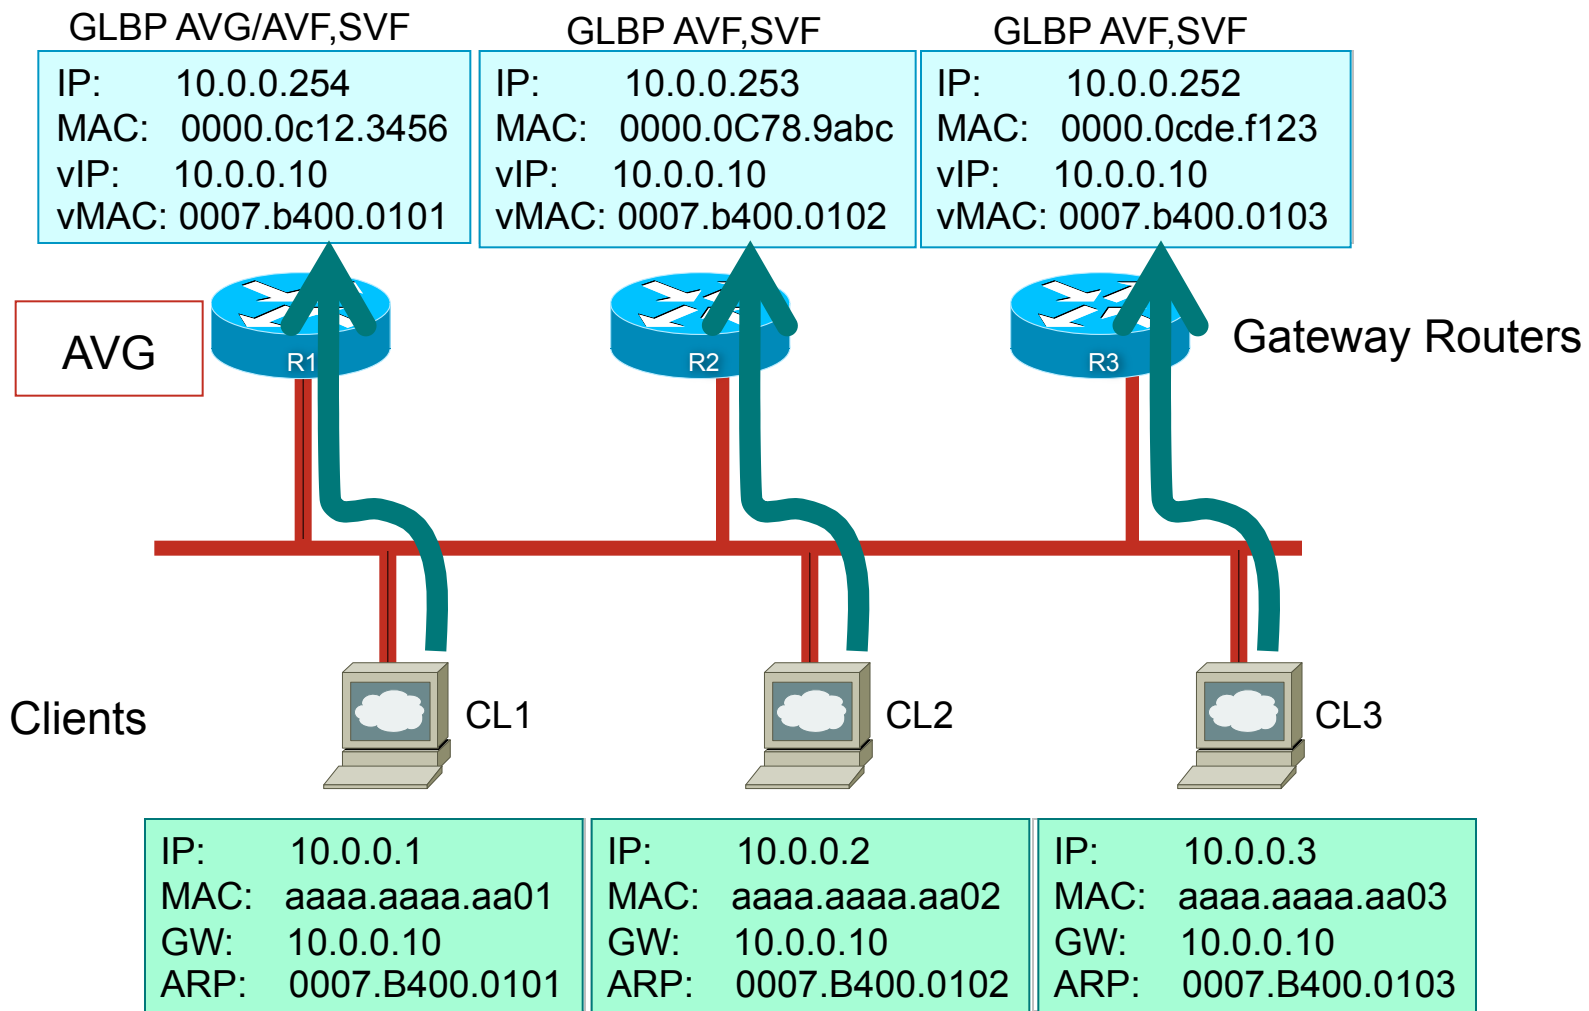
How GLBP Works

R1—AVG; R1, R2, R3 All Forward Traffic



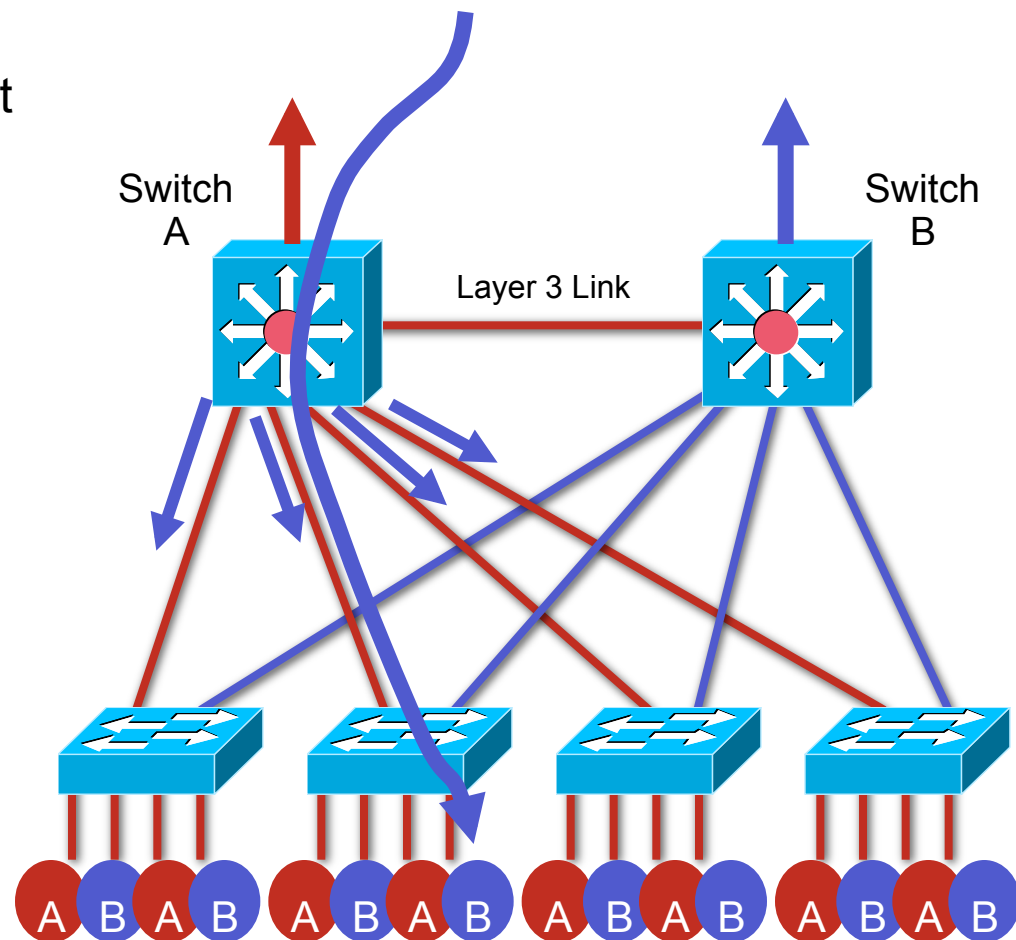
How GLBP Works

R1—AVG; R1, R2, R3 All Forward Traffic



What about Flooding?

- Traffic from 'B' devices may not be seen on switch A
- CAM aging may cause excessive flooding for asymmetric return traffic
- Mitigate by matching CAM aging timer with ARP cache timeout (default, 4 hours)
 - CAM aging > ARP cache timeout



GLBP – Protocol Details

- 'Hello' messages are exchanged between group members
 - AVG election by priority
 - vMAC distribution, learning of VF instances
- GLBP will use the following multicast destination for packets sent to all GLBP group members:
 - 224.0.0.102, UDP port 3222
- Virtual MAC addresses will be of the form:
 - 0007.b4yy.yyyy
 - where yy.yyyy equals the lower 24 bits; these bits consist of 6 zero bits, 10 bits that correspond to the GLBP group number, and 8 bits that correspond to the virtual forwarder number
 - 0007.b400.0102 : last 24 bits = 0000 0000 0000 0001 0000 0010 = GLBP group 1, forwarder 2
- Protocol allows for 1024 groups and 255 forwarders
 - Number of forwarders are capped at 4
 - Hardware restrictions limit actual number of groups and forwarders

GLBP Configuration Example

```
!  
interface GigabitEthernet2/0  
  ip address 10.88.49.1 255.255.255.0  
  duplex full  
  glbp 1 ip 10.88.49.10  
  glbp 1 priority 105  
  glbp 1 authentication text magicword  
  glbp 1 weighting 100 lower 95  
  glbp 1 weighting track 10 decrement 10  
  glbp 1 forwarder preempt delay minimum 0
```

Agenda

- Gateway Load-Balancing Protocol (GLBP)
- Network Virtualization with VRF-lite

What Is Network Virtualization?

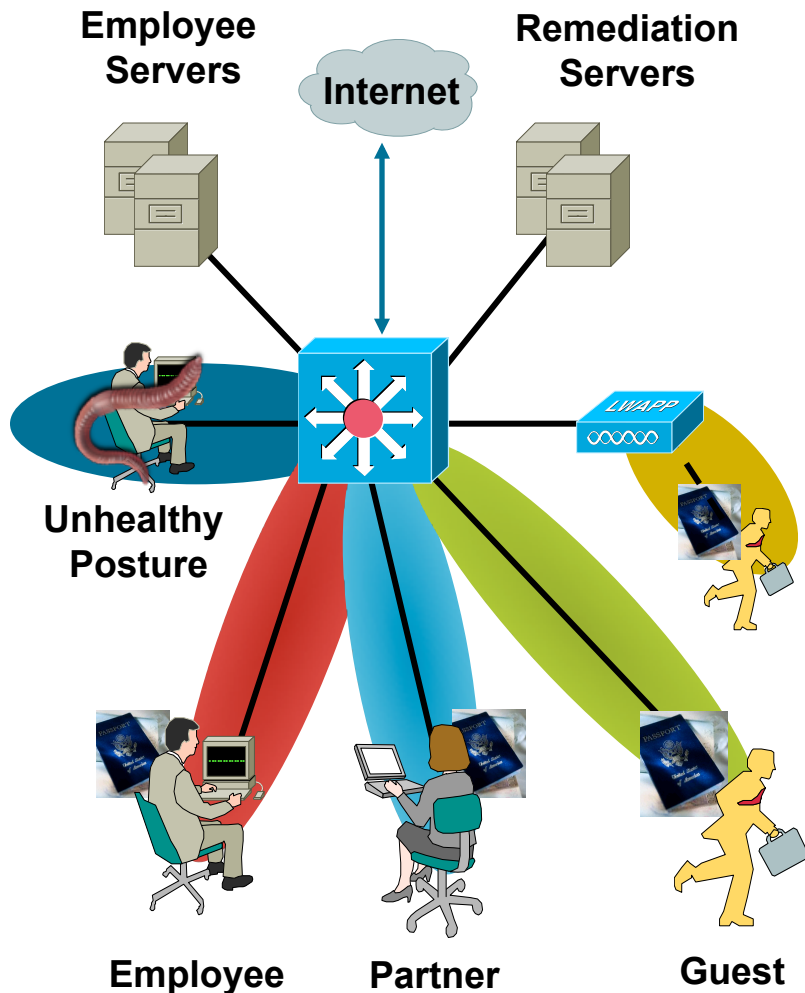
Network Virtualization Components

Deploying Network Virtualization in the Campus

Extending VRFs Across the MAN/WAN

Network Virtualization

Problem Definition



- NV provides an answer to multiple business problems

Communities of interest
NAC remediation
Regulatory compliance
...

- Closed user groups

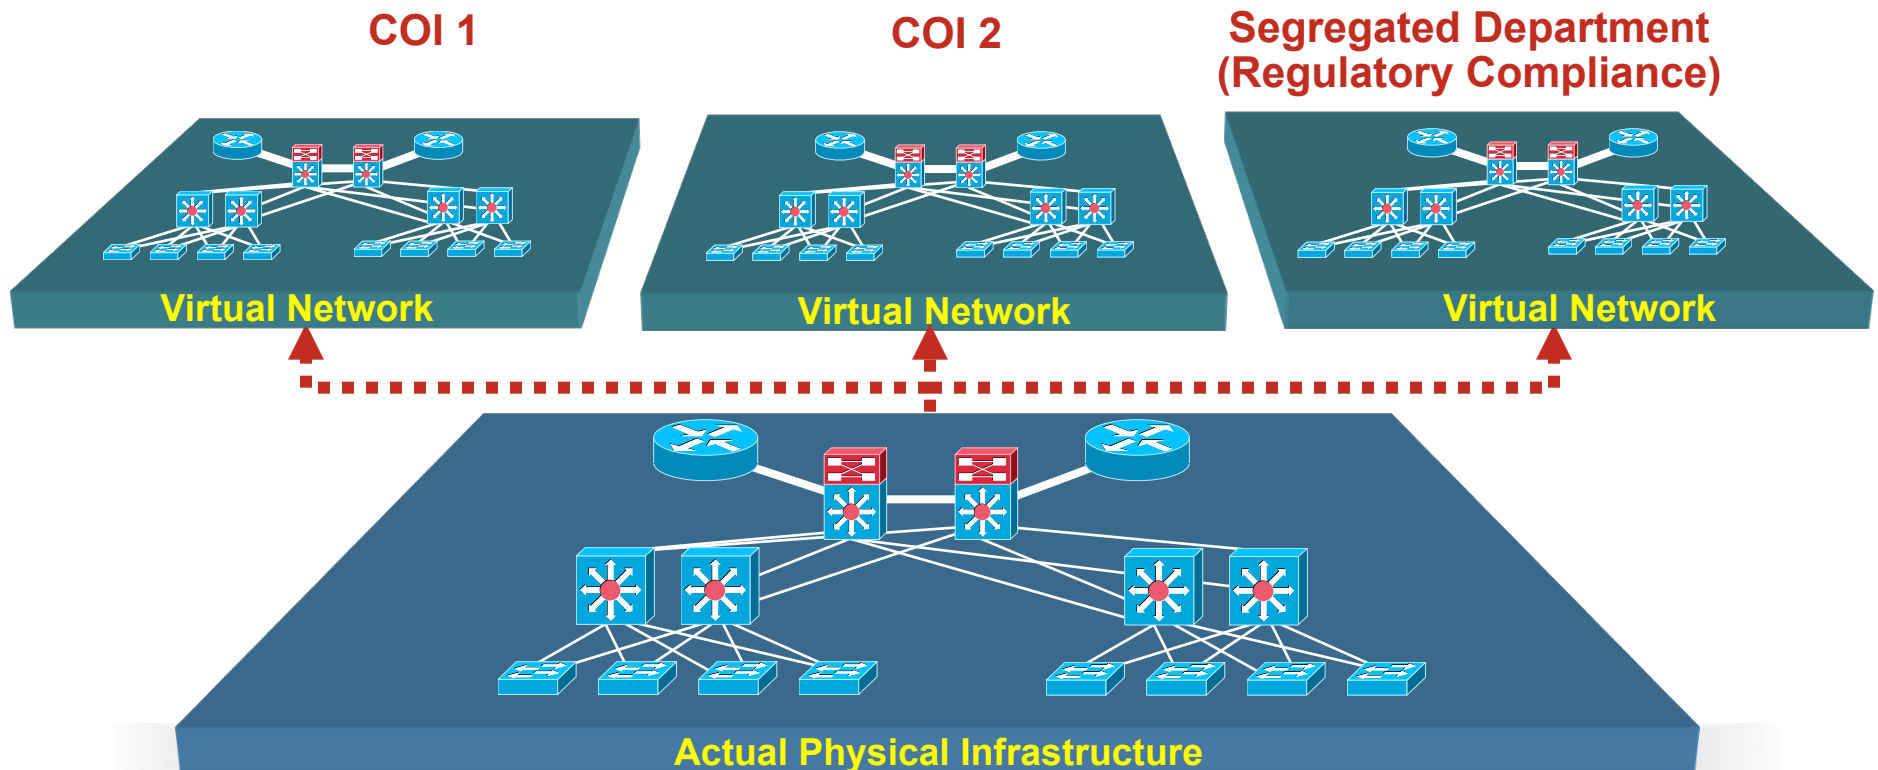
Private
Secure
Independent policies

- End-to-end shared infrastructure

Network Virtualization

Creation of Logical Partitions

- Virtualization: one-to-many (one network supports many virtual networks)
- End-user perspective is that of being connected to a dedicated network (security, independent set of policies, routing decisions...)
- Must have a rock-solid campus design in place before adding virtualization to the network



Agenda

- Gateway Load-Balancing Protocol (GLBP)
- Network Virtualization with VRF-lite

What Is Network Virtualization?

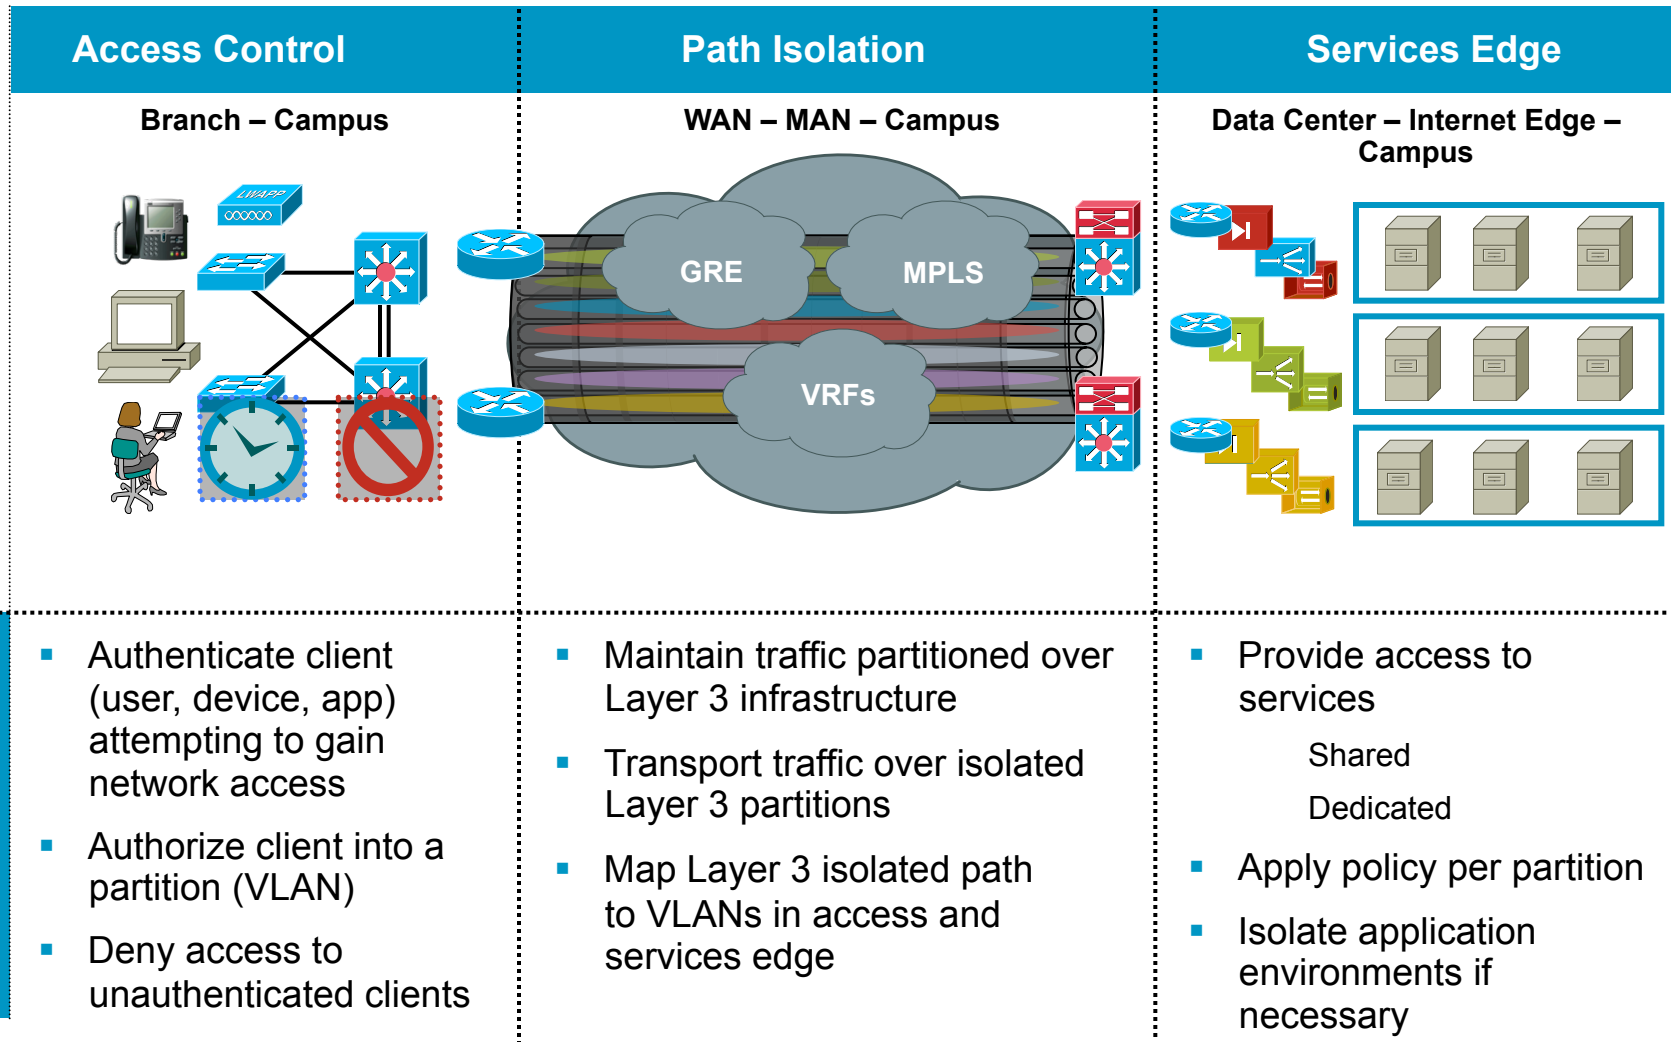
Network Virtualization Components

Deploying Network Virtualization in the Campus

Extending VRFs Across the MAN/WAN

Network Virtualization

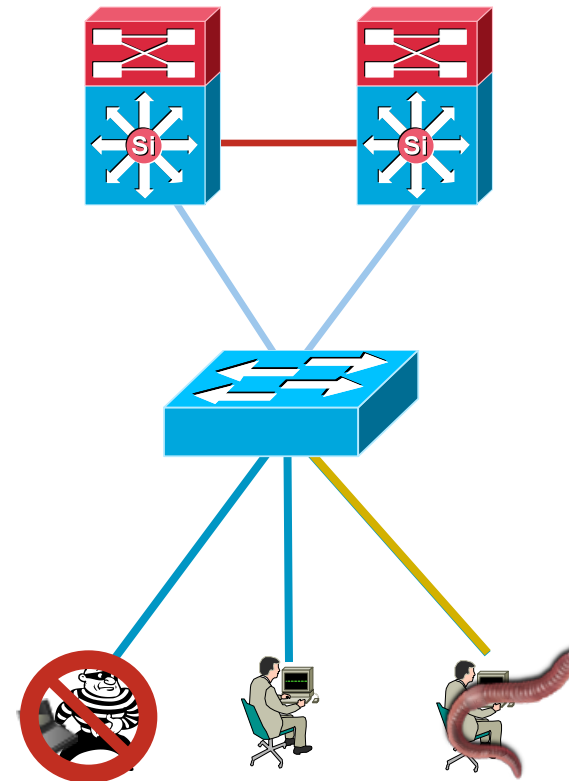
Functional Architecture



Access Control

Authentication, Authorization

- Authentication—Who/what is requesting access?
 - Holistic control—Client-based, infrastructure integrated— 802.1X
 - User-based control—Clientless— Web authentication
 - Device-specific control—MAC-address based
 - Static control—Physical security
- Authorization—Where/how is the access granted?
 - Allow access to the network from a particular VLAN




Edge Access Control


Path Isolation

Functional Components

- Device virtualization
 - Control plane virtualization
 - Data plane virtualization
 - Services virtualization



- Data path virtualization
 - Hop-by-Hop (VRF-Lite End-to-End) ↗
 - Multi-Hop (VRF-Lite+GRE, MPLS-VPN) ↘



VRF: Virtual Routing and Forwarding

Agenda

- Gateway Load-Balancing Protocol (GLBP)
- Network Virtualization with VRF-lite

What Is Network Virtualization?

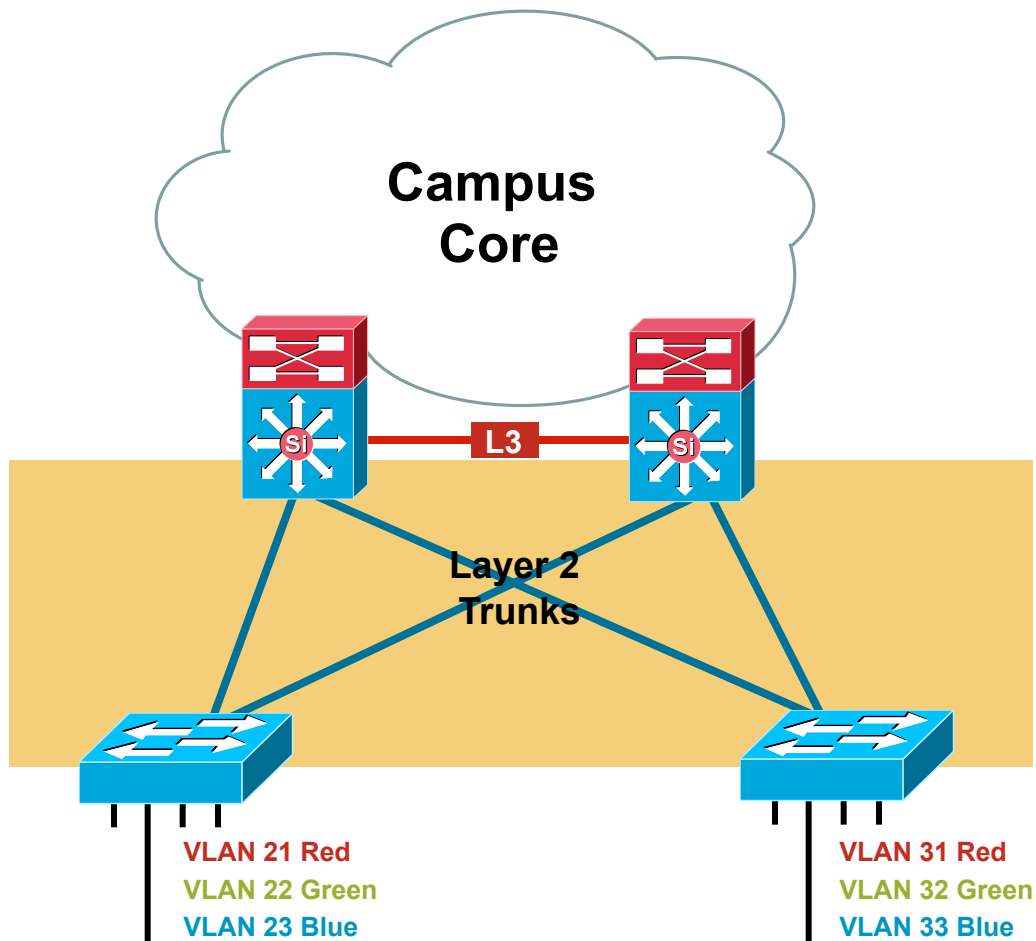
Network Virtualization Components

Deploying Network Virtualization in the Campus

Extending VRFs Across the MAN/WAN

Step 1: Definition of New VLANs

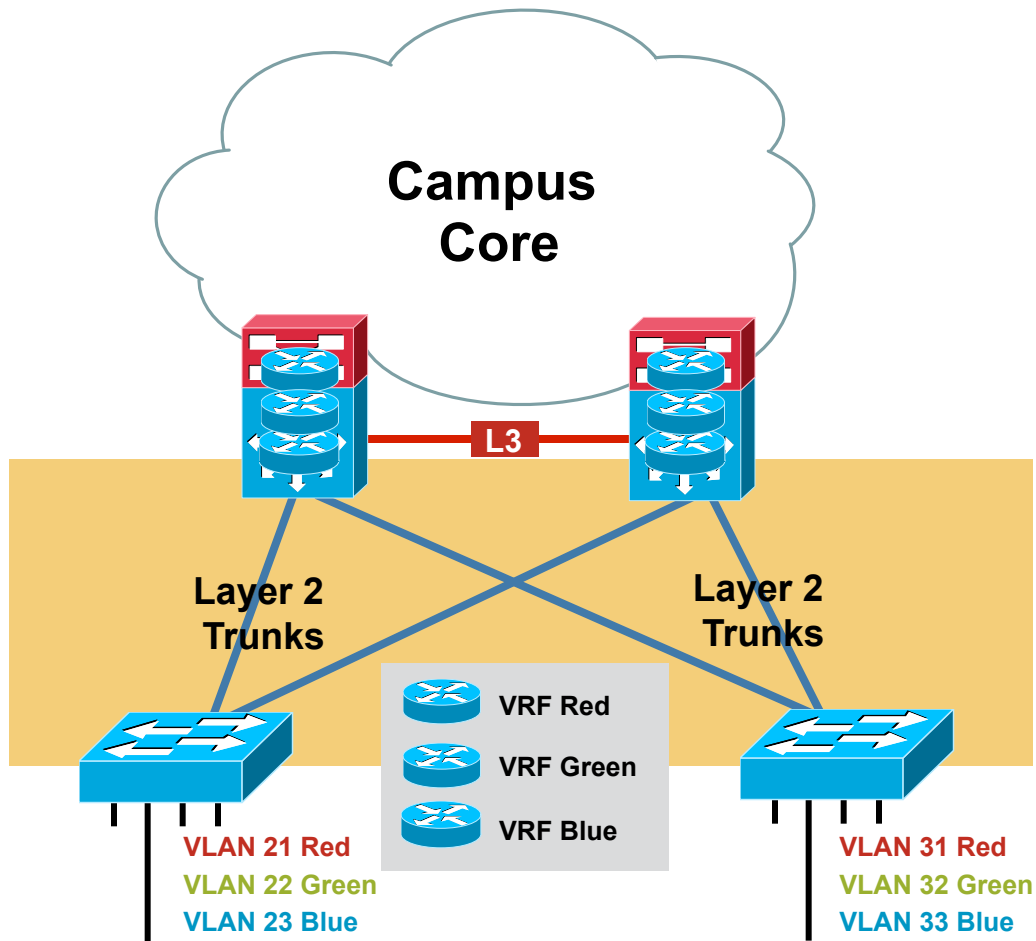
Multitier Deployment



- Campus best practice design is to keep VLAN IDs unique per access layer switch
- Total number of required VLANs is the product of the number of VRFs configured and the number of access layers switches
- Requirement to plan for new VLANs and IP subnets allocation
- Increase control plane load for protocols like STP, HSRP, etc.

Step 2: VLANs to VRF Mapping


Multitier Deployment



- Define VRFs on the distribution layer devices (first L3 hop in a campus multitier design)
- One VRF dedicated to each virtual network (“Red”, “Green”, etc.)
- Multiple VLANs defined at the access layer map to the same VRF
 - “Red” VLANs (21, 31) are mapped to the same “Red” VRF
- The chosen path Isolation technique is deployed from the distribution layer toward the routed core

Step 1: Definition of New VLANs

Routed Access Deployment




- Move the boundaries between L2 and L3 domains down to the access layer
- Same VLAN IDs can be used on each access layer switch
- Requirement to plan for new IP subnets allocation
- No increase on control plane load

No need for HSRP/GLBP/VRRP or STP between access and distribution layer devices

Step 2: VLANs to VRF Mapping

Routed Access Deployment



- Define VRFs on the access layer devices (first L3 hops in a campus routed access design)
- One VRF dedicated to each virtual network (“Red”, “Green”, etc.)
- Each VLAN defined at the Access Layer maps to the corresponding VRF
 - “Red” VLANs (21, 31) are mapped to the same “Red” VRF defined in the different switches
- The chosen path isolation technique must be deployed from the access layer devices

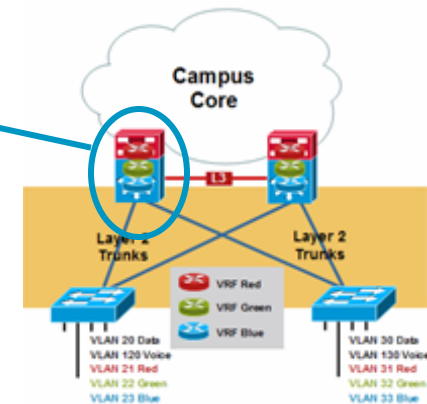
Example CLI

VLANs to VRF Mapping Configuration

```
ip vrf Red
  rd 1:1
!
ip vrf Green
  rd 2:2
!
vlan 21
  name Red_access_switch_1
!
vlan 22
  name Green_access_switch_1
!
interface Vlan21
  description Red on Access Switch 1
  ip vrf forwarding Red
  ip address 10.137.21.1 255.255.255.0
!
interface Vlan22
  description Green on Access Switch 1
  ip vrf forwarding Green
  ip address 10.137.22.1 255.255.255.0
```

Defining the VRFs


Defining the VLANs
(L2 and SVI) and Mapping
Them to the VRFs



VRF-Lite End-to-End

How Does It Work?

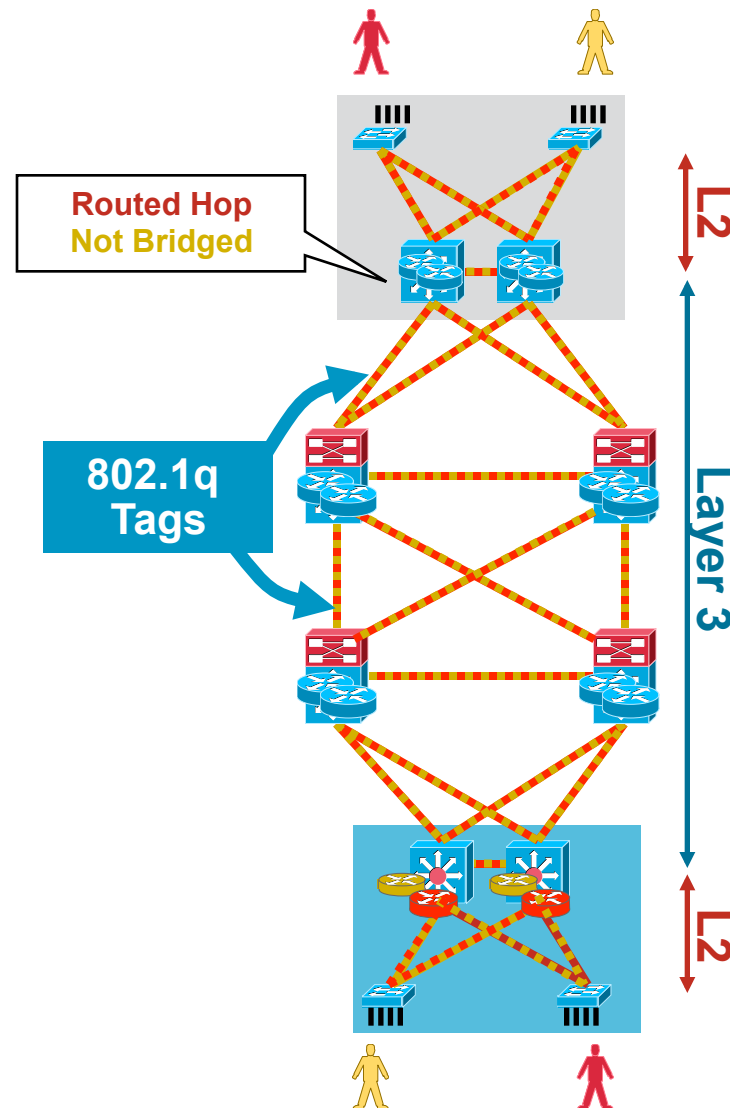
1. Create L2 VLANs and trunk them to the first L3 device
2. Define VRFs at the first L3 device and map the L2 VLANs to the proper VRF
3. Define VRFs on all the other L3 devices in the network
4. Configure as trunks all the physical links connecting the L3 devices in the network
Create VLAN interfaces or subinterfaces and map them to the corresponding VRF
5. Define unique VLANs on each trunk to be associated to each VRF
6. Enable a routing protocol in each VRF
7. Traffic is now carried end-to-end across the network maintaining logical isolation between the defined groups



VRF-Lite End-to-End

General Design Considerations

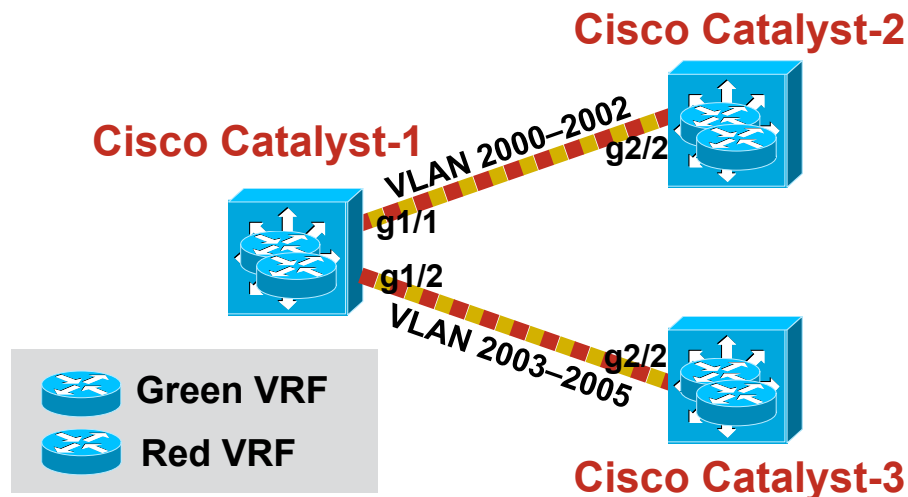
- VRF-lite **on all routed hops**: core and distribution (sometimes access)
 - VLANs are **not** extended across the Campus network
- Every physical link is virtualized to carry multiple logical routed links
 - 802.1q tags provide single hop data path virtualization
- These virtualized links **do not** extend VLANs throughout the campus
- The relationship of physical to logical networks is a matter of replication
 - Virtualization of every network device and every physical link connecting them



VRF-Lite End-to-End

Trunk with Switchports and SVIs

- Links between L3 devices defined as L2 trunks with switchports
- Unique VLANs used for global table, Green and Red traffic
- Logical SVIs mapped to the Green and Red VRFs



SVI: Switched Virtual Interface

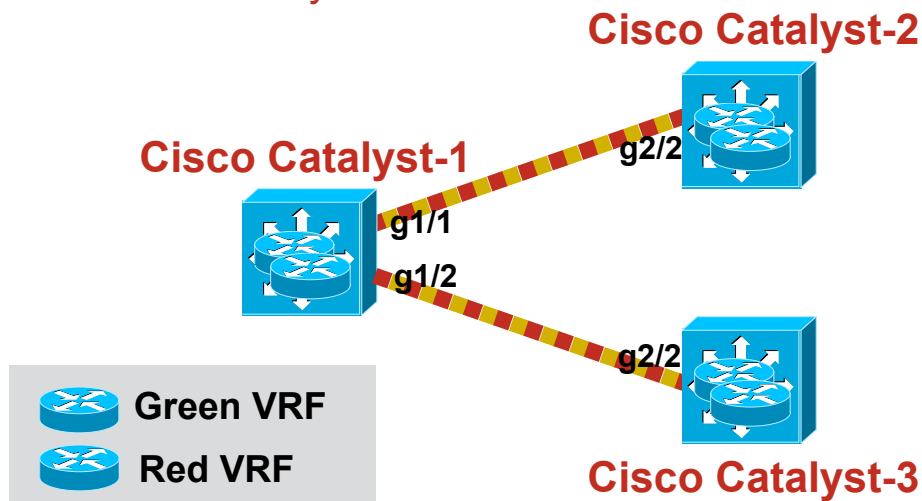
```
Catalyst-1
interface GigabitEthernet1/1
description --- Trunk to Catalyst-2 ---
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 2000-2002
switchport mode trunk
spanning-tree portfast trunk
!
interface Vlan2000
description --- Global table ---
ip address 10.1.1.1 255.255.255.252
!
interface Vlan2001
description --- Green VPN ---
ip vrf forwarding Green
ip address 11.1.1.1 255.255.255.252
!
interface Vlan2002
description --- Red VPN ---
ip vrf forwarding Red
ip address 12.1.1.1 255.255.255.252
```

```
Catalyst-2
interface GigabitEthernet2/2
description --- Trunk to Catalyst-1 ---
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 2000-2002
switchport mode trunk
spanning-tree portfast trunk
!
interface Vlan2000
description --- Global table ---
ip address 10.1.1.2 255.255.255.252
!
interface Vlan2001
description --- Green VPN ---
ip vrf forwarding Green
ip address 11.11.1.2 255.255.255.252
!
interface Vlan2002
description --- Red VPN ---
ip vrf forwarding Red
ip address 12.1.1.2 255.255.255.252
```

VRF-Lite End-to-End

Trunk with Routed Ports

- Links between L3 devices defined as routed port with subinterfaces
- Global table traffic is sent untagged
- Each additional subinterface associated to a unique VLAN and mapped to a separate VRF
- Easier migration: configuration on main interface (used for global traffic) remains unchanged
- Currently supported on Cisco Catalyst 6500 Series only



```
Catalyst-1
interface GigabitEthernet1/1
description --- Global table ---
ip address 10.1.1.1 255.255.255.252
!
interface GigabitEthernet1/1.2001
description --- Green VPN ---
encapsulation dot1q 2001
ip vrf forwarding Green
ip address 11.11.1.1 255.255.255.252
!
interface GigabitEthernet1/1.2002
description --- Red VPN ---
encapsulation dot1q 2002
ip vrf forwarding Red
ip address 12.1.1.1 255.255.255.252
```

```
Catalyst-2
interface GigabitEthernet2/2
description --- Global table ---
ip address 10.1.1.2 255.255.255.252
!
interface GigabitEthernet2/2.2001
description --- Green VPN ---
encapsulation dot1q 2001
ip vrf forwarding Green
ip address 11.1.1.2 255.255.255.252
!
interface GigabitEthernet1/1.2002
description --- Red VPN ---
encapsulation dot1q 2002
ip vrf forwarding Red
ip address 12.1.1.2 255.255.255.252
```

VRF-Lite End-to-End

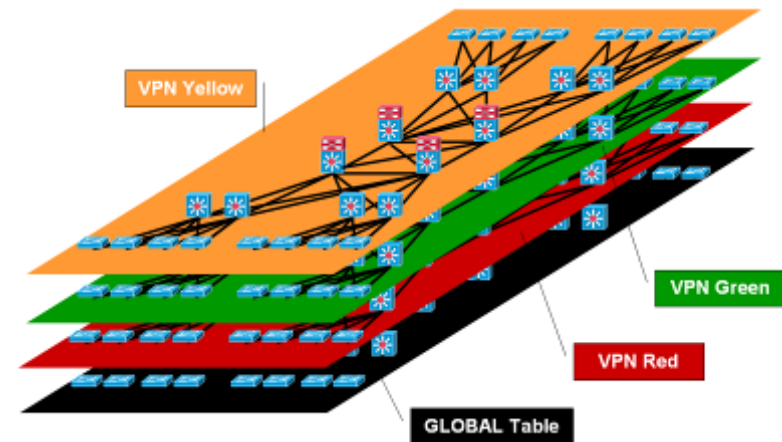
Virtualizing the Routing Protocol

- Recommendation is to use in each VRF the same routing protocol already leveraged in global table (usually EIGRP or OSPF)
- Routing design principles adopted in global table can simply be replicated in each virtual network

Summarization boundaries

IGP timer tuning

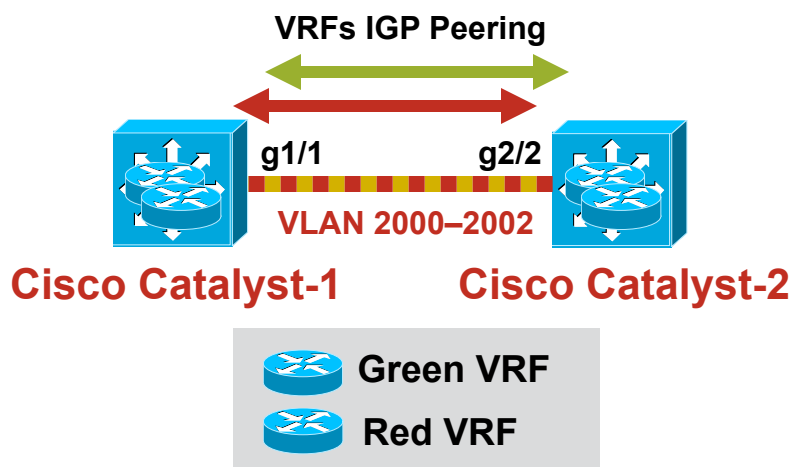
Areas definition for OSPF



VRF-Lite End-to-End

Virtual Routing Processes

- Each VRF instance needs a separate IGP process (OSPF) or address family (EIGRP, RIPv2)
 - Enabled on all L3 devices
- Devices peer over separate routing instances



```
router ospf 1
 network 10.0.0.0 0.255.255.255 area 0
 passive-interface default
 no passive-interface vlan 2000
 !
router ospf 100 vrf Green
 network 11.0.0.0 0.255.255.255 area 0
 no passive-interface vlan 2001
 !
router ospf 200 vrf Red
 network 12.0.0.0 0.255.255.255 area 0
 no passive-interface vlan 2002
```

```
router eigrp 100
 network 10.0.0.0 0.255.255.255
 passive-interface default
 no passive-interface vlan 2000
 no auto-summary
 !
address-family ipv4 vrf Green
 network 11.0.0.0 0.255.255.255
 no auto-summary
 exit-address-family
 !
address-family ipv4 vrf Red
 network 12.0.0.0 0.255.255.255
 no auto-summary
 exit-address-family
```

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
- Enterprise scale
- Supported on Catalyst 6500, 4500, 3700 families
- Supported on Nexus 7000

Application and Services


- Supports both wired and wireless networks
- Multiple VRF-aware Services available

Learning Curve

- Familiar routing protocols can be used
- IP Alternative to MPLS

Management

- Virtual Network Management (VNM) available with LMS 3.2 (Summer 2009)
- Provisioning, Troubleshooting and monitoring for VRF network




Agenda

- What Is Network Virtualization?
- Network Virtualization Components
- Deploying Network Virtualization in the Campus
- Extending VRFs Across the MAN/WAN

Extensibility over the MAN/WAN

Groups Must Be Extensible Over:

- The private MAN/WAN
- The Internet



MAN/WAN Extensibility

Different Options Available

- The virtual networks may need to be extended over the MAN/WAN
- There are several technical alternatives; some examples
 - MPLS over L2 service
 - DMVPN per VRF
 - RFC2547 over DMVPN
 - Carrier-supporting-carrier (where the service is available)
- The choice depends largely on the enterprise's MAN/WAN contracts and platform support
- **Next-generation MPLS VPN MAN/WAN design guide**
http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidances_list.html#anchor13

Trivia

Q and A

1) Question: In GLBP, which component answers ARPs from hosts – the AVF or the AVG?

Answer: The Active Virtual Gateway (AVG)

2) Question: What does the acronym VRF stand for?

Answer: Virtual Routing and Forwarding



