
TOMORROW starts here.
cisco
Cisco live!

Hot Topics and Capabilities for the Campus in 2014 and Beyond

BRKCRS-2663

Glenn Fullager

Systems Engineer

Cisco live!

Build a Foundation, th

Without it you're Asking for Trouble

Agenda

- The Principles
- The Basics
- The Cool Stuff
- The End

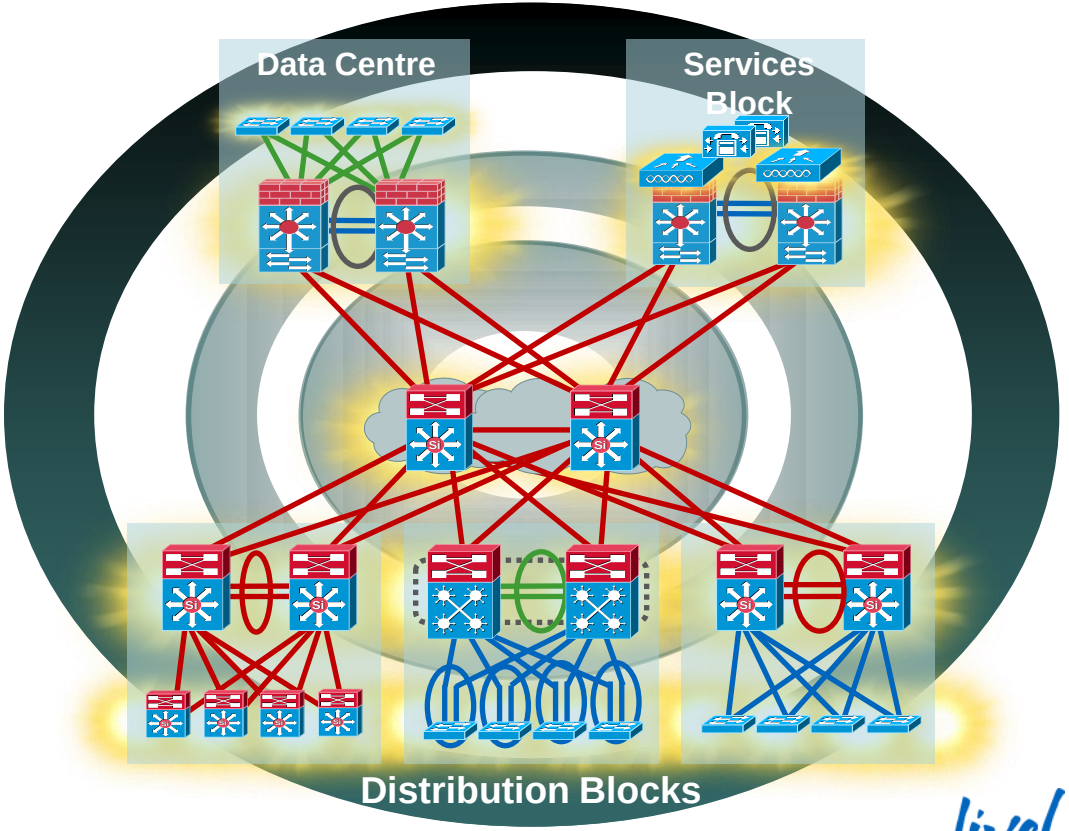

High-Availability Campus Design

Structure, Modularity, and Hierarchy

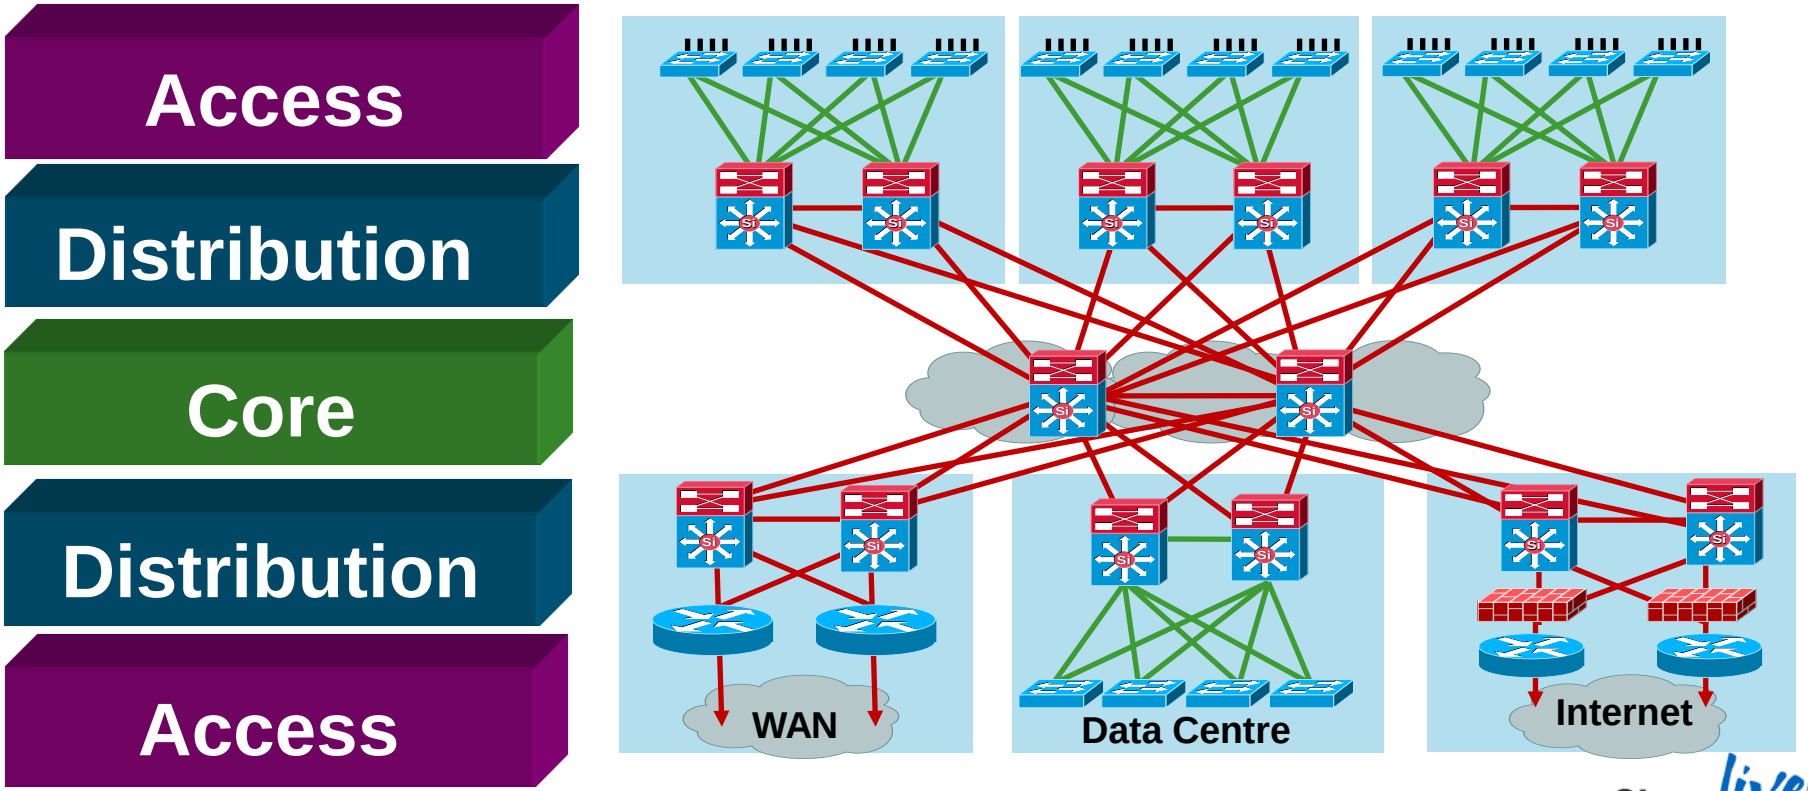

Hierarchical Network Design

Without a Rock Solid Foundation the Rest Doesn't Matter

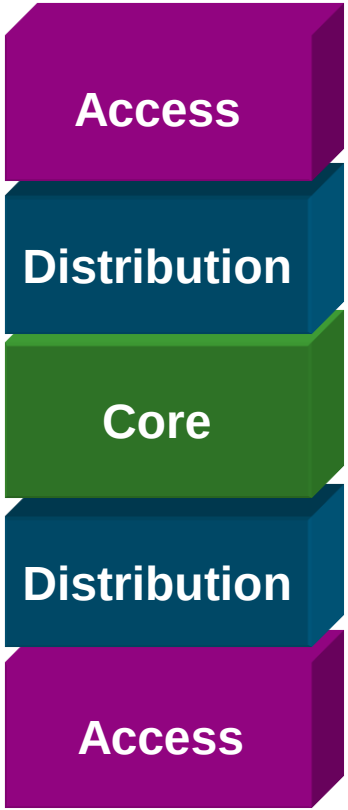

- Offers hierarchy—each layer has specific role
- Modular topology—building blocks
- Easy to grow, understand, and troubleshoot
- Creates small fault domains—clear demarcations and isolation
- Allows for implementation of new technologies per building block

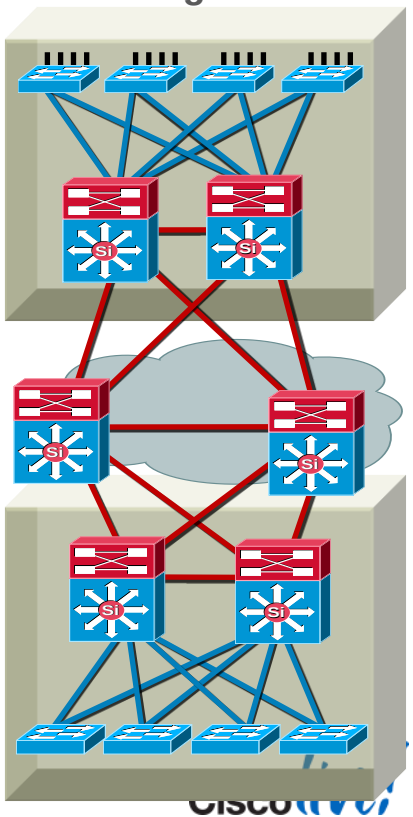

Hierarchical Campus Design

Structure, Modularity, and Hierarchy

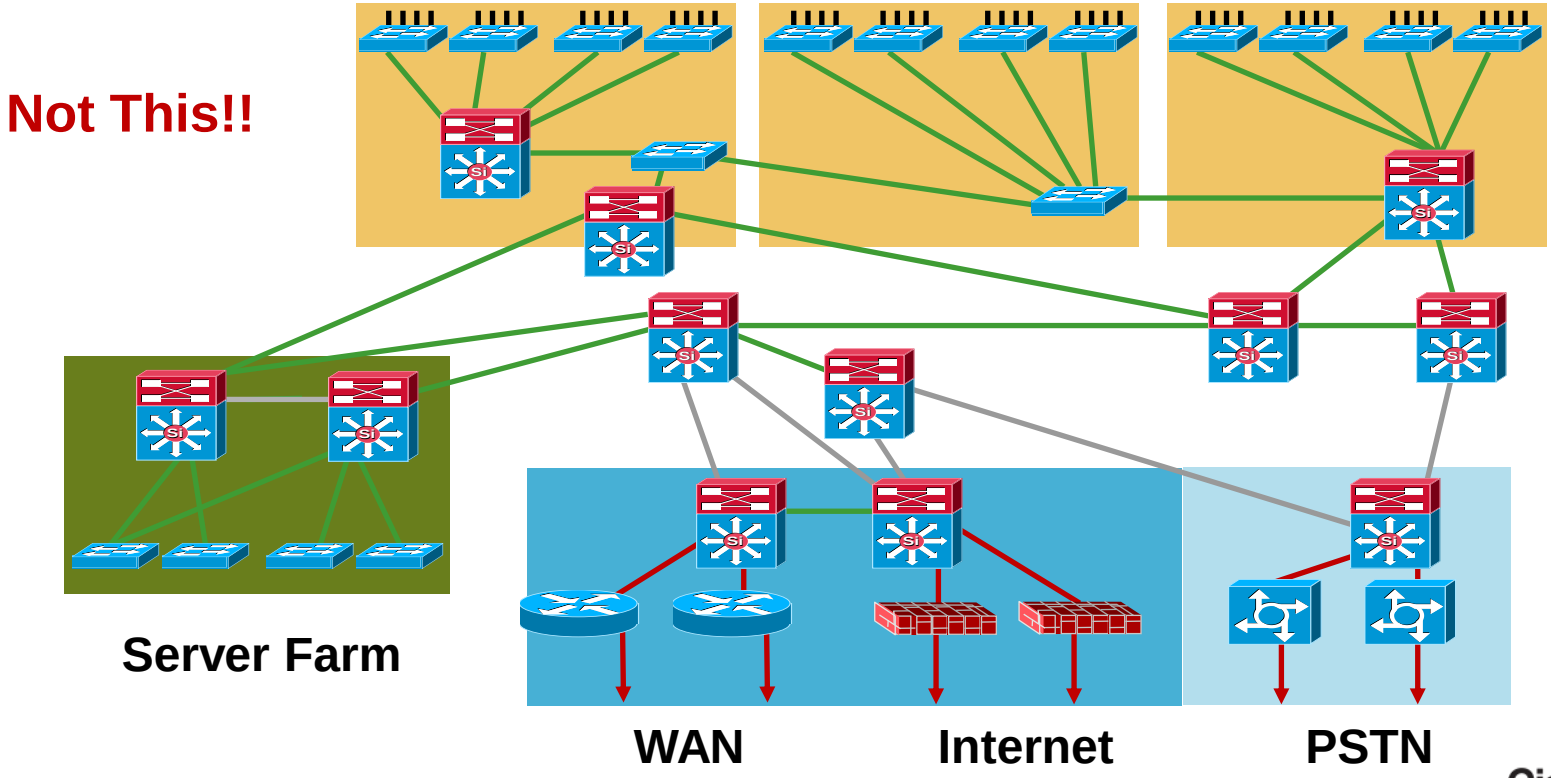

Access Layer
Feature Rich Environment

- It's not just about connectivity
- Layer 2/Layer 3 feature rich environment: convergence, HA, security, multicast
- Intelligent network services: QoS, trust boundary, broadcast suppression, IGMP snooping
- Intelligent network services: PVST+, Rapid PVST+, EIGRP, OSPF, DTP, PAgP/LACP, UDLD, FlexLink, etc.
- Cisco Catalyst® integrated security features IBNS (802.1x), (CISF): port security, DHCP snooping, DAI, IPSG, etc.
- Automatic phone discovery, conditional trust boundary, PoE, auxiliary VLAN, etc.
- Spanning tree toolkit: PortFast, UplinkFast, BackboneFast, LoopGuard, BPDU Guard, BPDU Filter, RootGuard, etc.

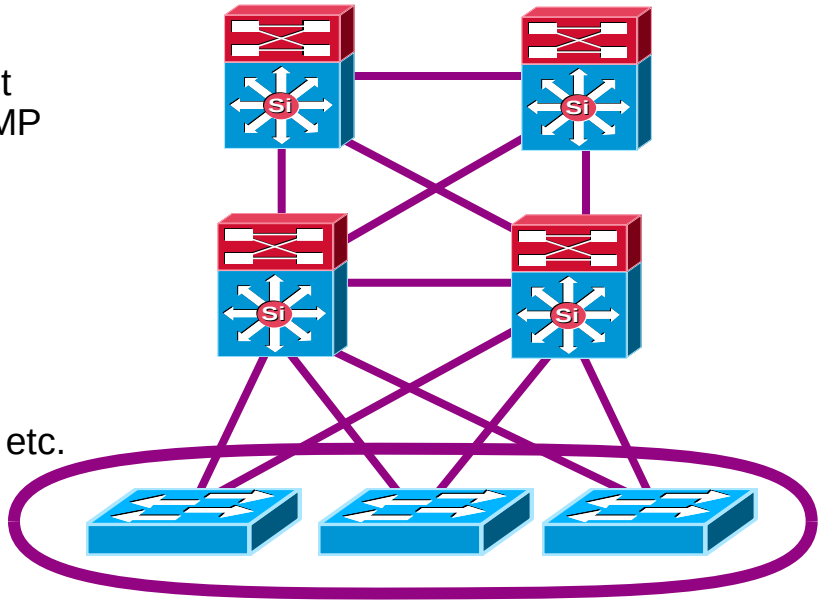

Distribution Layer

Policy, Convergence, QoS and High Availability

- Availability, load balancing, QoS and provisioning are the important considerations at this layer
- Aggregates wiring closets (access layer) and uplinks to core
- Protects core from high density peering and problems in access layer
- Route summarisation, fast convergence,
  redundant path load sharing
- HSRP or GLBP to provide first hop redundancy

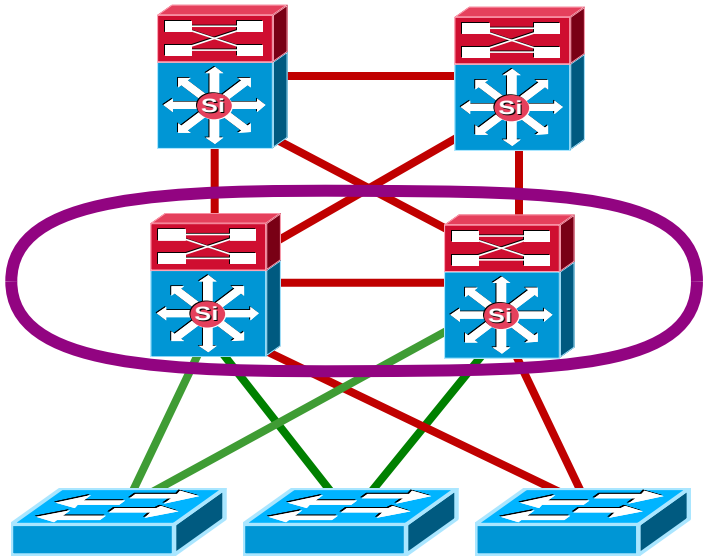

Core Layer

Scalability, High Availability and Fast Convergence

- Backbone for the network—connects network building blocks
- Performance and stability vs. complexity— less is more in the core
- Aggregation point for distribution layer
- Separate core layer helps in scalability during future growth
- Keep the design technology-independent

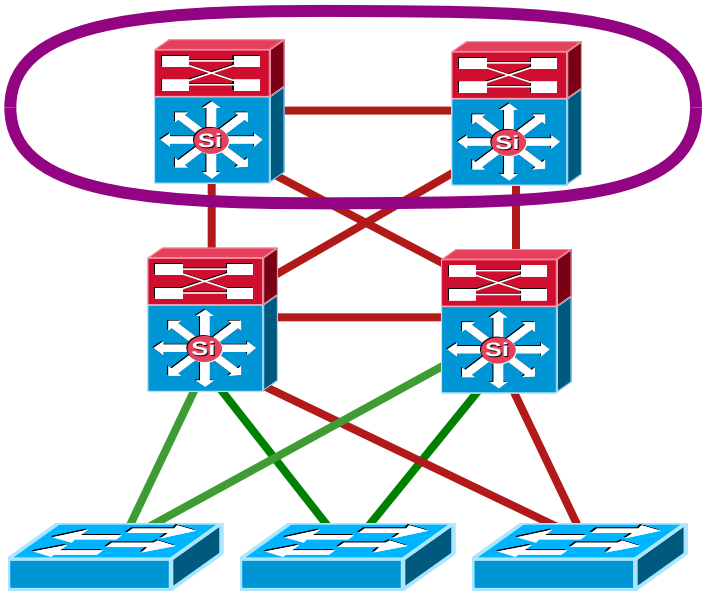

Agenda

- The Principles
- The Basics
- The Cool Stuff
- The End

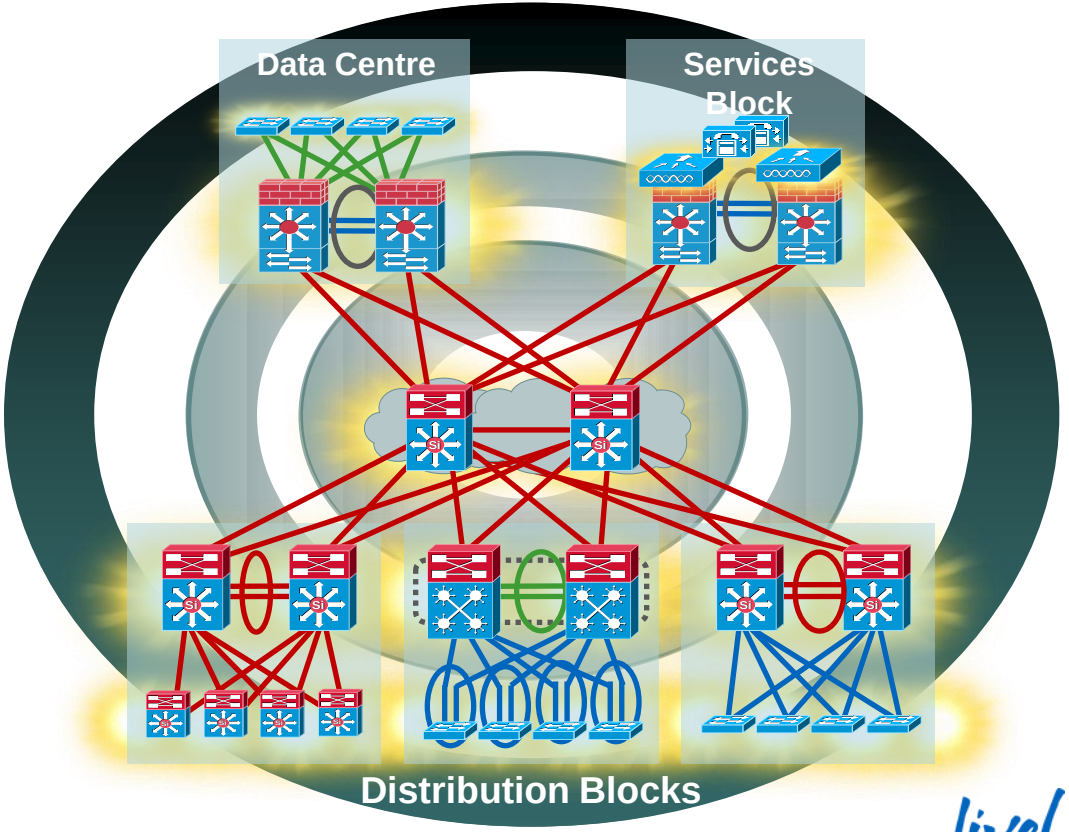

Multilayer Network Design
L2 Access with L3 Distribution

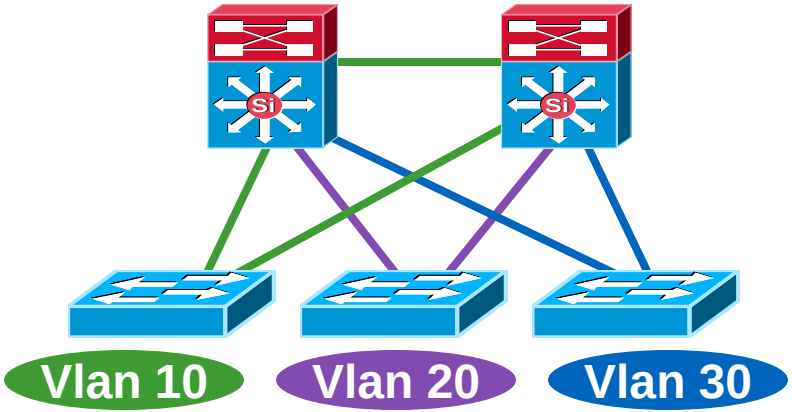

- Each access switch has unique VLANs
- No Layer 2 loops
- Layer 3 link between distribution
- No blocked links

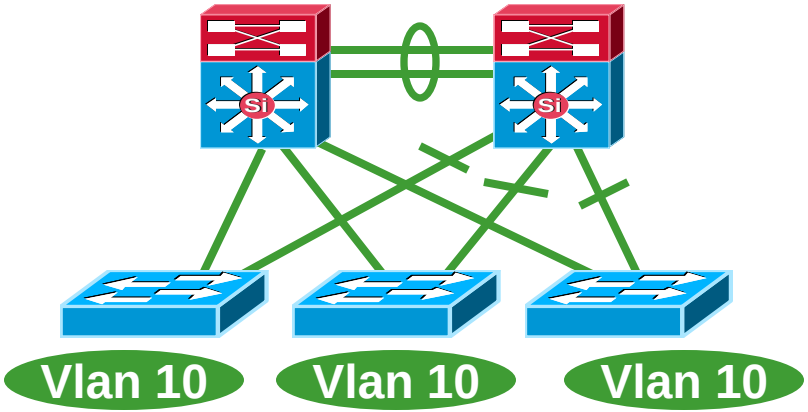

- At least some VLANs span multiple access switches
- Layer 2 loops, blocked links
- Layer 2 and 3 running over link between distribution

L2 Hardening

Spanning Tree should behave the way you expect

- Place the root where you want it
- The root bridge should stay where you put it
  - RootGuard
  - LoopGuard
  - UplinkFast
  - UDLD
- Only end-station traffic should be seen on an edge port
  - BPDU Guard
  - RootGuard
  - PortFast, PortSecurity

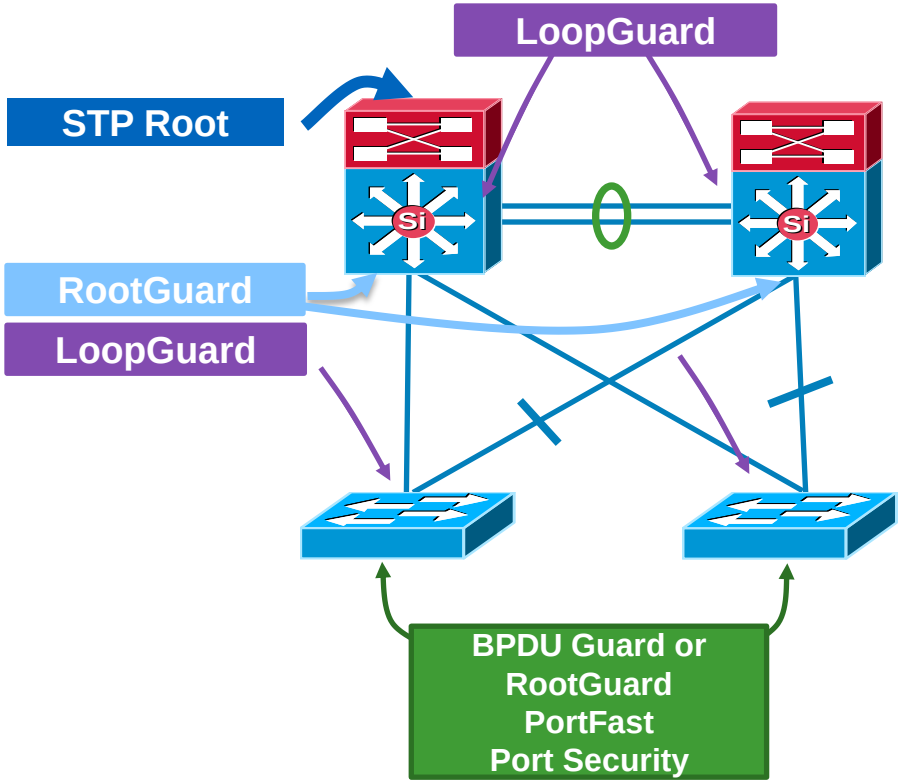

L3 Best Practice
Build Triangles not Squares

Triangles: Link/Box failure does not require routing protocol convergence

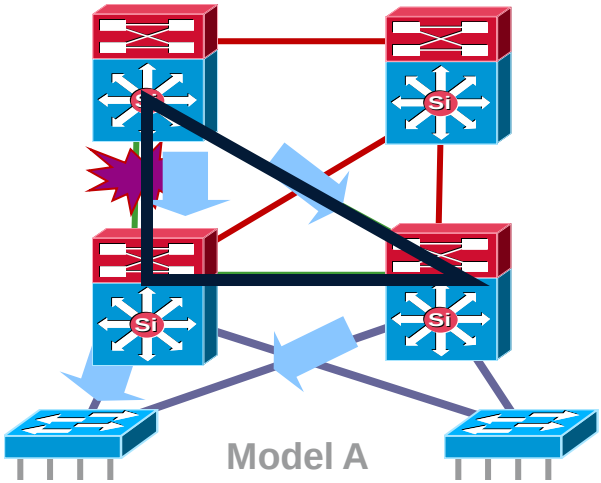

Squares: Link/Box failure requires routing protocol convergence

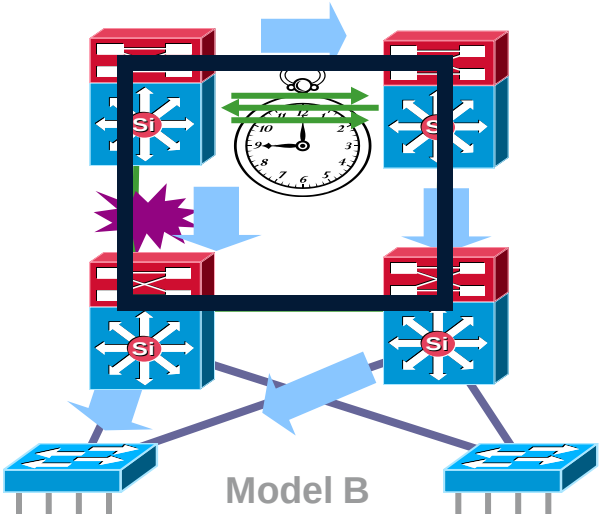

- Layer 3 redundant equal cost links support fast convergence
- Hardware based—fast recovery to remaining path
- Convergence is extremely fast (dual equal-cost paths: no need for OSPF or EIGRP to recalculate a new path)

L3 Summarise at the Distribution

Limit EIGRP Queries and OSPF LSA Propagation

- It is important to force summarisation at the distribution towards the core
- For return path traffic an OSPF or EIGRP re-route is required
- By limiting the number of peers an EIGRP router must query or the number of LSAs an OSPF peer must process we can optimise this reroute

EIGRP example:

```
interface Port-channel1
 description to Core#1
 ip address 10.122.0.34 255.255.255.252
 ip hello-interval eigrp 100 1
 ip hold-time eigrp 100 3
 ip summary-address eigrp 100 10.1.0.0
  255.255.0.0 5
```

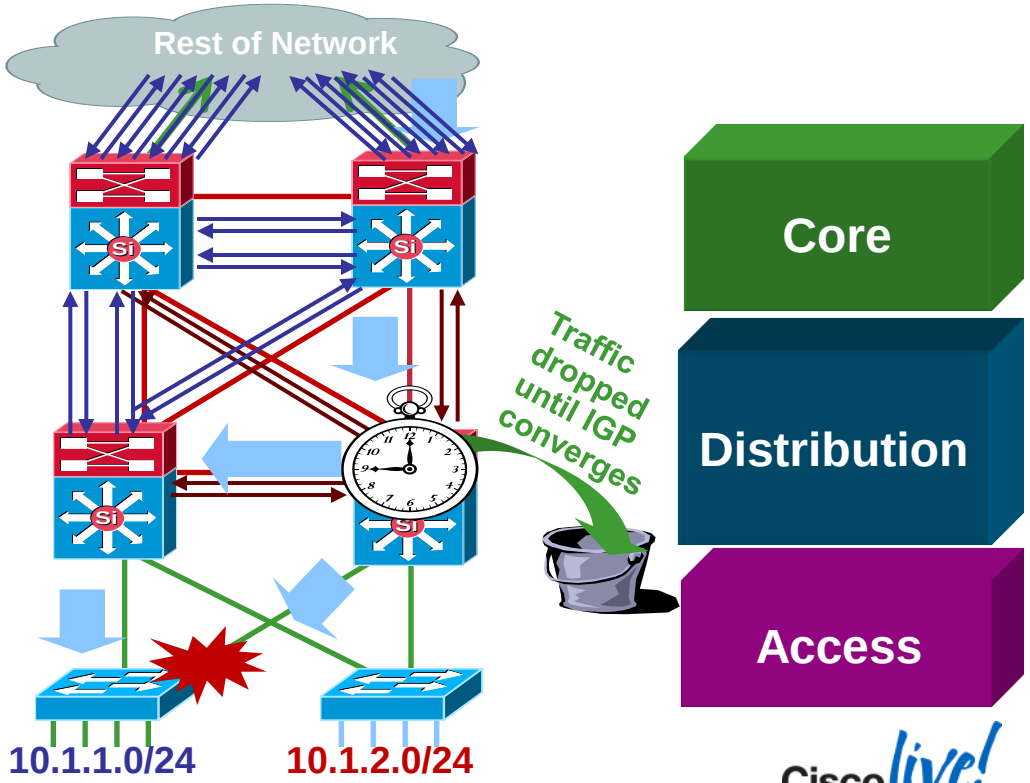

L3 Summarise at the Distribution

Reduce the Complexity of IGP Convergence

- It is important to force summarisation at the distribution towards the core
- For return path traffic an OSPF or EIGRP re-route is required
- By limiting the number of peers an EIGRP router must query or the number of LSAs an OSPF peer must process we can optimise his reroute
- For EIGRP if we summarise at the distribution we stop queries at the core boxes for an access layer flap
- For OSPF when we summarise at the distribution (area border or L1/L2 border) the flooding of LSAs is limited to the distribution switches; SPF now deals with one LSA not three

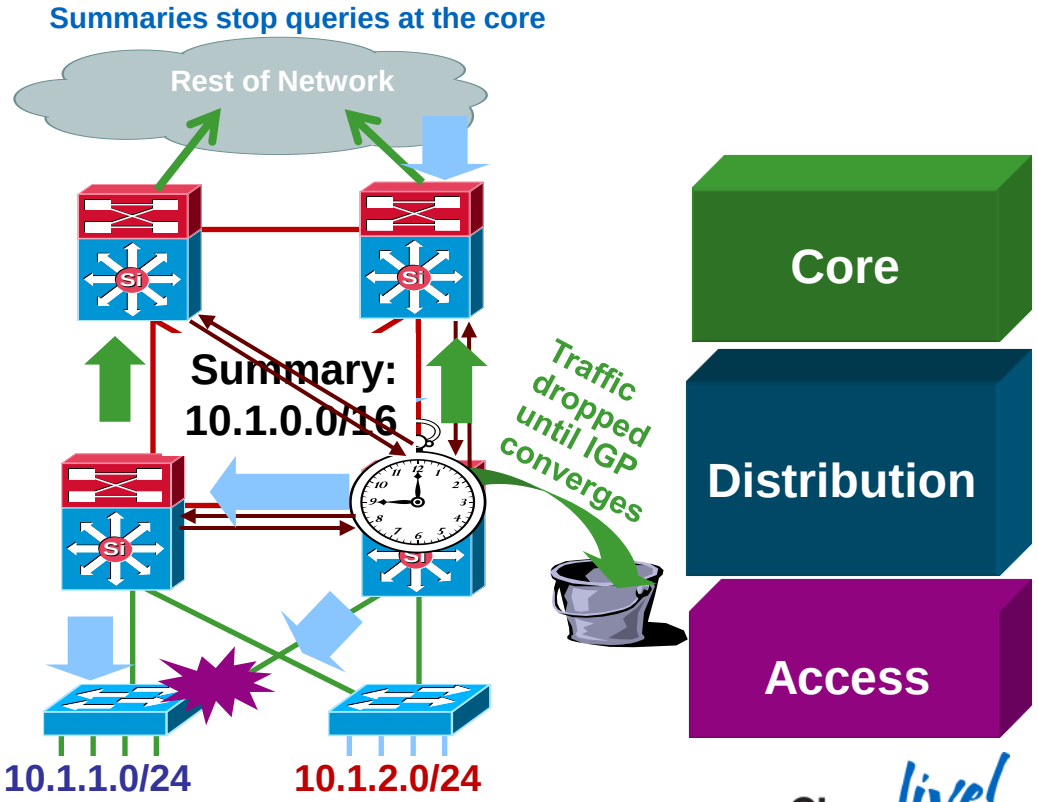

L3 Equal-Cost Multipath

Optimising CEF Load Sharing

- Depending on the traffic flow patterns and IP addressing in use, one algorithm may provide better load-sharing results than another
- Be careful not to introduce polarisation in a multi-tier design by changing the default to the same thing in all tiers/layers of the network

Catalyst 4500 Load-Sharing Options	
Original	Src IP + Dst IP
Universal*	**Src IP + Dst IP + Unique ID**
Include Port	Src IP + Dst IP + (Src or Dst Port) + Unique ID

Catalyst 6500 Load-Sharing Options	
Default*	**Src IP + Dst IP + Unique ID**
Full	Src IP + Dst IP + Src Port + Dst Port
Full Exclude Port	Src IP + Dst IP + (Src or Dst Port)
Simple	Src IP + Dst IP
Full Simple	Src IP + Dst IP + Src Port + Dst Port

* = Default Load-Sharing Mode

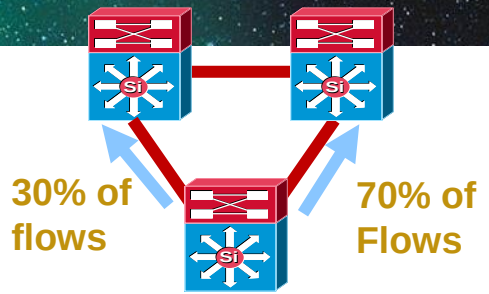

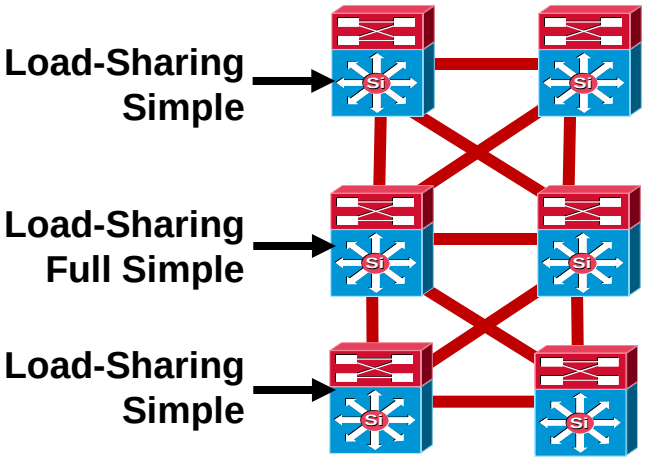

L3 CEF Load Balancing

Avoid Underutilising Redundant L3 Paths

Distribution
Default L3 Hash

Core
Default L3 Hash

Distribution
Default L3 Hash

Redundant Paths Ignored

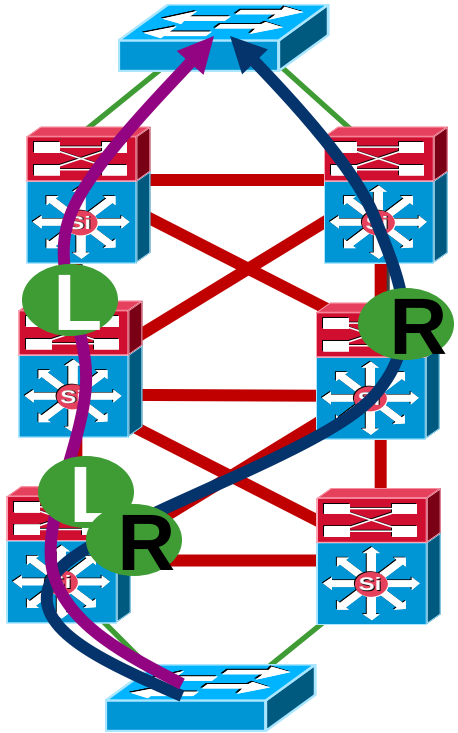

- CEF polarisation: without some tuning CEF will select the same path left/left or right/right
- Imbalance/overload could occur
- Redundant paths are ignored/underutilised
- The default CEF hash input is L3
- We can change the default to use L3 + L4 information as input to the hash derivation

L3 CEF Load Balancing

Avoid Underutilising Redundant L3 Paths

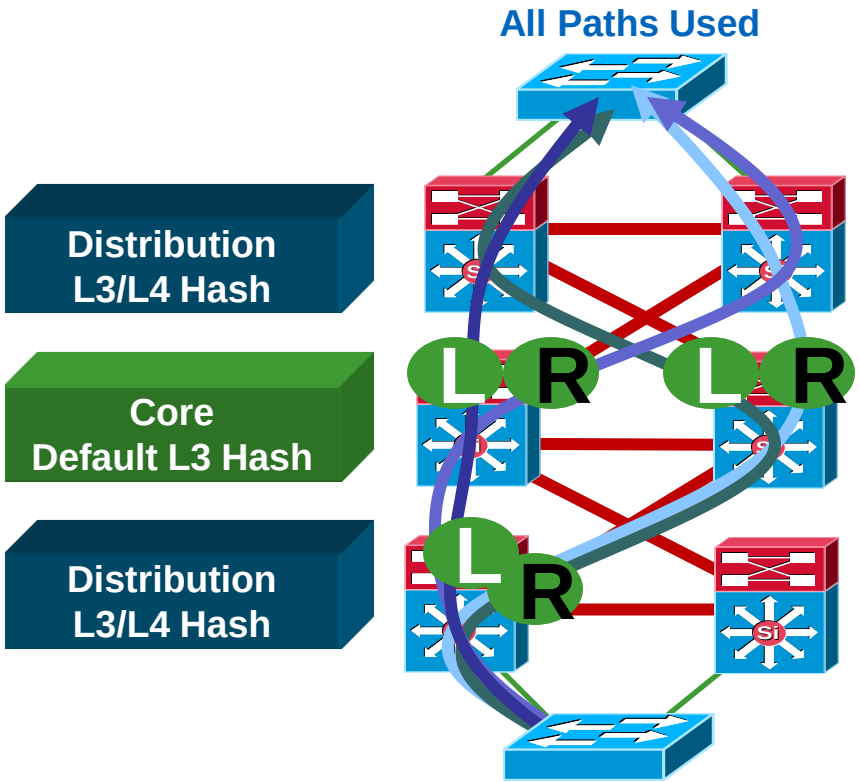

- Depending on IP addressing and flows, imbalance could occur
- Alternating L3/L4 hash and L3 hash will give us the best load balancing results
- Use **simple** in the core and **full simple** in the distribution to add L4 information to the algorithm at the distribution and maintain differentiation tier-to-tier

First Hop Redundancy
Why You Want HSRP Preemption

- Spanning tree root and HSRP primary aligned
- When spanning tree root is re-introduced, traffic will take a two-hop path to HSRP active
- HSRP preemption will allow HSRP to follow spanning tree topology

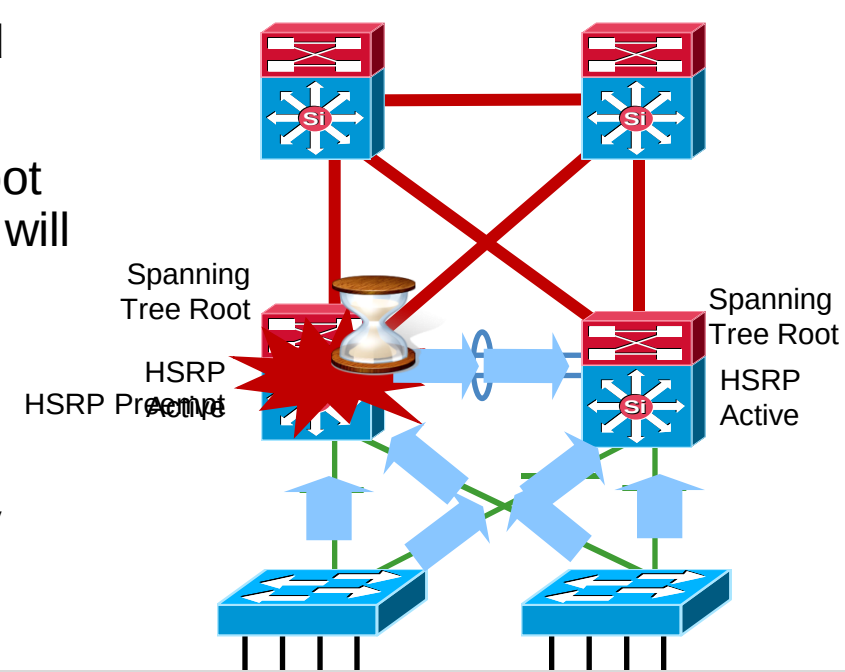

Without preempt delay HSRP can go active before box completely ready to forward traffic due to L1 (Boards), L2 (STP), L3 (IGP Convergence)

IOS (config-if)# **standby 1 preempt delay minimum 30**

What if you don't link the Distributions?

Black Holes and Multiple Transitions …

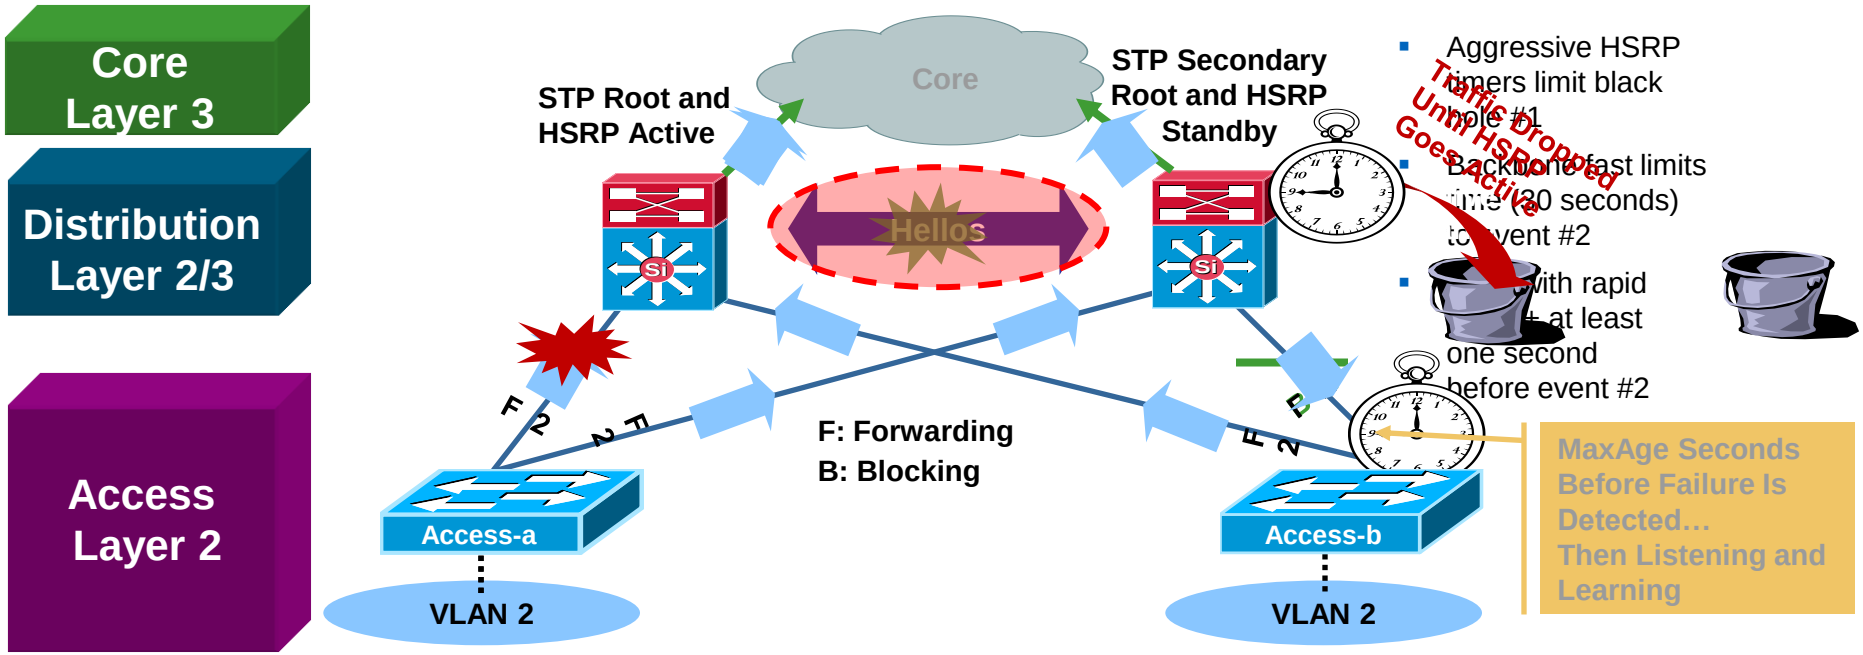

- **Blocking link on access-b will take 50 seconds to move to forwarding → traffic black hole until HSRP goes active on standby HSRP peer**
- **After MaxAge expires (or backbone fast or Rapid PVST+) converges HSRP preempt causes another transition**
- **Access-b used as transit for Access-a's traffic**

What if you don't link the Distributions?

Return Path Traffic Black-holed

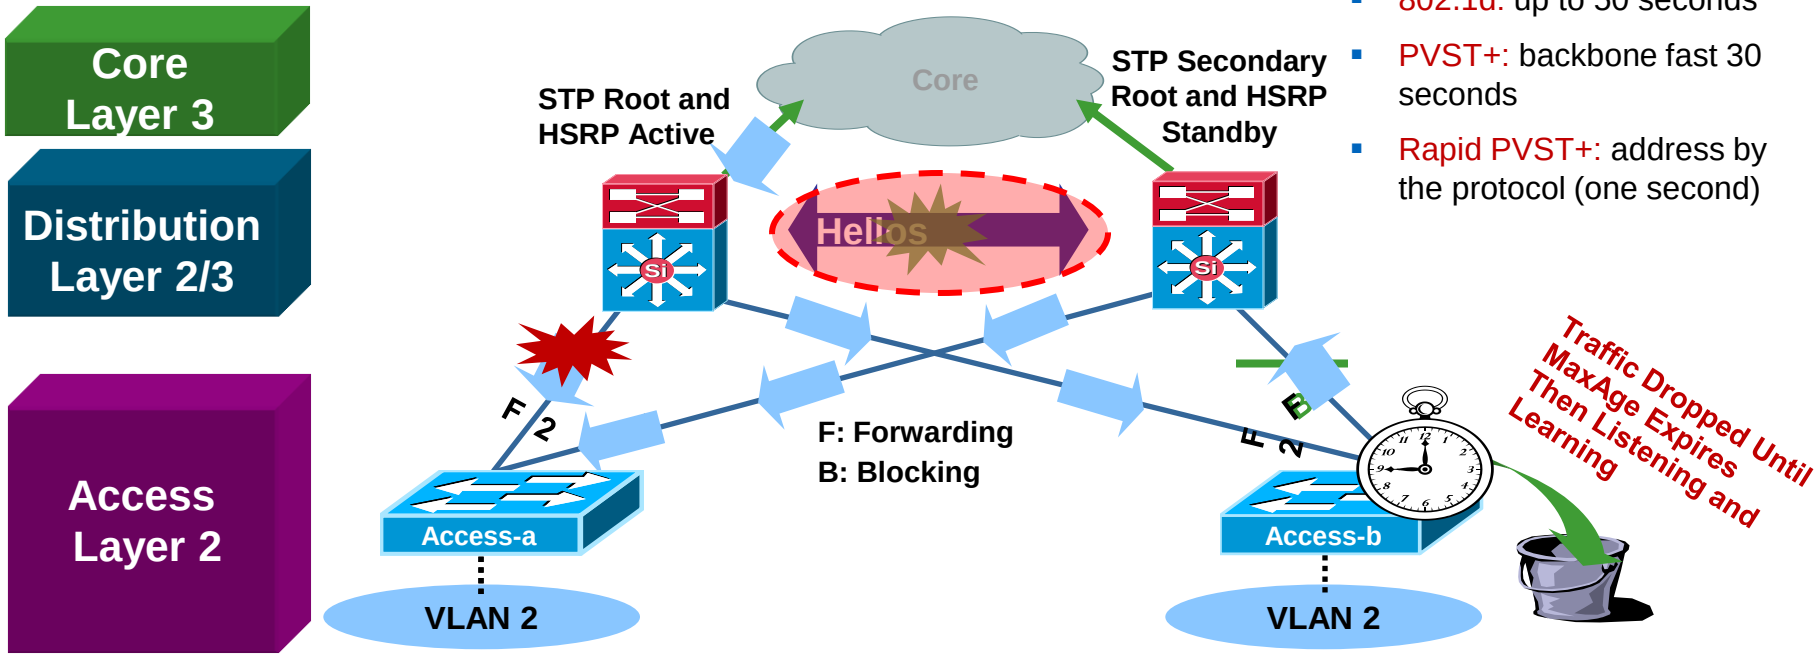

- 802.1d: up to 50 seconds
- PVST+: backbone fast 30 seconds
- Rapid PVST+: address by the protocol (one second)

- Blocking link on access-b will take 50 seconds to move to forwarding → return traffic black hole until then

Asymmetric Routing

- Affects redundant topologies with shared L2 access
- One path upstream and two paths downstream
- CAM table entry ages out on standby HSRP
- Without a CAM entry packet is flooded to all ports in the VLAN

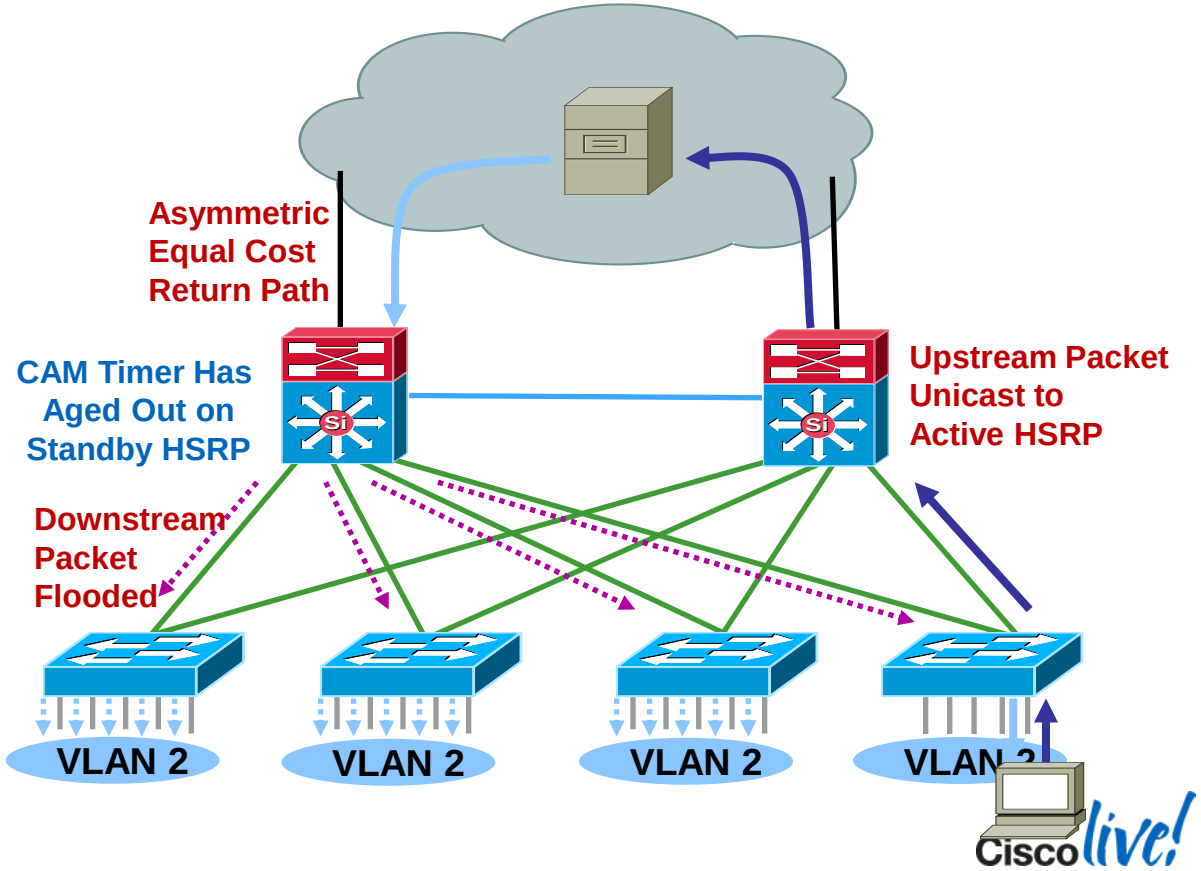

Asymmetric Routing

Best Practice to Prevent Excessive Flooding

- Assign one unique data and voice VLAN to each access switch
- Traffic is now only flooded down one trunk
- Access switch unicasts correctly; no flooding to all ports
- If you have to:
  - Tune ARP and CAM aging timers; CAM timer exceeds ARP timer
  - Bias routing metrics to remove equal cost routes

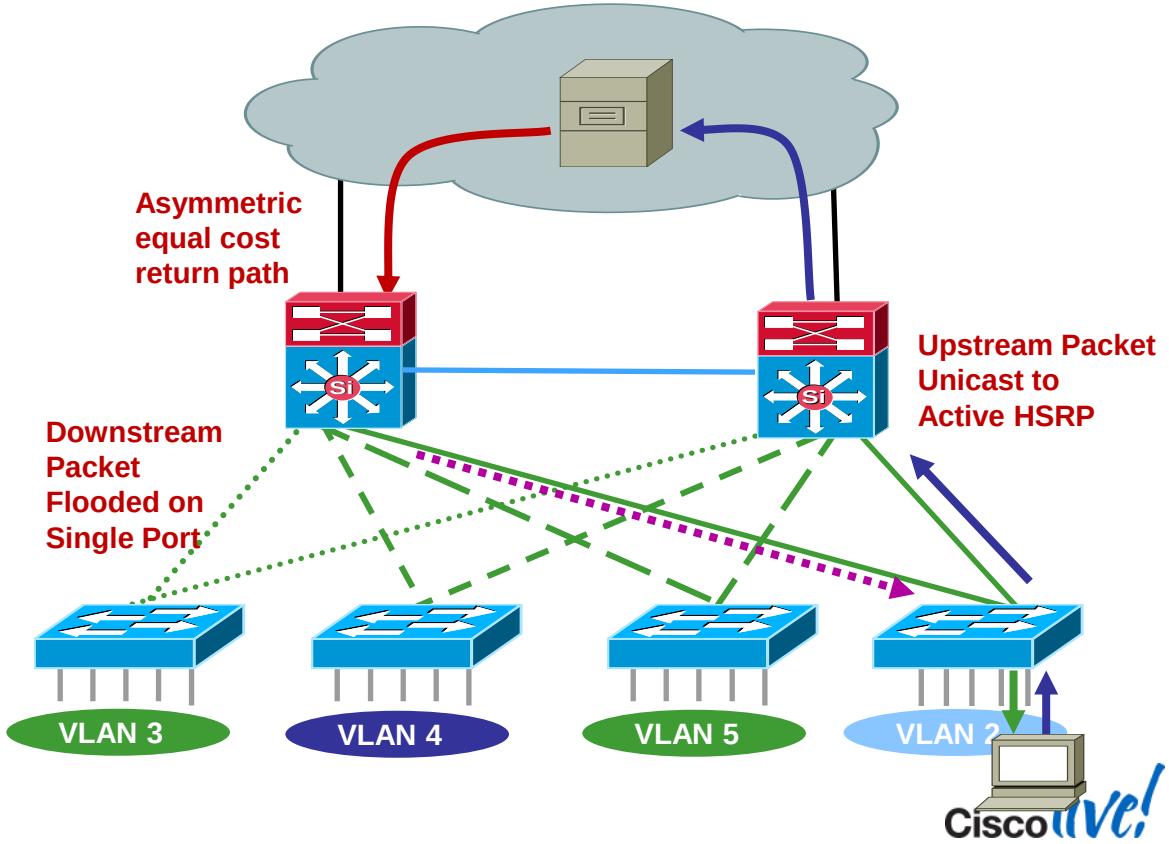

Agenda

- The Principles
- The Basics
- The Cool Stuff
- The End

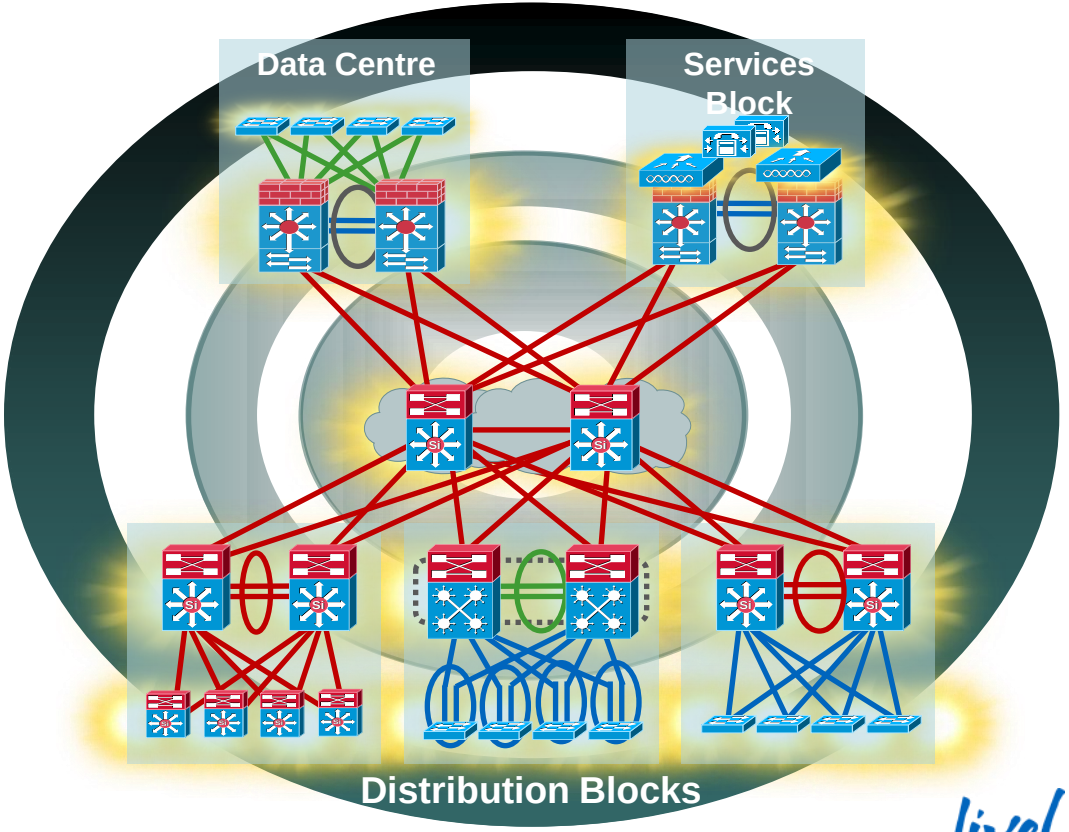

Virtualisation (VSS)

Cisco Virtual Switching System (VSS)

Traditional Campus Design

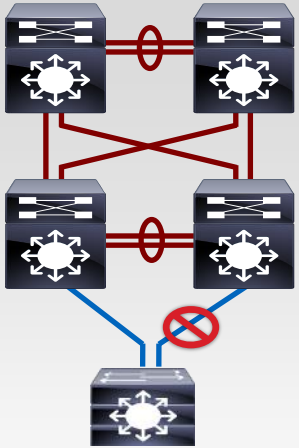

- Complex Network Design and Operation
- Underutilise Network Resource
- Sub-Optimal Application and Network Performance

Optimised Network

VSS Campus Design

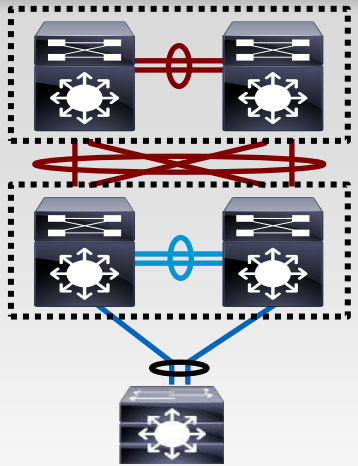

- Optimised Network Design
- Double Switching Capacity
- Deterministic Application and Network Performance

Simplified Operation

VSS Campus Design

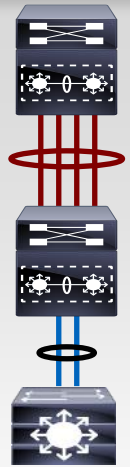

- Simplified System Operation
- Single Neighbour and Network Per Layer
- Simplified and Highly Redundant Network Topologies

Unified Access/Converged Access/Instant Access

Campus Deployment Models

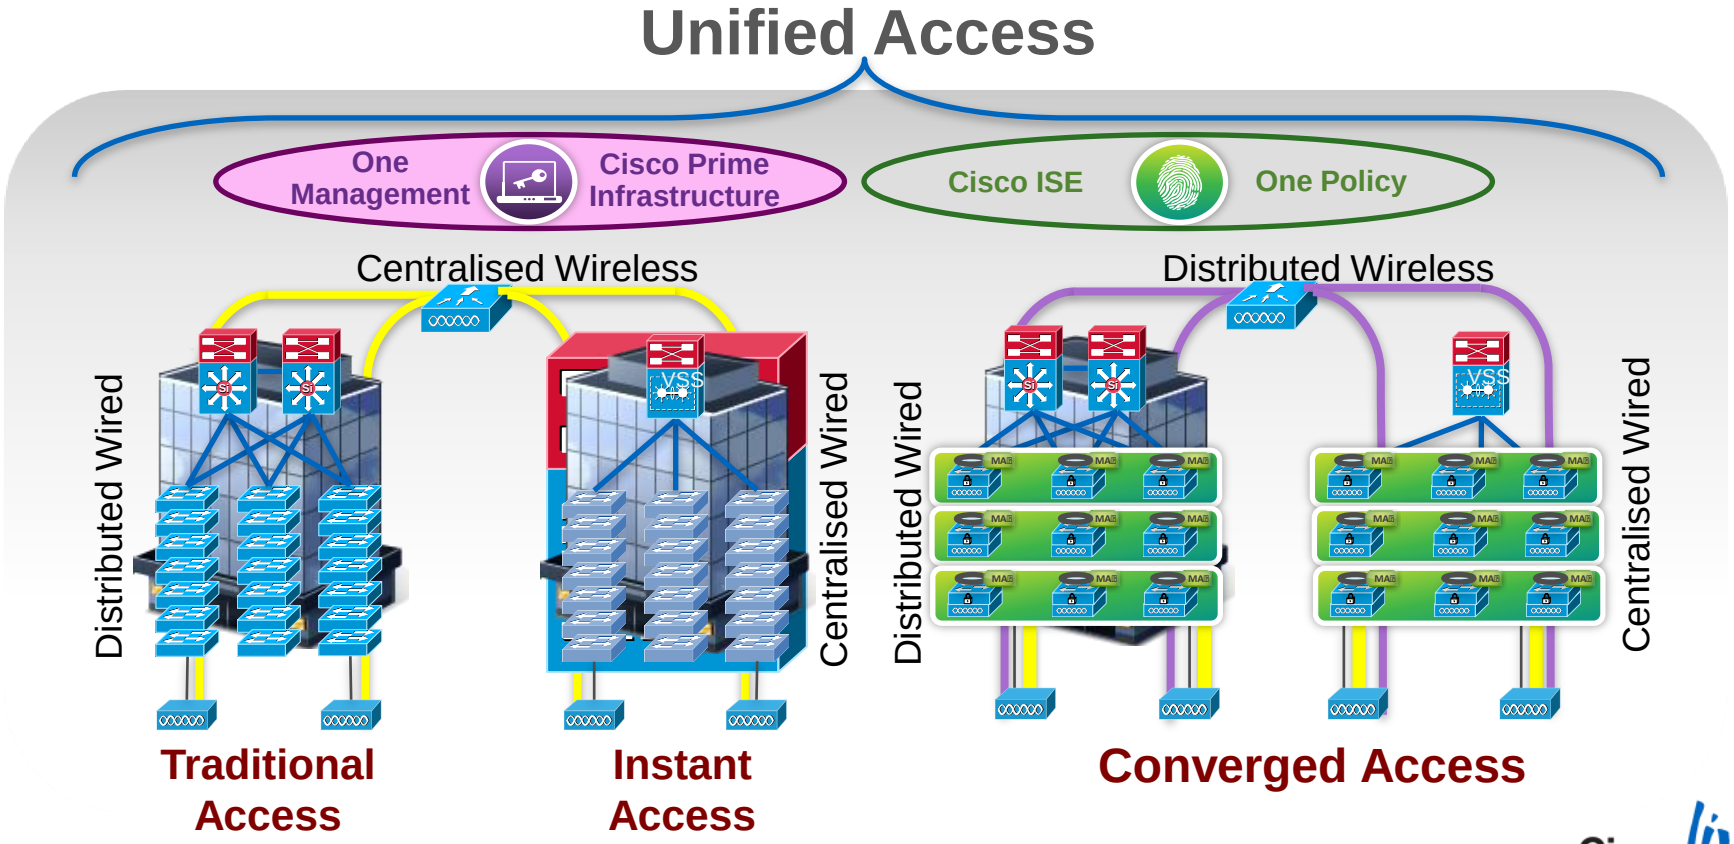

Converged Access

Centralised Wireless

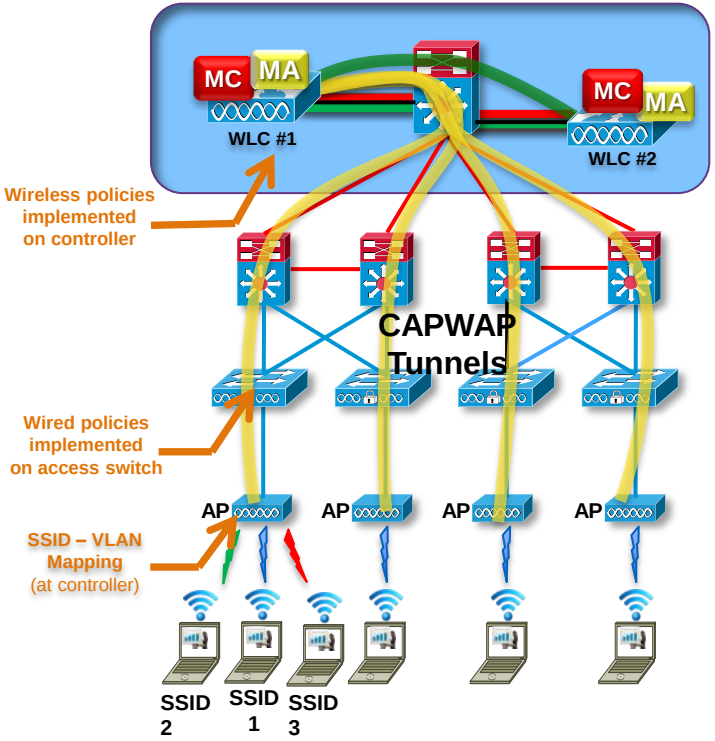

VS.

Distributed Wireless

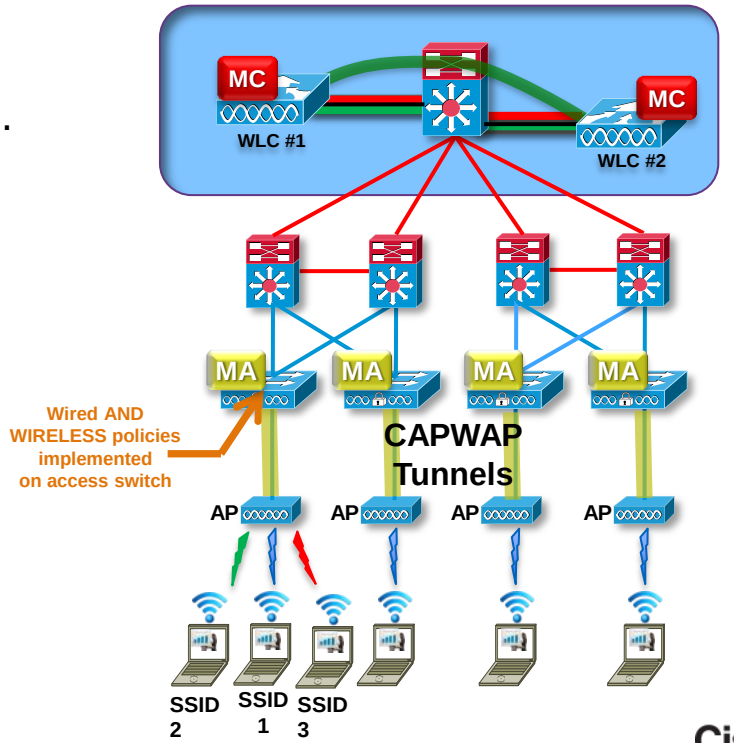

Converged Access (cont.)

Centralised Wireless

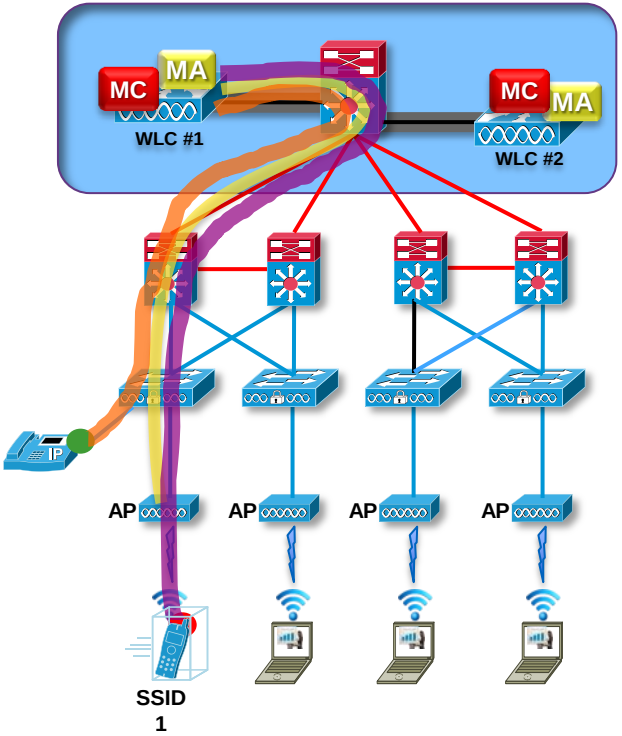

VS.

Distributed Wireless

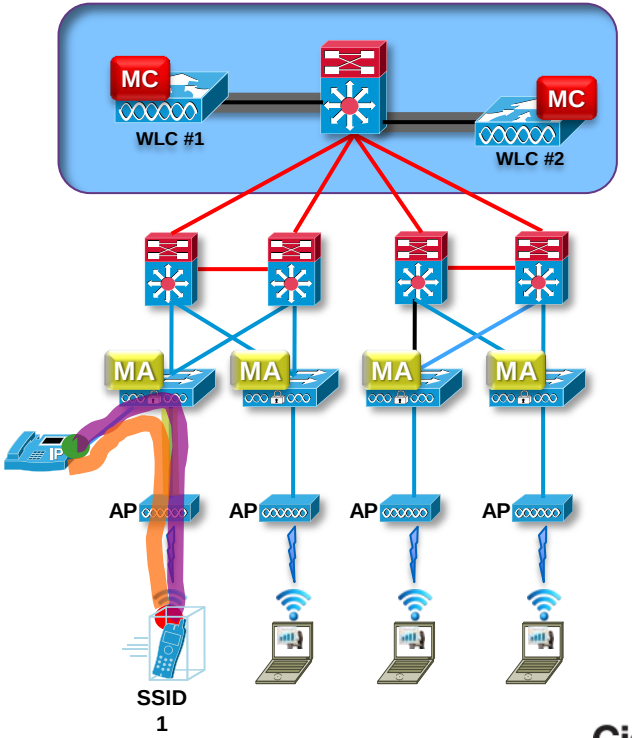

Instant Access

Managed Devices = 2

Example: 1000 User-Port Campus Distribution POD

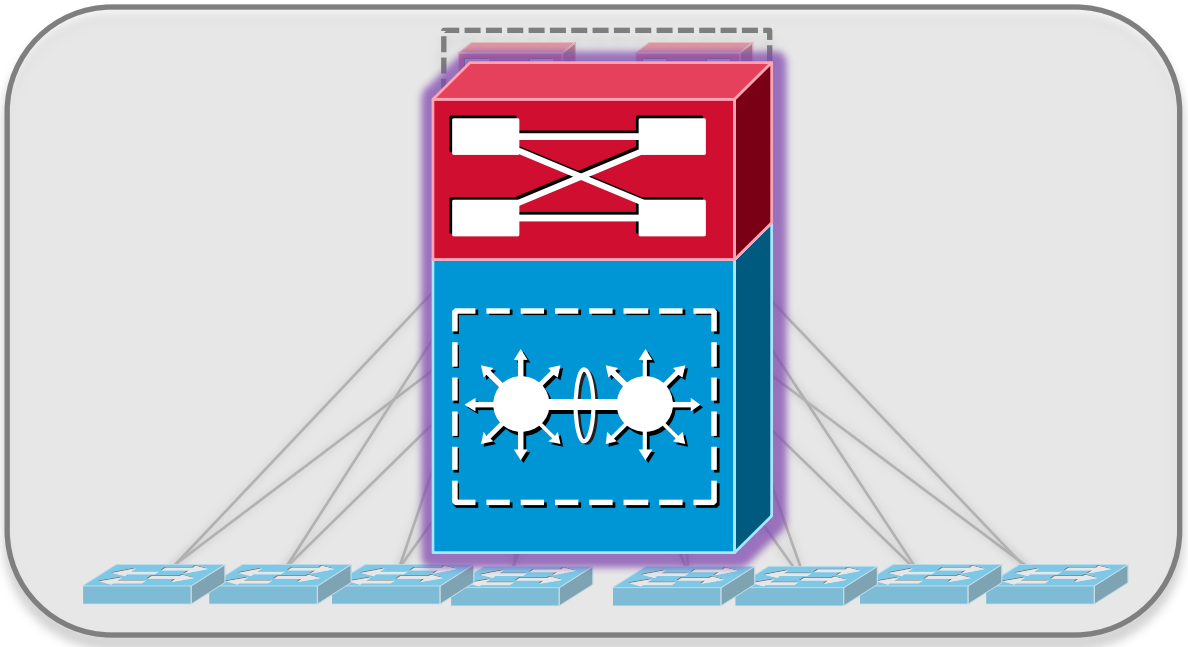

- Single point of management
- Single configuration
- Single IOS image
- No spanning-tree

Instant Access components

Parent switch:

- Sup2T or Sup10 (6800) VSS
- WS-X6904-40G and FourX adapter
- 15.1(2)SY, IP Services or higher

Client switch:

- 6800-IA switch
  - PoE or non-PoE
  - Stackable (3 max at Phase 1)
- No local processing of packets

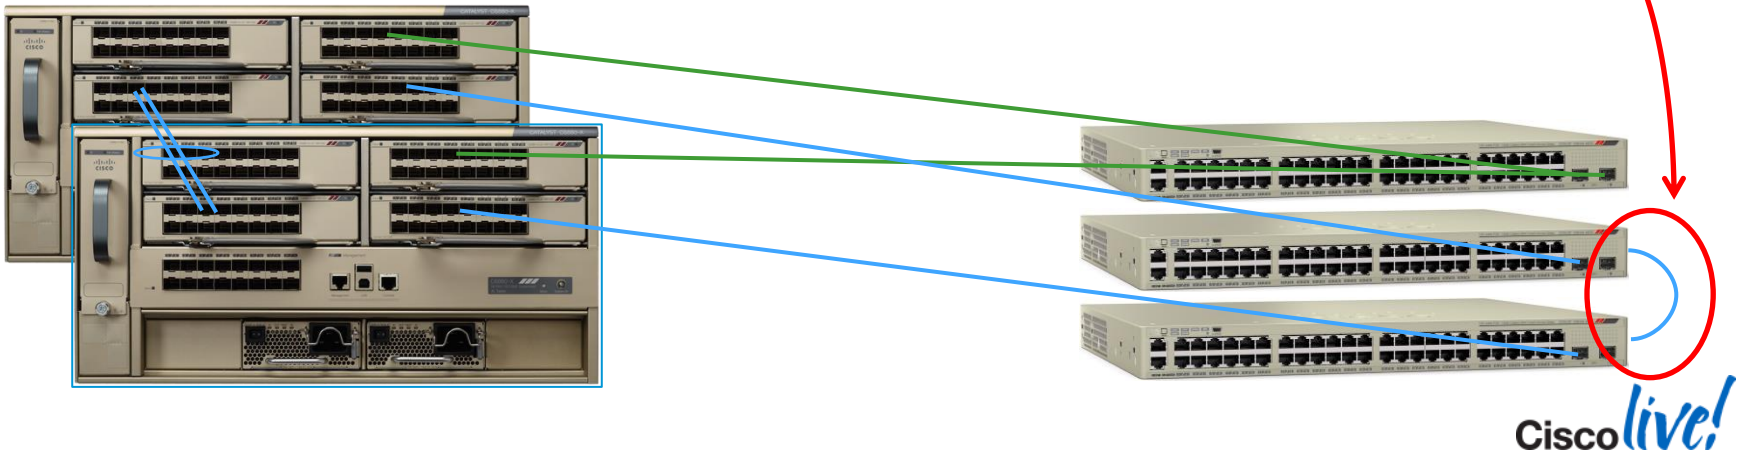

Unified/Converged/Instant Access

For more information:

BRKARC-2665 Converged Access Architecture, Design and Deployment
Thursday 2:00pm – 4:00pm

BRKARC-3465 Cisco Catalyst 6500 Instant Access Solution - Design and Migration Case Studies
Thursday 8:30am – 10:30am

Smart Operations

Cisco Catalyst SmartOperations

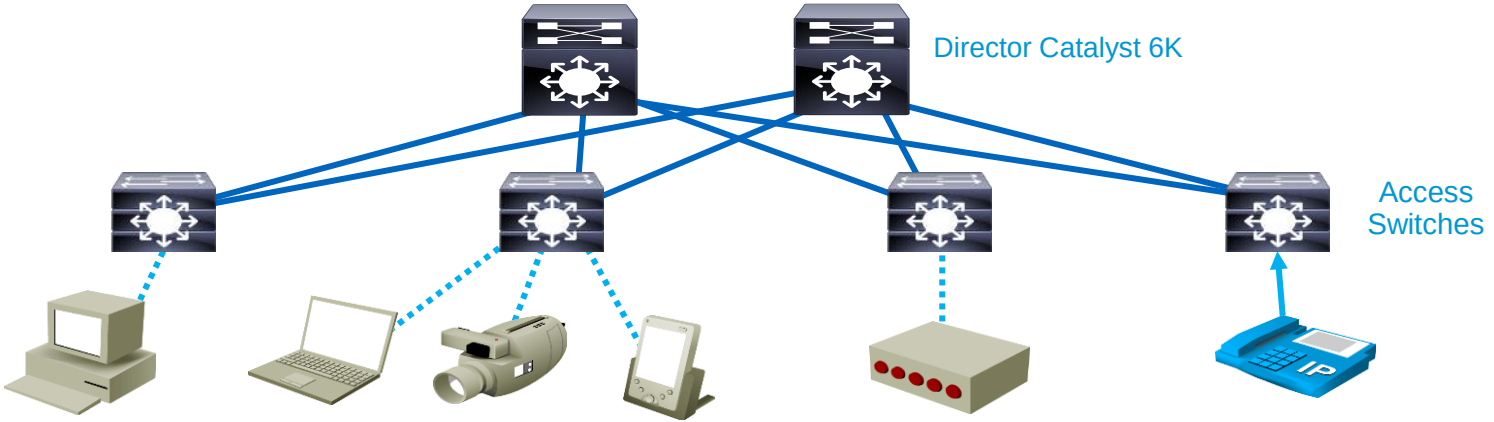

Auto-QoS
Automatically Creates Relevant QoS Configuration

New Configuration

- No in-depth QoS knowledge needed
- VoIP feature simplifies QoS implementation
- Can use existing Cisco commands to modify the automatically generated configuration

Auto Smartports
Plug and Play for End Devices

New Device Attached

- Port configuration: Applied
- QoS policy: Enforced
- Security policy: Enforced

Smart Install
Zero-Touch Deployments and Maintenance

New Switch Connected

- Software image downloaded
- Configuration automatically applied

AutoQoS

```
Switch(config-if)# auto qos ?

classify Classify untrusted traffic
trust    Trust the DSCP/CoS marking
voip     Configure AutoQoS for VoIP
video    Configure AutoQos for video
```

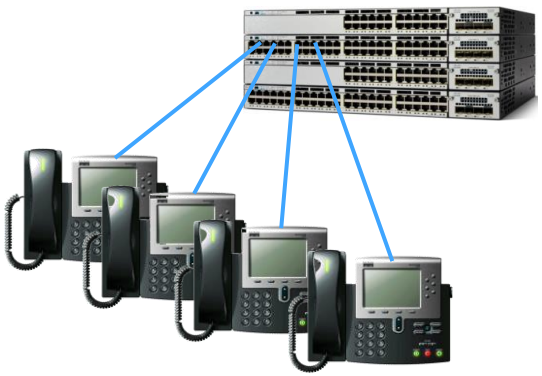

```
Switch(config)# default fastethernet 0/48
Switch(config)# interface fa0/48
Switch(config-if)# auto qos voip cisco-phone
Switch(config-if)# do show run interface fa0/48
interface FastEthernet0/48
 switchport mode access
 mls qos trust device cisco-phone
 mls qos trust cos
 auto qos voip cisco-phone
 wrr-queue bandwidth 10 20 70 1
 wrr-queue min-reserve 1 5
 wrr-queue min-reserve 2 6
 wrr-queue min-reserve 3 7
 wrr-queue min-reserve 4 8
 wrr-queue cos-map 1 0 1
 wrr-queue cos-map 2 2 4
 wrr-queue cos-map 3 3 6 7
 wrr-queue cos-map 4 5
 priority-queue out
 auto qos voip cisco-phone
 service-policy input AutoQoS-CiscoPhone-Policy
end
```

SmartPorts - Predefined Configurations

```
Switch# show parser macro brief
   default global   : cisco-global
   default interface: cisco-desktop
   default interface: cisco-phone
   default interface: cisco-switch
   default interface: cisco-router
   default interface: cisco-wireless
```

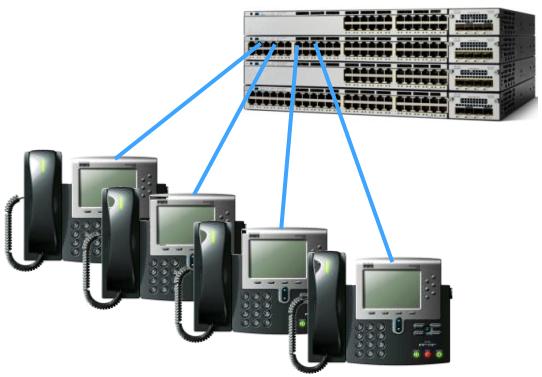

```
Switch(config)# default fastethernet 0/48
Switch(config)# int fa0/48
Switch(config-if)# macro apply cisco-phone $ACCESS_VLAN 20
                   $VOICE_VLAN 10

Switch# show run int fa0/48
switchport access vlan 20
switchport mode access
switchport voice vlan 10
switchport port-security maximum 2
switchport port-security
switchport port-security aging time 2
switchport port-security violation restrict
srr-queue bandwidth share 10 10 60 20
srr-queue bandwidth shape  10  0  0  0
mls qos trust device cisco-phone
mls qos trust cos
 macro description cisco-phone
 auto qos voip cisco-phone
 spanning-tree portfast
 spanning-tree bpduguard enable
end
```

SmartPorts – Automatic Configurations of ports


```
Switch (config)# macro auto global processing
Switch (config-if)# no macro auto global processing
```

Enable auto-smartports

Or disable per-port

```
Switch (config)# macro auto device phone ACCESS_VLAN=10 VOICE_VLAN=20
```

Change macro defaults

If necessary

```
Switch (config-if)# macro auto sticky
```

Make applied config permanent

“Last resort” macro applied if device not known (no CDP/LLDP or pre-defined OUI/MAC list matches.

Smart Install - Components

Client	Receives image and configuration from Director
Groups	Collection of Clients with same image and configuration
Director	Manages Client image installation and configuration
DHCP and TFTP Servers	Centrally located and shared across network

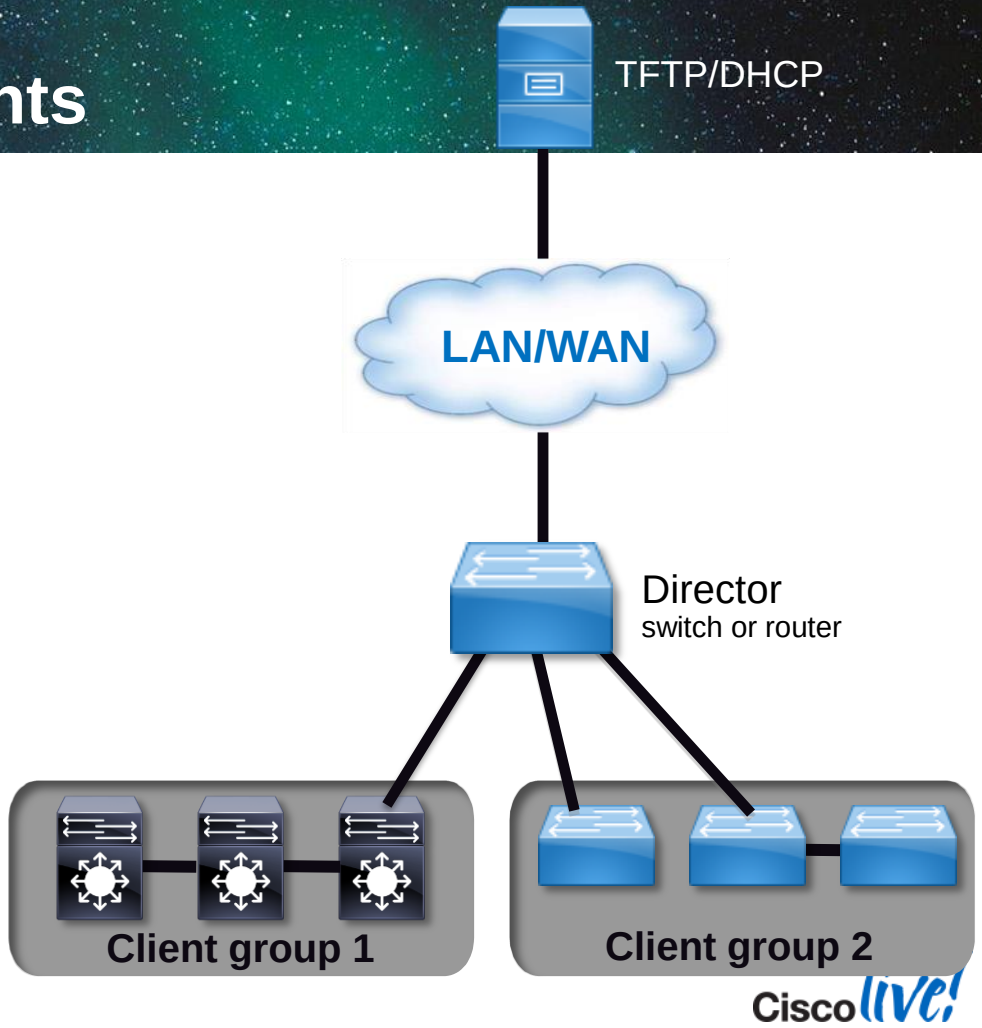

Smart Install – How it Works

1. Director discovers client via CDP
2. New switch issues DHCP discover
3. Director adds options to DHCP offer
4. Client retrieves image, config via TFTP
5. Client reboots with new configuration and image

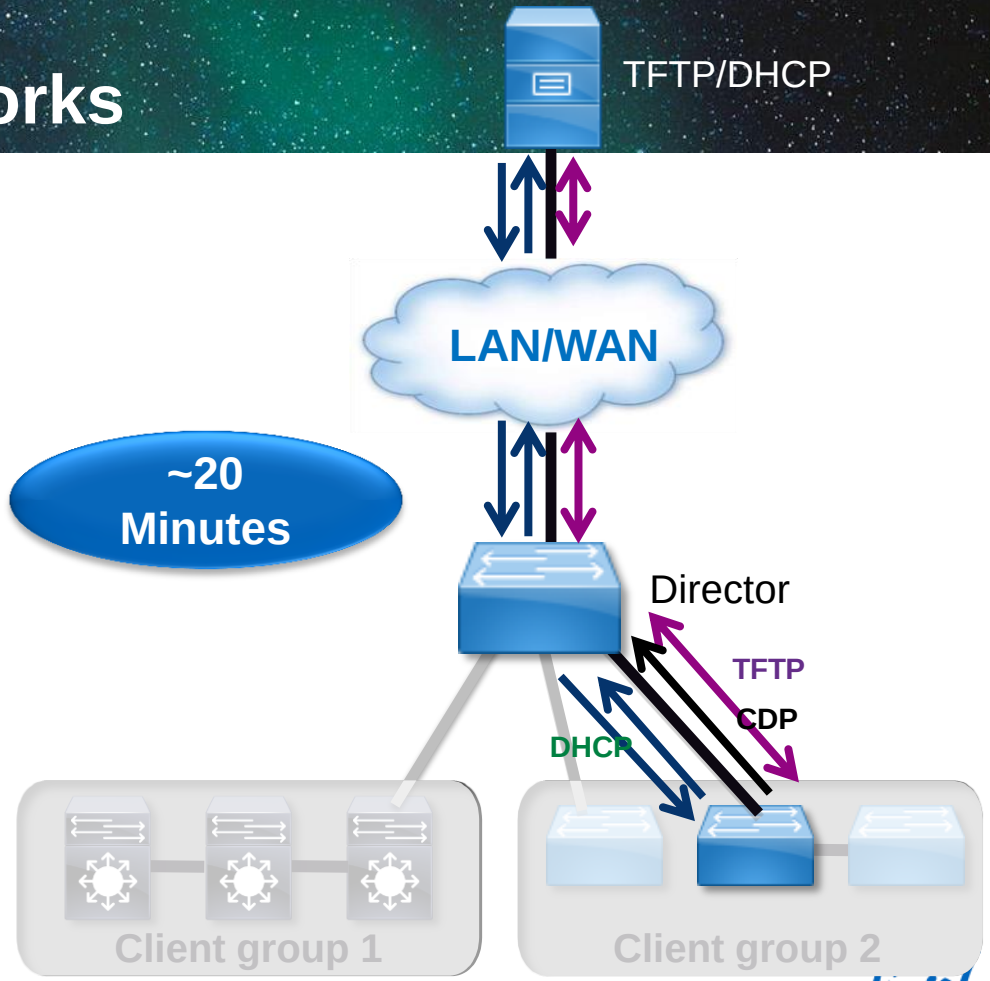

Catalyst Smart Operations

For more information:

BRKCRS-3090 Implementing Network Automations - Power Tools for Catalyst Switching Network Operations

Thursday 4:30pm – 6:30pm

Virtualisation (EVN)

Why Network Virtualisation?

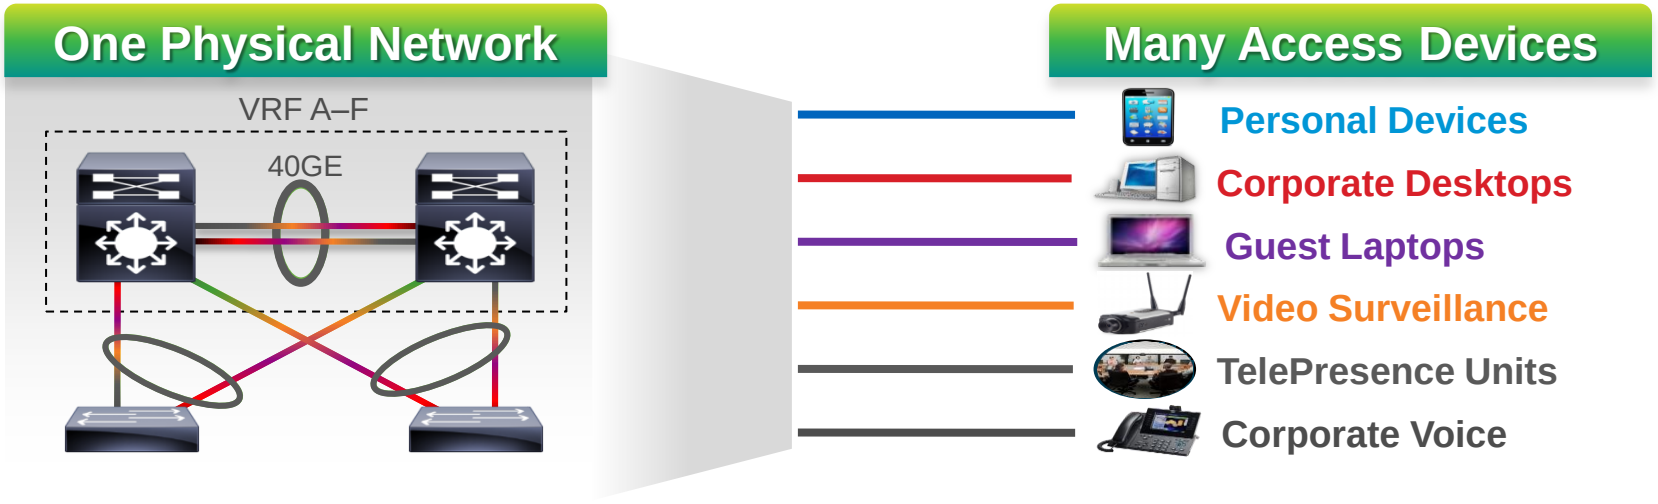

Simplified Network Design via MPLS, VRF-Lite and **EVN**

Enhanced Security**,** Group Segregation, and Shared Services via **Virtualised Firewalls**

Better Monitoring and Operations with VRF-Aware Services

What is Multi-VRF CE (VRF-Lite)?

- It is a device virtualisation technique to virtualise Layer 3 routing and forwarding.
- It allows the switch to maintain multiple routing and forwarding tables.
- Each VRF has its own interfaces.
- It allows overlapping address spaces, and complete Layer 2 and Layer 3 traffic isolation: virtual networks.

```
192.168.1.0/32 is subnetted, 1 subnets
C       192.168.1.102 is directly connected, Loopback12
```

```
192.168.1.0/32 is subnetted, 1 subnets
C       192.168.1.102 is directly connected, Loopback11
```

```
192.168.255.0/32 is subnetted, 1 subnets
C       192.168.255.253 is directly connected, Loopback0
```

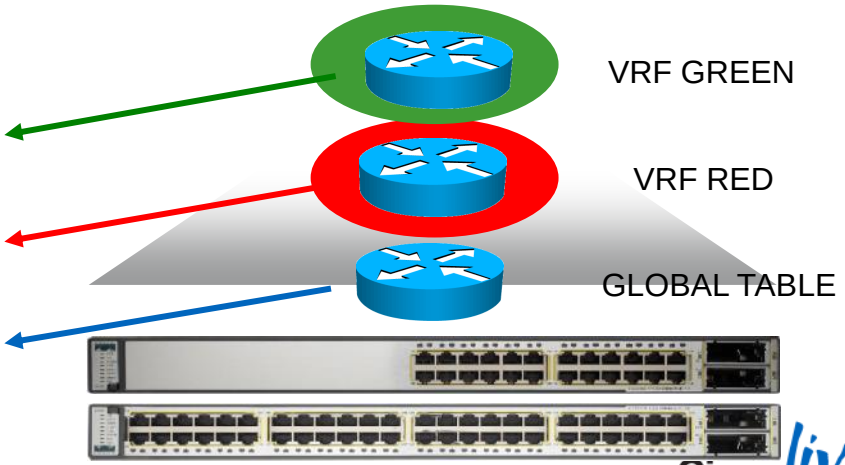

What is Easy Virtual Network (EVN)?

Multi-VRF Network

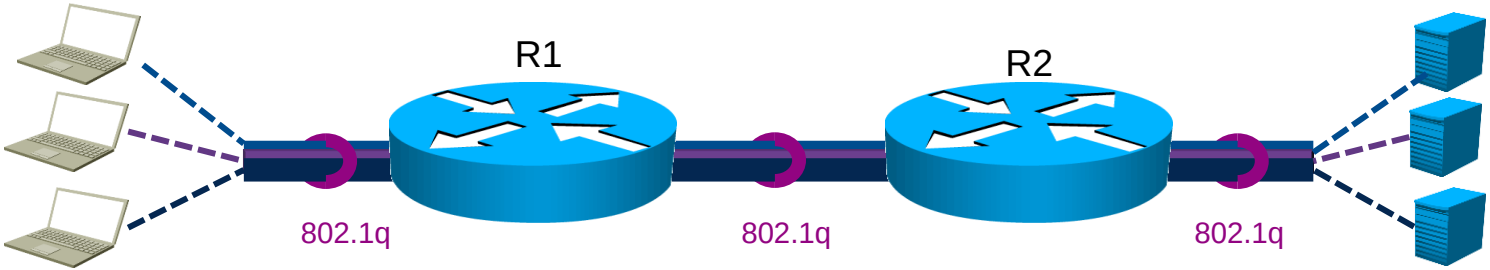

Easy Virtual Network

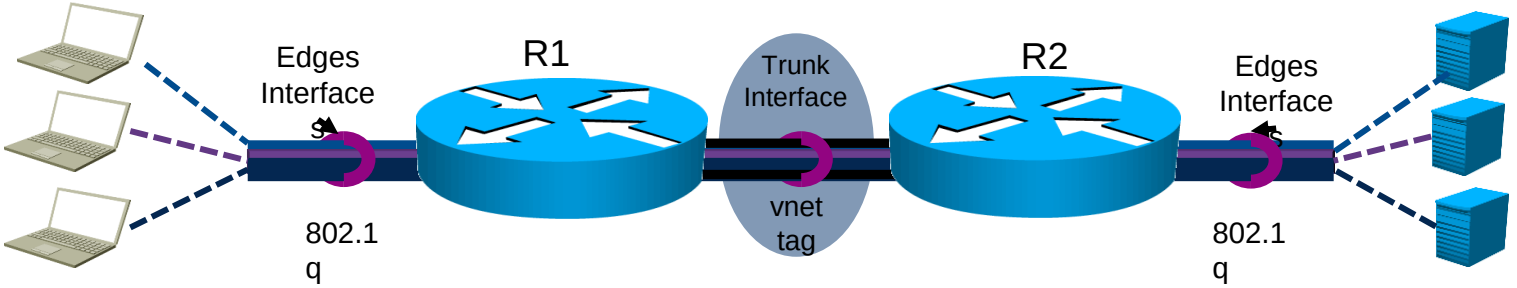

Trunk Configuration Comparison

VRF-lite end-to-end example

```
vrf definition RED
 address-family ipv4
vrf definition GREEN
 address-family ipv4
vrf definition BLUE
 address-family ipv4
!
interface GigabitEthernet0/0
 description Trunk interface
!
interface GigabitEthernet0/0.100
 vrf forwarding RED
 encapsulation dot1Q 100
 ip address 10.100.1.1 255.255.255.0
!
interface GigabitEthernet0/0.101
 vrf forwarding GREEN
 encapsulation dot1Q 101
 ip address 10.101.1.1 255.255.255.0
!
interface GigabitEthernet0/0.102
 vrf forwarding BLUE
 encapsulation dot1Q 102
 ip address 10.102.1.1 255.255.255.0
```

EVN example

```
vrf definition RED
 vnet tag 100
 address-family ipv4
vrf definition GREEN
 vnet tag 101
 address-family ipv4
vrf definition BLUE
 vnet tag 102
!
interface GigabitEthernet0/0
 description Trunk interface
 ip address 10.1.1.1 255.255.255.0
 vnet trunk
```

New command

Automatically creates sub-interfaces for each VRF.
"show derived-config gig0/0.101" will show sub-interface config.

Can filter using "list" option pointing to a "vrf-list"

Shared Service Configuration Comparison

VRF-lite end-to-end example

```
ip vrf SHARED
 rd 3:3
 route-target export 3:3
 route-target import 1:1
 route-target import 2:2
!
ip vrf RED
 rd 1:1
 route-target export 1:1
 route-target import 3:3
!
ip vrf GREEN
 rd 2:2
 route-target export 2:2
 route-target import 3:3
!
!
router bgp 65001
 bgp log-neighbor-changes
 !
 address-family ipv4 vrf SHARED
 ……
```

EVN example

```
vrf definition SHARED
  address-family ipv4
   route-replicate from vrf RED
                  unicast all route-map red-map
   route-replicate from vrf GREEN
                  unicast all route-map green-map

vrf definition RED
  address-family ipv4
   route-replicate from vrf SHARED
         unicast all

vrf definition GREEN
  address-family ipv4
   route-replicate from vrf SHARED
         unicast all
```

CLI Comparison

```
Router# show ip route vrf RED
……
Gateway of last resort is not set

O    10.0.6.0/24 [200/0] via 10.10.10.6,.
O    10.1.6.0/24 [200/0] via 10.10.10.6,.
C    10.0.4.0/24 is directly connected, ..


Router# ping vrf RED 10.0.6.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.6.1:
!!!!!
Success rate is 100 percent (5/5),
  round-trip min/avg/max = 176/264/576 ms
```

VRF-lite end-to-end example

```
Router# routing-context vrf RED

Router%RED# show ip route
……
Gateway of last resort is not set

O    10.0.6.0/24 [200/0] via 10.10.10.6,.
O    10.1.6.0/24 [200/0] via 10.10.10.6,.
C    10.0.4.0/24 is directly connected, ..


Router%RED# ping 10.0.6.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.6.1:
!!!!!
Success rate is 100 percent (5/5),
  round-trip min/avg/max = 176/264/576 ms
```

EVN example

Virtual Networking

For more information:

BRKRST-2045 Network Virtualisation Design Concepts over the WAN
Thursday 2:00pm – 4:00pm




Campus Lisp

LISP 101

Locator/ID Separation

A routing protocol that separates routable IP addresses of networking devices from endpoint IP addresses of hosts

- **EID (Endpoint Identifier)** is the IP address of a host – just as it is today
- **RLOC (Routing Locator)** is the IP address of the LISP router for the host
- **EID-to-RLOC mapping** is the distributed architecture that maps **EIDs** to **RLOCs**

LISP 101

Traditional Routing in Today's Campus

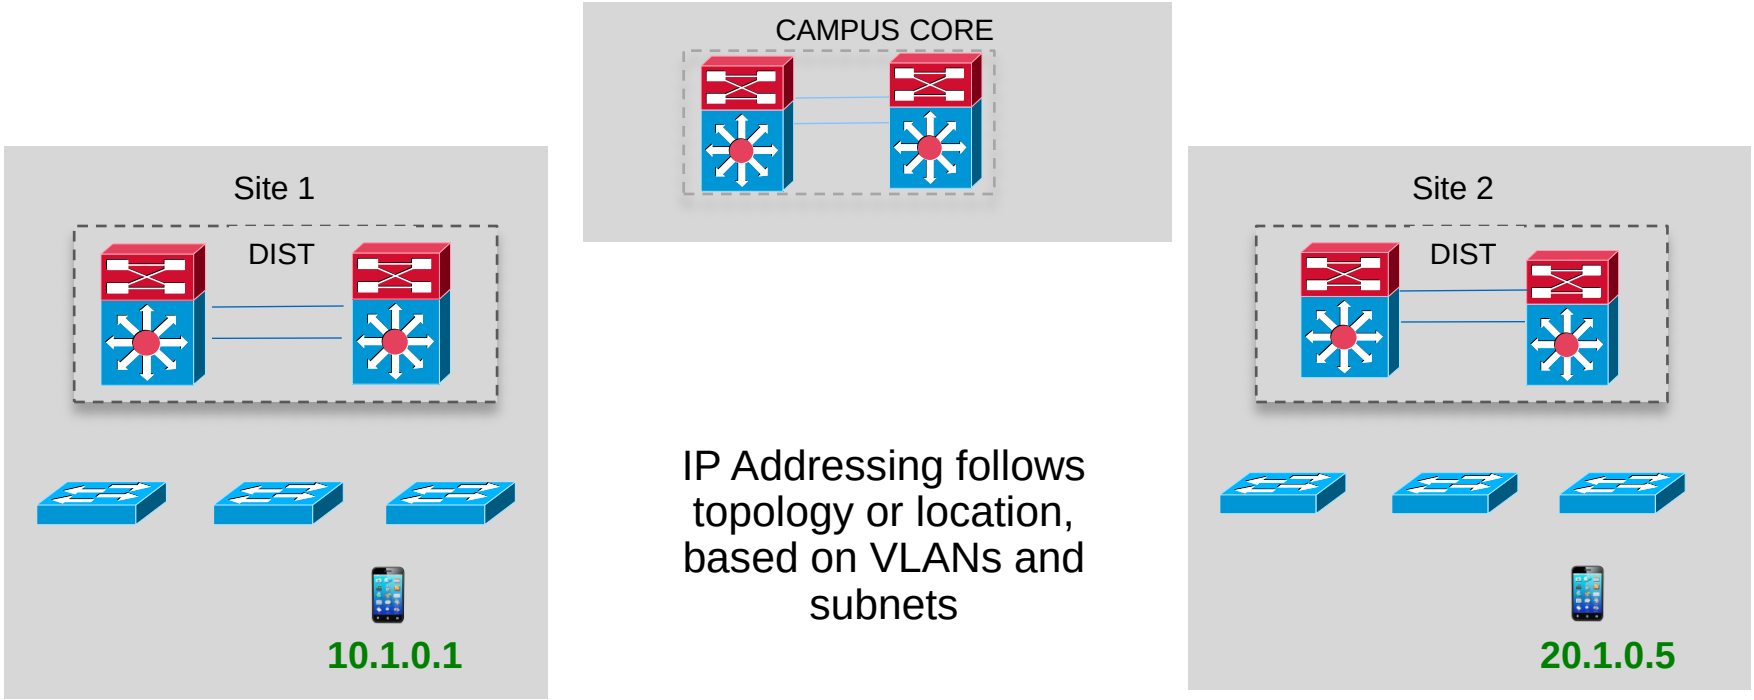

LISP 101

Implementing Locator/ID Separation

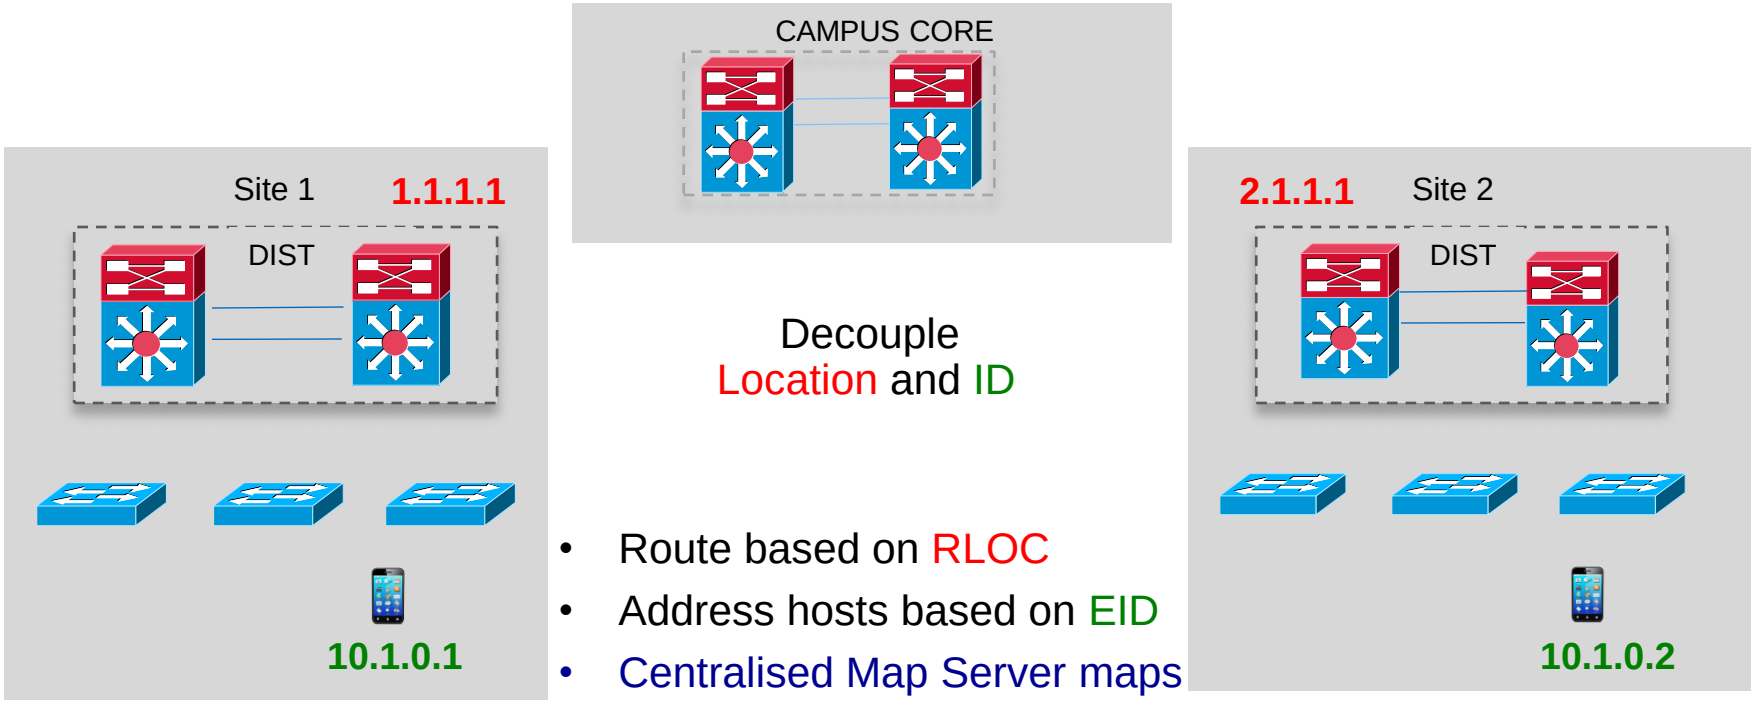

Decouple
Location and ID

- Route based on RLOC
- Address hosts based on EID
- Centralised Map Server maps EID to RLOC

LISP Packet Forwarding

Unicast IPv4 – Data Plane

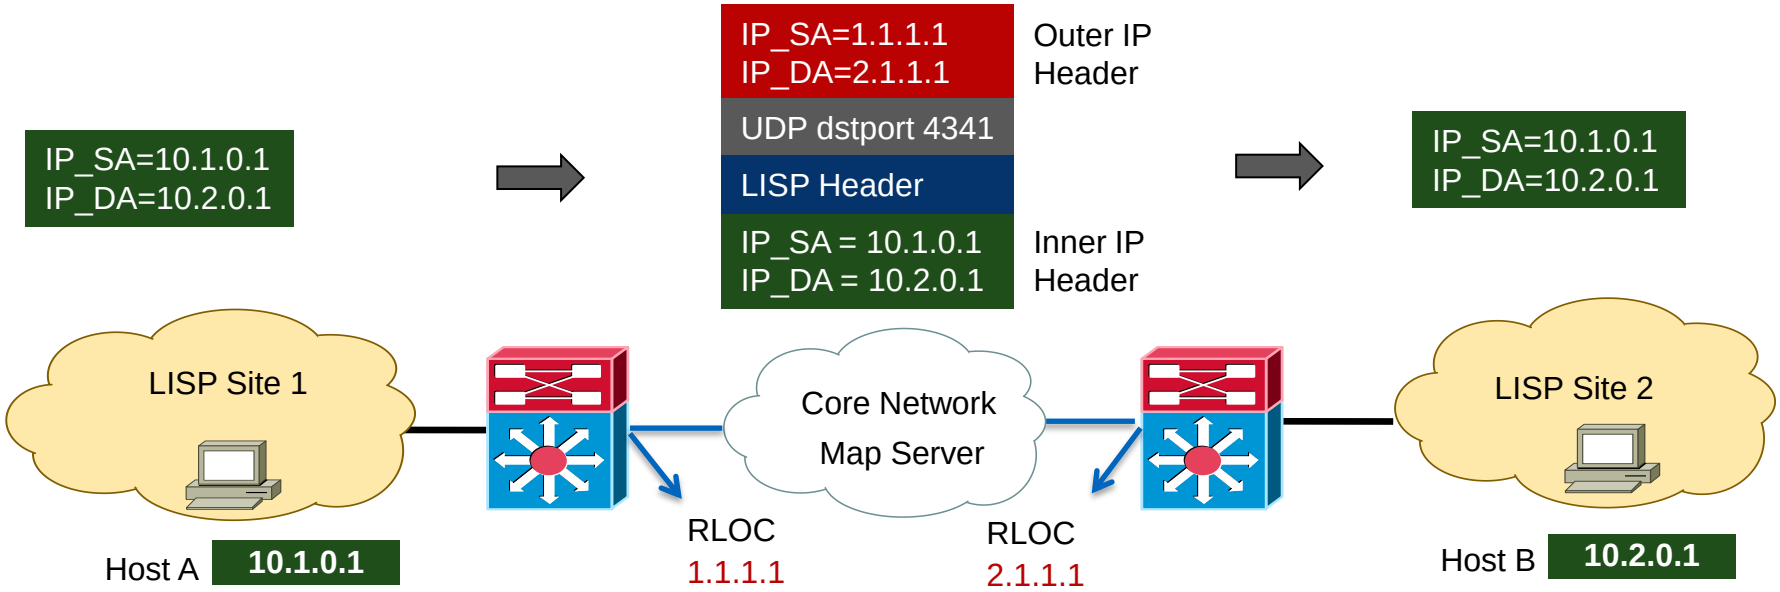

Mapping Database:
EID Prefix 10.1.0.0/24 -> RLOC 1.1.1.1

Map Cache:
10.2.0.0/24 -> RLOC 2.1.1.1

Mapping Database:
EID Prefix 10.2.0.0/24 -> RLOC 2.1.1.1

LISP Packet Forwarding

Unicast IPv4 with Multihoming – Data Plane

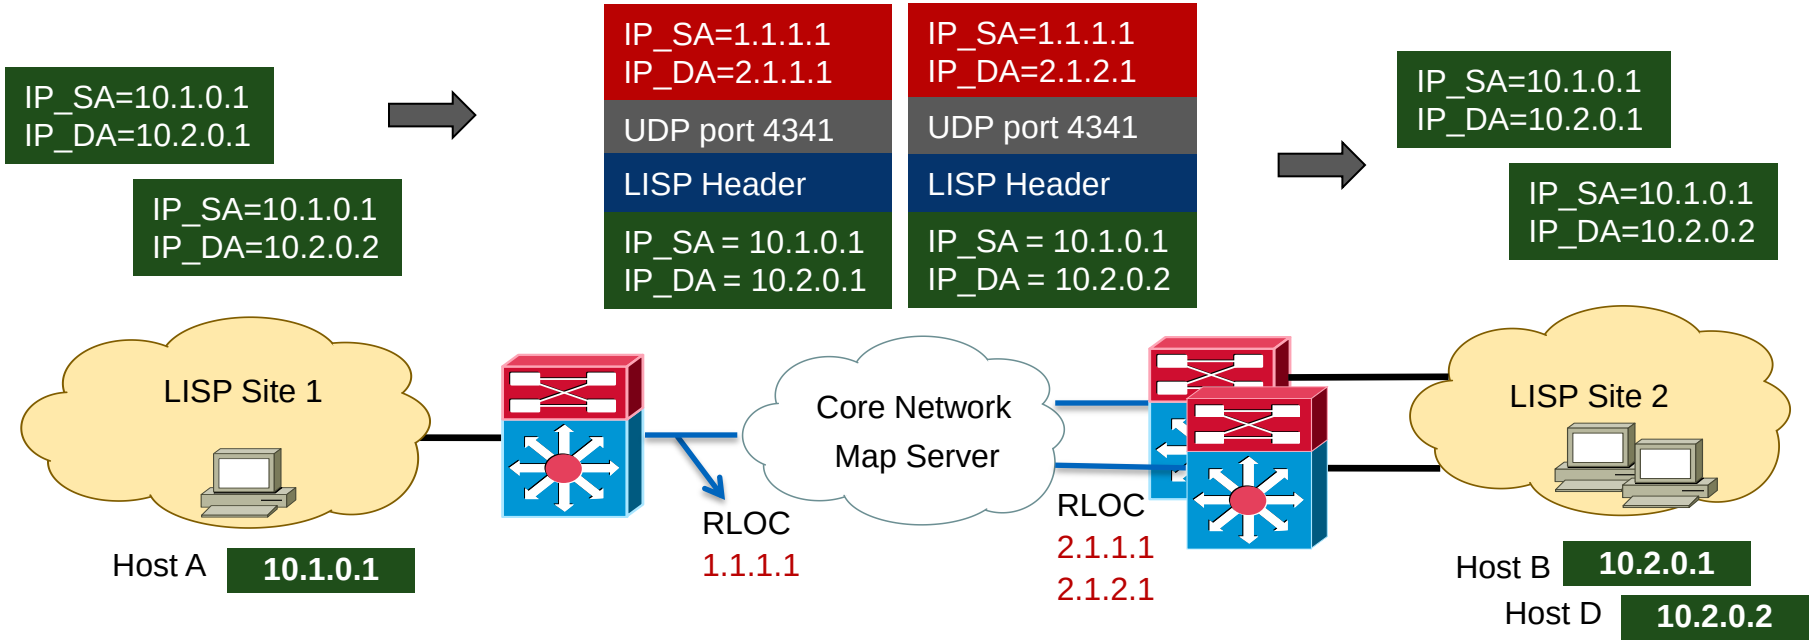

Mapping Database:
EID Prefix 10.1.0.0/24 -> RLOC 1.1.1.1

Map Cache:
10.2.0.0/24 -> RLOC 2.1.1.1 and 2.1.2.1(priority and weight)

Mapping Database:
EID Prefix 10.2.0.0/24 -> RLOC 2.1.1.1 and 2.1.2.1

LISP Features

LISP is a routing architecture, not a feature

- LISP enables IP address portability (using EIDs)
- LISP enables pull versus push routing (using mapping)
- LISP is address-family agnostic (IPv6 deployment)
- LISP has inherent advantages in multihoming and virtualisation
- LISP is an open standard (approved RFC in experimental section)

Campus LISP

For more information:

BRKCRS-3510 LISP in Campus Networks
Thursday 11:00am– 12:30am

Software Defined Networking (SDN)

SDN 101

Since time began, network devices have included a control plane and a data plane.

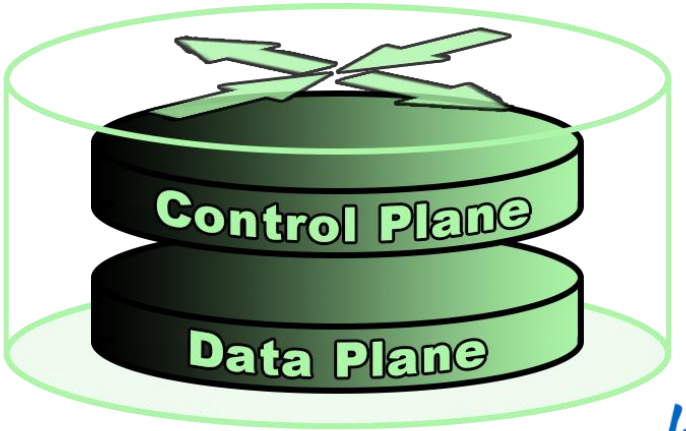

SDN 101

Control Plane:

- Runs on device CPU
- 1000’s of packets per second
- Routing, STP, AAA, syslog, CLI, etc

Data Plane:

- Runs on dedicated HW ASICs
- Millions or billions of pps
- L2 switching, L3 forwarding,
- QoS marking/classification/policing

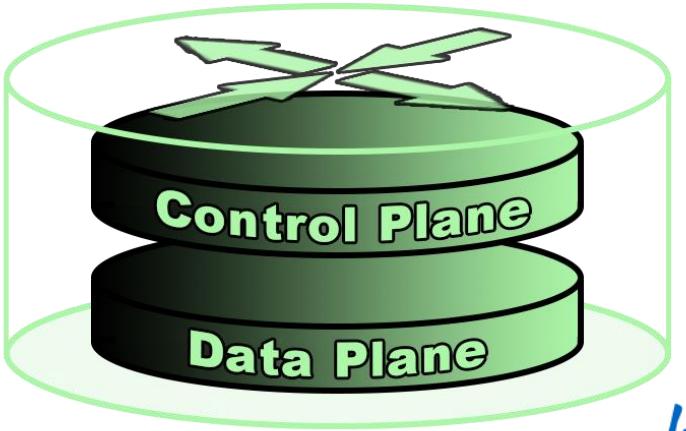

SDN 101

SDN is an approach to building computer networks that separates and abstracts elements of these systems

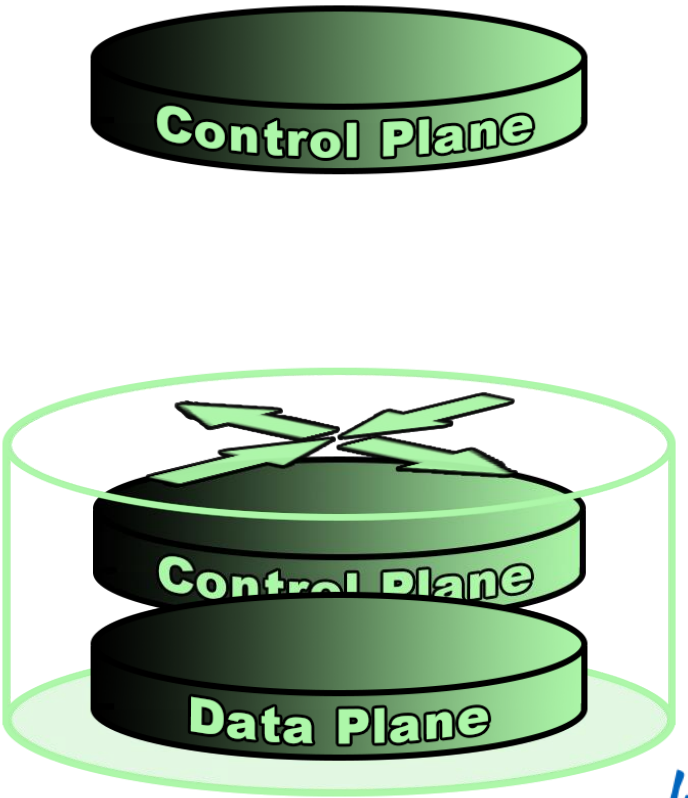

Cisco ONE 101

Open Network Environment – Cisco's Strategy for SDN

Platform APIs

onePK

Comprehensive Developer Kit IOS, IOS-XR and NX-OS

Controllers & Agents

SDN Controller (Nexus 9000, APIC)

OpenFlow Agent

Overlay Virtual Networks

Nexus 1000V

OpenStack REST API

Multi-Hypervisors VXLAN Gateway

Cisco ONE 101

Most companies focusing on SDN in the data centre.

How will SDN affect the campus?

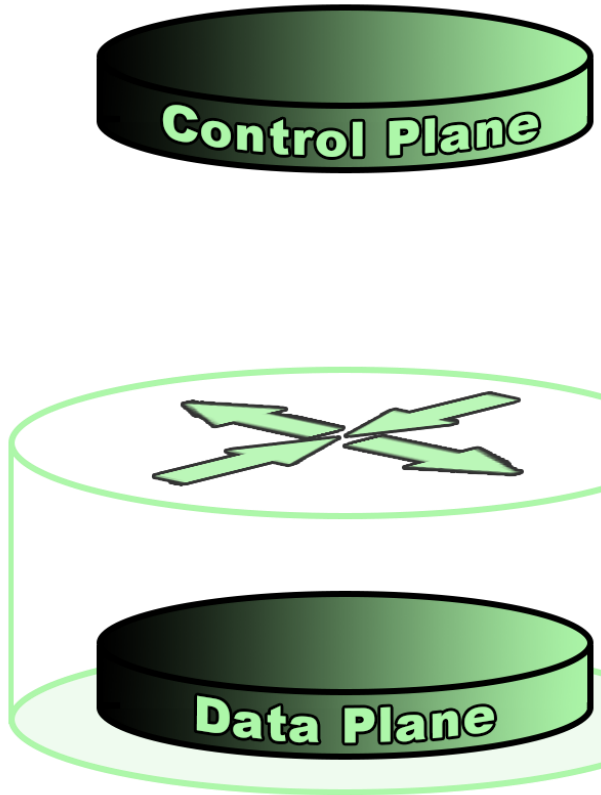

Cisco ONE 101

Cisco ONE 101

DC Module

- Focusing on DC
- Uses Nexus 9000 spine and leaf nodes
- Application centric policy creation

Enterprise Module

- Focusing on WAN and branch
- Catalyst, ISR and ASR devices
- Centralised policy creation for:
  - Security (ACL's, threat detection and mitigation)
  - Network-wide QoS
  - Path optimisation (PfR)

Software Defined Networking

For more information:

BRKAPP-9000 Introduction to Application Centric Infrastructure
Thursday 8:30am – 10:30am

BRKAPP-9001 Policy Driven Data Centre Design
Thursday 11:00am – 12:30pm

BRKDCT-3640 Nexus 9000 Architecture
Friday 2:00pm – 4:00pm

Agenda

- The Principles
- The Basics
- The Cool Stuff
- The End

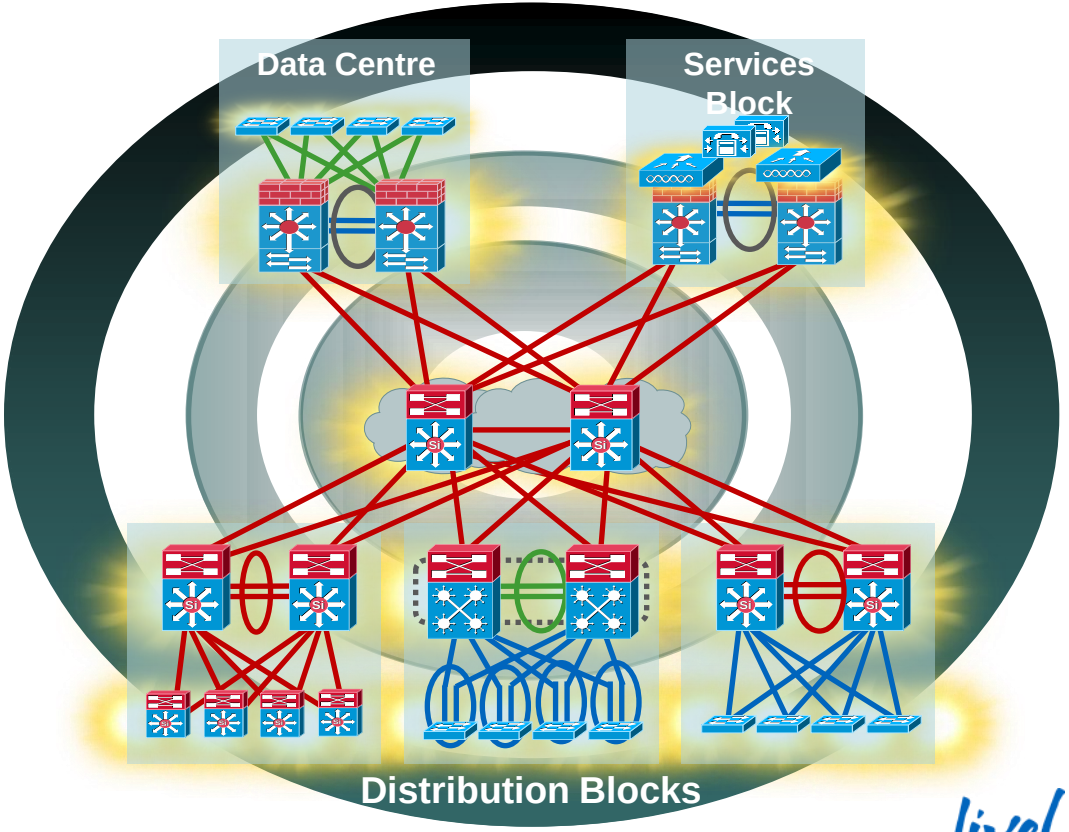

Summary

- Offers hierarchy—each layer has specific role
- Modular topology—building blocks
- Easy to grow, understand, and troubleshoot
- Creates small fault domains— clear demarcations and isolation
- Allows for implementation of new technologies per building block

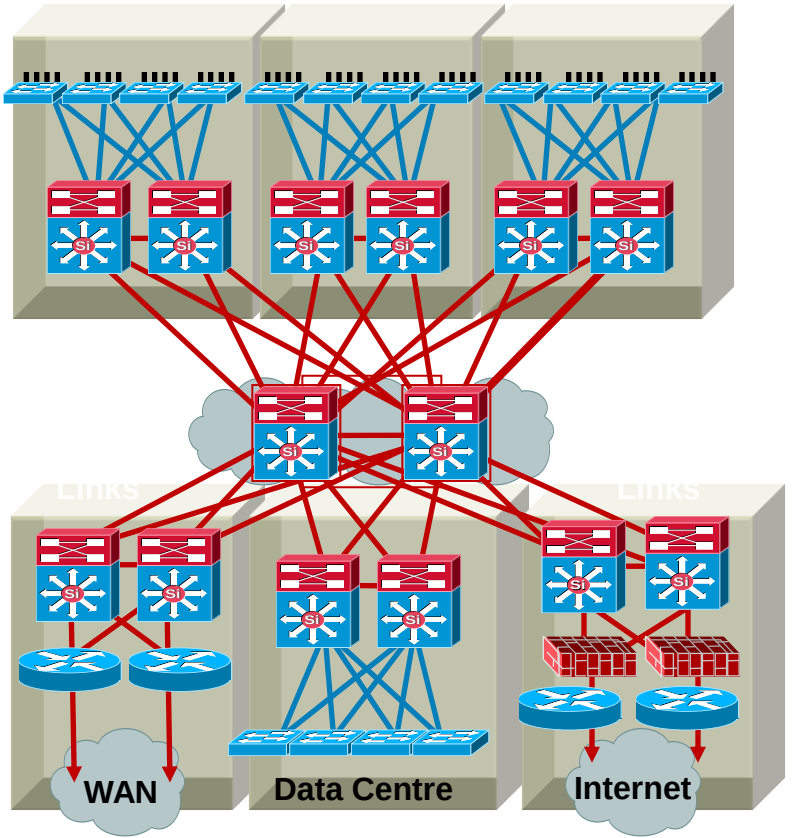

Access
Distribution
Core
Distribution
Access

Hierarchical Network Design

Without a Rock Solid Foundation the Rest Doesn't Matter

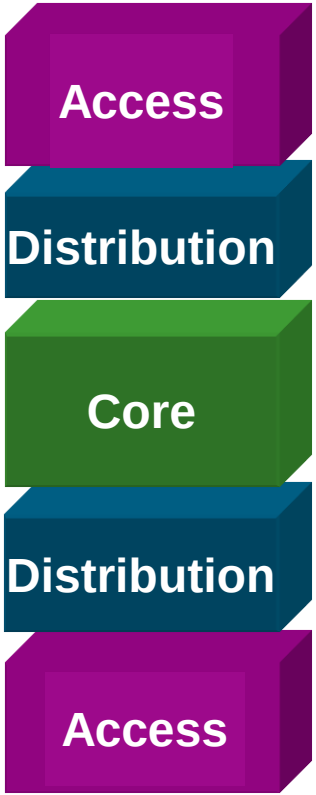

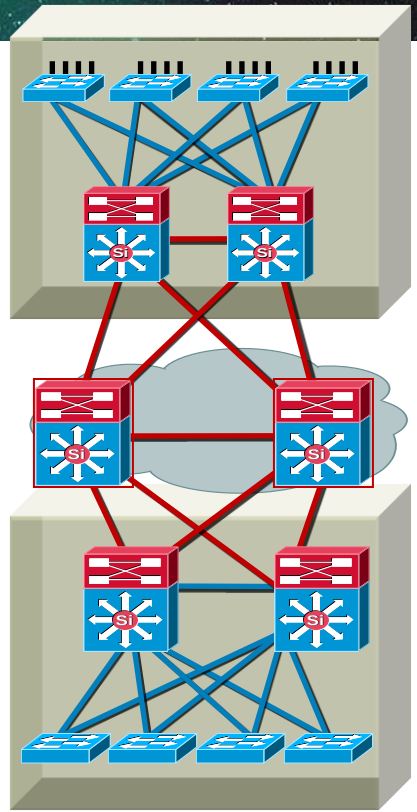




Q & A

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