

Experience Today the
Network of Tomorrow.

Cisco Expo
2009

Enterprise Multilayer and Routed Access Campus Design


Yaman Hakmi
Systems Engineer



Welcome to the Human Network.




- Multilayer Campus Design Principles
- Latest Cisco Campus Networking Portfolio
 - Catalyst 6500
 - Nexus 7000
- Routed Access Campus Design
- Summary



Mitigating the Exposure

Most Common Causes of Downtime



**Common Causes of
Enterprise Network Downtime***

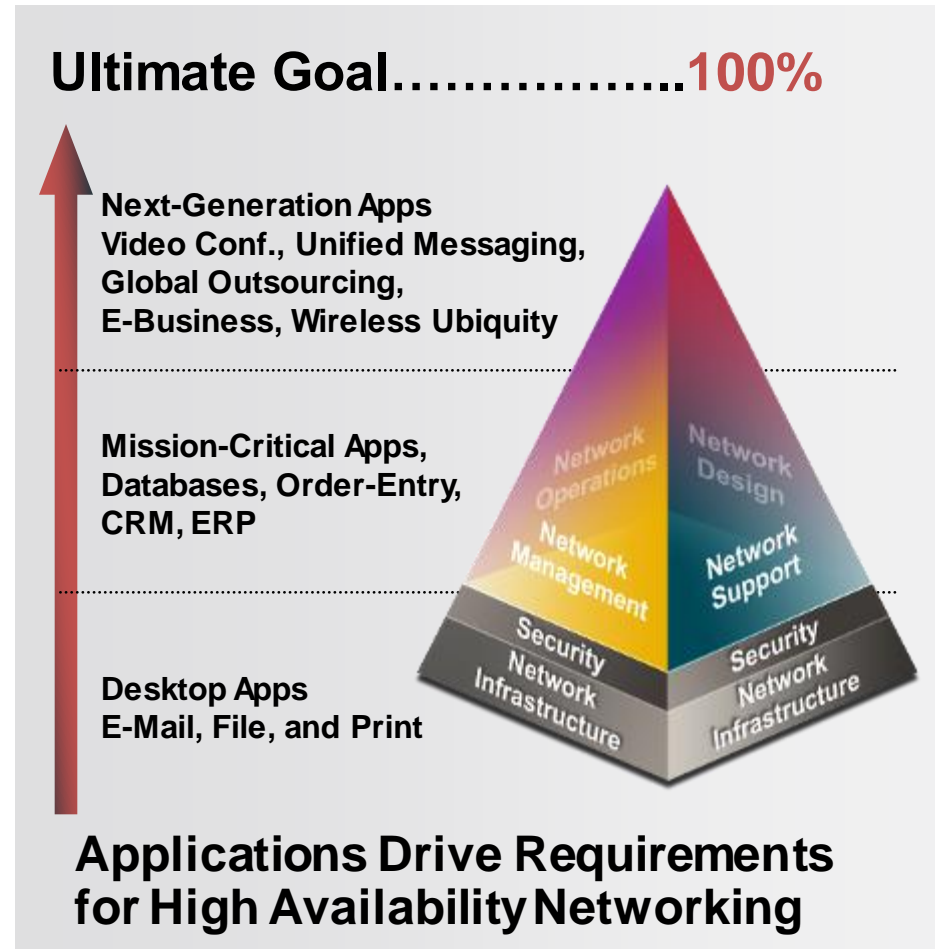
*Source: Yankee Group, The Road to Five-Nines Network

Enterprise Class Availability

Resilient Campus Communication Fabric


Systems Design Approach to High Availability

- VOIP availability is the baseline for the enterprise networks
Human ear notices the difference in voice within 150–200 msec, which translates only ten consecutive packet loss with G711 codec
- Video loss is even more noticeable and it is rapidly becoming new frontier for jitter and delay requirements
- 200 msec end-to-end campus convergence is the design goal



High Availability Campus Design


Structure, Modularity, and Hierarchy



Hierarchical Campus Network


Structure, Modularity and Hierarchy

Not This !!




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
- Promotes load balancing and redundancy
- Promotes deterministic traffic patterns
- Incorporates balance of both Layer 2 and Layer 3 technology, leveraging the strength of both
- Utilizes Layer 3 Routing for load balancing, fast convergence, scalability, and control



Access Layer

Feature Rich Environment


- It's not just about connectivity
- Layer 2/Layer 3 feature rich environment; convergence, HA, security, QoS, IP multicast, etc.
- Intelligent network services: QoS, trust boundary, broadcast suppression, IGMP snooping
- Intelligent network services: PVST+, Rapid PVST+, EIGRP, OSPF, PAgP/LACP, UDLD, FlexLink, etc
- Cisco Catalyst integrated security features (802.1x, CISF): port security, DHCP snooping, DAI, IPSG, etc.
- Automatic phone discovery, conditional trust boundary, power over Ethernet, auxiliary VLAN, etc.
- Spanning tree toolkit: Portfast, UplinkFast, BackboneFast, LoopGuard, BPDUGuard, 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 summarization, 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




Do I Need a Core Layer?

It's really a question of:
Scale, Complexity, and Convergence

No Core

- Fully meshed distribution layers
- Physical cabling requirement
- Routing complexity




Do I Need a Core Layer?

It's Really a Question of
Scale, Complexity, and Convergence

No Core

- Fully meshed distribution layers
- Physical cabling requirement
- Routing complexity

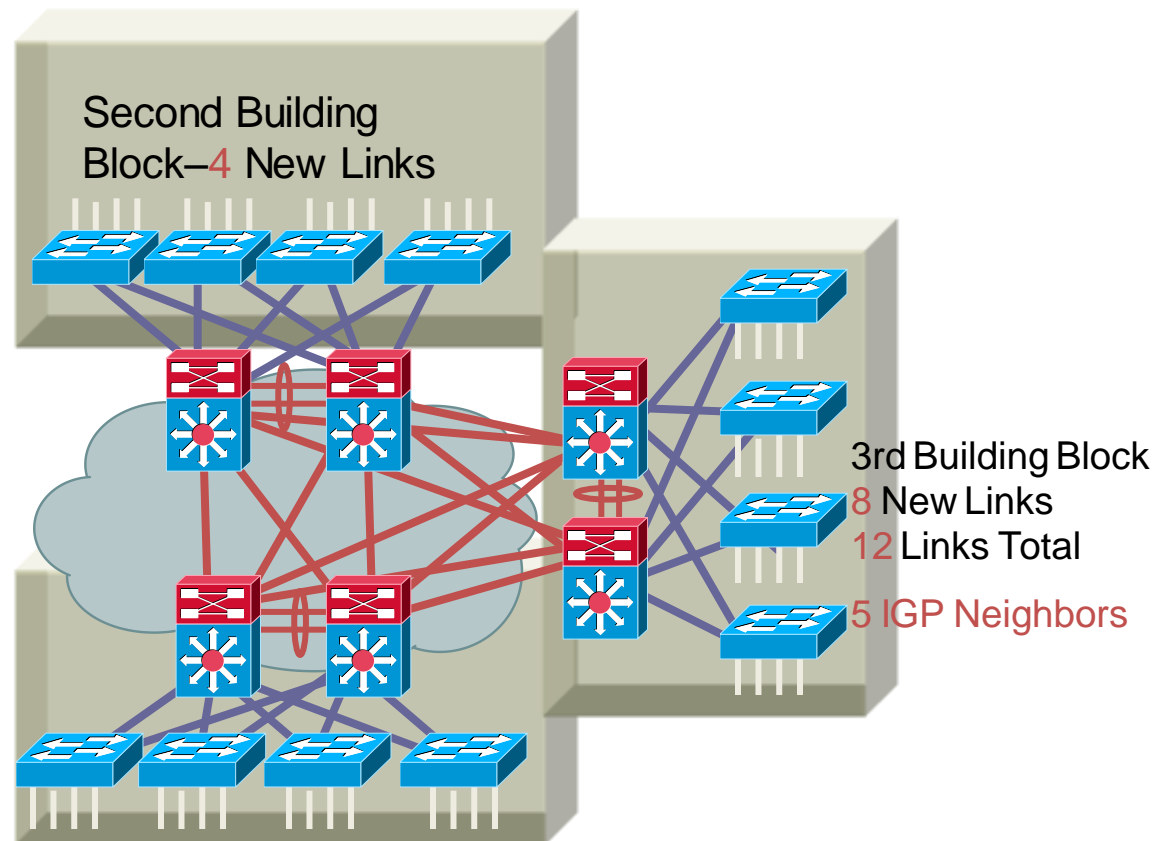


Do I Need a Core Layer?

It's Really a Question of
Scale, Complexity, and Convergence

No Core

- Fully meshed distribution layers
- Physical cabling requirement
- Routing complexity

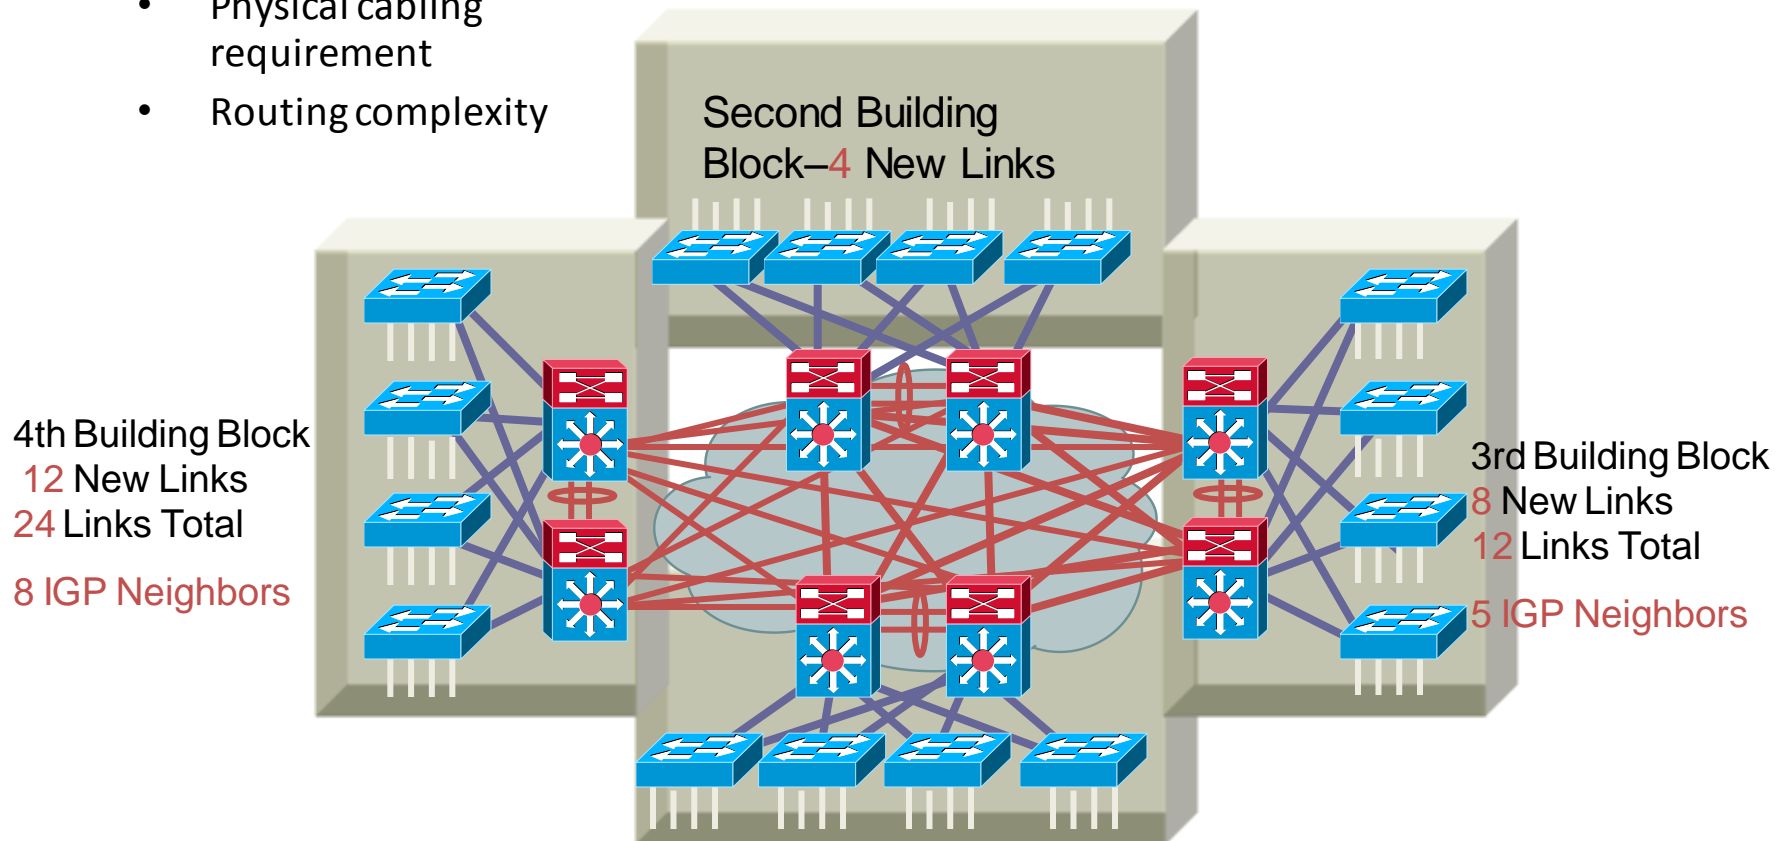


Do I Need a Core Layer?

It's Really a Question of
Scale, Complexity, and Convergence

No Core

- Fully meshed distribution layers
- Physical cabling requirement
- Routing complexity




Do I Need a Core Layer?


It's Really a Question of
Scale, Complexity, and Convergence

Dedicated Core Switches

- Easier to add a module
- Fewer links in the core
- Easier bandwidth upgrade
- Routing protocol peering reduced
- Equal cost Layer 3 links for best convergence

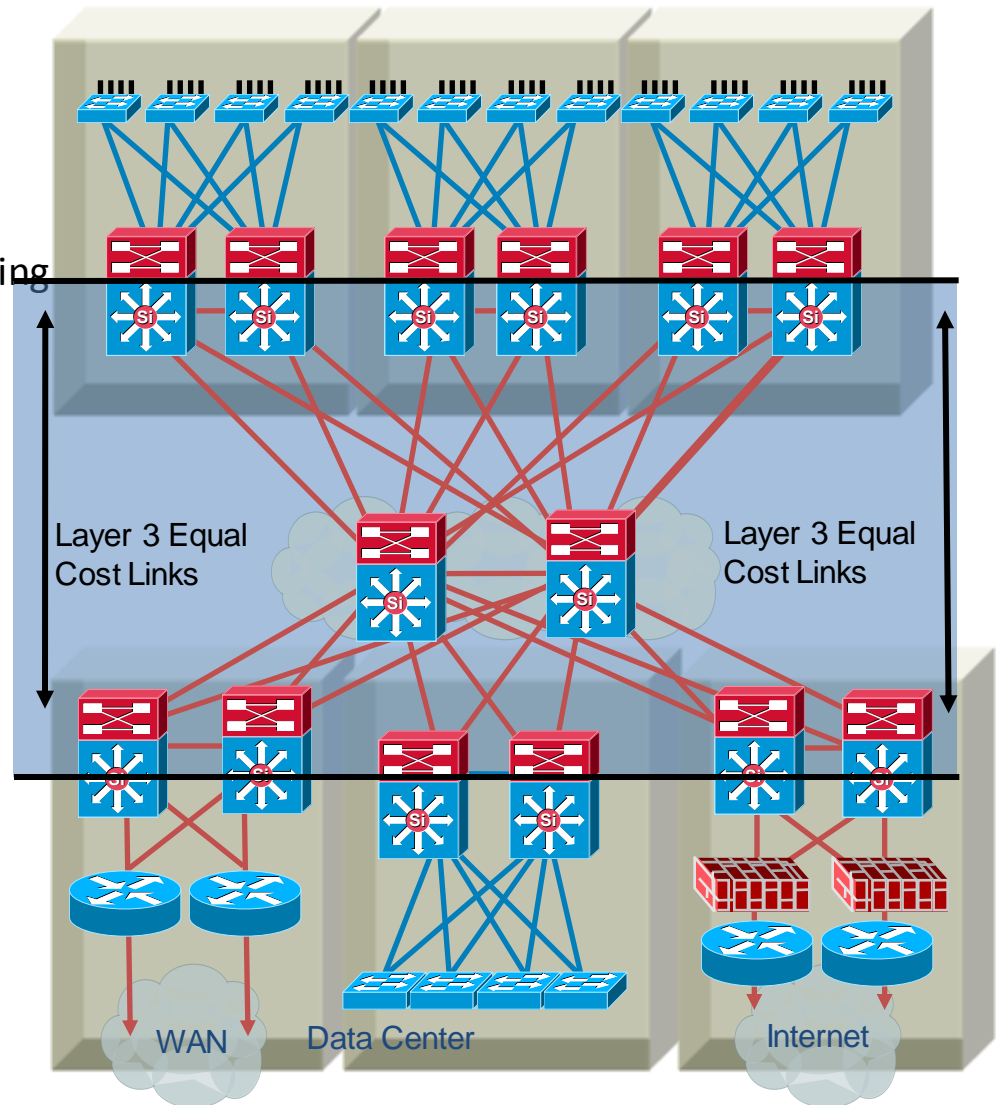


Design Alternatives come within a Building (or Distribution) Block



Core Layer: L3 Routing


- Typically deployed in distribution to core, and core to core interconnections
- Used to quickly re-route around failed node/links while providing load balancing over redundant paths
- Build triangles not squares for deterministic convergence
- Only peer on links that you intend to use as transit
- Insure redundant L3 paths to avoid black holes
- Summarize distribution to core to limit EIGRP query diameter or OSPF LSA propagation
- Tune CEF L3/L4 load balancing hash to achieve maximum utilization of equal cost paths (CEF polarization)



Layer 3 Distribution Interconnection

Reference Design—No VLANs Span Access Layer


- Tune CEF load balancing
- Match IOS EtherChannel settings and tune load balancing
- Summarize routes towards core
- Limit redundant IGP peering
- STP Root and HSRP primary tuning or GLBP to load balance on uplinks
- Set trunk mode on/negotiate
- Disable EtherChannel unless needed
- Set Port Host on access layer ports:
 - Disable Trunking
 - Disable EtherChannel
 - Enable PortFast
- RootGuard or BPDU-Guard
- Use security features



Layer 2 Distribution Interconnection

Some VLANs Span Access Layer


- Tune CEF load balancing
- Match IOS EtherChannel settings and tune load balancing
- Summarize routes towards core
- Limit redundant IGP peering
- STP Root and HSRP primary or GLBP and STP port cost tuning to load balance on uplinks
- Set trunk mode on/nonegotiate
- Disable EtherChannel unless needed
- RootGuard on downlinks
- LoopGuard on uplinks
- Set port host on access Layer ports:
 - Disable trunking
 - Disable EtherChannel
 - Enable PortFast
- RootGuard or BPDU-Guard
- Use security features



Multilayer Network Design

Many Moving Parts

- Evolved due to historical pressures
 - Speed of routing vs. switching
 - Flexibility with spanning VLANs
 - Nonroutable protocols
- Well understood optimization of interaction between the various control protocols and the topology
 - STP root and HSRP primary tuning to load balance on uplinks
 - Spanning tree toolkit (RootGuard, LoopGuard...)
- Many moving parts increase a chance of instability



Multilayer Network Design

Good Solid Design Option, But...

- Utilizes multiple Control Protocols

Spanning Tree (802.1w, ...), FHRP (HSRP, ...), Routing Protocol (EIGRP, ...)

- Convergence is dependent on multiple factors


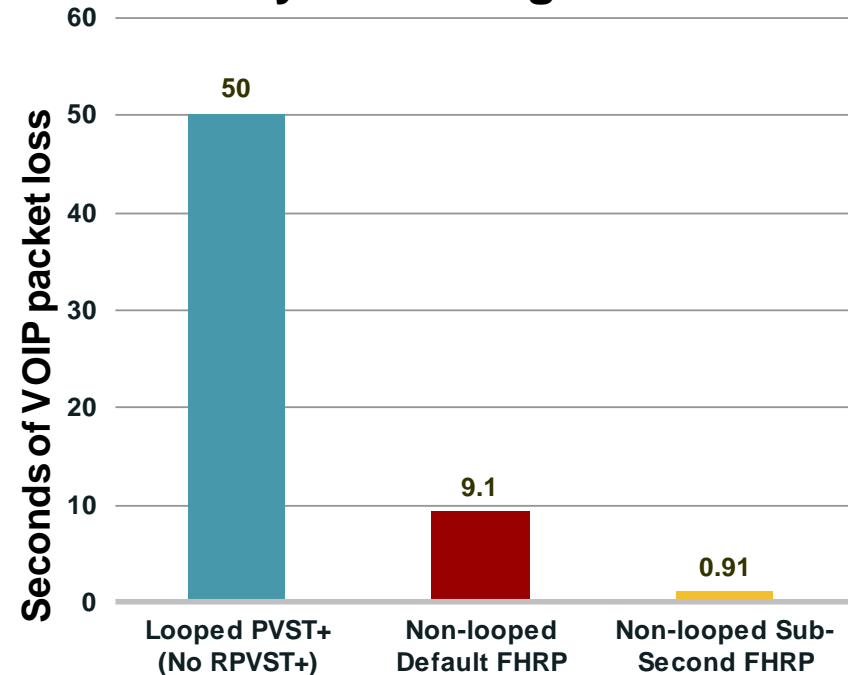
FHRP—900msec to 9 seconds

Spanning Tree—Upto 50 seconds

Poor load balancing—single uplink, asymmetric routing etc

- STP, if it breaks badly, no inherent mechanism to stop the loop

Multi-Layer Convergence



Latest Campus Networking Technologies



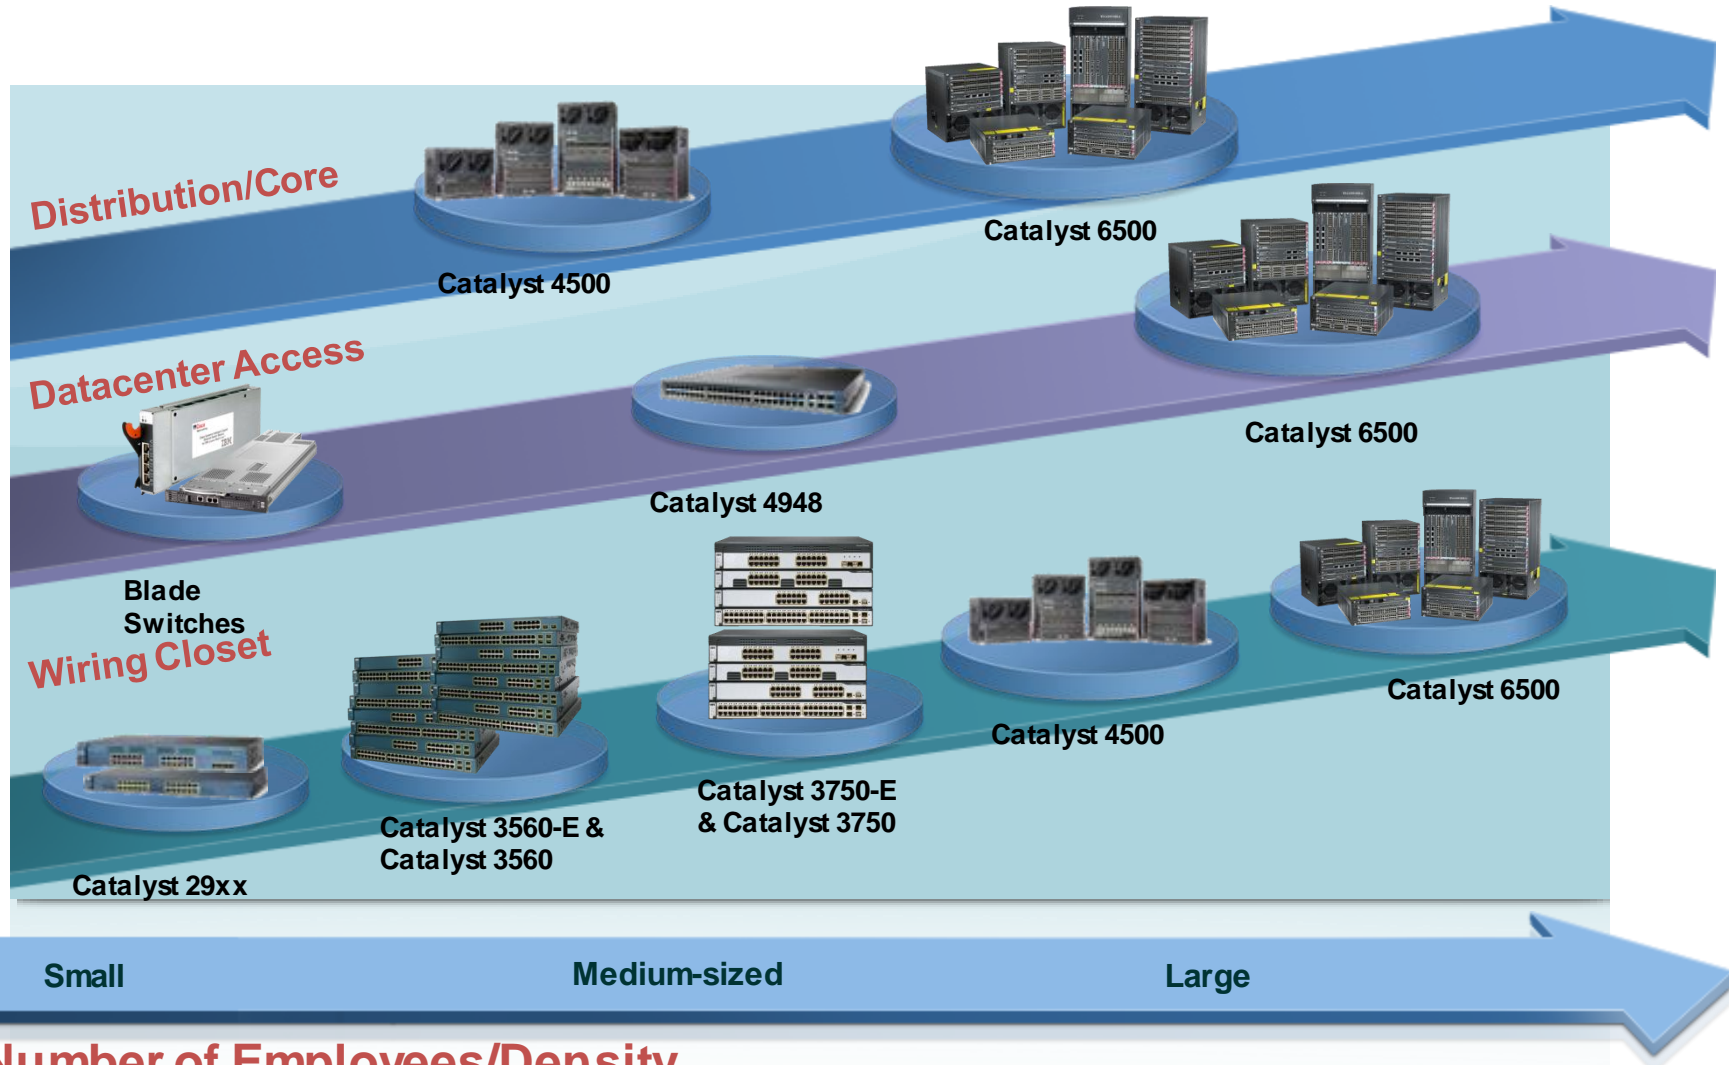
Cisco Expo
2009

Welcome to the Human Network.



Catalyst Switching Portfolio

Features, Scalability, Longevity



Number of Employees/Density

Catalyst 6500 Linecards

Latest Core, Distribution, and Data Center Portfolio



WS-X6708-10GE

8 Port 10GE X2 (2x20G Fabric Enabled Card)

ER, LR, LX4, SR, CX4 optics; DFC included; Queuing: TX - 1p7q4t, RX - 8q8t; Jumbo frame support: 200MB/port Buffer



WS-X6704-10GE

4 Port 10GE XENPAK (2x20G Fabric Enabled Card)

ER, LR, LX4, SR, CX4, ZR optics; Optional DFC; Queuing: TX - 1p7q4t, RX - 1q8t or 8q8t (w/DFC); Jumbo frame support



WS-X6748-SFP

48 Port GE SFP (2x20G Fabric Enabled Card)

SX, LX, ZX, Tx, CWDM SFPs; Optional DFC; Queuing: TX - 1p3q8t, RX - 1q8t or 2q8t (w/DFC); Jumbo frame support



WS-X6724-SFP

24 Port GE SFP (1x20G Fabric Enabled Card)

SX, LX, ZX, Tx, CWDM SFPs; Optional DFC; Queuing: TX - 1p3q8t, RX - 1q8t or 2q8t (w/DFC); Jumbo frame support



WS-X6748-GE-TX

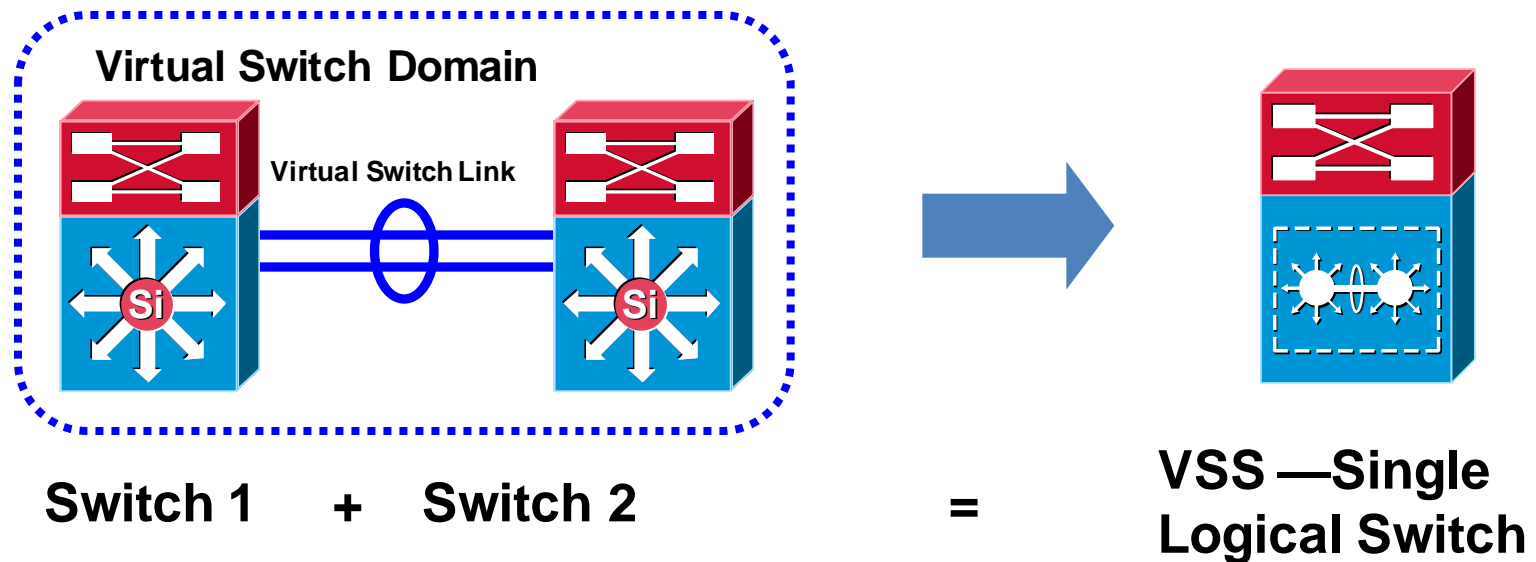
48 Port 10/100/1000 (2x20G Fabric Enabled Card)

Supports TDR; Optional Distributed Forwarding Card (DFC); Queuing: TX - 1p3q8t, RX - 1q8t; Jumbo frame support

Virtual Switch

Virtual Switching System 1440 (VSS)

- Virtual Switching System consists of two Cisco Catalyst 6500 Series defined as members of the same virtual switch domain
- Single control plane with dual active forwarding planes
- Design to increase forwarding capacity while increasing availability by eliminating STP loops
- Reduced operational complexity by simplifying configuration




Virtual Switching System

Hardware and Software Requirements

- **Software Support**
Native and modular Cisco IOS are supported
Minimum IOS required is 12.2(33)SXH1, however current recommendation is 12.2(33)SXH2(a)
- **Supervisor—VS-S720-10G-3C/XL**
PFC3C/XL contains new hardware support to forward traffic across multiple physical chassis and lookup enhancements
- **Virtual switch link**
VS header encapsulation requires new port ASIC
VS-S720-10G-3C/XL Supervisor 10G port or WS-X6708-10G-3C/XL
10 Gigabit Ethernet only
- **WS-X6716: 16-port 10Gbps line card**

Catalyst 6500 16p 10GBASE-T line card



VS-S720-10G-3C/XL




WS-X6708-10G-3C/XL


Virtual Switching System

Multichassis EtherChannel (MEC)

- MEC is an advanced EtherChannel technology extending link aggregation to two separate physical switches
- MEC enables the VSS to appear as a single logical device to devices connected to VSS, thus significantly simplifying campus topology
- Traditionally spanning VLANs over multiple closets would create STP looped topology, MEC with VSS eliminates these loops in the campus topology
- MEC replaces spanning tree as the means to provide link redundancy and thus doubling bandwidth available from access
- MEC is supported only with VSS



Multichassis EtherChannel




BW Capacity in Non-MEC and MEC Topology

Cisco Nexus 7000


- Nexus Family: Increasingly seen in the enterprise campus network
- High density
 - 256 10G interfaces per system
 - 384 1G interfaces per systems
- High performance
 - 64 non-blocking 10G ports
 - 1.2Tbps system bandwidth at initial release
 - 80Gbps per slot
- Future proof
 - Initial fabric provides up to 4.1Tbps
 - Product family scaleable to 15+Tbps
- **Nexus 7018: Second chassis in Nexus 7000 family**
- Ultra-high density
 - 512 10G interfaces per system
 - 768 1G interfaces per system
- High performance
 - 128 non-blocking 10G ports
 - 2.5Tbps system bandwidth at initial release; 80G/slot
- Future proof
 - Initial fabric provides up to 7.8Tbps
 - Chassis scaleable to 17.6Tbps



Recall: Campus Network Design




Campus Network Design with the Nexus 7000



Campus Network Design with the Nexus 7000

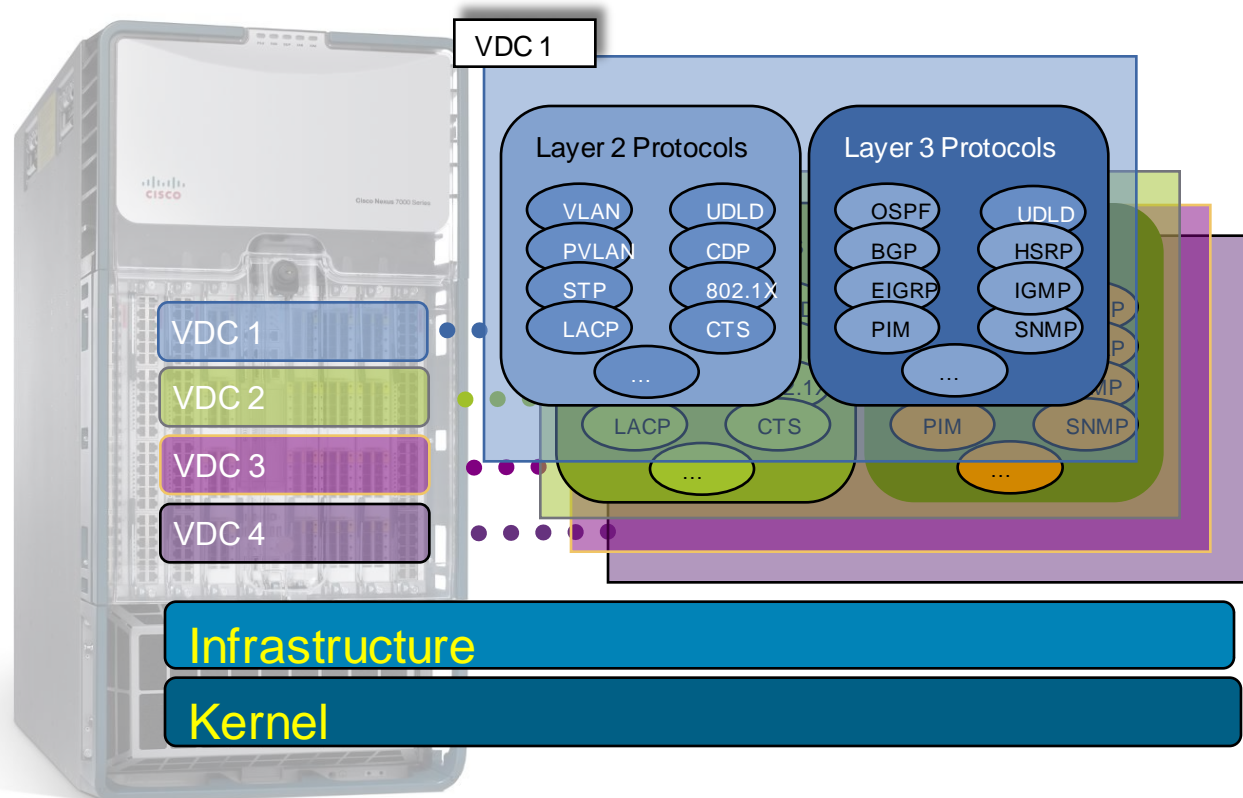
Drive for Nexus 7000 in the enterprise core:

- Higher 10 Gbps port density
- Increased backbone switching and routing capacity
- Enhanced network resiliency with SSO/ISSU




Virtual Device Contexts

Virtual Device Contexts provides virtualization at the device level allowing multiple instances of the device to operate on the same physical switch at the same time.



Network Segmentation


- 1 to Many: One network supports many virtual networks



Virtual Port-Channel (vPC)


Feature Overview

- Allow a single device to use a port channel across two upstream switches
- Eliminate STP blocked ports
- Uses all available uplink bandwidth
- Dual-homed server operate in active-active mode
- Provide fast convergence upon link/device failure
- Reduce CAPEX and OPEX
- Available in NX-OS 4.1 with current and future hardware (10G card only)





How does vPC help with STP?

- Before vPC
 - STP blocks redundant uplinks
 - VLAN based load balancing
 - Re-convergence relies on STP
 - Protocol Failure → 💣
- With vPC
 - No blocked uplinks
 - Lower oversubscription
 - EtherChannel load balancing (hash)
 - Convergence sub-second
 - Reduced STP logical port count



Virtual Port Channels

Multi-Chassis Etherchannel (MEC)



Virtual Switching System (VSS)

Virtual Port Channel (vPC)

- Both VSS-MEC and vPC are a Port-channeling concept extending link aggregation to two separate physical switches
- Allows the creation of resilient L2 topologies based on Link Aggregation.
- Eliminates the dependence on STP in the L2 access-distribution Layer
- Enable seamless VM Mobility, Server HA Clusters
- Scale Available Layer 2 Bandwidth
- Simplify Network Design

Routed Access Campus Design




Cisco Expo
2009

Welcome to the Human Network.



Routed Access

Layer 3 Distribution with Layer 3 Access




- Move the Layer 2/3 demarcation to the network edge
- Upstream convergence times triggered by hardware detection of light lost from upstream neighbor
- Beneficial for the right environment

Routed Access

Advantages, Yes in the Right Environment

- EIGRP converges in **< 200 msec**
- OSPF with subsecond tuning converges in **< 200 msec**
- Ease of implementation, less to get right
 - No matching of STP/HSRP/GLBP priority
 - No L2/L3 multicast topology inconsistencies
- Single control plane and well known tool set
 - traceroute, show ip route, show ip eigrp neighbor, etc.
- Convergence times dependent on GLBP/HSRP tuning


Both L2 and L3 Can Provide Subsecond Convergence



Routed Access

Simplified Network Recovery

- Routed access network recovery is dependent on L3 reroute
- Time to restore upstream traffic flows is based on ECMP (Equal Cost Multi-Path) reroute
 - Time to detect link failure
 - Process the removal of the lost routes from the software CEF (Cisco® Express Forwarding)
 - Update the hardware CEF table
- Time to restore downstream flows is based on a full routing protocol reroute
 - Time to detect link failure
 - Time to determine new route
 - Process the update of the software RIB and FIB
 - Update the hardware FIB
- Except uplink failure, all other fault recovery is ECMP-based i.e., consistent and predictable

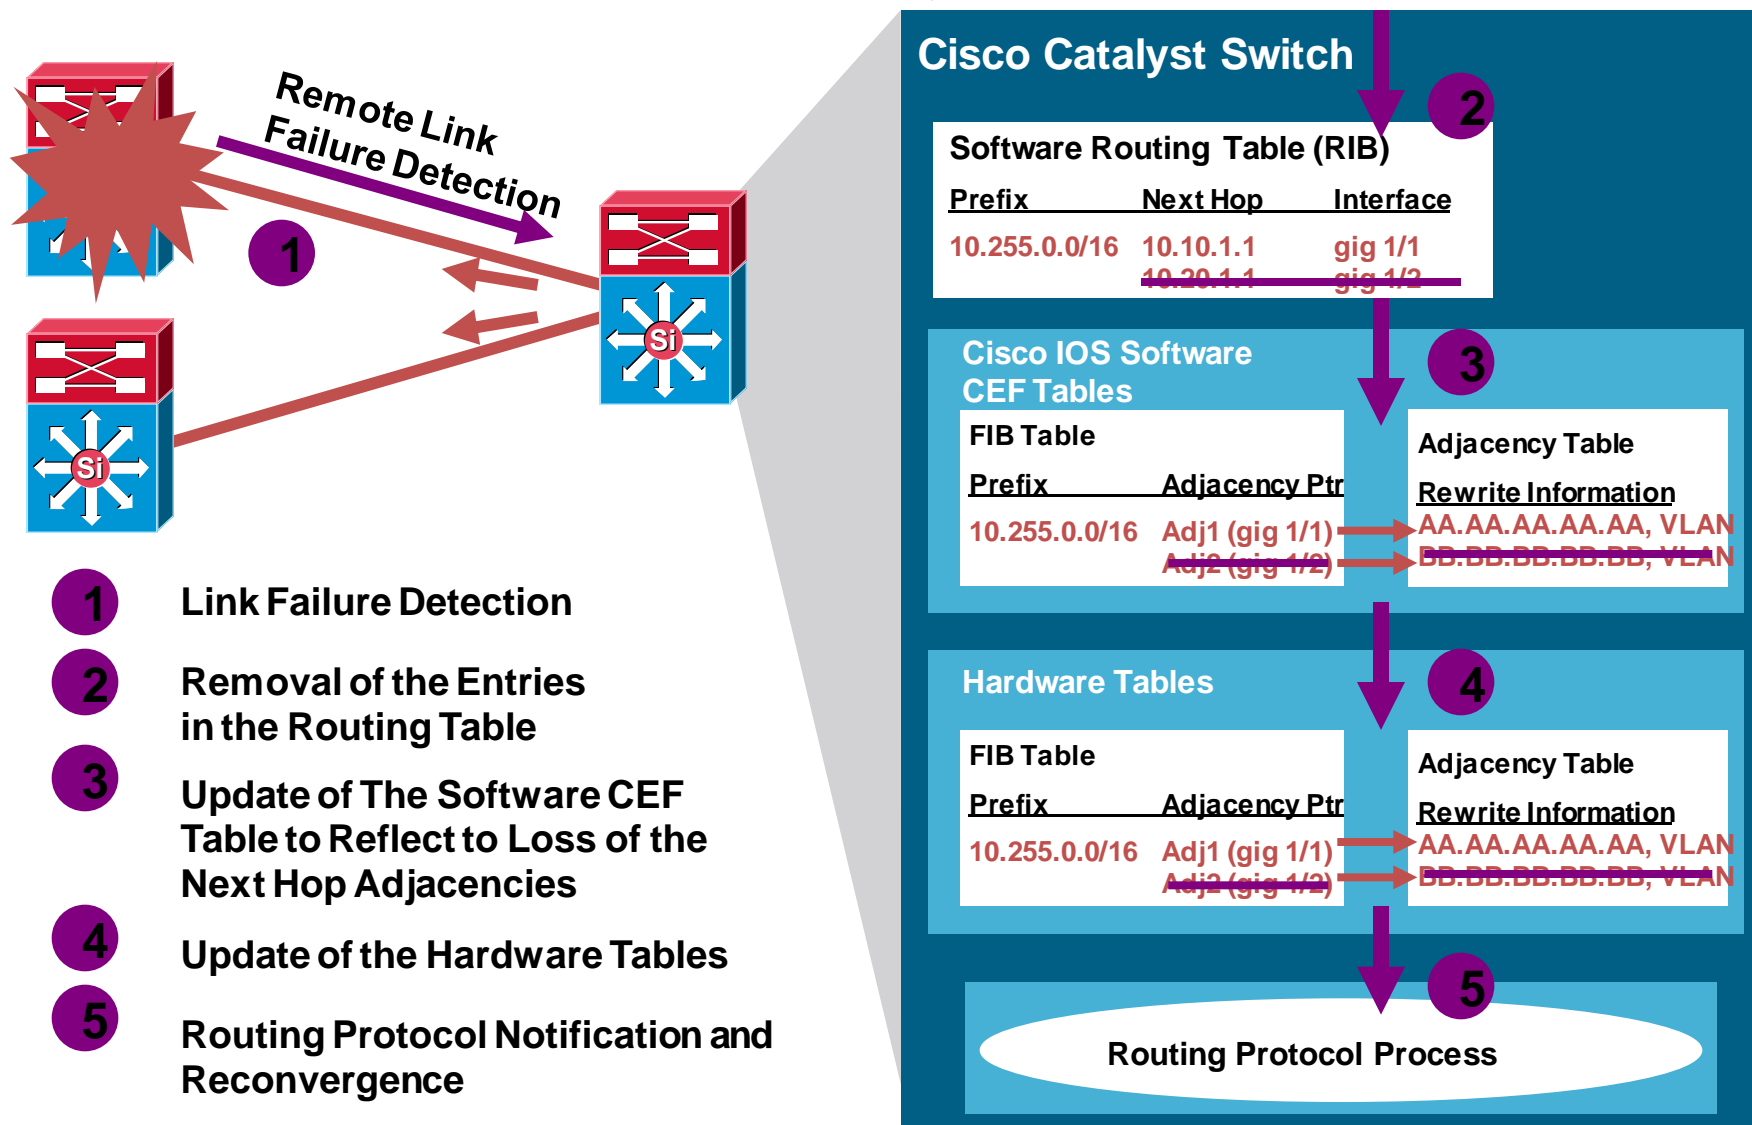


Upstream: ECMP Recovery

Downstream: Routing Protocol Recovery

Routed Access


Time to Recovery CEF Paths




Equal Cost Multipath

Optimizing CEF Load-Sharing

- CEF load-sharing options influence the traffic utilization of Equal Cost Multi-Path (ECMP) links
- Criteria for optimizing CEF load-sharing options would depend on IP address plan and application flows (ports, location, etc.)



CEF Polarization




Unequal Load-Sharing

- CEF polarization and unequal load-sharing are two different possible side effects when using ECMP
- CEF polarization can only happen with multiple stages of CEF load balancing decisions
- If each stage uses the same decision criteria some links can be totally unused

EIGRP Design Rules for Routed Access

Leverage the Tools Provided

- The greatest advantages of EIGRP are gained when the network has a structured addressing plan that allows for use of summarization and stub routers
- EIGRP provides the ability to implement multiple tiers of summarization and route filtering
- Relatively painless to migrate to a L3 access with EIGRP if network addressing scheme permits
- Able to maintain a deterministic convergence time in very large L3 topology



EIGRP Design Rules for HA Campus

High-Speed Campus Convergence


- EIGRP convergence is largely dependent on query response times
- Minimize the number and time for query response to speed up convergence
- Summarize distribution block routes upstream to the core
- Configure all access switches as EIGRP stub routers
- Filter routes sent down to access switches

```
interface TenGigabitEthernet < Uplinks to the core >  
  ip summary-address eigrp 100 10.120.0.0 255.255.0.0 5  
<No summary routes on links between distributions>
```

```
router eigrp 100  
  network 10.0.0.0  
  distribute-list Default out <downstream links>
```

```
ip access-list standard Default  
  permit 0.0.0.0
```


```
router eigrp 100  
  network 10.0.0.0  
  eigrp stub connected
```



OSPF Design Rules for Routed Access

Where Are the Areas?

- Area size/border is bounded by the same concerns in the campus as the WAN
- In campus the lower number of nodes and stability of local links could allow you to build larger areas
- Area design also based on address summarization
- Area boundaries should define buffers between fault domains
- Keep area 0 for core infrastructure and do not extend to the access routers




Utilize Totally Stubby Areas

Reduce SPF and LSA Load in Distribution Area

- ABR for a regular area forwards
 - Summary LSAs (Type 3)
 - ASBR summary (Type 4)
 - Specific externals (Type 5)
- Stub area ABR forwards
 - Summary LSAs (Type 3)
 - Summary default (0.0.0.0)
- A totally stubby area ABR forwards
 - Summary default (0.0.0.0)

```
router ospf 100
  area 120 stub no-summary
  area 120 range 10.120.0.0 255.255.0.0 cost 10
  network 10.120.0.0 0.0.255.255 area 120
  network 10.122.0.0 0.0.255.255 area 0
```




Summarization Distribution to Core

Reduce SPF and LSA Load in Area 0

- Summarize routes from the distribution block upstream into the core
- Minimize the number of LSAs and routes in the core
- Reduce the need for SPF calculations due to internal distribution block changes
- Incremental SPF (iSPF) is a mechanism to reduce the computational load of larger OSPF areas but is more applicable to WAN than campus environments


```
router ospf 100
 area 120 stub no-summary
 area 120 range 10.120.0.0 255.255.0.0 cost 10
 network 10.120.0.0 0.0.255.255 area 120
 network 10.122.0.0 0.0.255.255 area 0
```



OSPF Design Rules for HA Campus

OSPF SPF and LSA Throttling

- OSPF has an SPF throttling timer designed to dampen route recalculation
 - 12.2S OSPF enhancements let us tune this timer to milliseconds; prior to 12.2S one second was the minimum
 - After a failure, the router waits for the SPF timer to expire before recalculating a new route
-
- By default, there is a 500 ms delay before generating router and network LSAs; the wait is used to collect changes during a convergence event and minimize the number of LSAs sent
 - Propagation of a new instance of the LSA is limited at the originator
 - Acceptance of a new LSAs is limited by the receiver
 - Make sure lsa-arrival < lsa-hold





```
timers throttle spf 10 100 5000
timers throttle lsa all 10 100 5000
timers lsa arrival 80
```

```
timers throttle spf 10 100 5000
timers throttle lsa all 10 100 5000
timers lsa arrival 80
```


Convergence

ECMP Convergence Is Dependent on Number of Routes

- Time to update of switch HW FIB is linearly dependent on the number of entries (routes) to be updated
- Summarization will serve to decrease RP load as well as speed up convergence



Routed Access Design Considerations

Design Requirements

- VLANs are localized to a single wiring closet switch
- IP addressing—do you have an address allocation plan to support a routed access design?
- Platform requirements
 - Requires a Cisco Catalyst® 3550 Series, Cisco Catalyst 3560 Series or above
 - Requires Cisco IOS® (native or hybrid)
 - Cisco Catalyst 6500 Series requires a Supervisor with an MSFC
 - Cisco Catalyst IOS feature set considerations
 - IP base feature set for EIGRP-Stub and PIM
 - IP services feature set for OSPF and PIM
 - Cisco Catalyst 3000 Series require 12.2(37)SE for PIM stub in IP base



Routed Access Design Considerations

Design Motivations

- Simplified Control Plane
 - No STP feature placement (root bridge, loopguard, ...)
 - No default gateway redundancy setup/tuning
 - No matching of STP/HSRP priority
 - No L2/L3 multicast topology inconsistencies
- Ease of Troubleshooting (leverage well know toolset)
 - Show ip route
 - Traceroute
 - Ping and extended pings
 - Extensive protocol debugs
 - End to end consistent troubleshooting: Campus, WAN to Data Center
- Failure differences
 - Routed topologies fails safe—i.e. loss of routing protocol neighbor disables path
 - Layer 2 topologies fail open—i.e. loss of spanning tree BPDU's can open a blocked link and cause flooding



VSS Enabled Access Layer Campus Design



Cisco Expo
2009


Welcome to the Human Network.



VSS Enabled Campus Design

Control Plane Simplification


- VSS makes the network loop-free in normal topology; do **not** disable spanning tree to safeguard against possible loop introduced at the edge due to user error and daisy chaining
- Simplifies the topology allowing VLANs to span to increase flexibility in design options
- Ease of implementation, less to get it right
 - No need for HSRP, GLBP, or VRRP
 - No reliance on subsecond FHRP timers
 - No asymmetric forwarding
- A single logical multicast router on the access subnets simplifies the multicast topology resulting in elimination of non-RPF traffic
- Redundant supervisors provide resiliency via SSO-enabled protocols; consistent recovery during the failover of nodes at the distribution



VSS Enabled Campus Design

Impact to the Campus Topology


- Physical network topology does not change
 - Still have redundant chassis
 - Still have redundant links
- Logical topology is simplified as we now have a single control plane
 - No unicast flooding
 - Single configuration management
- Allows the design to replace traditional topology control plane with multichassis EtherChannel (MEC)
 - Enables loopfree topology, thus doubles the link capacity
 - Convergence and load balancing are based on EtherChannel



VSS Enabled Campus Design

MEC Configuration

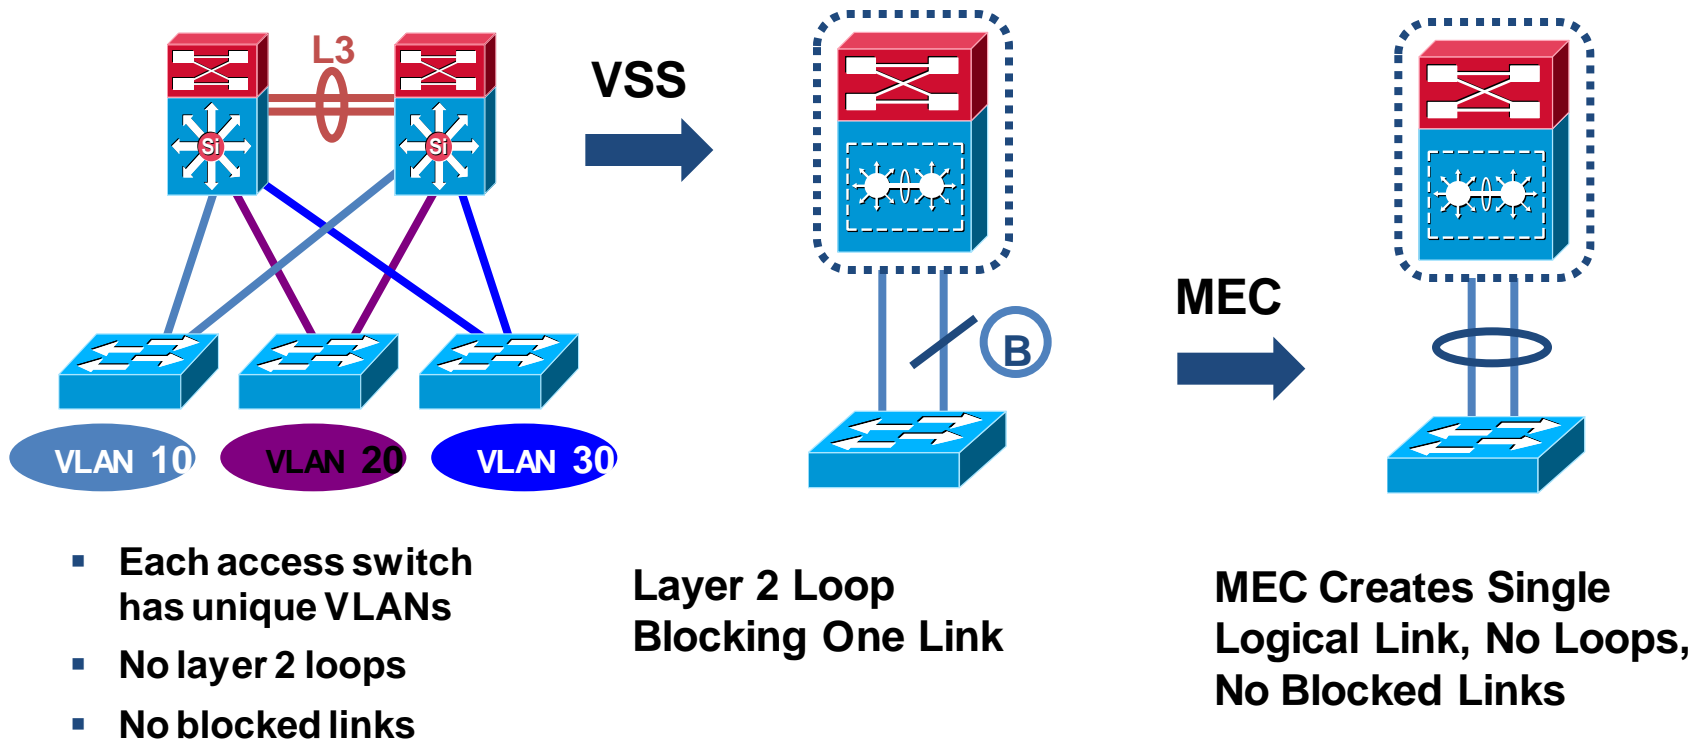
- MEC links on both switches are managed by PAgP or LACP running on the ACTIVE switch via internal control messages
 - All the rules and properties of EtherChannel applies to MEC such as negotiation, link characteristics (port-type, trunk), QoS, etc.
- Do not use “on” and “off” options with PAgP or LACP protocol negotiation
 - PAgP—Run Desirable-Desirable with MEC links
 - LACP—Run Active-Active with MEC links
- L2 MEC enables loop free topology and doubles the uplink bandwidth as no links are blocked
- L3 MEC provides reduced neighbor counts, consistent load-sharing (L2 and L3) and reduced VSL link utilization for multicast flows



VSS Enabled Campus Design

Multilayer Topology


- Optimized multilayer topology uses “V” shape design where VLANs do not span closets
- Deploying VSS in such topology without MEC reintroduces STP loops in the networks
- Use of MEC is recommended any time two L2 links from the same devices connected to VSS



VSS Enabled Campus Design


STP Optimization

- Make sure VSS remains root of all VLANs
- Do not use loop guard as it will disable the entire MEC channel on fault detection
- Use root guard at the edge port to protect external switch introducing superior BPDUs, e.g., temporary connectivity
- BPDU guard and root guard are mutually exclusive
- PortFast and BPDU guard is still necessary at the edge switch to prevent accidental loop introduce either due to user error or topology change




VSS Enabled Campus Design

End-to-End VSS Design Option



**STP-Based
Redundant Topology**

 = STP Blocked Link




**Fully Redundant
Virtual Switch Topology**

VSS Enabled Campus Design

Summary

- VSS enables highly available campus with sub-second convergence without the complexity of managing dual node at distribution layer
- Eliminates FHRP configuration
- Must use L2 MEC to create loop free topology, STP should remained enabled
- Use of L3 MEC significantly improves convergence for multicast traffic
- Enabled NSF in adjacent routed devices for better convergence
 - Use default Hello and Hold timers for EIGRP & OSPF
- Use STP tool kits guidance applicable to loop free “V” shape design



Summary



Cisco Expo
2009

Welcome to the Human Network.



Next Generation Campus Design

Evolving the Campus Foundation Architecture

	Multitier Access	Routed Access	Virtual Switch
Access Distribution Control Plane Protocols	Spanning Tree (PVST+, Rapid-PVST+ or MST)	EIGRP or OSPF	PAgP, LACP
Spanning Tree Required	STP Required for Network Redundancy and to Prevent L2 Loops	No	No
Network Recovery Mechanisms	Spanning Tree and FHRP (HSRP, GLBP, VRRP)	EIGRP or OSPF	Multichassis EtherChannel (MEC)
VLAN Spanning Wiring Closets	Supported (Not Desirable Design)	No	Supported

Next Generation Campus Design


Evolving the Campus Foundation Architecture

	Multitier Access	Routed Access	Virtual Switch
Layer 2/3 Demarcation	Distribution	Access	Distribution (Could Be Access)
First Hop Redundancy Protocol	HSRP, GLBP, VRRP Required	Not Required	Not Required
Load Balancing	Per Subnet or Host	Per Flow—ECMP	Per Flow—MEC
Convergence	900 msec—50 Seconds (Dependent on STP Topology and FHRP Tuning)	50–600 msec	50–600 msec

Next Generation Campus Design

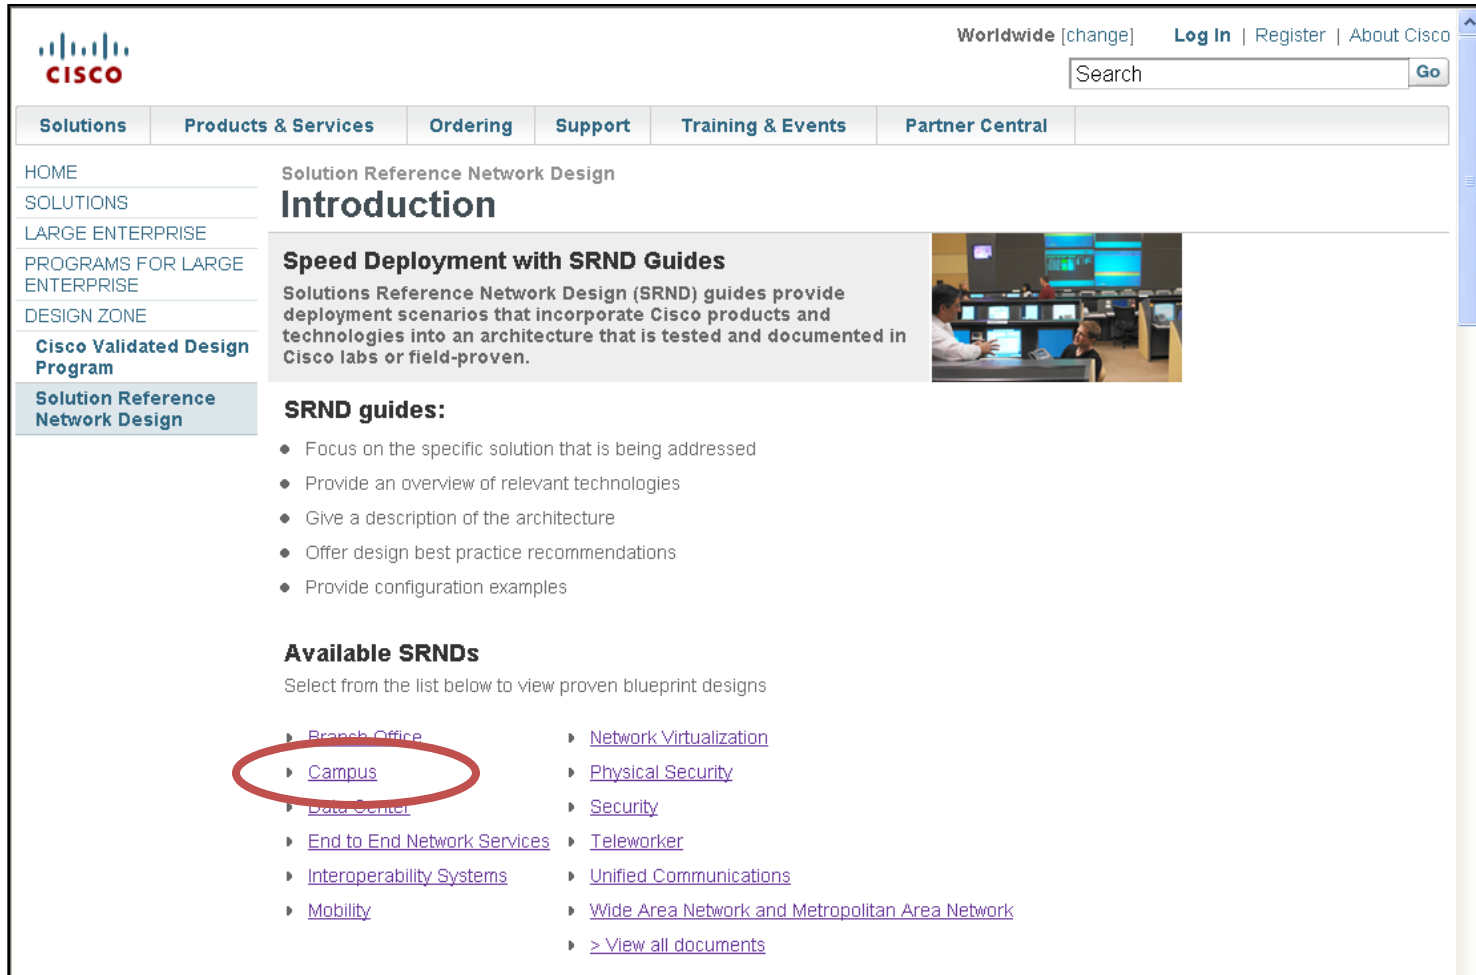
Evolving the Campus Foundation Architecture

- Traditional Layer 2 designs can evolve to VSS
- Evolving architectures provide:
 - Simplified control plane: remove dependence on STP
 - Increased capacity: provide flow-based load balancing
 - High availability: 200 msec or better recovery
- Flexibility to provide for the right implementation for each network requirement



Campus Design Guidance

Where to Go for More Information



CISCO

Worldwide [change] Log In | Register | About Cisco

Search Go

Solutions | **Products & Services** | **Ordering** | **Support** | **Training & Events** | **Partner Central**


HOME
SOLUTIONS
LARGE ENTERPRISE
PROGRAMS FOR LARGE ENTERPRISE
DESIGN ZONE
Cisco Validated Design Program
Solution Reference Network Design

Solution Reference Network Design

Introduction

Speed Deployment with SRND Guides

Solutions Reference Network Design (SRND) guides provide deployment scenarios that incorporate Cisco products and technologies into an architecture that is tested and documented in Cisco labs or field-proven.



SRND guides:

- Focus on the specific solution that is being addressed
- Provide an overview of relevant technologies
- Give a description of the architecture
- Offer design best practice recommendations
- Provide configuration examples

Available SRNDs

Select from the list below to view proven blueprint designs

- ▶ [Branch Office](#)
- ▶ [Campus](#)
- ▶ [Data Center](#)
- ▶ [End to End Network Services](#)
- ▶ [Interoperability Systems](#)
- ▶ [Mobility](#)
- ▶ [Network Virtualization](#)
- ▶ [Physical Security](#)
- ▶ [Security](#)
- ▶ [Teleworker](#)
- ▶ [Unified Communications](#)
- ▶ [Wide Area Network and Metropolitan Area Network](#)
- ▶ [View all documents](#)

<http://www.cisco.com/go/srnd/> and <http://www.cisco.com/go/cvd>

Q & A



Yaman Hakmi

Cisco Expo
2009

Welcome to the Human Network.





CISCO