



Deploying a Fully Routed Campus Network (Routing in the Access Layer) Advanced Design

BRKCAM-3004



Scott Van de Houten

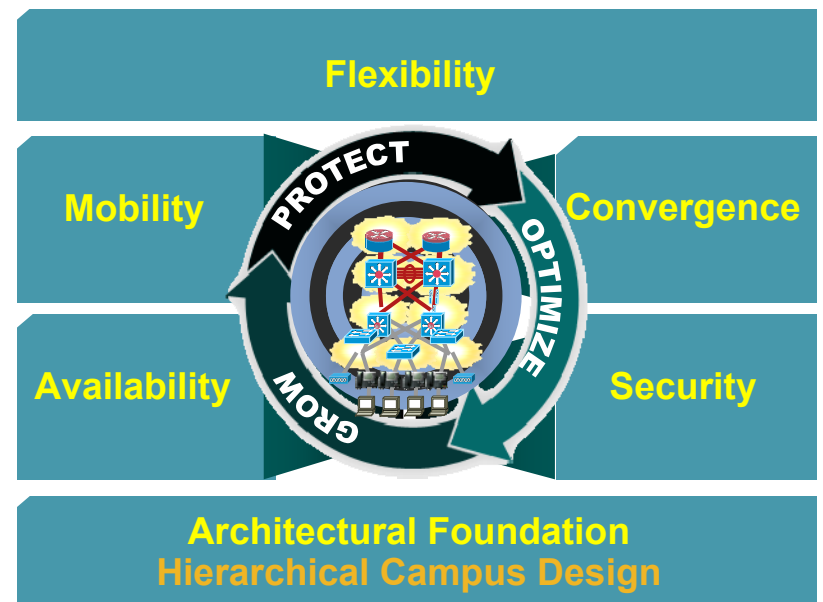
Cisco Networkers
2007

HOUSEKEEPING

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

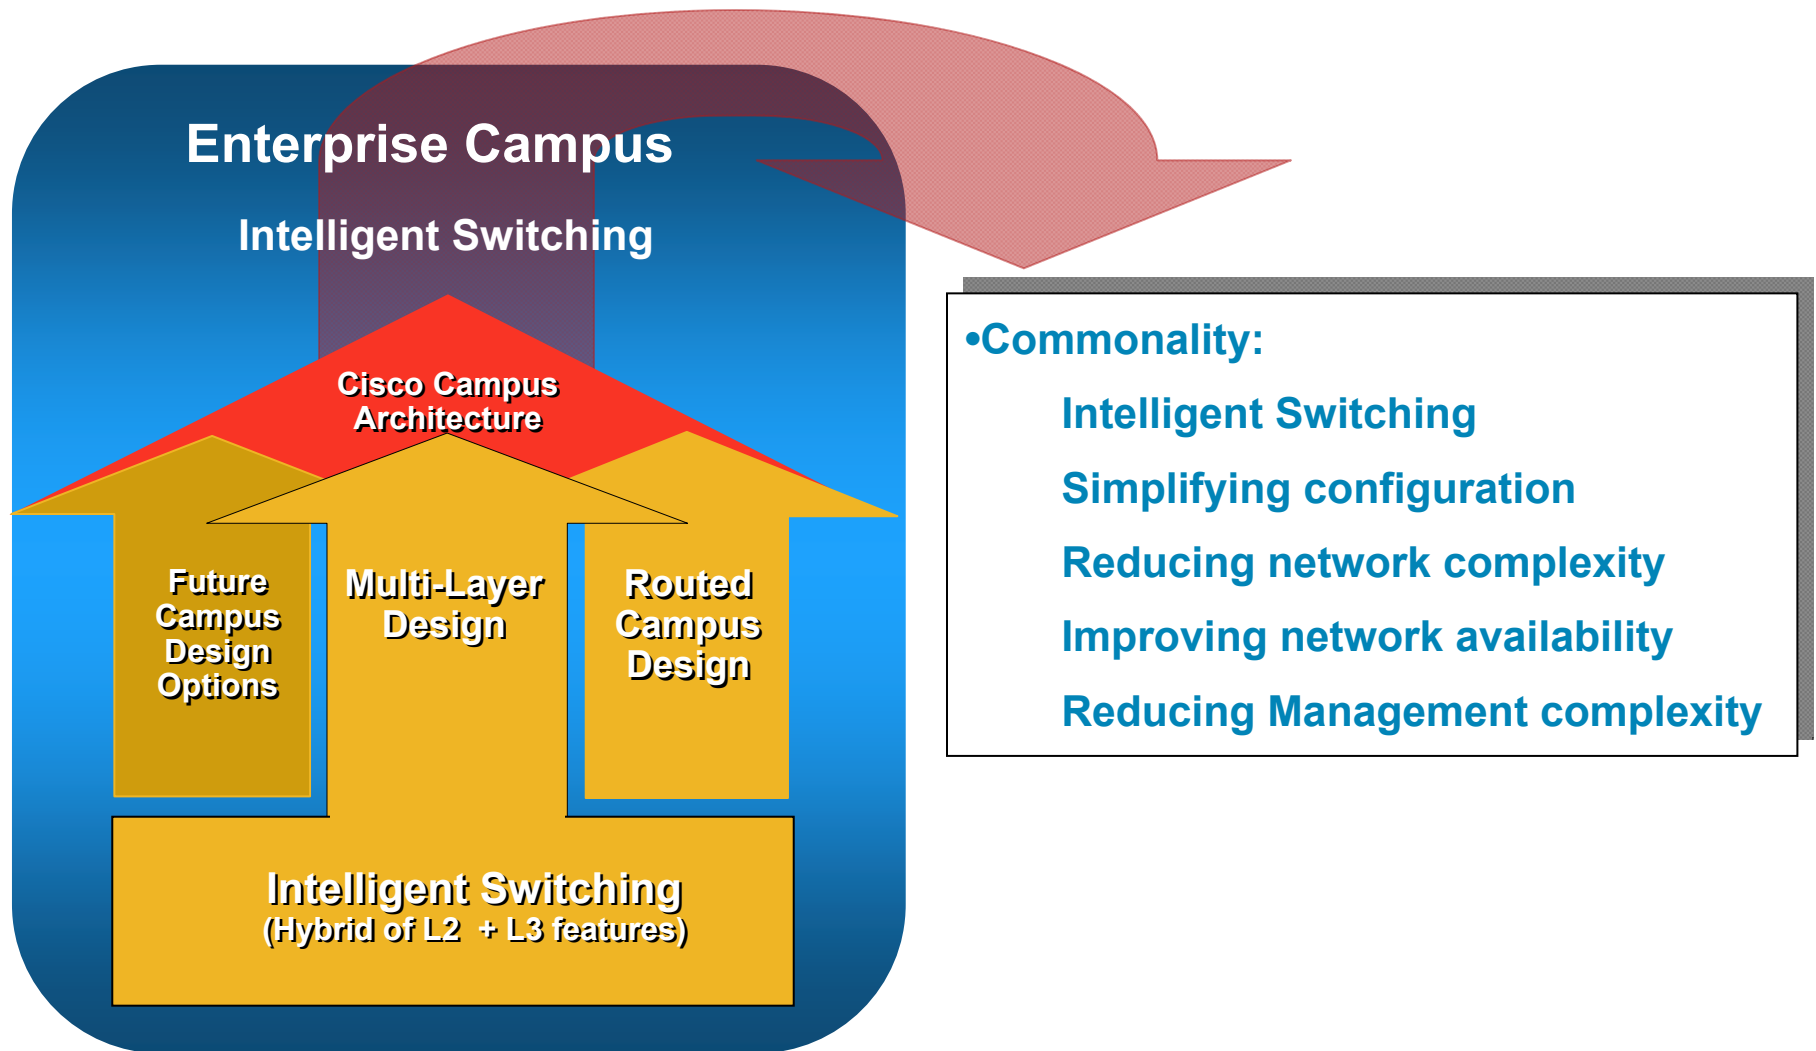
Agenda

- Cisco Campus Architecture
- Campus Network Resiliency
- Routed Campus Design
 - EIGRP Design Details
 - OSPF Design Details
 - PIM Design Details
- Impact on Advanced Technologies and Services
- Summary



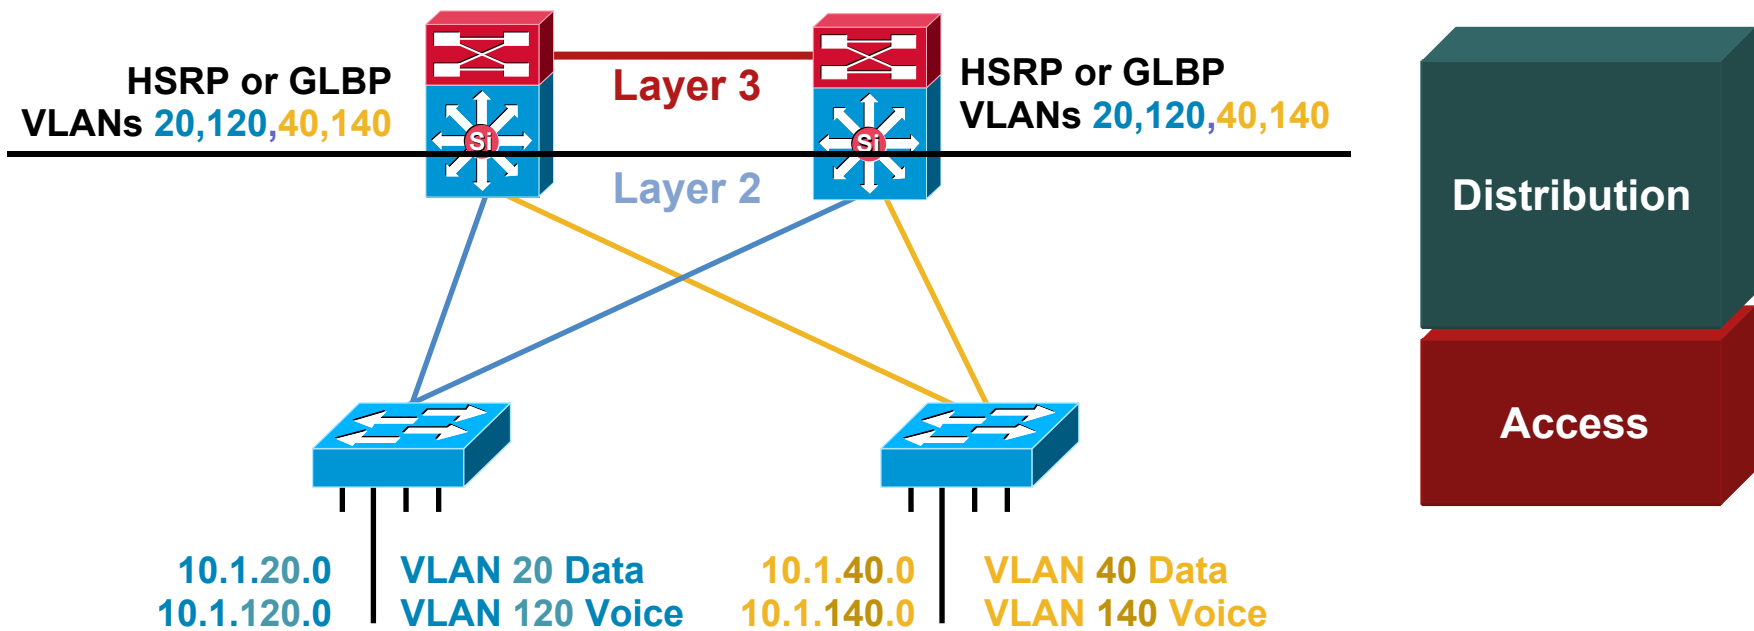
Cisco Campus Architecture

One Architecture with Multiple Design Options



Multilayer Reference Design

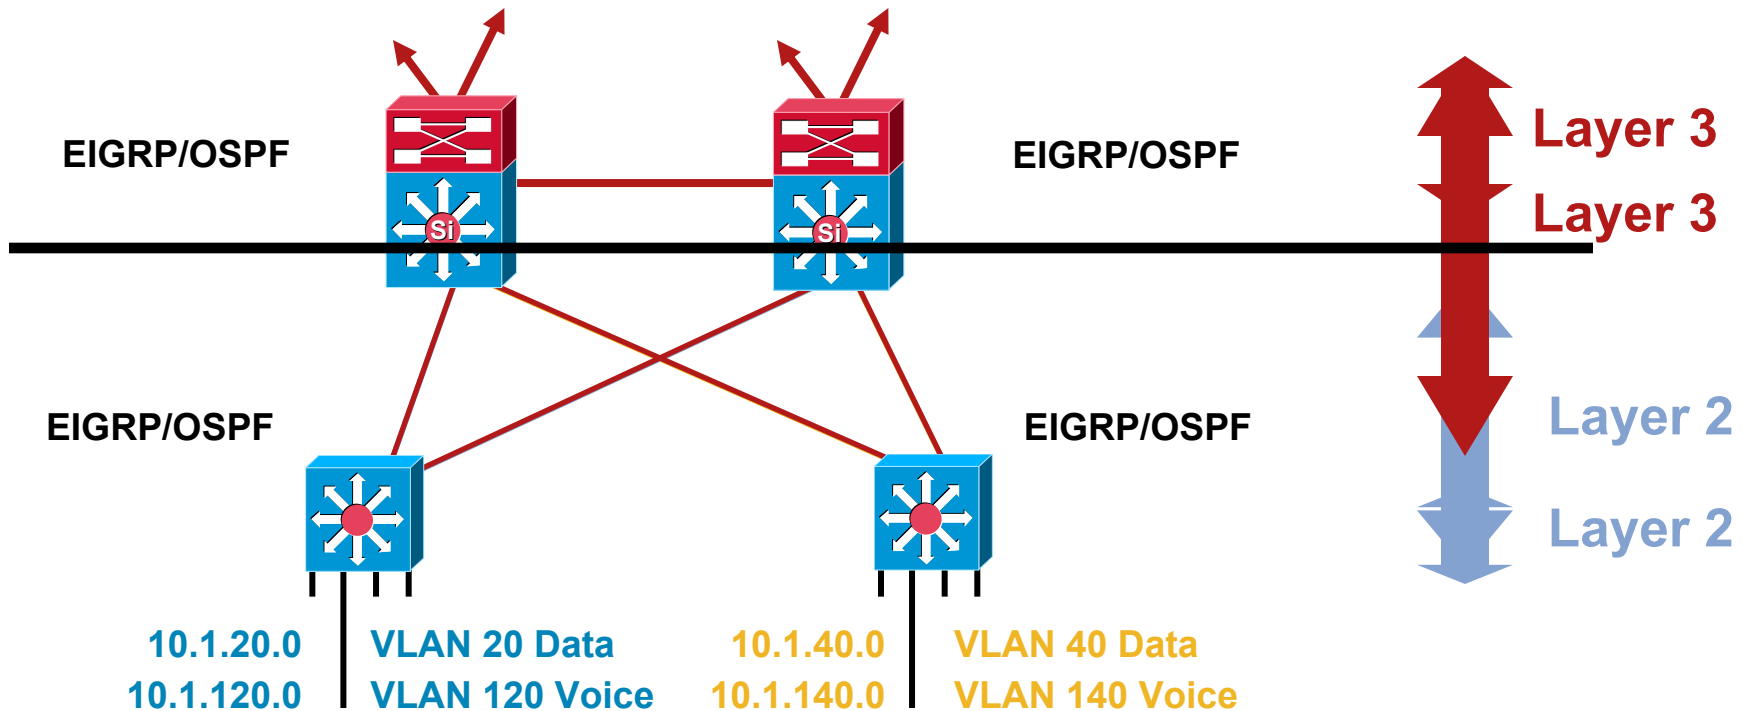
Layer 2/3 Distribution with Layer 2 Access



- Consider fully utilizing uplinks via GLBP
- Distribution-to-distribution link required for route summarization
- STP convergence not required for uplink failure/recovery
- Map L2 VLAN number to L3 subnet for ease of use/management
- Can easily extend VLANs across access layer switches **if required**

Routed Campus Design


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 link lost from upstream neighbor
- Beneficial for the right environment

Hierarchical Campus Design


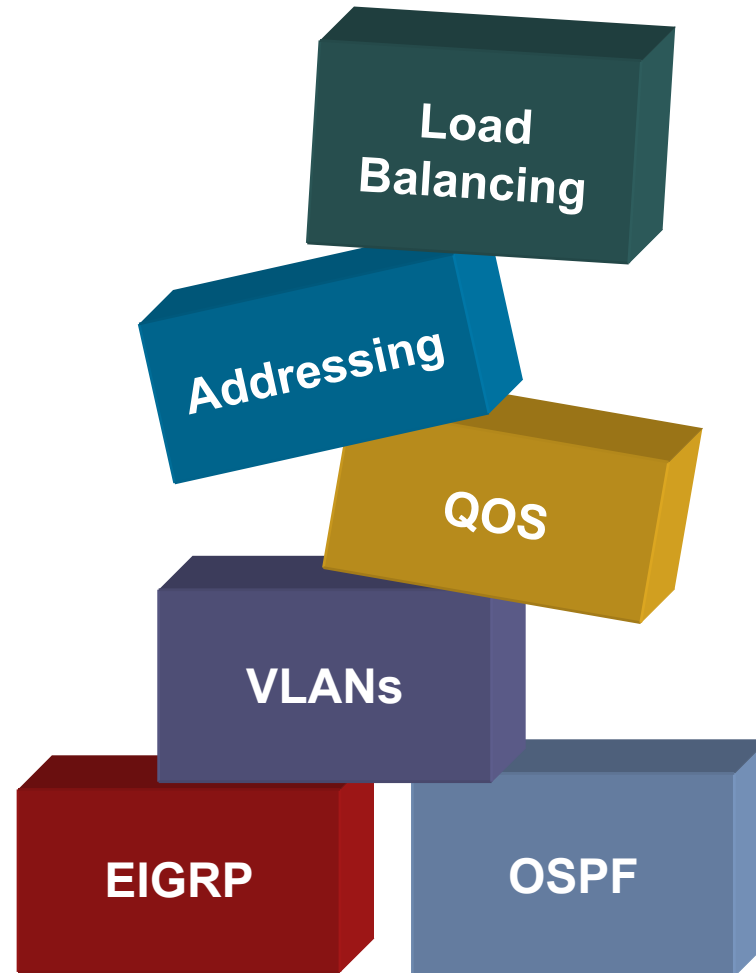

Building Blocks



- 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
- Can be applied to both the **multilayer** and **routed** campus designs




Hierarchical Campus Design— Without a Rock Solid Foundation the Rest Doesn't Matter




Agenda

- Cisco Campus Architecture
- **Campus Network Resiliency**
- Routed Campus Design
 - EIGRP Design Details
 - OSPF Design Details
 - PIM Design Details
- Impact on Advanced Technologies and Services
- Summary



What Is High Availability? And, Why Does It Matter?

Availability	Downtime Per Year (24 x 365)		
99.000%	3 Days	15 Hours	36 Minutes
99.500%	1 Day	19 Hours	48 Minutes
99.900%		8 Hours	46 Minutes
99.950%		4 Hours	23 Minutes
99.990%			53 Minutes
99.999%			5 Minutes
99.9999%			30 Seconds

To Achieve 5–9s or Better
Seconds Count

More Than Just Revenue Impacted

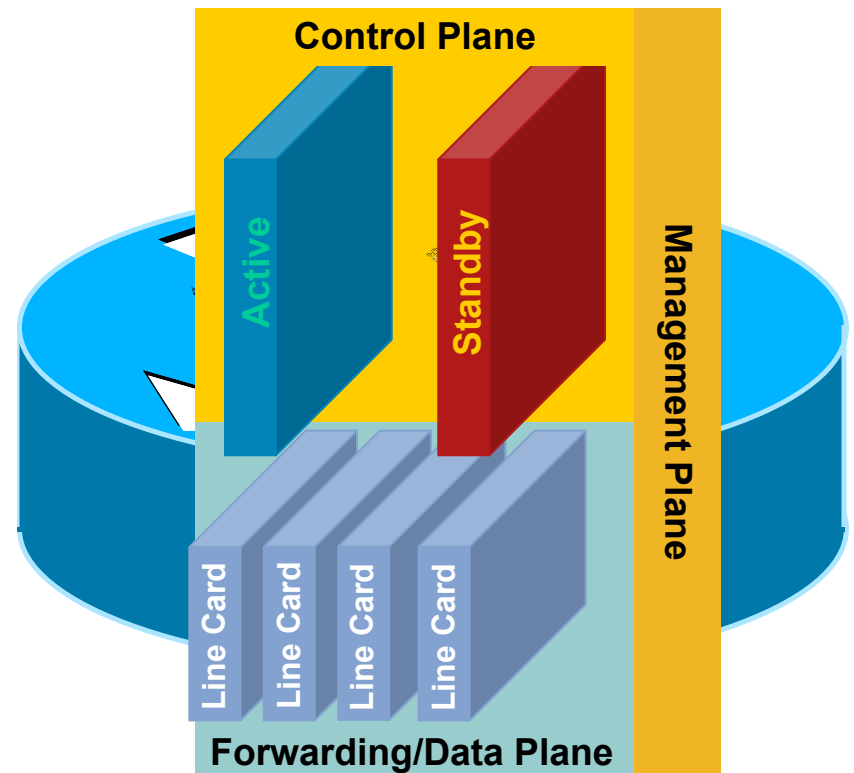
- Revenue loss
- Productivity loss
- Impaired financial performance
- Damaged reputation
- Employee frustration

Industry Sector	Revenue/ Hour	Revenue/ Employee-Hour
Energy	\$2,817,846	\$569
Telecommunications	\$2,066,245	\$186
Manufacturing	\$1,610,654	\$134
Financial Institution	\$1,495,134	\$1,079
Insurance	\$1,202,444	\$370
Retail	\$1,107,274	\$244
Transportation	\$668,586	\$107
Average	\$1,010,536	\$205

System Level Resiliency Overview

Eliminate Single Points of Failure for Hardware and Software Components

- Control/data plane resiliency
 - Separation of control and forwarding plane
 - Fault isolation and containment
 - Seamless restoration of Route Processor control and data plane failures
- Link resiliency
 - Reduced impact of line card hardware and software failures
- Planned outages
 - Seamless software and hardware upgrades



Network Level Resiliency Overview

Hierarchical Network Design


- Scalability to expand/shrink without affecting network behavior
- Predictable performance under normal conditions and failure conditions

Convergence and Self-Healing

- Reduce convergence times for major network protocols—EIGRP, OSPF, IS-IS, BGP
- Leverage in network wherever redundant paths exist

Intelligent Protocol Fabric


- Embed NSF intelligence network-wide in service provider and enterprise networks



Campus Network Resilience

Sub-Second Convergence

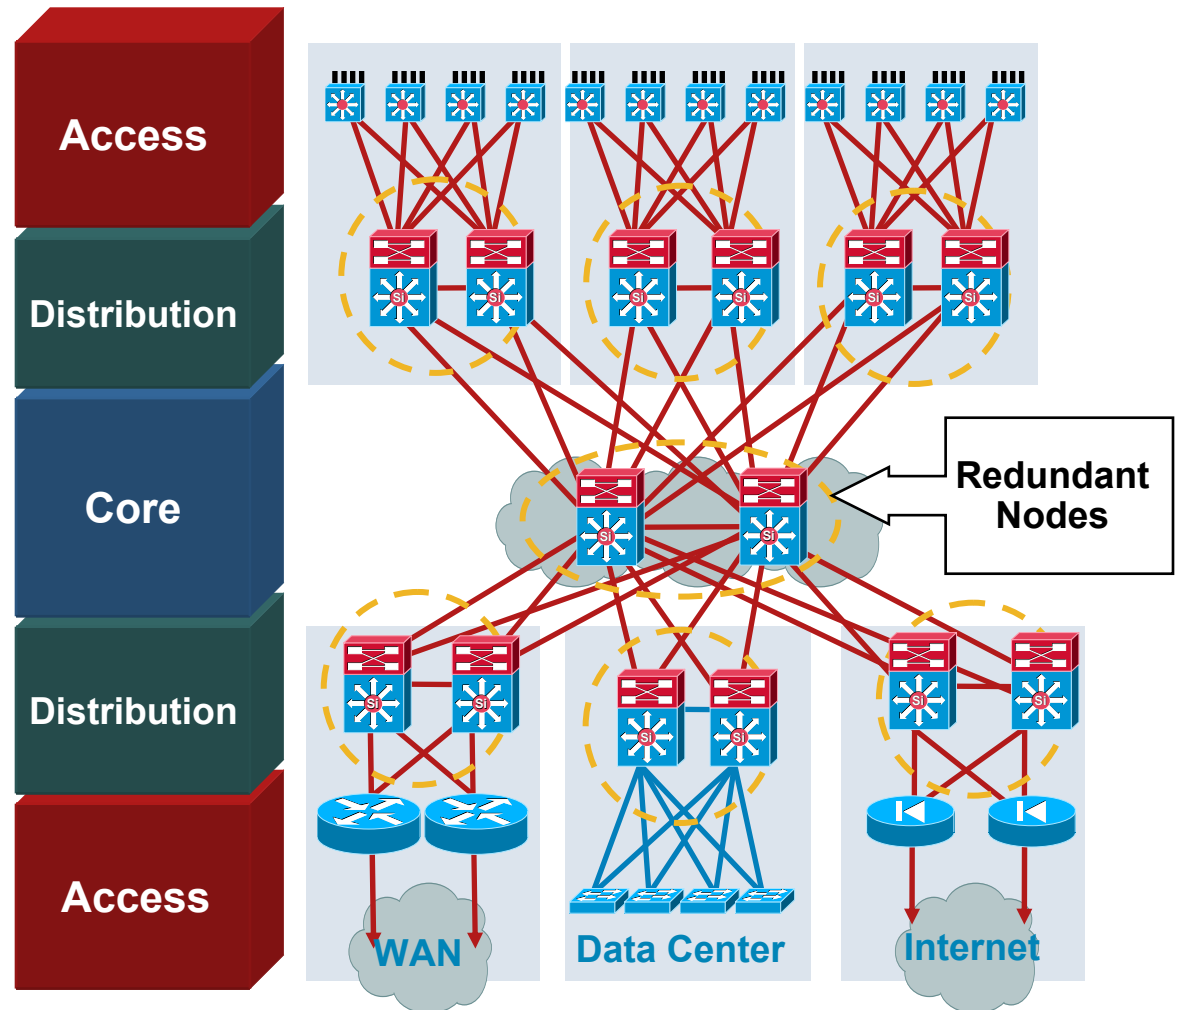
Convergence Times for Campus Best Practice Designs



* OSPF Results Require Sub-Second Timers


Optimal Redundancy

- Core and distribution engineered with redundant nodes and links to provide maximum redundancy and optimal convergence
- Network bandwidth and capacity engineered to withstand node or link failure
- Sub-second converge around most failure events



Single Points of Termination

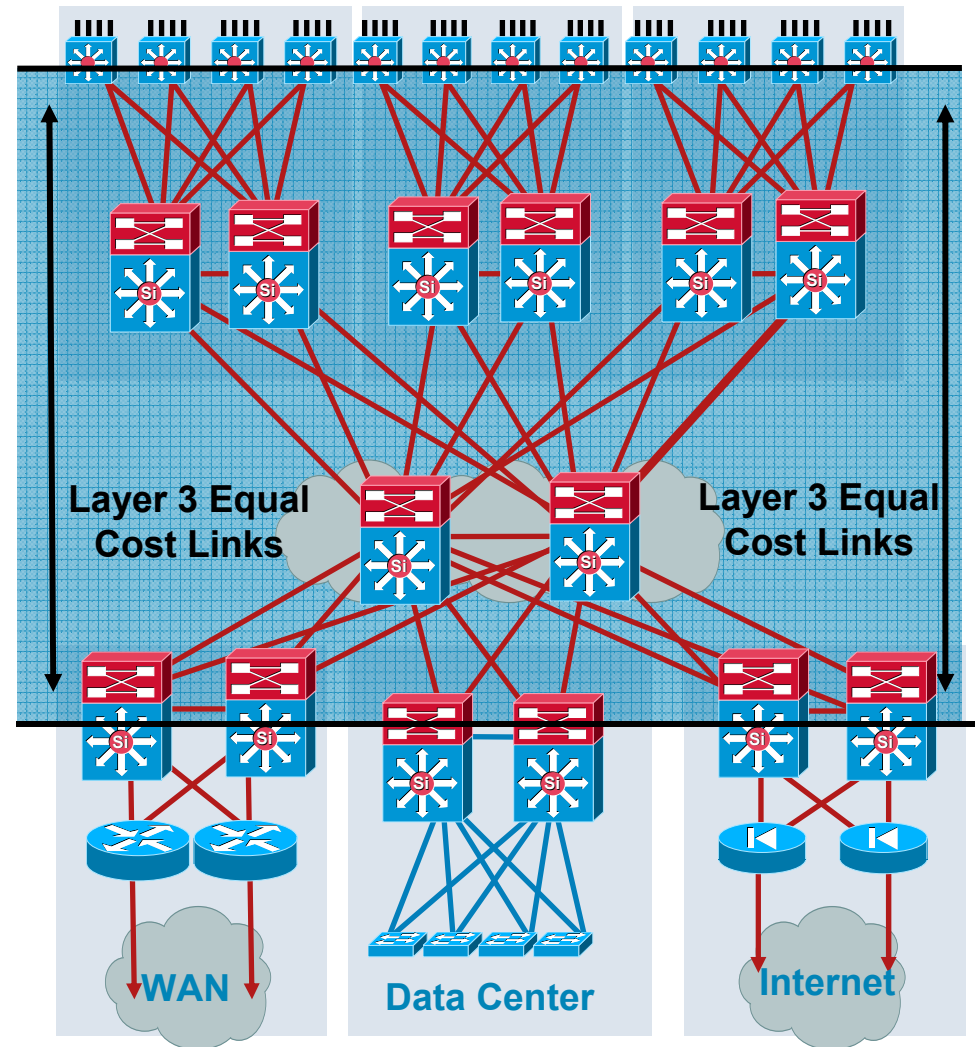
NSF/SSO Avoiding Total Network Outage



- The access layer and other single points of failure are candidates for supervisor redundancy
- L2 access layer SSO
- L3 access layer NSF and SSO
- Single points of failure risk network outage until physical replacement or reload vs. NSF zero to 1.8 seconds

Best Practices—Layer 3 Routing Protocols

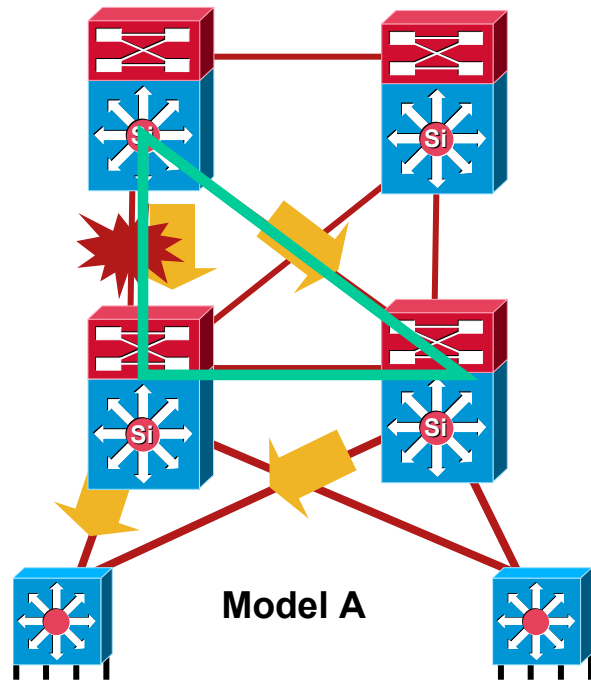
- Used to quickly re-route around failed node/links while providing load balancing over redundant paths
- **Build triangles** not squares for deterministic convergence
- Only peer on links that you intend to use as transit
- Insure redundant L3 paths to avoid black holes
- **Summarize** distribution 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)



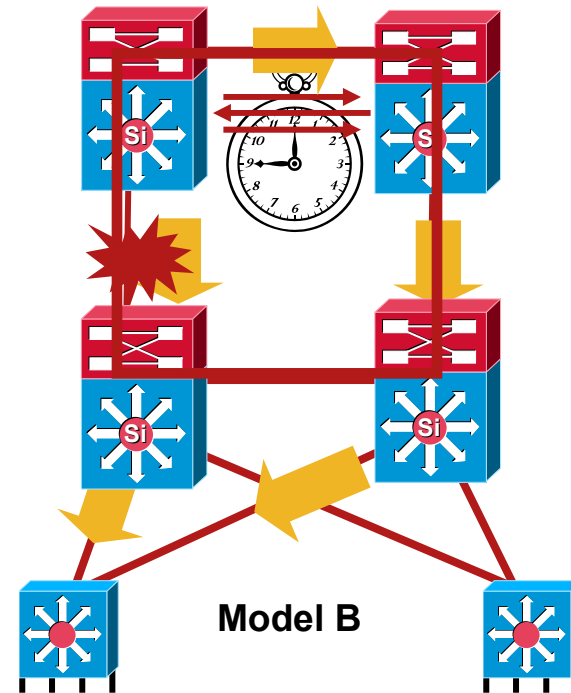
Best Practice—Build Triangles Not Squares

Deterministic vs. Non-Deterministic

Triangles: Link/Box Failure Does **Not** Require Routing Protocol Convergence



Squares: Link/Box Failure Requires Routing Protocol Convergence

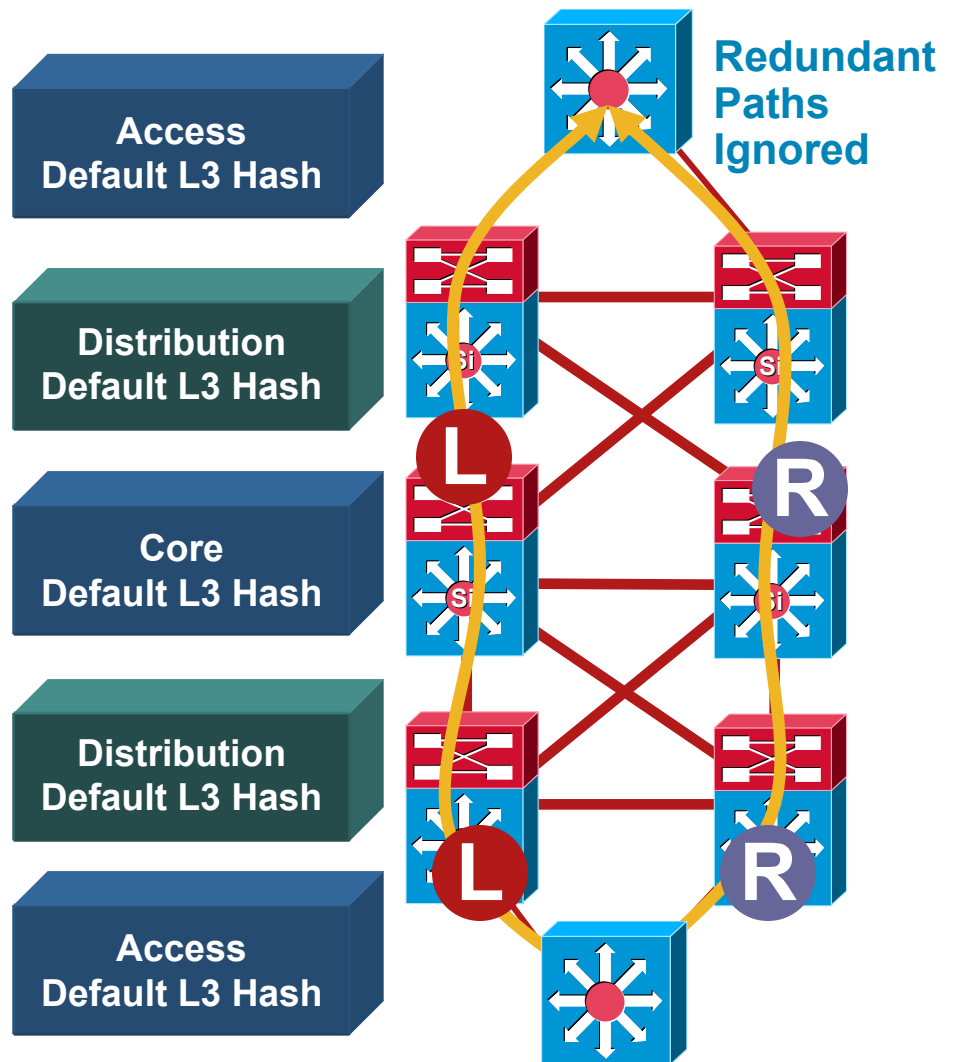


- Layer 3 redundant equal cost links support fast convergence
- Hardware based—fast recovery to remaining path
- Convergence is extremely fast (dual equal-cost paths: no need for OSPF or EIGRP to recalculate a new path)

CEF Load Balancing

Avoid Underutilizing Redundant Layer 3 Paths

- The default CEF hash 'input' is L3
- **CEF polarization:** in a multihop design, CEF could select the same left/left or right/right path
- Imbalance/overload could occur
- Redundant paths are ignored/underutilized



CEF Load Balancing

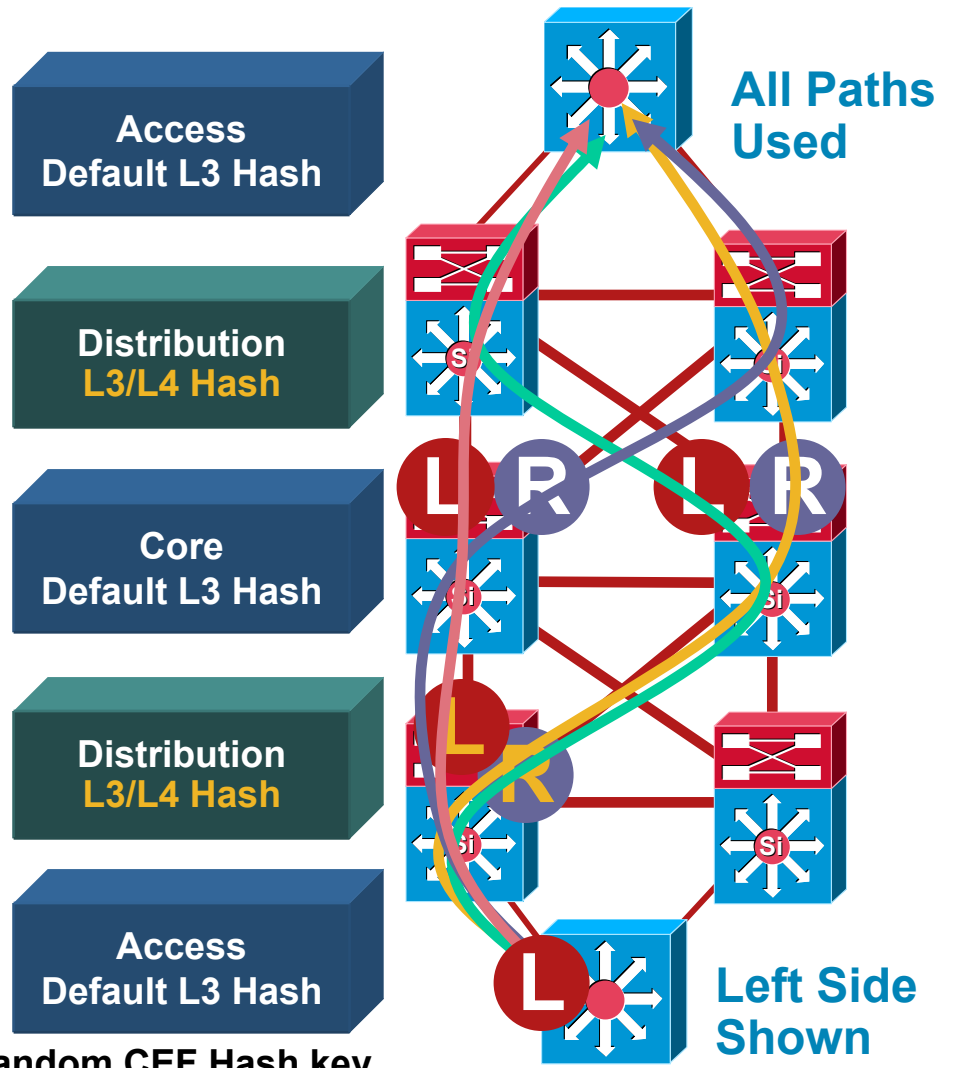
Avoid Underutilizing Redundant Layer 3 Paths

- With defaults, CEF could select the same left/left or right/right paths and ignore some redundant paths
- Alternating L3/L4 hash and default L3 hash will give us the best load balancing results
- The default is L3 hash—no modification required in core or access
- In the distribution switches use:

```
mls ip cef load-sharing full
```

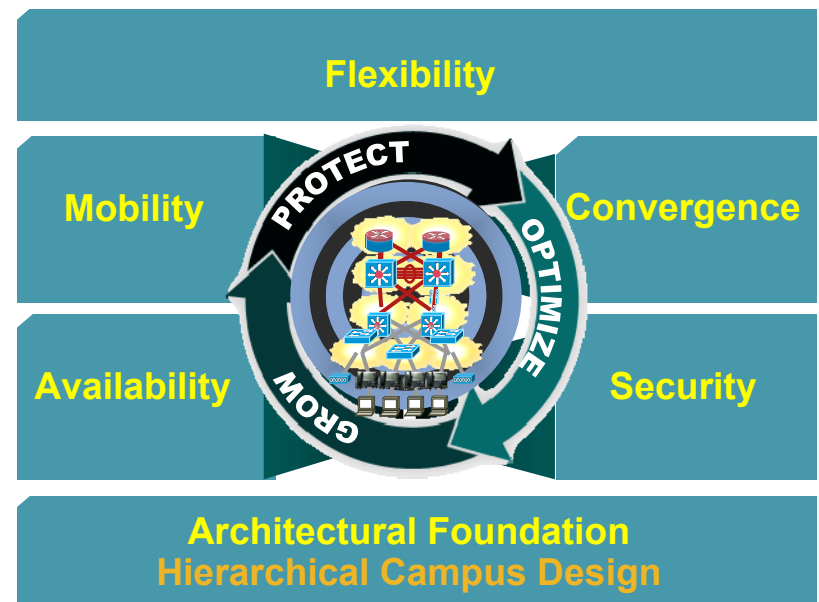
to achieve better redundant path utilization*

* Cisco Catalyst 6500 Sup720 & Sup32 have random CEF Hash key
Cisco Catalyst® 3000s Do Not Support L3/L4 Hash




Agenda

- Cisco Campus Architecture
- Campus Network Resiliency
- **Routed Campus Design**
 - EIGRP Design Details
 - OSPF Design Details
 - PIM Design Details
- Impact on Advanced Technologies and Services
- Summary



Why Routed Access Campus Design?



- Most Cisco Catalysts support L3 switching today
- IGP enhancements; stub router/area, fast reroute, etc.
- EIGRP/OSPF routing preference over spanning tree
- Ease of implementation and troubleshooting
- Single control plane and well known tool set
Traceroute, show ip route, sho ip eigrp neighbor, etc.

Routed Access Considerations

- VLANs are localized to a single wiring closet switch
- IP addressing—do you have and address allocation plan to support a routed access design?
- Platform requirements;
 - Requires a Cisco Catalyst 3550, Catalyst 3560 or above
 - Requires Cisco IOS® (Native or Hybrid)
 - Cisco Catalyst 6500 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 3000s require IP Services for PIM, PIM Stub for IP Base in 2007.



Ease of Implementation

- Fewer control plane options
- No STP feature placement
 - LoopGuard
 - RootGuard
 - STP Root
- No default gateway redundancy setup/tuning
- No matching of STP/HSRP priority
- No L2/L3 multicast topology inconsistencies



Ease of Troubleshooting

- Routing troubleshooting tools

 - Show ip route

 - Traceroute

 - Ping and extended pings

 - Extensive protocol debugs

 - Consistent troubleshooting: access, dist, core

- Bridging troubleshooting tools

 - Show ARP

 - Show spanning-tree, standby, etc.

 - Multiple show CAM dynamics to find a host

- Failure differences


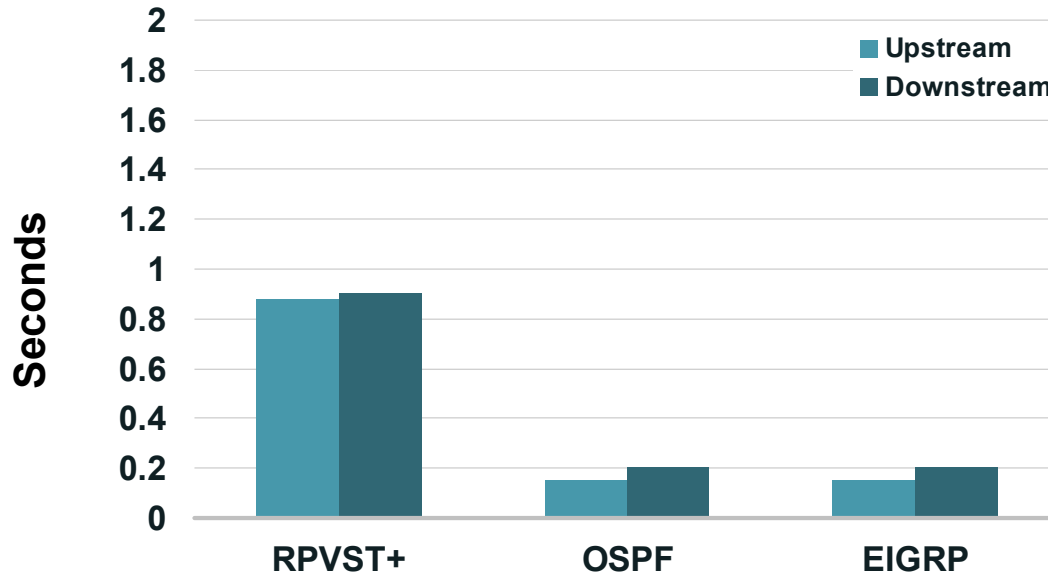
 - Routed topologies fail closed—i.e. neighbor loss

 - Layer 2 topologies fail open—i.e. broadcast and unknowns flooded



Routing to the Edge

Resiliency Advantages? Yes, with a Good Design



- EIGRP and OSPF converge in <200 msec
- OSPF convergence times dependent on timer tuning
- RPVST+ convergence times dependent on GLBP/HSRP tuning

EIGRP or OSPF as Your Campus IGP

DUAL or Dijkstra—They Are Both Good

- **Convergence:**

Within the campus environment, both EIGRP and OSPF provide extremely fast convergence

EIGRP requires summarization

OSPF requires summarization and timer tuning for fast convergence

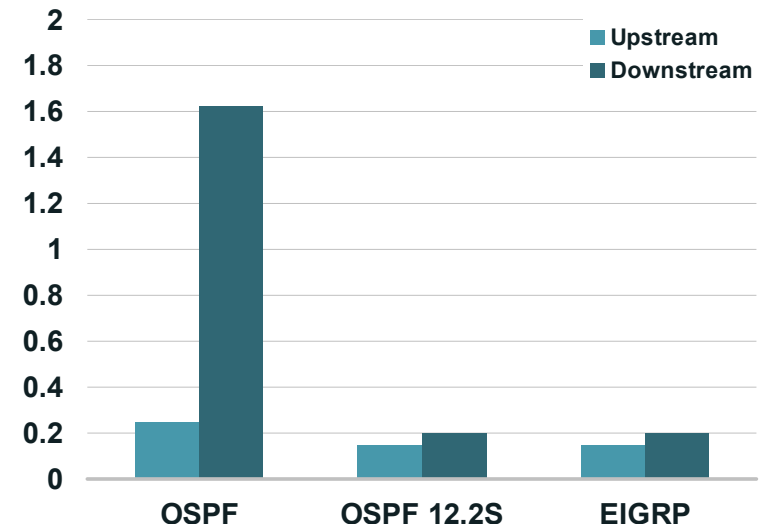
- **Flexibility:**

EIGRP supports multiple levels of route summarization and route filtering which simplifies migration from the traditional multilayer L2/L3 campus design

OSPF area design restrictions need to be considered

- **Scalability:**


Both protocols can scale to support very large enterprise network topologies



Routed Access Design



High-Speed Campus Convergence

- Convergence is the time needed for traffic to be rerouted to the alternative path after the network event
- Network convergence requires all affected routers to process the event and update the appropriate data structures used for forwarding
- Network convergence is the time required to:
 - Detect the event
 - Propagate the event
 - Process the event
 - Update the routing table/FIB



High-Speed Campus Convergence— Event Detection

- When physical interface changes state, the routing process is notified
 - This should happen in the ms range
- Some events are detected by the routing protocol hellos
 - L2 switch between L3 devices is a typical example
 - Neighbor is lost, but interface is UP/UP
- To improve failure detection
 - Use routed interfaces between L3 switches
 - Decrease interface carrier-delay to 0s
 - Decrease IGP hello timers without NSF
 - EIGRP: hellos = 1, hold-down = 3
 - OSPF: hellos = 250ms
- Bidirectional Forwarding Detection (BFD)*
 - Rapid (sub-second) failure detection with low overhead
 - Link/device/protocol failure detection at any protocol layer and over any media



***Verify Cisco IOS Release Availability, Performance Gains Not Yet Verified**

ESE Campus Solution Test Bed

Verified Design Recommendations


Total of 68 Access Switches,
2950, 2970, 3550, 3560, 3750,
4507 SupII+, 4507 SupIV, 6500
Sup2, 6500 Sup32, 6500 Sup720
and 40 APs (1200)

Three Distribution Blocks
6500 with Redundant Sup720
4507 with Redundant SupV

6500 with Redundant Sup720s


Three Distribution Blocks
6500 with Redundant Sup720s
7206VXR NPEG1

4500 SupII+, 6500 Sup720,
FWSM, WISM, IDSM2, MWAM



Agenda

- Cisco Campus Architecture
- Campus Network Resiliency
- Routed Campus Design
 - EIGRP Design Details
 - OSPF Design Details
 - PIM Design Details
- Impact on Advanced Technologies and Services
- Summary




Strengths of EIGRP

- Advanced distance vector
- Maps easily to the traditional multilayer design
- 100% loop free
- Fast convergence
- Easy configuration
- Incremental update
- Supports VLSM and discontinuous network
- Classless routing
- Protocol independent
 - IPv6, IPX and AppleTalk
- Unequal cost paths load balancing
- Flexible topology design options

EIGRP Design Rules for HA Campus

Similar to WAN Design, But—


- EIGRP design for the campus follows all the same best practices as you use in the WAN with a few differences
 - No BW limitations
 - Lower neighbor counts
 - Direct fiber interconnects
 - Lower cost redundancy
 - HW switching
- WAN → stability and speed
- Campus → stability, redundancy, load sharing, and high speed



EIGRP in the Campus

Conversion to an EIGRP Routed Edge

- The greatest advantages of extending EIGRP to the access 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 Protocol Fundamentals

Metric:

- Metric = $[K1 \times BW + (K2 \times BW)/(256 - \text{Load}) + K3 \times \text{Delay}] \times [K5/(\text{Reliability} + K4)] \times 256$
By Default: $K1 = 1, K2 = 0, K3 = 1, K4 = K5 = 0$
- Delay is sum of all the delays along the path
Delay = Delay/10
- Bandwidth is the lowest bandwidth link along the path
Bandwidth = $10000000/\text{Bandwidth}$

Packets:

- Hello: establish neighbor relationships
- Update: send routing updates
- **Query**: ask neighbors about routing information
- **Reply**: response to query about routing information
- Ack: acknowledgement of a reliable packet

EIGRP Neighbors

Event Detection

- EIGRP neighbor relationships are created when a link comes up and routing adjacency is established
- When physical interface changes state, the routing process is notified



Carrier-delay should be set as a rule because it varies based upon the platform

- Some events are detected by the routing protocol
Neighbor is lost, but interface is UP/UP

- To improve failure detection
Use routed interfaces and not SVIs
Decrease interface carrier-delay to 0
Decrease EIGRP hello and hold-down timers

Hello = 1
Hold-down = 3


```
interface GigabitEthernet3/2
 ip address 10.120.0.50 255.255.255.252
 ip hello-interval eigrp 100 1
 ip hold-time eigrp 100 3
 carrier-delay msec 0
```



EIGRP Design Rules for HA Campus

Limit Query Range to Maximize Performance

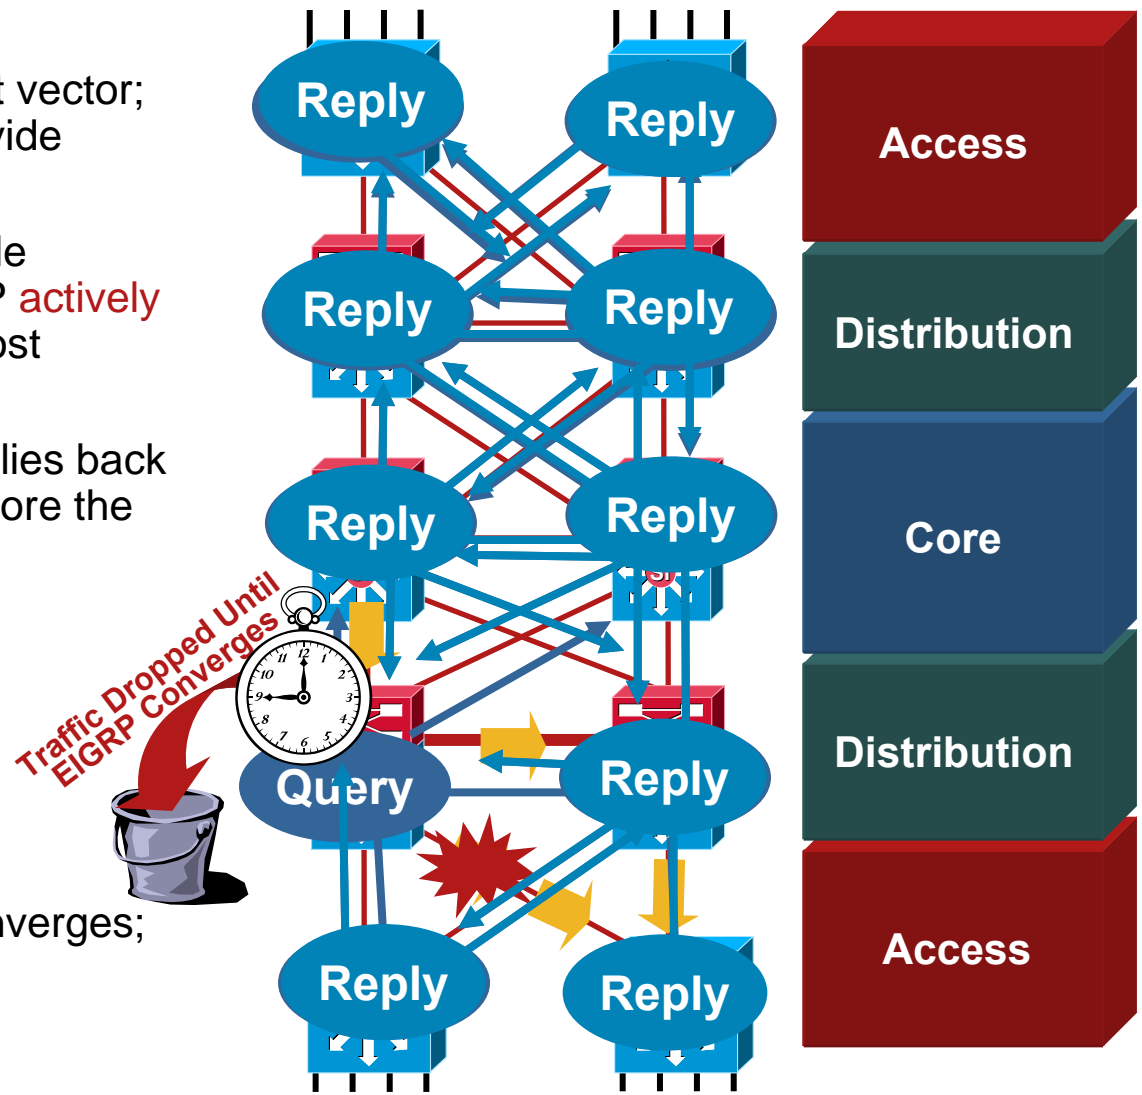
- EIGRP convergence is largely dependent on **query response times**
- **Minimize the number of queries** to speed up convergence
- **Summarize** distribution block routes upstream to the core
 - Upstream queries are returned immediately with infinite cost
- Configure all access switches as **EIGRP stub routers**
 - No downstream queries are ever sent



EIGRP Query Process

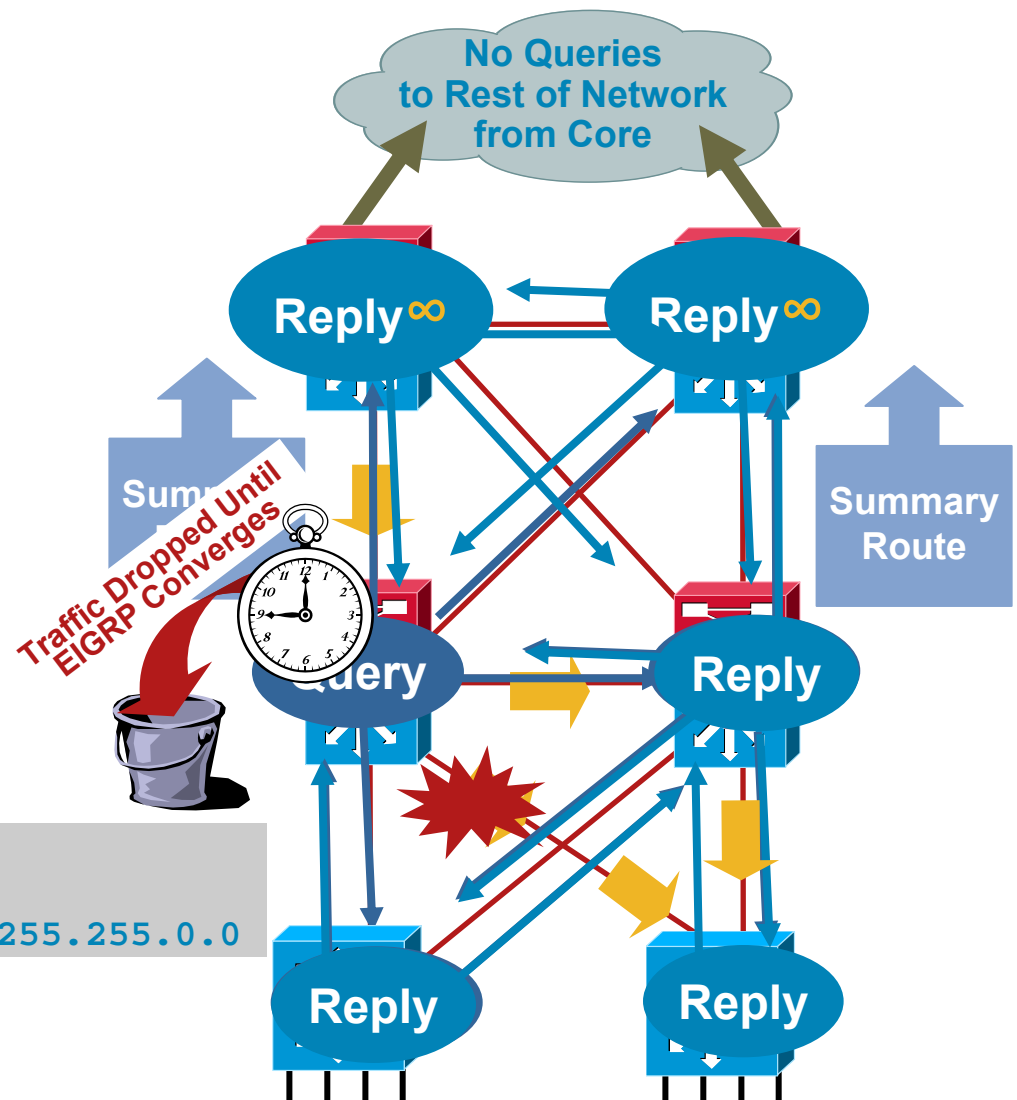
Queries Propagate the Event

- EIGRP is an advanced distant vector; it relies on its neighbor to provide routing information
- If a route is lost and no feasible successor is available, EIGRP **actively** queries its neighbors for the lost route(s)
- The router will have to get replies back from **all** queried neighbors before the router calculates successor information
- If any neighbor fails to reply, the queried route is **stuck in active** and the router resets the neighbor that fails to reply
- The fewer routers and routes queried, the faster EIGRP converges; **solution is to limit query range**



Limiting the EIGRP Query Range with Summarization

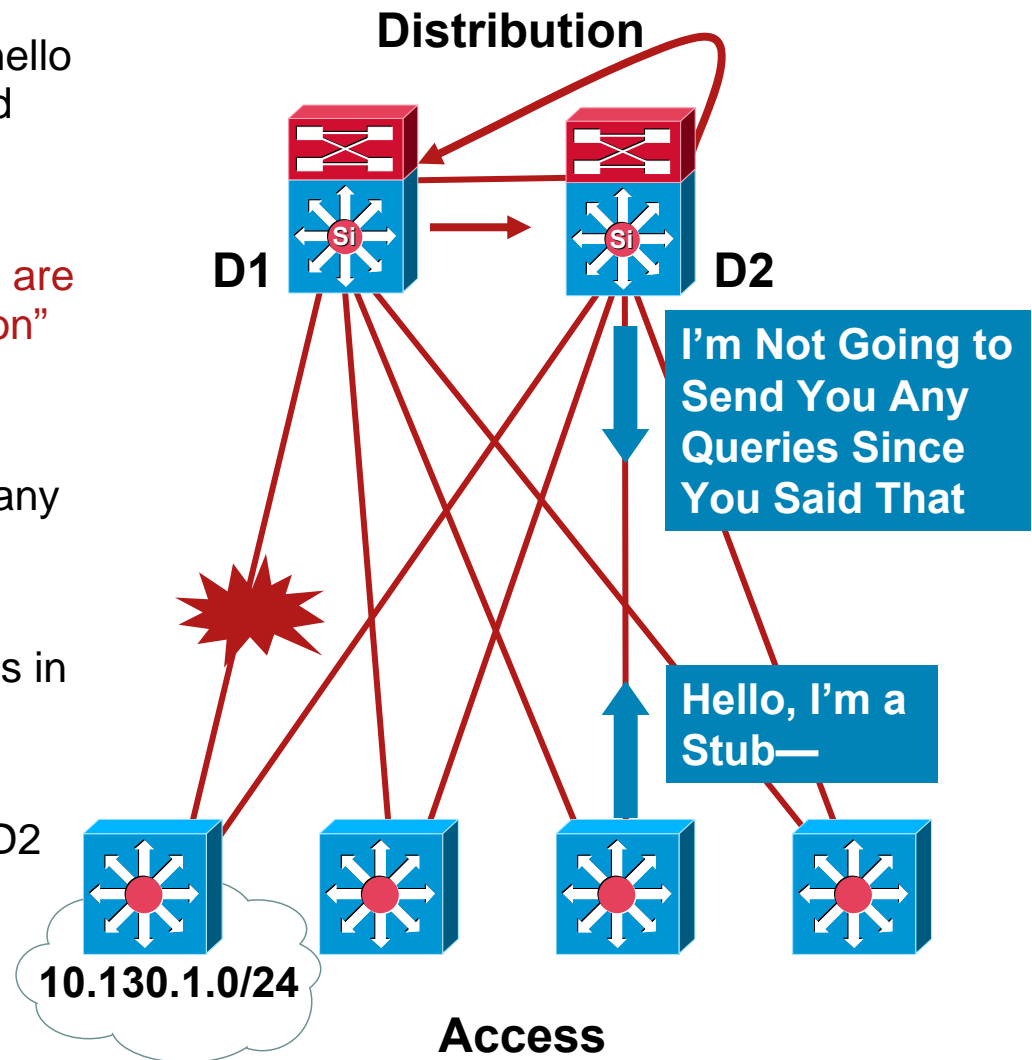
- When we summarize from distribution to core for the subnets in the access we can limit the upstream query/reply process
- In a large network this could be significant because queries will now stop at the core; no additional distribution blocks will be involved in the convergence event
- The access layer is still queried



```
interface gigabitethernet 3/1
ip address 10.120.10.1 255.255.255.252
ip summary-address eigrp 1 10.130.0.0 255.255.0.0
```

Limiting the EIGRP Query Range with Stub Routers

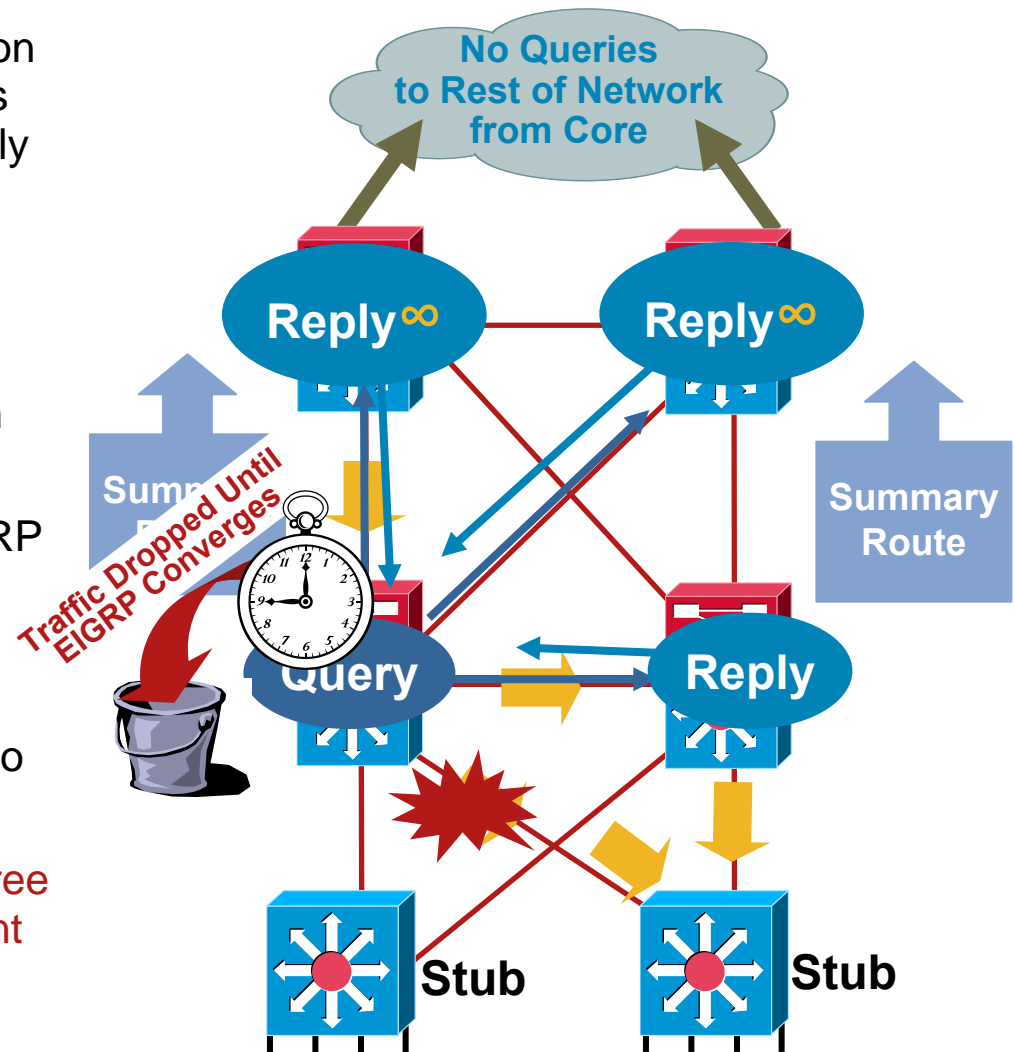
- A stub router signals (through the hello protocol) that it is a stub and should not transit traffic
- Queries that would have been generated towards the stub routers are marked as if a “No path this direction” reply had been received
- D1 will know that stubs cannot be transit paths, so they will not have any path to 10.130.1.0/24
- D1 simply will not query the stubs, reducing the total number of queries in this example to one
- These stubs will not pass D1’s advertisement of 10.130.1.0/24 to D2
- D2 will only have one path to 10.130.1.0/24



EIGRP Query Process

With Summarization and Stub Routers

- When we summarize from distribution to core for the subnets in the access we can limit the upstream query/reply process
- In a large network this could be significant because queries will now stop at the core; no additional distribution blocks will be involved in the convergence event
- When the access switches are EIGRP stubs we can further reduce the query diameter
- Non-stub routers do not query stub routers—so no queries will be sent to the access nodes
- **No secondary queries—and only three nodes involved in convergence event**




EIGRP Route Filtering in the Campus

Control Route Advertisements

- Bandwidth is not a constraining factor in the campus but it is still advisable to control number of routing updates advertised
- Remove/filter routes from the core to the access and inject a default route with distribute-lists
- Smaller routing table in access is simpler to troubleshoot
- Deterministic topology

```
ip access-list standard Default
  permit 0.0.0.0


router eigrp 100
  network 10.0.0.0
  distribute-list Default out <mod/port>
```



EIGRP Routed Access Campus Design

Overview


- Detect the event:
 - Set hello-interval = 1 second and hold-time = 3 seconds to detect soft neighbor failures
 - Set carrier-delay = 0
- Propagate the event:
 - Configure all access layer switches as stub routers to limit queries from the distribution layer
 - Summarize the access routes from the distribution to the core to limit queries across the campus
- Process the event:
 - Summarize and filter routes to minimize calculating new successors for the RIB and FIB



**For More Discussion on EIGRP:
BRKIPM-3008 - Advances in EIGRP**

Agenda

- Cisco Campus Architecture
- Campus Network Resiliency
- Routed Campus Design
 - EIGRP Design Details
 - OSPF Design Details
 - PIM Design Details
- Impact on Advanced Technologies and Services
- Summary



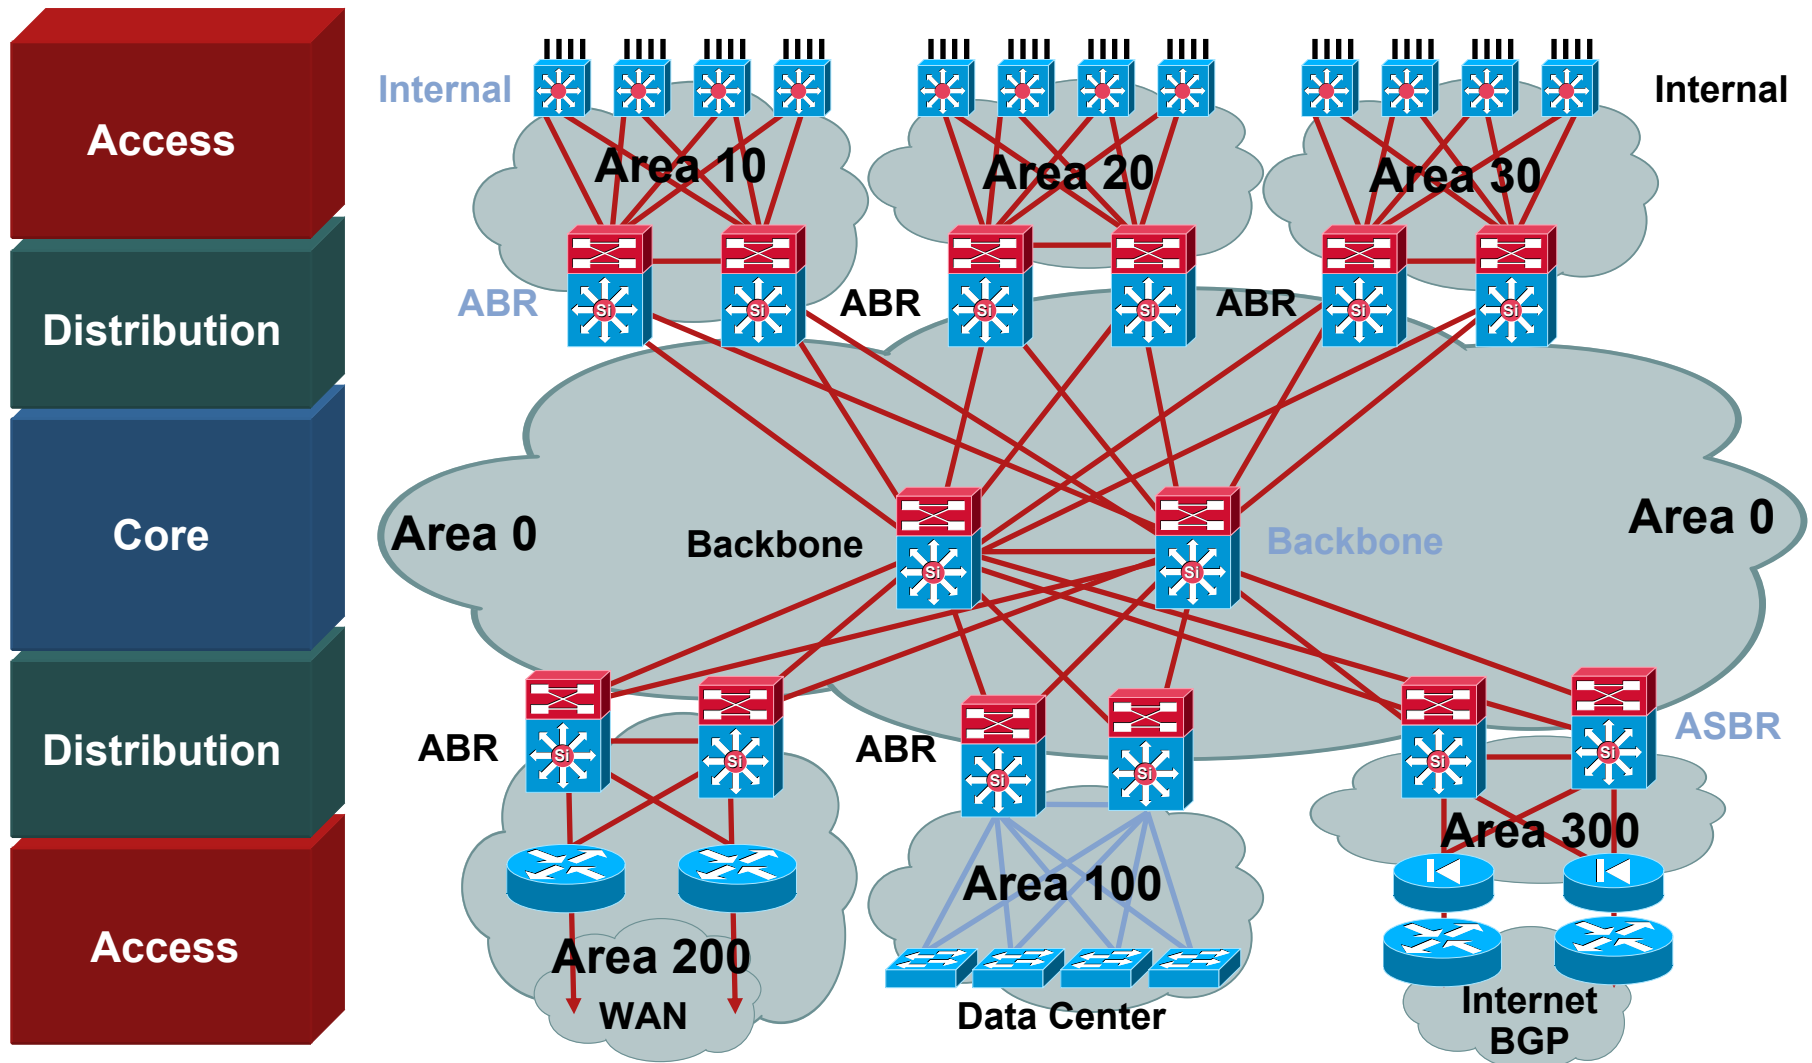
Open Shortest Path First (OSPF)

Overview

- OSPFv2 established in 1991 with RFC 1247
- Goal—a link state protocol more efficient and scaleable than RIP
- Dijkstra Shortest Path First (SPF) algorithm
- Metric—path cost
- Fast convergence
- Support for CIDR, VLSM, authentication, multipath and IP unnumbered
- Low steady state bandwidth requirement
- OSPFv3 for IPv6 support

Hierarchical Campus Design

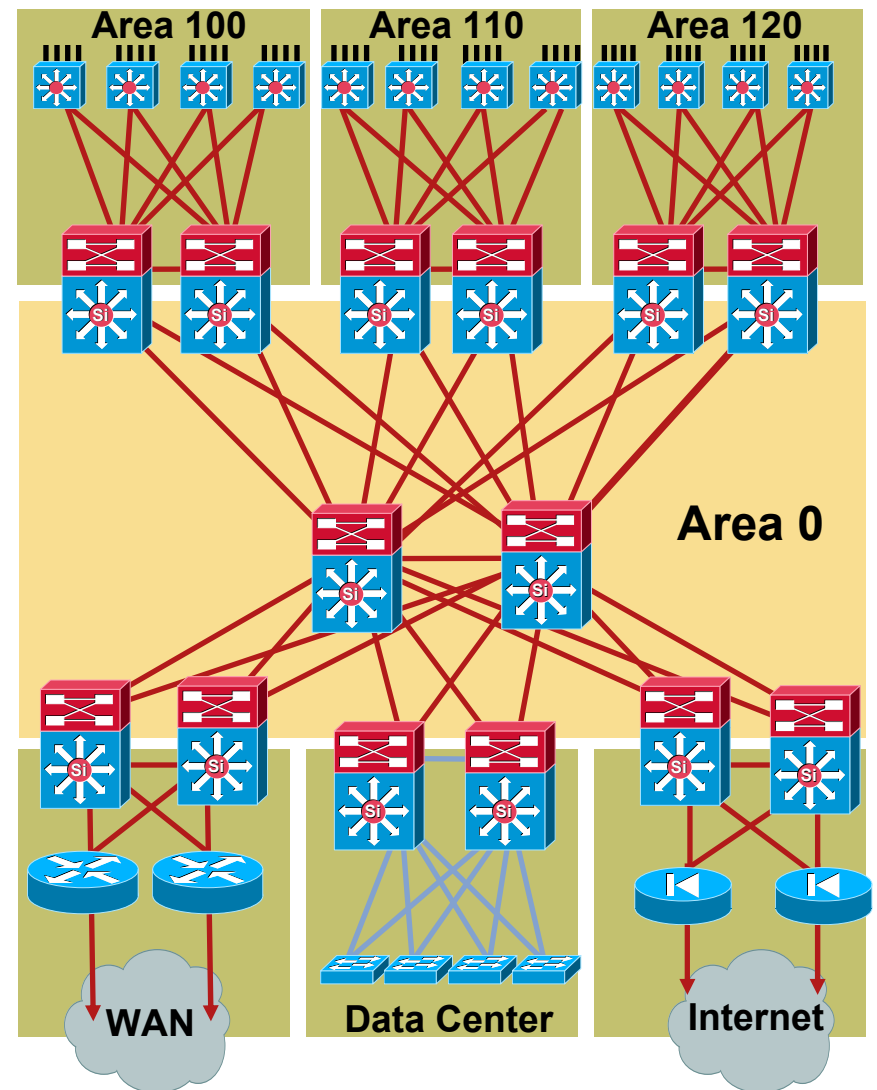
OSPF Areas with Router Types



OSPF Design Rules for HA Campus

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 however—
- Area design also based on address summarization
- Area boundaries should define buffers between fault domains
- Keep area 0 for core infrastructure do not extend to the access routers



OSPF in the Campus


Conversion to an OSPF Routed Edge

- OSPF designs that utilize an area for each campus distribution building block allow for straight forward migration to Layer 3 access
- Converting L2 switches to L3 within a contiguous area is reasonable to consider as long as new area size is reasonable
- How big can the area be?
- It depends
 - Switch type(s)
 - Number of links
 - Stability of fiber plant



When a Link Changes State

Router 1, Area 1



LSA



ACK



Router 2, Area 1

Link State Table

Dijkstra Algorithm

Old Routing Table

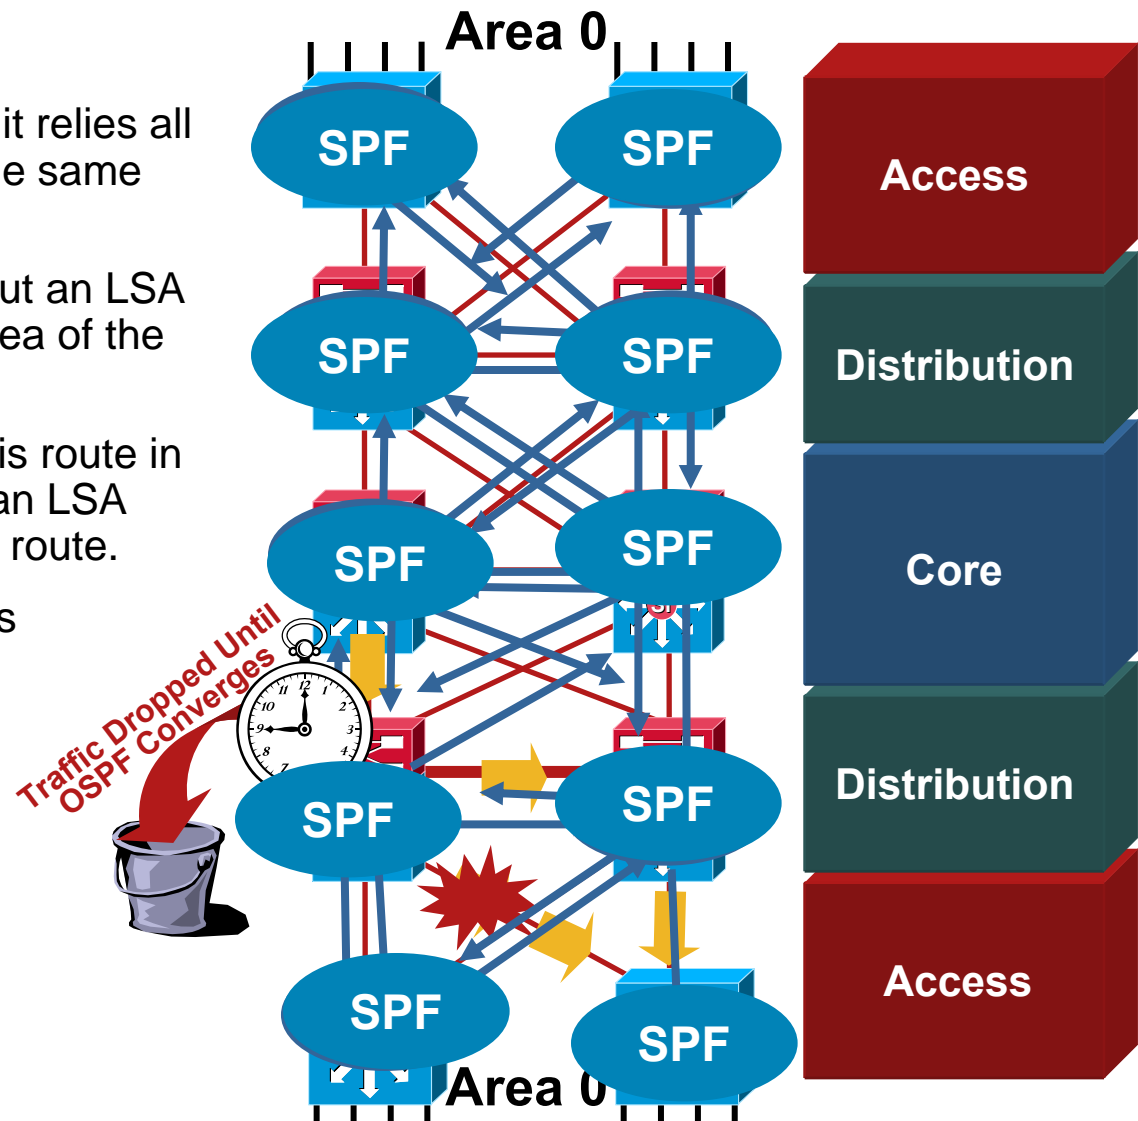
New Routing Table

- Every router in area hears a specific link LSA
- Each router computes shortest path routing table

OSPF LSA Process

LSAs Propagate the Event


- OSPF is a Link State protocol; it relies all routers within an area having the same topology view of the network.
- If a route is lost, OSPF sends out an LSA to inform it's peers within the area of the lost route.
- All routers with knowledge of this route in the OSPF network will receive an LSA and run SPF to remove the lost route.
- The fewer the number of routers with knowledge of the route, the faster OSPF converges;
- Solution is to limit LSA propagation range**



Regular Area

ABRs Forward All LSAs from Backbone

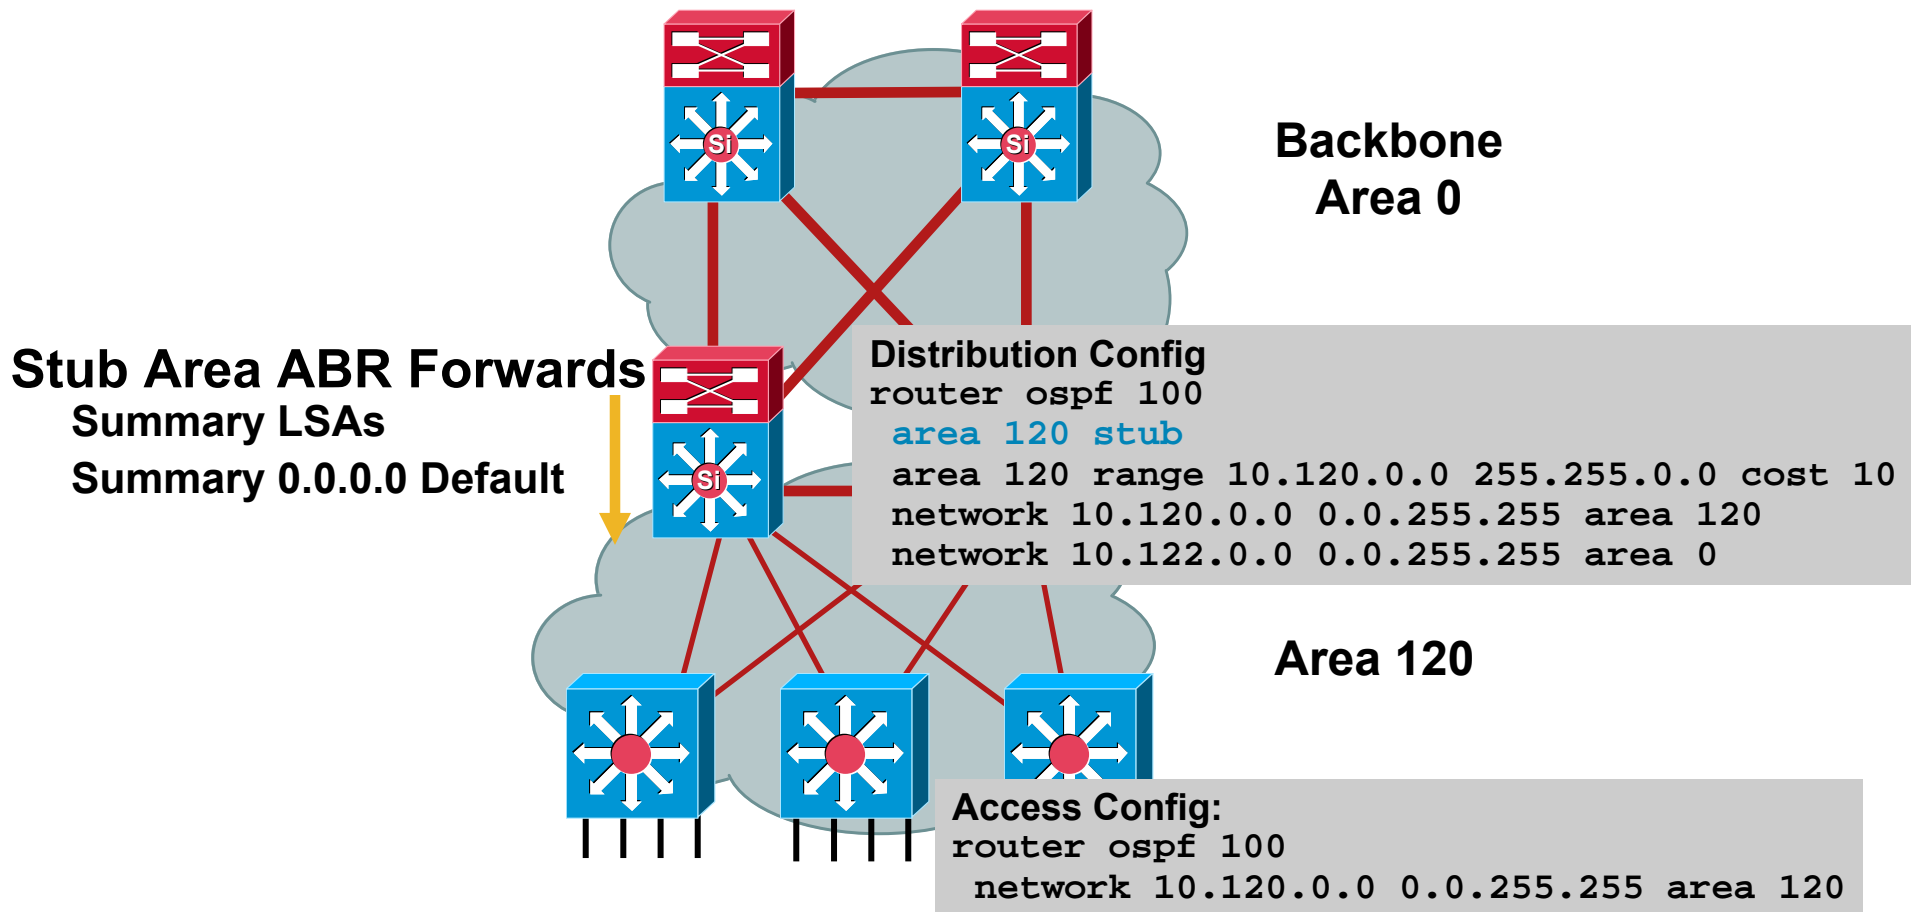
External Routes/LSA Present in Area 120



Stub Area

Consolidates Specific External Links—Default 0.0.0.0

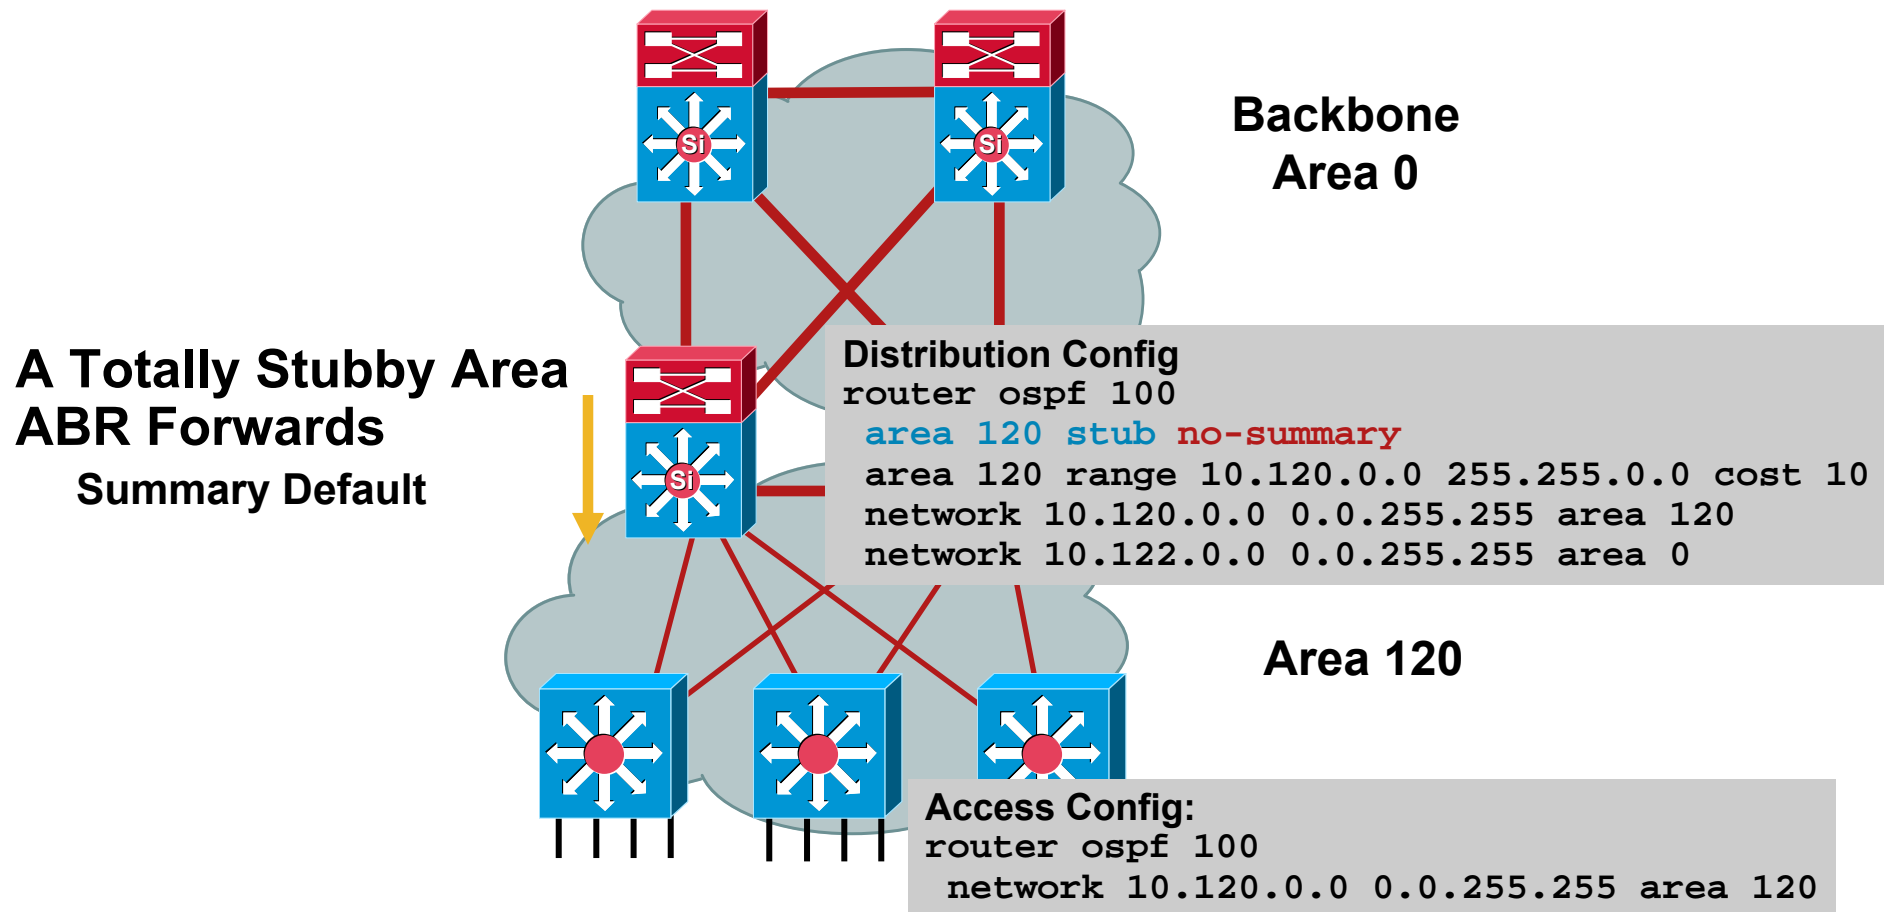
Eliminates External Routes/LSA Present in Area (Type 5)



Totally Stubby Area

Use This for Stable—Scalable Internetworks

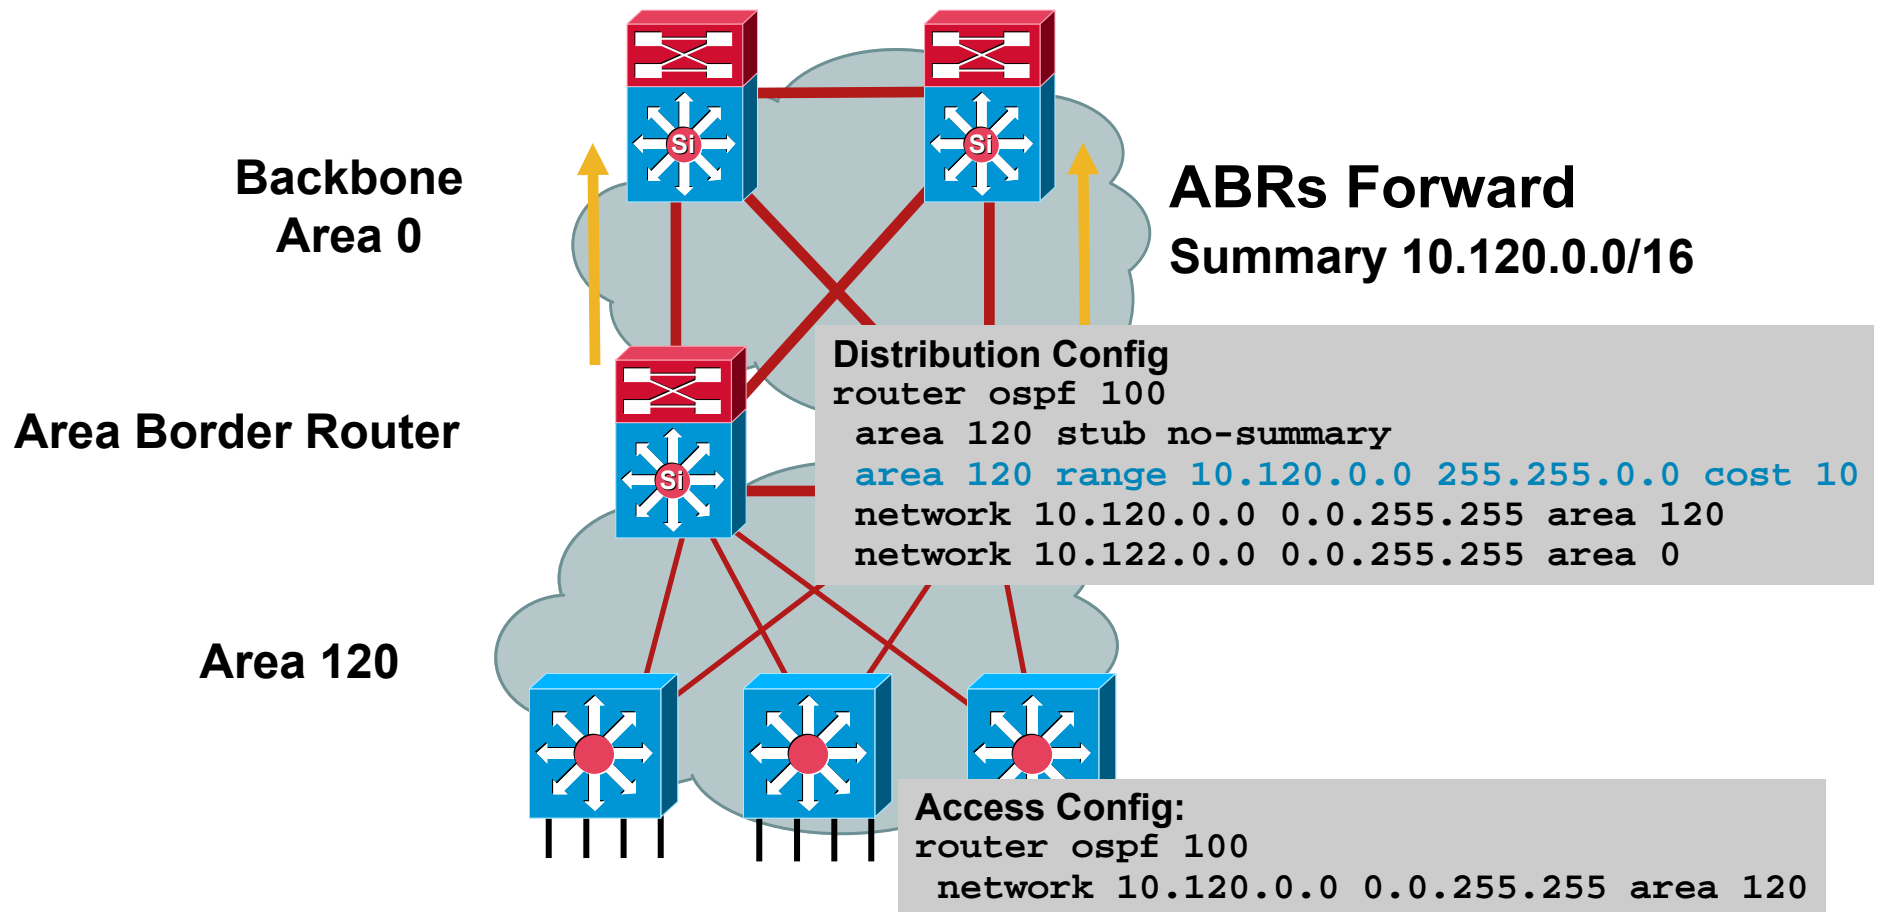
Minimize the Number of LSAs and the Need for Any External Area SPF Calculations



Summarization Distribution to Core

Reduce SPF and LSA Load in Area 0

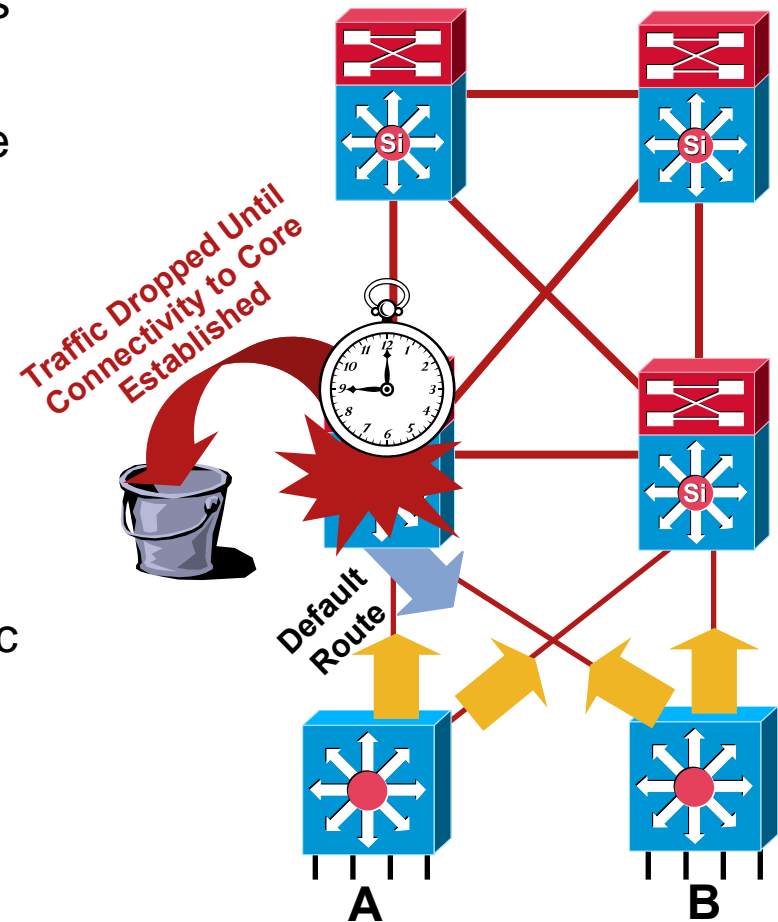
Minimize the Number of LSAs and the Need for Any SPF Recalculations at the Core



OSPF Totally Stubby Area

Default Route Restoration

- Reintroducing Distribution ABR black holes traffic
- Totally stubby areas are used to isolate the access layer switches from route calculations due to events in other areas
- This means that the ABR (the distribution switch) will send a default route to the access layer switch when the neighbor relationship is established
- The default route is sent regardless of the distribution switches ability to forward traffic on to the core (area 0)
- Traffic could be black holed until connectivity to the core is established



Note: Solution to This Anomaly Is Being Investigated

OSPF Timer Tuning

High-Speed Campus Convergence

- OSPF by design has a number of throttling mechanisms to prevent the network from thrashing during periods of instability
- Campus environments are candidates to utilize OSPF timer enhancements

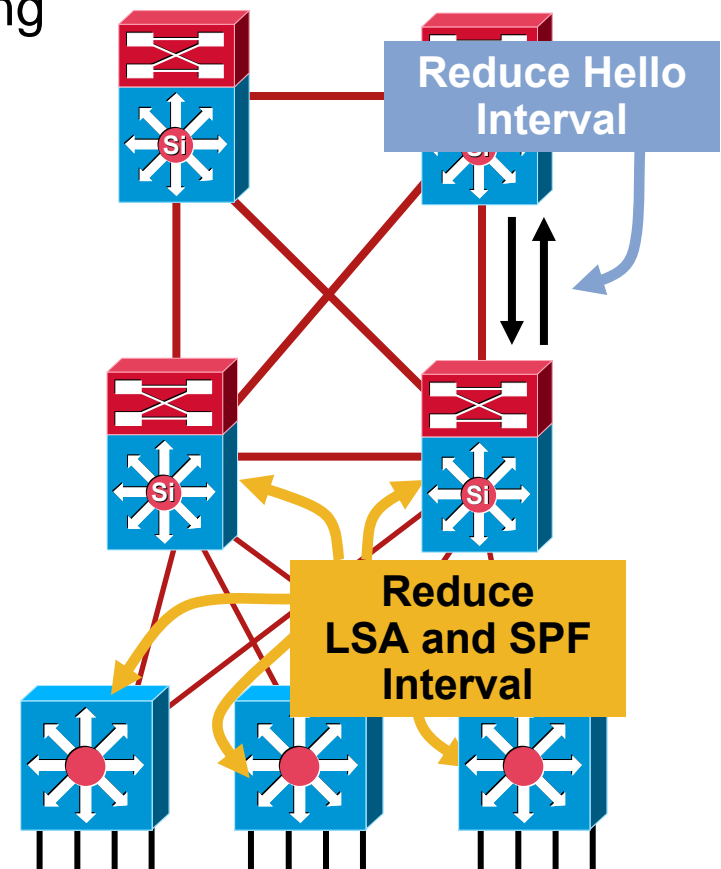
Sub-second hellos

Generic IP (interface) dampening mechanism

Back-off algorithm for LSA generation

Exponential SPF backoff

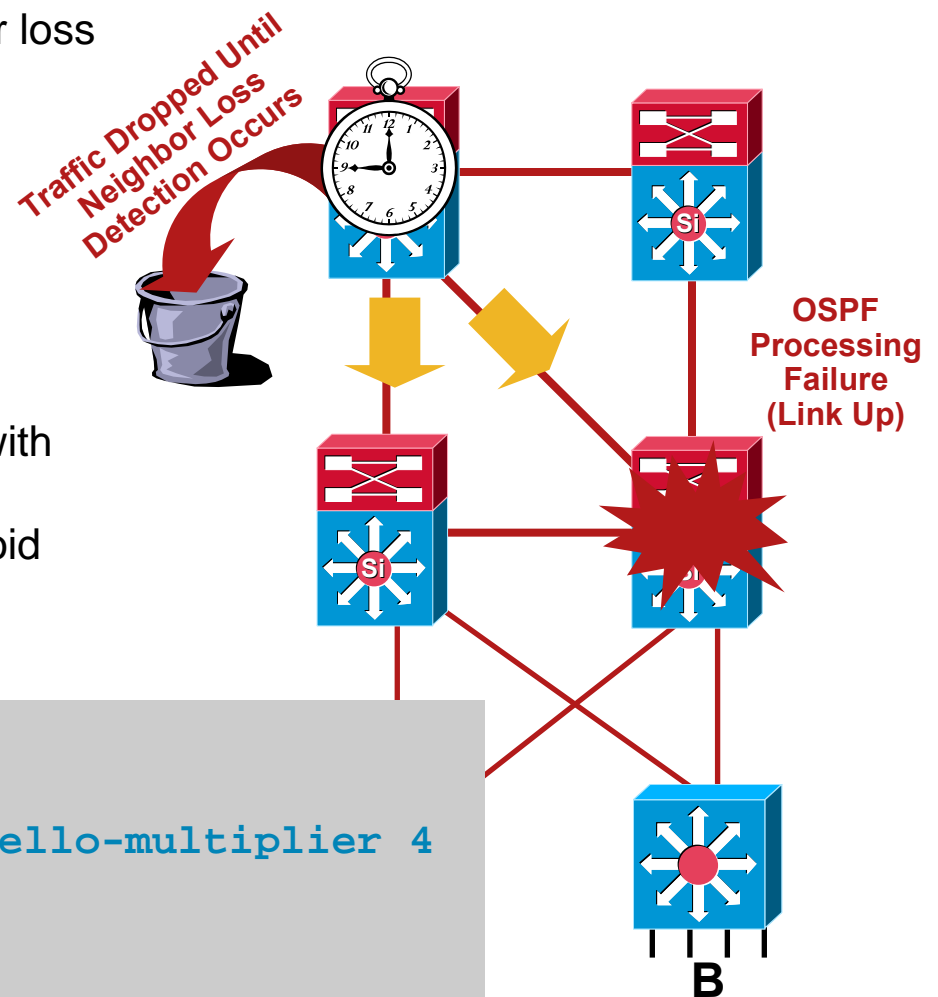
Configurable packet pacing



Subsecond hellos

Neighbor Loss Detection—Physical Link Up

- OSPF hello/dead timers detect neighbor loss in the absence of physical link loss
- Useful in environments where an L2 device separates L3 devices (Layer 2 core designs)
- Aggressive timers quickly detect neighbor failure
- **Not recommended with NSF/SSO**
- Interface **dampening** is recommended with sub-second hello timers
- OSPF point-to-point network type to avoid designated router (DR) negotiation.



Access Config:

```
interface GigabitEthernet1/1
dampening
ip ospf dead-interval minimal hello-multiplier 4
ip ospf network point-to-point
```

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


OSPF LSA Throttling

- By default, there is a 500ms 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

```
timers throttle lsa all <start-interval>  
    <hold-interval> <max-interval>
```


- Acceptance of a new LSAs is limited by the receiver

```
timers lsa arrival <milliseconds>
```

Access Config:

```
interface GigabitEthernet1/1  
    dampening  
    ip ospf dead-int min hello-multiplier 4  
    ip ospf network point-to-point
```

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




OSPF SPF Throttling

- OSPF has an SPF throttling timer designed to dampen route recalculation (preserving CPU resources) when a link bounces
- 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; SPF timer was one second

Access Config:

```
interface GigabitEthernet1/1
  dampening
  ip ospf dead-int min hello-multiplier 4
  ip ospf network point-to-point
```


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



OSPF Routed Access Campus Design

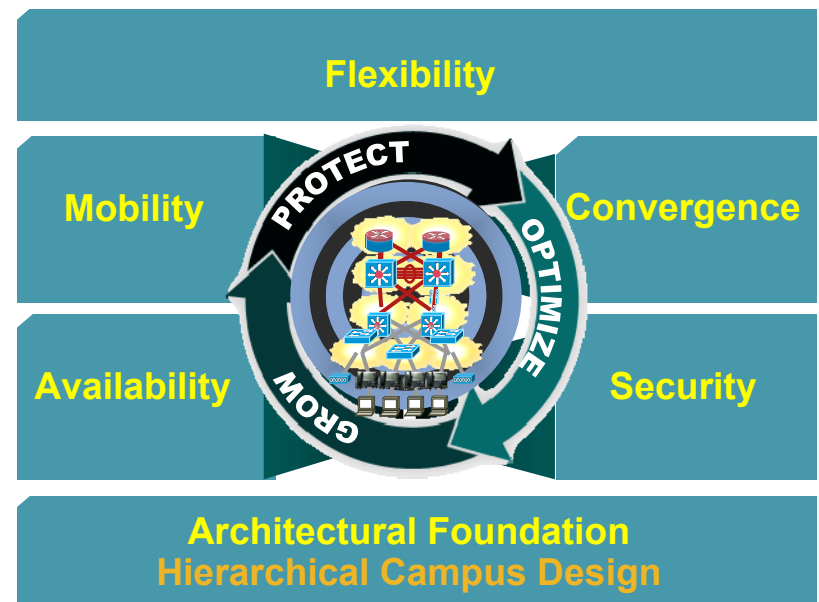
Overview—Fast Convergence

- Detect the event:
 - Decrease the hello-interval and dead-interval to detect soft neighbor failures
 - Enable interface dampening
 - Set carrier-delay = 0
- Propagate the event:
 - Summarize routes between areas to limit LSA propagation across the campus
 - Tune LSA timers to minimize LSA propagation delay
- Process the event:
 - Tune SPF throttles to decrease calculation delays



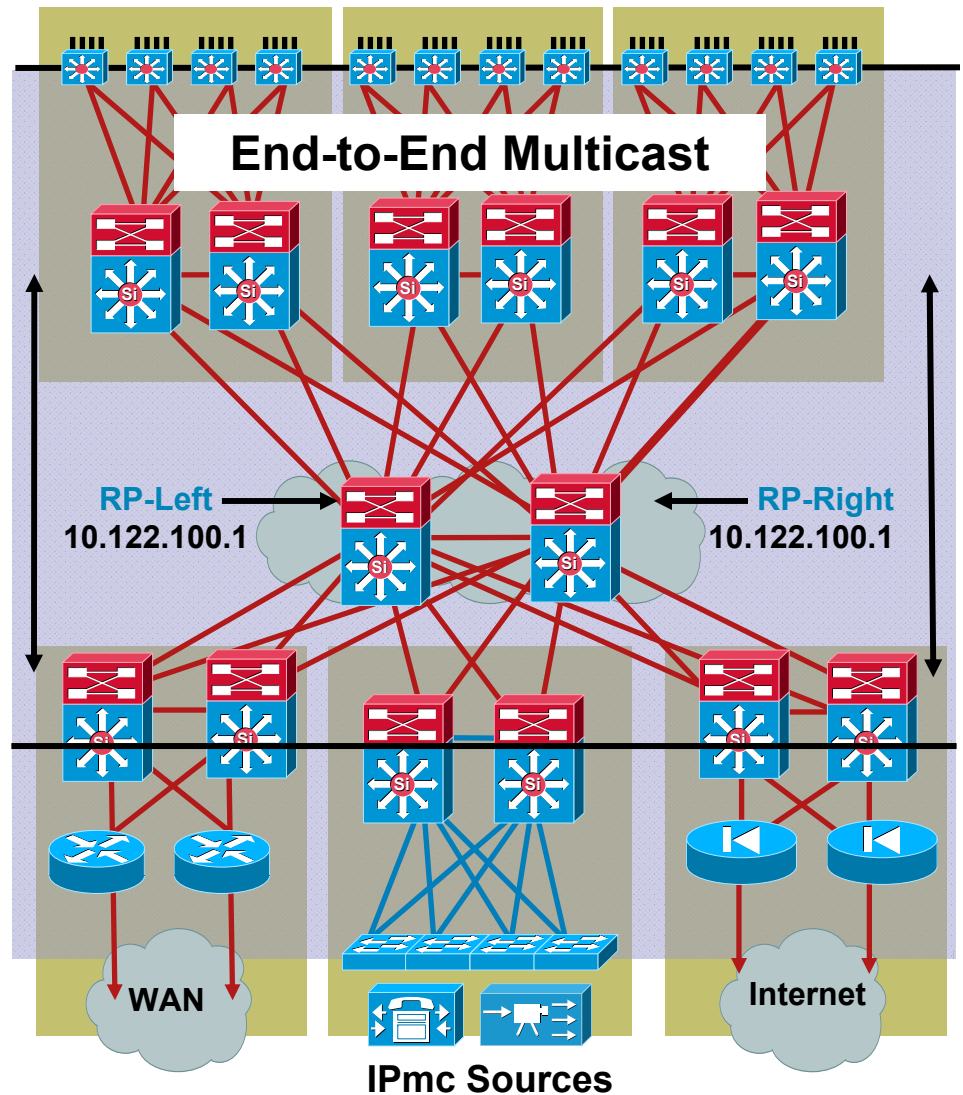
Agenda

- Cisco Campus Architecture
- Campus Network Resiliency
- Routed Campus Design
 - EIGRP Design Details
 - OSPF Design Details
 - PIM Design Details**
- Impact on Advanced Technologies and Services
- Summary



Best Practices— Multicast Deployment

- Use PIM sparse mode
- Use IGMP snooping capable hardware in the access
- Enable PIM sparse mode on routing nodes (core, distribution, and access)
- **Enable PIM on ALL interfaces**
- Use Anycast RP for RP redundancy and fast convergence
- **With OSPF, define the Router-id** so that it does not conflict with the Anycast RP address.
- There are other combinations of RP redundancy; see the multicast session for details



Which PIM Mode—Sparse or Dense?

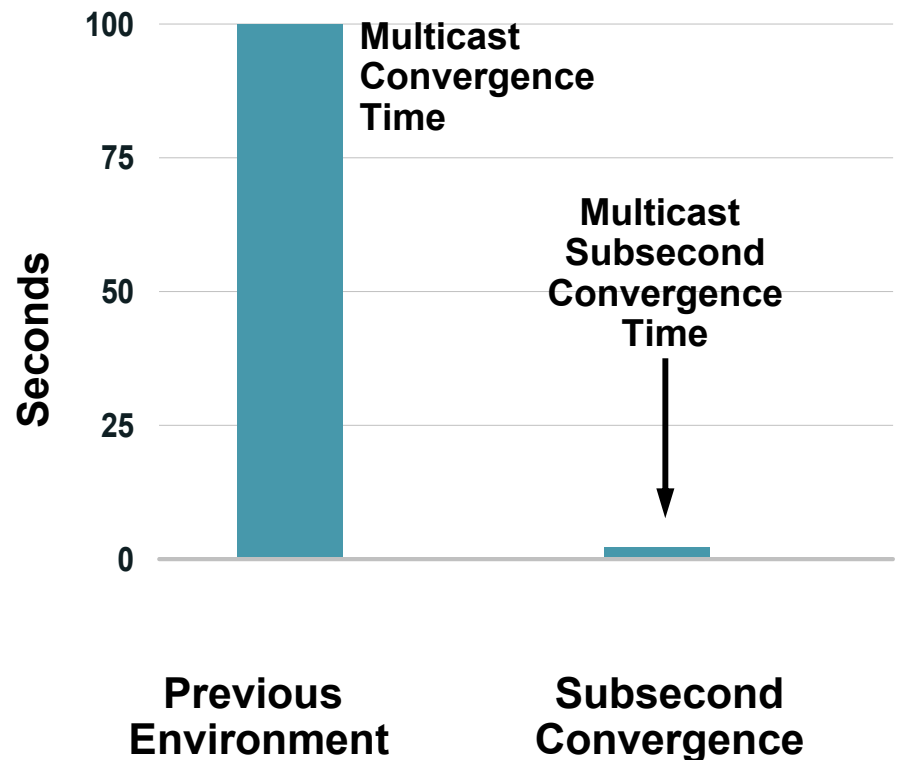
“Sparse mode Good! Dense mode Bad!”

**Source: “The Caveman’s Guide to IP Multicast”,
©2000, R. Davis**

Multicast Sub-Second Convergence

Really a Series of Enhancements

- **Triggered RPF checks following unicast convergence**
Now, as soon as unicast is converged, it causes an instantaneous start of RPF checks (old default was five seconds)
- **Join/prune aggregation**
PIM used to send one packet per (S,G) or (*,G) entry after a Rendezvous Point failover
PIM now aggregates these into only a few PIM packets with multiple entries
- **PIM HELLO option (new)**
Added option which advertises the HOLDDTIME in milliseconds
Allows subsecond failover of designated router



Multicast Subsecond Convergence Provides Almost Instantaneous Recovery of Multicast Paths Following Unicast Routing Recovery

Multicast in the Campus

```
interface loopback 0
 ip address 10.0.0.1 255.255.255.255

interface loopback 1
 ip address 10.0.0.3 255.255.255.255
!
 ip msdp peer 10.0.0.2 connect-source loopback 1
 ip msdp originator-id loopback 1
!
interface TenGigabitEthernet M/Y
 ip address 10.122.0.X 255.255.255.252
 ip pim sparse-mode
!
 ip pim rp-address 10.0.0.1
```

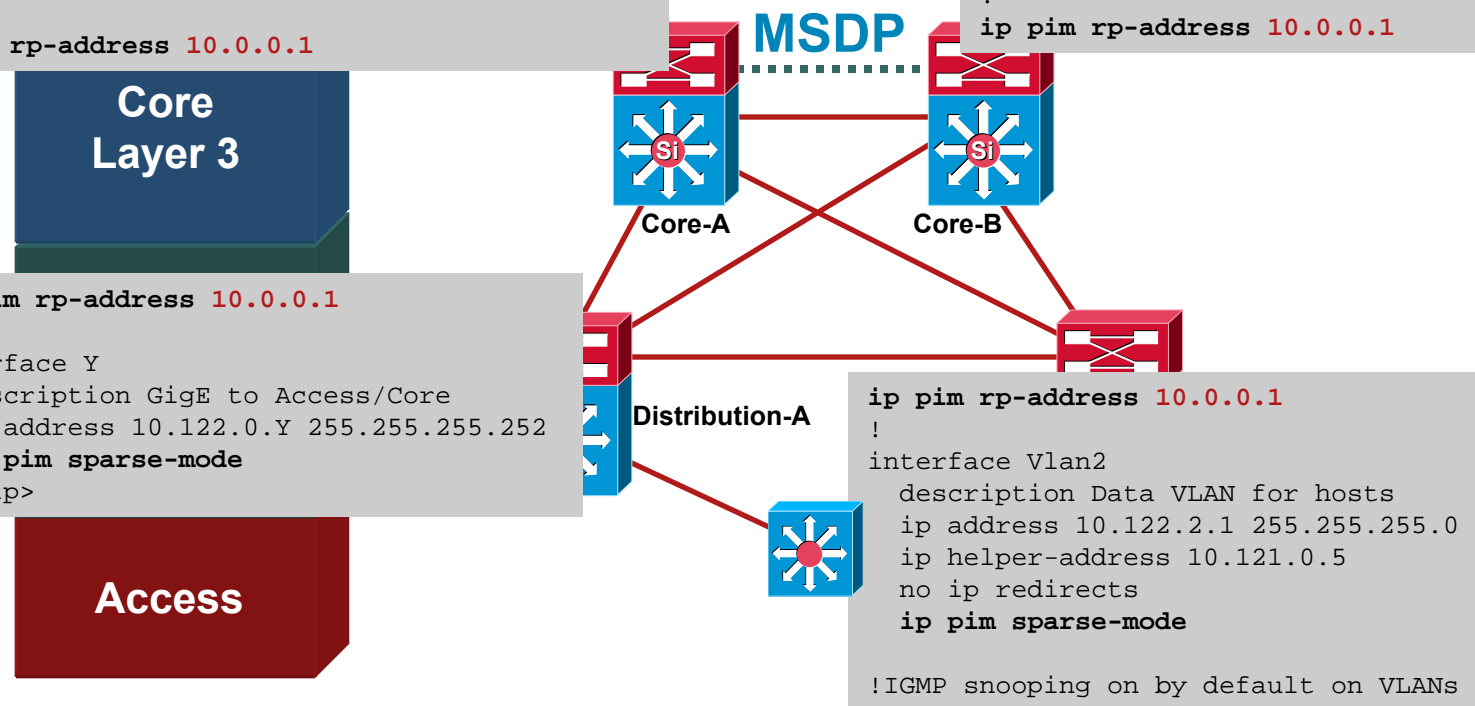
```
interface loopback 0
 ip address 10.0.0.1 255.255.255.255

interface loopback 1
 ip address 10.0.0.2 255.255.255.255
!
 ip msdp peer 10.0.0.3 connect-source loopback 1
 ip msdp originator-id loopback 1
!
interface TenGigabitEthernet M/Y
 ip address 10.122.0.X 255.255.255.252
 ip pim sparse-mode
!
 ip pim rp-address 10.0.0.1
```

```
ip pim rp-address 10.0.0.1
!
interface Y
 description GigE to Access/Core
 ip address 10.122.0.Y 255.255.255.252
 ip pim sparse-mode
!<snip>
```

```
ip pim rp-address 10.0.0.1
!
interface Vlan2
 description Data VLAN for hosts
 ip address 10.122.2.1 255.255.255.0
 ip helper-address 10.121.0.5
 no ip redirects
 ip pim sparse-mode

!IGMP snooping on by default on VLANs
```



What Are The Considerations?

- When using loopback interfaces—you can effect Router IDs (RID) for routing protocols like OSPF and BGP
- Duplicate IP addresses for Anycast RP redundancy is good
- Duplicate IP addresses as RID for OSPF is bad
- Set your RID don't let the router pick one
- There are many other multicast implementation and deployment considerations for details see the Multicast Breakouts and Techtorial


Multicast Sessions:
BRKIPM-3018 - Advanced Multicast Concepts
BRKIPM-2019 - Multicast Security

Multicast in the Routed Access Design

- There is only one default gateway per VLAN/subnet;
 - No PIM designated router timer tuning for fast convergence
 - Asymmetrical unicast and multicast traffic patterns are not a concern when troubleshooting.
- IGMP snooping is limited to VLANs in the wiring closet
- Troubleshooting—PIM show and debug commands on all layers of the network including the wiring closet

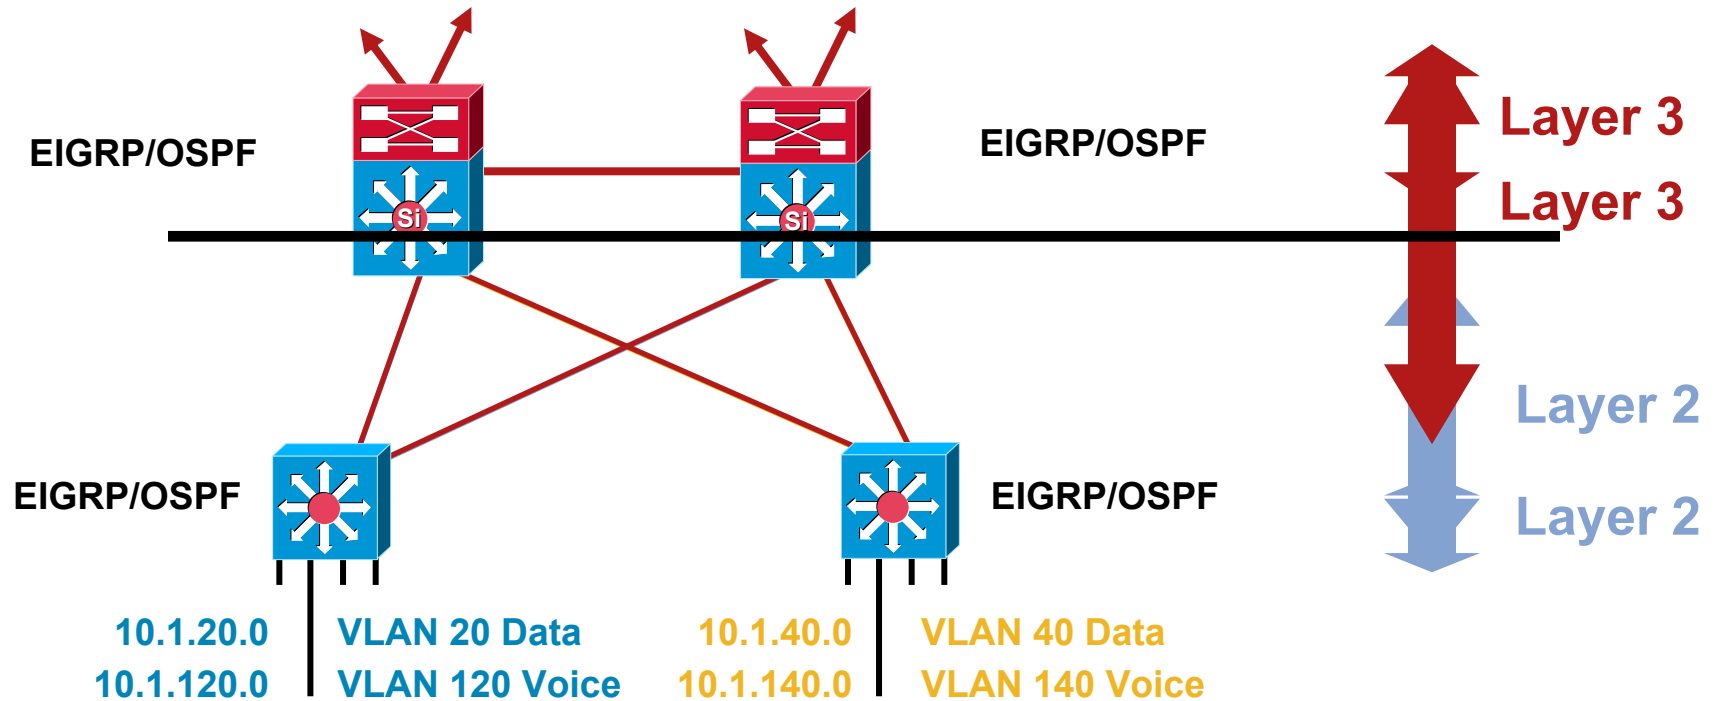
Agenda

- Cisco Campus Architecture
- Campus Network Resiliency
- Routed Campus Design
 - EIGRP Design Details
 - OSPF Design Details
 - PIM Design Details
- **Impact on Advanced Technologies and Services**
- Summary



Routed Campus Design

Impact of Moving the L2/L3 Demark



- The Layer 2/3 demarcation is closer to the network edge
- VLAN design in the access is the same
- Some distribution layer features move to the access layer
- Moving the L2/L3 demark is transparent to most applications


Wireless integration into the Campus

Layer 2 Wireless Roaming

- Layer-2 roaming requires spanning at least two VLANs between wiring closet switches
 - Common 'Trunk' or native VLAN for APs to communicate to WDS
 - The Wireless Voice VLAN
- Adding VLANs to the uplinks in a routed access design eliminates most of the advantages of this design

Spanning Tree, VLAN trunks, slower convergence, complexity

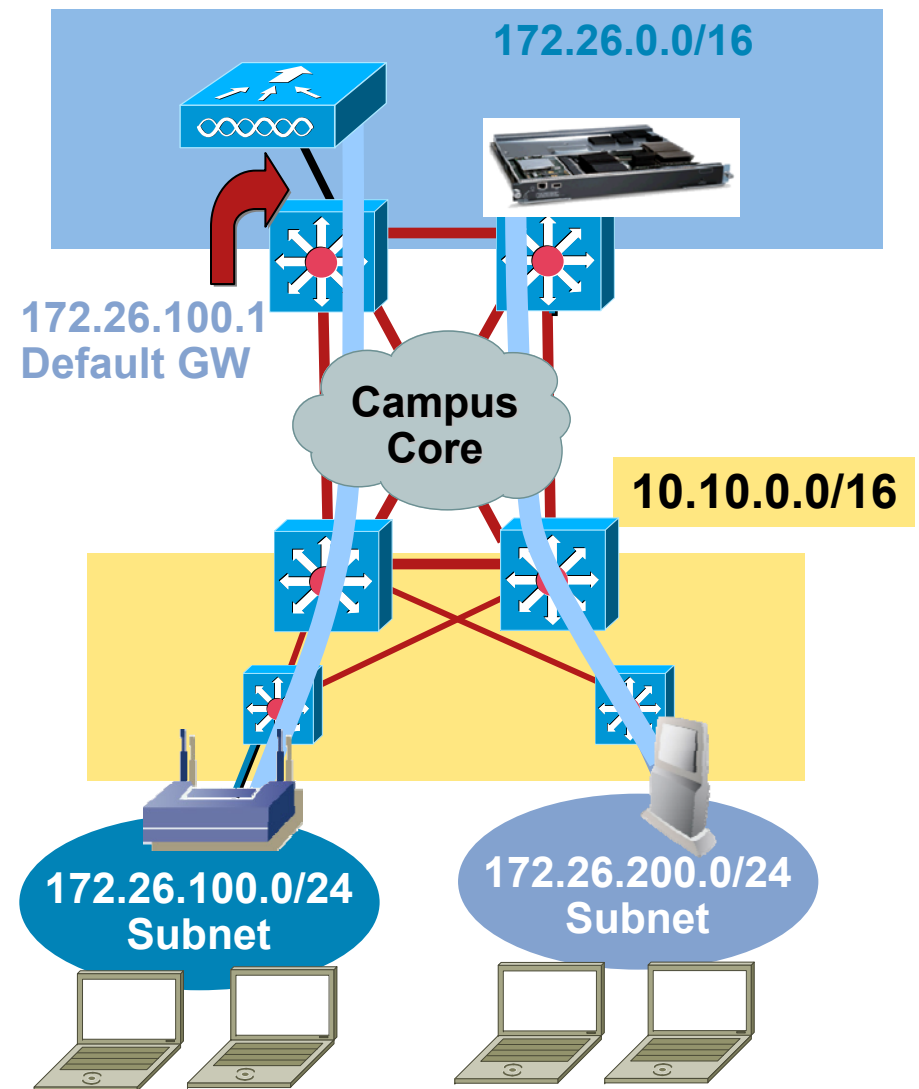
- Layer-2 wireless roaming is **not** recommended for the routed access design



Wireless Design Considerations

Controllers and IP Addressing


- Wireless controllers provide centralized management and fast roaming between APs
- Endpoints keep the same IP address as they move between APs
- The wireless mobile endpoints are addressed out of the **summary range** as defined by the location of the WLAN controller
- The default gateway for all wireless endpoints when using a WLAN controller is the adjacent Cisco Catalyst switch
- Communication between a wired client on an access switch and a wireless client is via the core



Design Considerations for NSF/SSO

NSF and Hello Timer Tuning?


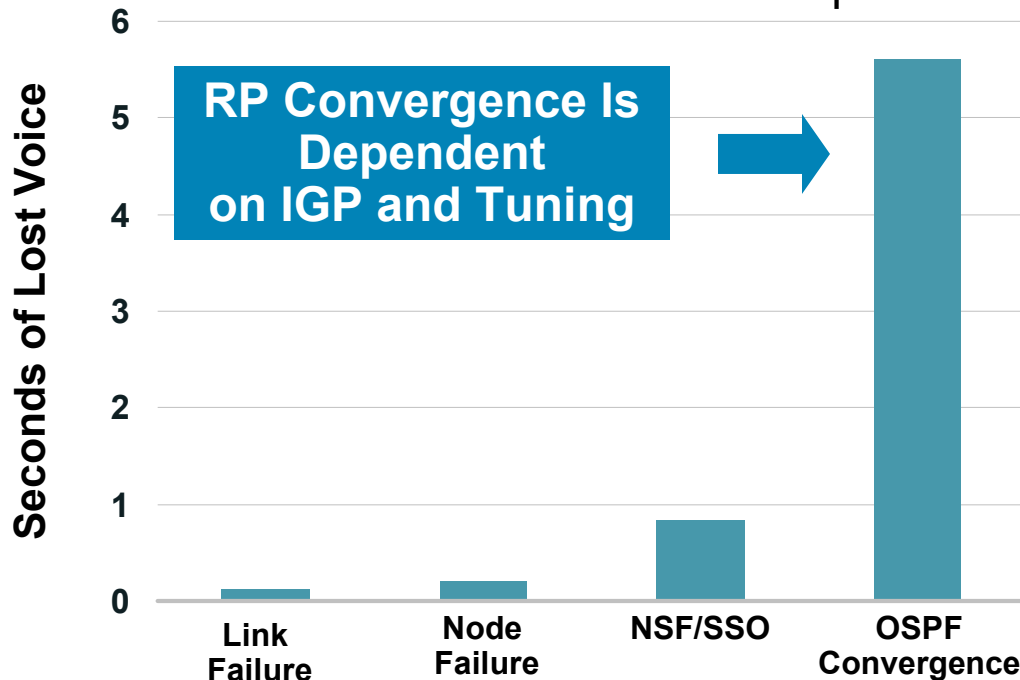
- **NSF** is intended to provide availability through **route convergence avoidance**
- **Fast IGP timers** are intended to provide availability through **fast route convergence**
- In an NSF environment dead timer must be greater than SSO Recovery + RP restart + time to send first hello
- Switches running Native IOS
 - OSPF 2/8 seconds for hello/dead
 - EIGRP 1/4 seconds for hello/hold
- Switches running Hybrid
 - OSPF 3/12 seconds for hello/dead
 - EIGRP 2/8 seconds for hello/hold



Design Considerations for NSF/SSO

Where Does It Make Sense?


- Redundant topologies with equal cost paths provide sub-second convergence
- NSF/SSO provides superior availability in environments with non-redundant paths



See
[BRKCAM-2001 – Multilayer Architecture Principals and Foundational Design Guidelines](#)
[BRKCAM-3005 – Advanced Enterprise Campus High Availability](#)

Cisco Catalyst Integrated Security Features

Hardening Layer 2/3



- **Port Security** prevents MAC flooding attacks
- **DHCP Snooping** prevents client attack on the switch and server
- **Dynamic ARP Inspection** adds security to ARP using DHCP snooping table
- **IP Source Guard** adds security to IP source address using DHCP snooping table

Applied Only in the Access Layer
in the Routed Access Design

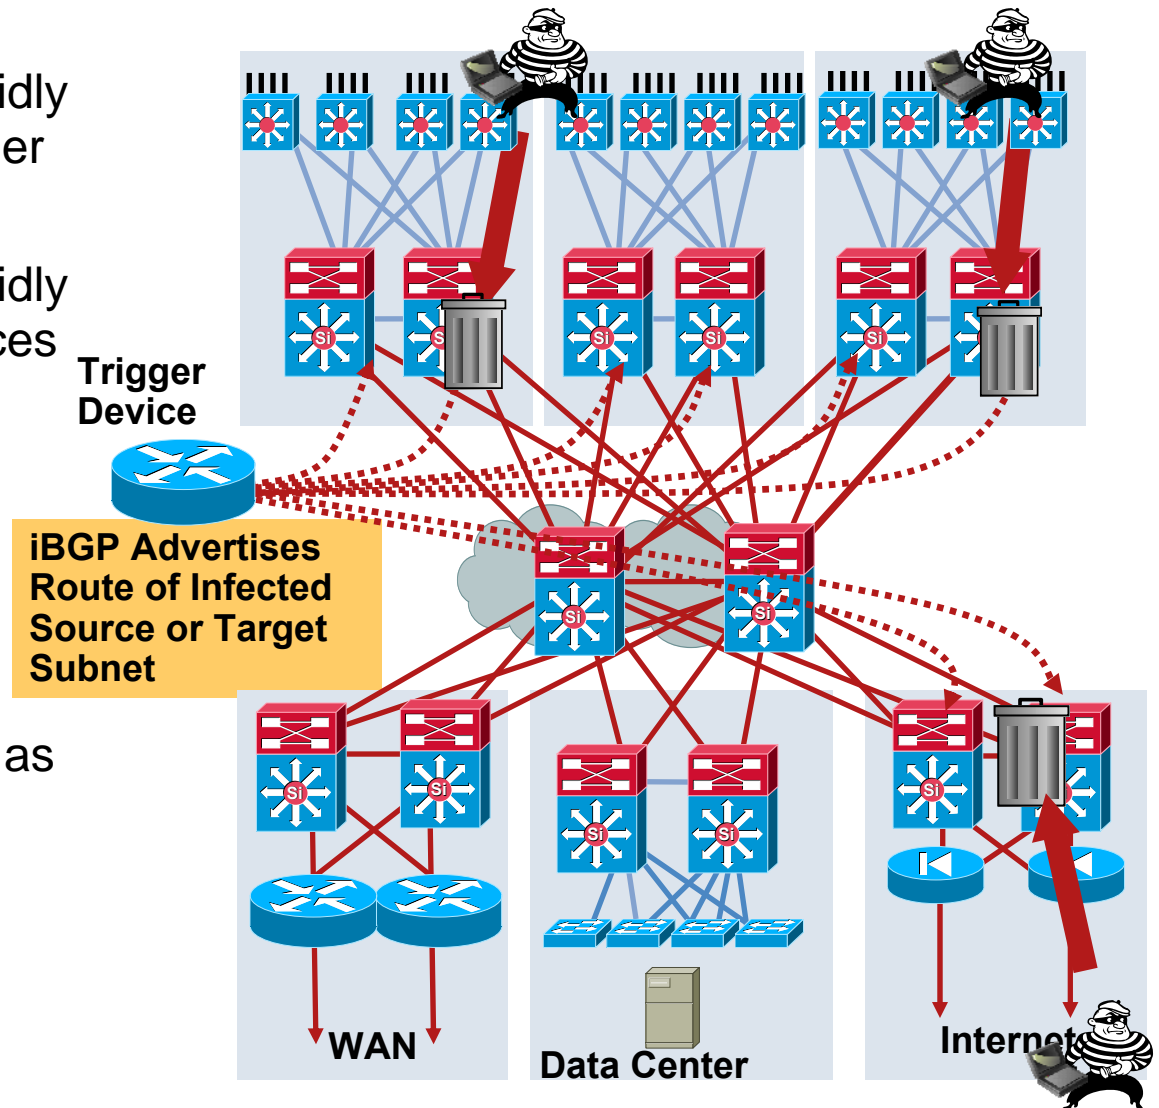
```
ip dhcp snooping
ip dhcp snooping vlan 2-10
ip arp inspection vlan 2-10
!
interface fa3/1
switchport port-security
switchport port-security max 3
switchport port-security violation restrict
switchport port-security aging time 2
switchport port-security aging type inactivity
ip arp inspection limit rate 100
ip dhcp snooping limit rate 100
!
interface gigabit1/1
ip dhcp snooping trust
ip arp inspection trust
```

See
BRKCAM-2008 – Understanding and Preventing Layer 2 attacks

Worm Containment

Remote Triggered Black Holes (RTBH)

- Need a scalable method to rapidly block traffic **to** destinations under attack (DDoS)
- Need a scalable method to rapidly block traffic **from** infected sources (Worm)
- RTBH allows you to do both
- Using iBGP push message to choke points to discard the attack packets
- Does NOT require to use BGP as your routing protocol
- uRPF requires Cisco Catalyst 6500 and Sup720 or Sup32



Worm Containment

Source and Destination Based RTBH

1. Establish iBGP peering between the **trigger device** and each edge router

Create a static route **on each edge** device pointing to Null0 a special (unreachable) IP address (**192.0.2.1**) to be the **sink hole** route

2. Add a static route to the **target** IP address (**10.2.2.2**) on the trigger device.

Enable BGP to advertise this static route to each edge device; set the sink hole IP address (**192.0.2.1**) **as the next-hop**

3. Because the edge routers have a static route for the **sink hole** IP address = Null0, the final FIB entry for the **target** IP address = Null0

Packets **directed** to target will be dropped

4. Enable **uRPF in loose mode**: packets **sourced** from the **target** address (**10.2.2.2**) will be drop

iBGP peering Trigger device ↔ edge routers

ip route **192.0.2.1** 255.255.255.255 Null0 at edge

1

Setup

ip route **10.2.2.2** 255.255.255.255 Null0

BGP update **10.2.2.2**, Next-hop **192.0.2.1**

2

Trigger Device

10.2.2.2 = 192.0.2.1 = Null0

Next-Hop for **10.2.2.2** Is Now Set to Null0

3

Edge Router

Enabling uRPF in Loose Mode Will Discard Packets Sourced from **10.2.2.2**

4

Edge Router

Worm Containment

Source and Destination Based RTBH—Configuration

```
! Apply uRPF in loose mode for source dropping
interface FastEthernet2/0
```

```
ip verify unicast source reachable-via any
```

```
! Define the static route pointing to Null0
```

```
ip route 192.0.2.1 255.255.255.255 Null0
```

```
edge-6500-1#sh ip route 10.2.2.2
```

```
Routing entry for 10.2.2.2/32
```

```
Known via "bgp 400", distance 200, metric 0,
type internal
```

```
Last update from 192.0.2.1 00:00:21 ago
```

```
router bgp 400
```

```
redistribute static route-map STATIC-TO-BGP
```

```
! Define route-map
```

```
route-map STATIC-TO-BGP permit 10
```

```
match tag 66
```

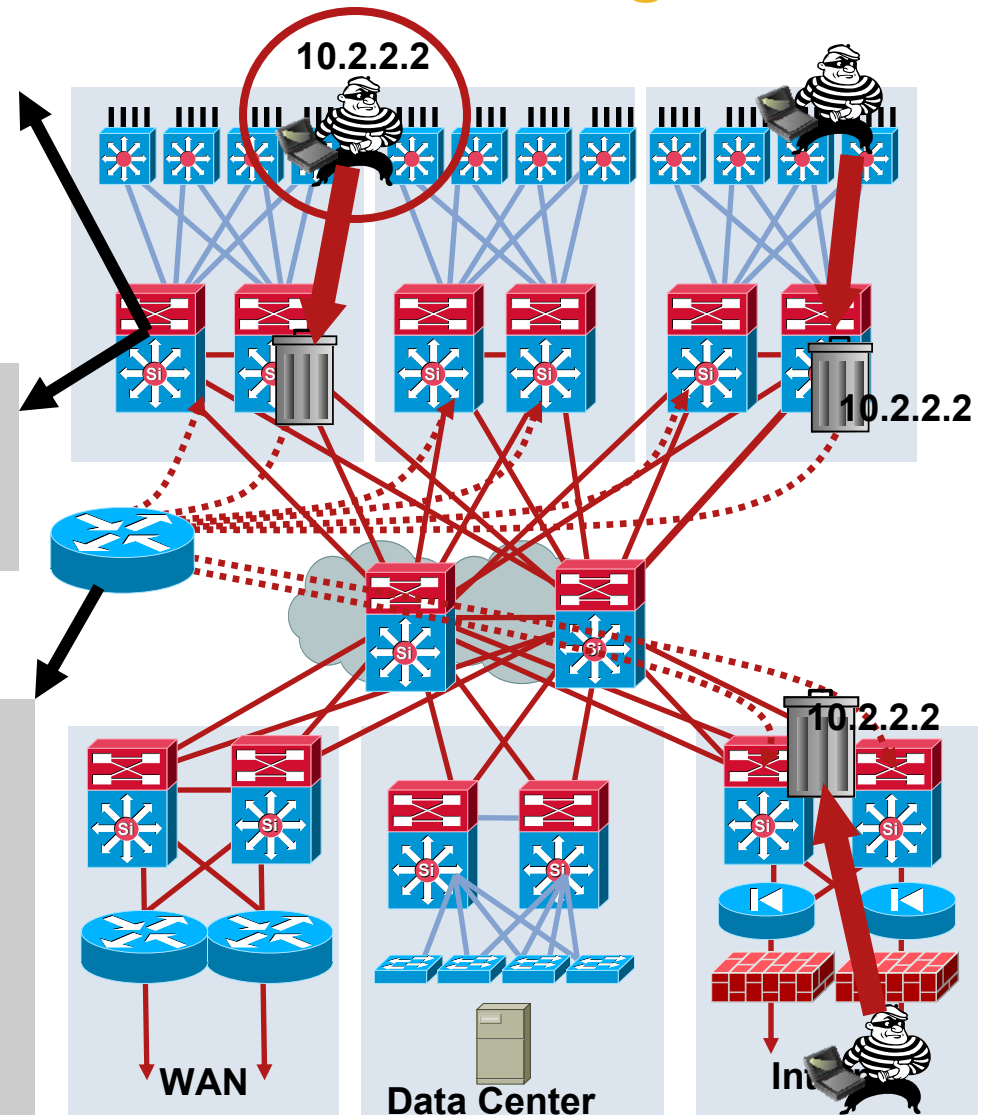
```
set ip next-hop 192.0.2.1
```

```
!
```

```
route-map STATIC-TO-BGP deny 20
```

```
! Define the static route (src or dest address)
```

```
ip route 10.2.2.2 255.255.255.255 Null0 Tag 66
```




IBNS (802.1x) and NAC

Access and Policy Control

- Identity-Based Networking Services (IBNS)
 - Identifies and authenticates the user or device on the network and ensures access to correct network resources
- Network Access Control (NAC)
 - Performs posture validation to ensure that machines not compliant with software posture, and therefore vulnerable to infection, can be isolated to a segment of the network where remediation can take place
- 802.1x provides port-based access control and operates at L2
- MAC Authentication Bypass (MAB) provide port-based access control and also operates at L2
- NAC provides posture assessment and device containment at L3 or L2
- Complimentary functions

See
BRKCAM-2007 for more details on IBNS
BRKSEC-2011 for more details on NAC




Edge Access Control

802.1x Access Control


Dynamic VLAN Assignment & Port Based ACLs

- 802.1x defines an access method for LAN switch ports
- You can permit or deny access based on authorization behavior
- After authentication, user access can be controlled by either Port Based ACLs or permissions of the assigned VLAN.
- The Cisco IBNS solution can dynamically assign a Port Based ACL or dynamically assign a VLAN to control access.
- **The Routed Access Campus Design can support either 802.1x Port Based ACLs or Dynamic VLAN Assignment.**

Port Based ACL Assignment



**Engineering VLAN,
Guest Default VLAN,
MAB Assigned VLAN or,
Auth-Fail Assigned VLAN**




Campus Design for 802.1x and NAC

What Do I Do with Them Once I Have Them in a VLAN?

- 802.1x and NAC with 802.1x control network access based on VLAN assignment
- Once they are assigned to a specific VLAN the network infrastructure needs to keep the traffic isolated
- Potential solutions
 - Distributed ACLs
 - VRF with GRE
 - VRF-lite end-to-end
 - RFC4364 over MPLS
- All provide some form of **network compartmentalization**
- **VLAN/Subnet expansion**

* RFC4364 Obsoletes RFC2547



Virtualized Network Devices

VRF (Virtual Routing and Forwarding)

- VRF allows for the creation of multiple logical forwarding tables

Distinct Routing Information Base (RIB)


Distinct Forwarding Information Base (FIB)

- It is possible to associate with each VRF a group of unique logical data paths, e.g.

802.1q VLANs

GRE tunnels

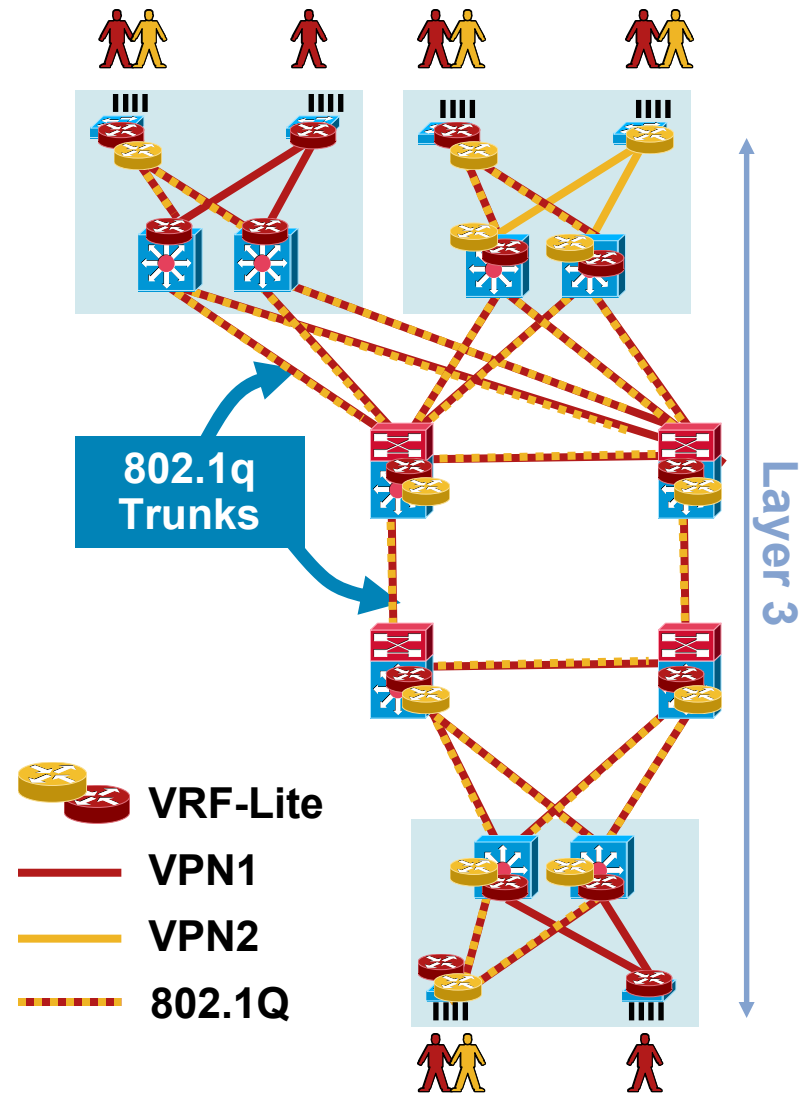
Label Switched Paths (LSPs)



Segmentation and Virtualization

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

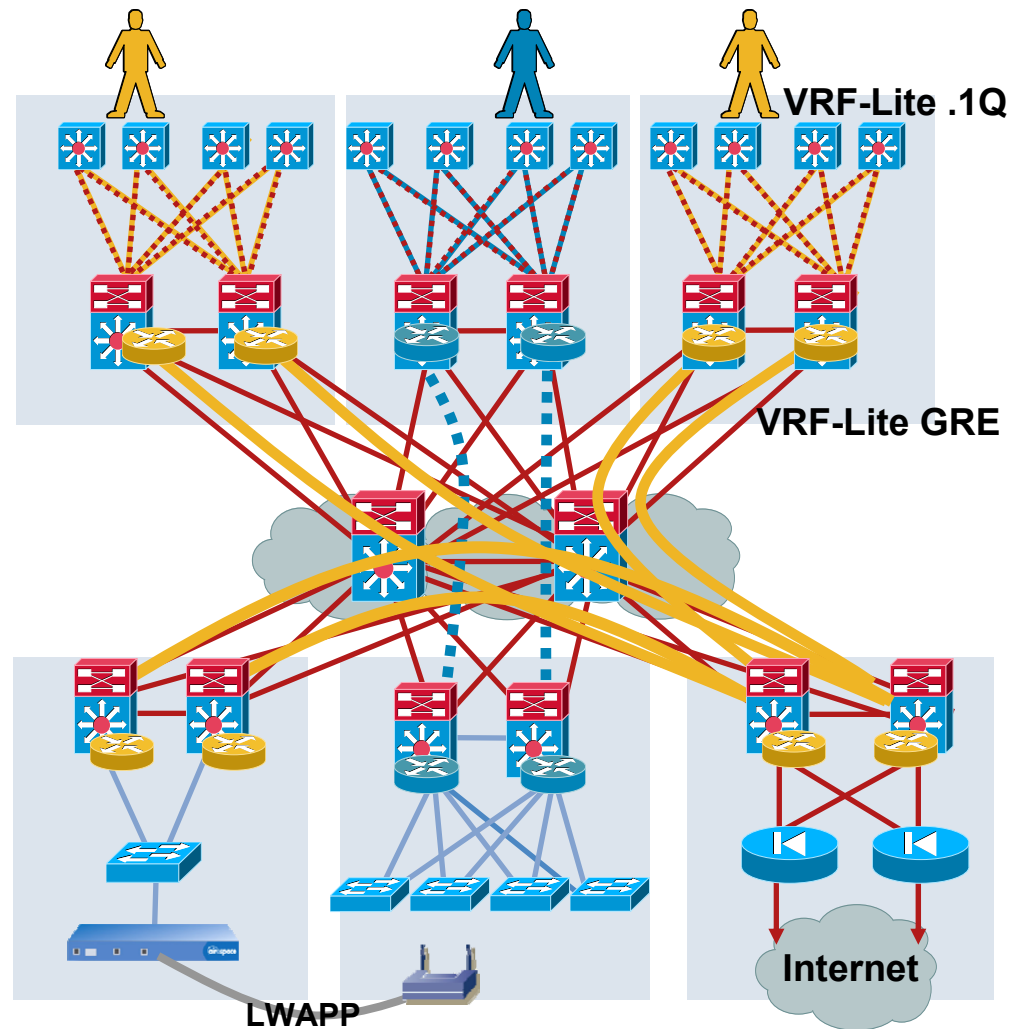
- In the RA design, VRF-lite needs to be extended to the wiring closet
- Every link is a 802.1q trunk
- **VLAN trunks converge slower than pt-2-pt links (250ms+100ms/SVI)**
- No BGP or MPLS
- VRF-lite on all routed hops: access, distribution and core
- 802.1q tags provide single hop data path virtualization
- These trunks do not extend VLANs throughout the campus
- Every physical link carries multiple logical routed links
- Limited scalability: very large configurations in the distribution and core; every VRF on every link



Segmentation and Virtualization

VRF-Lite and GRE

- Create separate logical overlay networks for guest access and NAC remediation traffic
- Limited GRE support on Cisco Catalyst platforms
- Use VRF-lite .1Q between access and distribution, with VRF-lite GRE tunnels from the distribution across the core
- Requires VLAN trunk between access and distribution layer
- Limited scalability; .1Q configuration and GRE tunnel limitations




MPLS-VPN – RFC4364 VPNs

Any-to-Any Connectivity

- Routed Access (CE) with VRF-lite to the Distribution
- Distribution (PE) with VRF-Lite and MPLS to the Core
- Any to any connectivity per user group
- Requires MPLS and BGP in the campus
- Limited Cisco Catalyst MPLS support

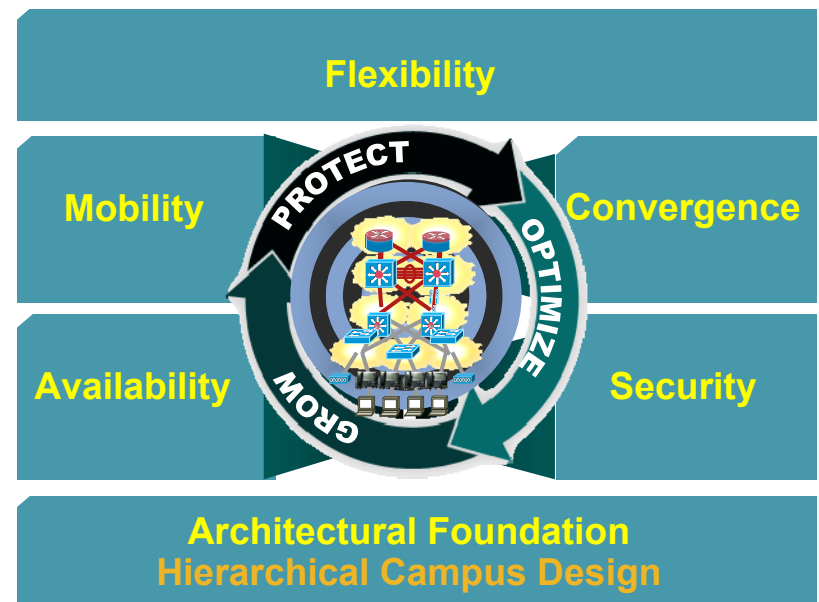
See **BRKCAM-3002**
Techniques for Enterprise Network Virtualization: Using the Infrastructure to Ease Policy Implementation and Enforcement

* RFC4364 Obsoletes RFC2547




Agenda

- Cisco Campus Architecture
- Campus Network Resiliency
- Routed Campus Design
 - EIGRP Design Details
 - OSPF Design Details
 - PIM Design Details
- Impact on Advanced Technologies and Services
- **Summary**



Campus High Availability

Non-Stop Application Delivery




Hierarchical, Systematic Approach


- System level resiliency for switches and routers
- Network resiliency with redundant paths
- Supports integrated services and applications
- Embedded management

Routed Campus Design

Resiliency Advantages? Yes, with a Good Design




- Sub-200 msec convergence for EIGRP and OSPF
- Ease of implementation; fewer things to get right
- Troubleshooting; well known protocols and tools
- Simplified IP multicast deployment
- Considerations; spanning VLANs, IP addressing, IGP selection



Routed Access Design

Summary



- EIGRP or OSPF routed links between access and distribution
- Routed interfaces, not VLAN trunks, between switches
- Equal cost multi path to load balance traffic across network
- Route summarization at distribution with stub routers/areas
- Single (IGP) control plan to configure/manage/troubleshoot

Meet the Experts

Campus and Wireless Evolution

- Mark Montanez
Corporate Dev Consulting Engineer



- Tim Szigeti
Technical Leader



- Sujit Ghosh
Technical Mktg Eng



- Victor Moreno
Technical Leader




- Mike Herbert
Technical Leader



Recommended Reading

BRKCAM -3004

- Routing TCP/IP, Volume I
- OSPF Network Design Solutions
- CCNP Self-Study: Building Scalable Cisco Internetworks (BSCI)
- Optimal Routing Design
- Campus Network Design Fundamentals
- Building Resilient IP Networks



Available in the Cisco Company Store





Backup Slides



Cisco Networkers
2007

EIGRP

Core Layer Configuration

6k-core configuration	
<pre>interface TenGigabitEthernet3/1 description 10GigE to Distribution 1 ip address 10.122.0.29 255.255.255.252 ip pim sparse-mode ip hello-interval eigrp 100 1 ip hold-time eigrp 100 3 ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp carrier-delay msec 0 mls qos trust dscp ! interface TenGigabitEthernet3/2 description 10GigE to Distribution 2 ip address 10.122.0.37 255.255.255.252 ip pim sparse-mode ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp carrier-delay msec 0 mls qos trust dscp</pre>	<pre>! router eigrp 100 network 10.0.0.0 no auto-summary</pre>

EIGRP

Distribution Layer Configuration

6k-distribution configuration

```
interface GigabitEthernet3/2
  description typical link to Access neighbor
  ip address 10.120.0.50 255.255.255.252
  ip pim sparse-mode
  ip hello-interval eigrp 100 1
  ip hold-time eigrp 100 3
  ip authentication mode eigrp 100 md5
  ip authentication key-chain eigrp 100 eigrp
  carrier-delay msec 0
  mls qos trust dscp
!
interface TenGigabitEthernet4/3
  description 10GigE to Distribution neighbor
  ip address 10.120.0.22 255.255.255.252
  ip pim sparse-mode
  ip hello-interval eigrp 100 1
  ip hold-time eigrp 100 3
  ip authentication mode eigrp 100 md5
  ip authentication key-chain eigrp 100 eigrp
  mls qos trust dscp
```

```
interface TenGigabitEthernet4/2
  description 10 GigE to Core neighbor
  ip address 10.122.0.38 255.255.255.252
  ip pim sparse-mode
  ip authentication mode eigrp 100 md5
  ip authentication key-chain eigrp 100 eigrp
  ip summary-address eigrp 100 10.120.0.0
  255.255.0.0 5
  mls qos trust dscp
!
router eigrp 100
  network 10.0.0.0
  distribute-list Default out GigabitEthernet3/1
  distribute-list Default out GigabitEthernet3/2
  ...
  distribute-list Default out
  GigabitEthernet9/15
  no auto-summary
!
ip access-list standard Default
  permit 0.0.0.0
  permit 10.0.0.0
```

EIGRP

Access Layer Configuration

Catalyst 4507 configuration	
<pre>interface GigabitEthernet2/1 description cr3-6500-2 Distribution no switchport ip address 10.120.0.53 255.255.255.252 ip hello-interval eigrp 100 1 ip hold-time eigrp 100 3 ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp ip pim sparse-mode carrier-delay msec 0 qos trust dscp tx-queue 3 priority high ! interface FastEthernet3/5 description Host port w/ IP Phone switchport access vlan 4 switchport mode access switchport voice vlan 104 qos trust cos tx-queue 3 priority high spanning-tree portfast spanning-tree bpduguard enable</pre>	<pre>interface Vlan4 ip address 10.120.4.1 255.255.255.0 ip helper-address 10.121.0.5 no ip redirects ip pim sparse-mode ip igmp snooping fast-leave ! interface Vlan104 ip address 10.120.104.1 255.255.255.0 ip helper-address 10.121.0.5 no ip redirects ip pim sparse-mode ip igmp snooping fast-leave ! router eigrp 100 passive-interface default no passive-interface GigabitEthernet1/1 no passive-interface GigabitEthernet2/1 network 10.0.0.0 no auto-summary eigrp stub connected</pre>

OSPF

Core Layer Configuration

6k-core configuration

```
interface Port-channel1
  description Channel to Peer Core node
  dampening
  ip address 10.122.0.19 255.255.255.254
  ip pim sparse-mode
  ip ospf network point-to-point
  ip ospf dead-interval minimal hello-multipl 4
  load-interval 30
  carrier-delay msec 0
  mls qos trust dscp
!

interface TenGigabitEthernet3/1
  description 10GigE to Distribution 1
  dampening
  ip address 10.122.0.20 255.255.255.254
  ip pim sparse-mode
  ip ospf network point-to-point
  ip ospf dead-interval minimal hello-multipl 4
  load-interval 30
  carrier-delay msec 0
  mls qos trust dscp
!
```

```
router ospf 100
  router-id 10.122.10.2
  log-adjacency-changes
  timers throttle spf 10 100 5000
  timers throttle lsa all 10 100 5000
  timers lsa arrival 80
  passive-interface Loopback0
  passive-interface Loopback1
  passive-interface Loopback2
  network 10.122.0.0 0.0.255.255 area 0.0.0.0
!
```

OSPF

Distribution Layer Configuration

6k-dist-left configuration	
<pre>interface GigabitEthernet3/2 description 3750 Access Switch dampening ip address 10.120.0.8 255.255.255.254 ip pim sparse-mode ip ospf network point-to-point ip ospf dead-interval minimal hello-multipl 4 load-interval 30 carrier-delay msec 0 mls qos trust dscp ! interface TenGigabitEthernet4/1 description 10 GigE to Core 1 dampening ip address 10.122.0.26 255.255.255.254 ip pim sparse-mode ip ospf network point-to-point ip ospf dead-interval minimal hello-multipl 4 load-interval 30 carrier-delay msec 0 mls qos trust dscp</pre>	<pre>router ospf 100 router-id 10.122.102.1 log-adjacency-changes area 120 stub no-summary area 120 range 10.120.0.0 255.255.0.0 timers throttle spf 10 100 5000 timers throttle lsa all 10 100 5000 timers lsa arrival 80 network 10.120.0.0 0.0.255.255 area 120 network 10.122.0.0 0.0.255.255 area 0</pre>

OSPF

Access Layer Configuration

3750-Access configuration	
<pre>interface GigabitEthernet1/0/1 description Uplink to Distribution 1 no switchport dampening ip address 10.120.0.9 255.255.255.254 ip pim sparse-mode ip ospf network point-to-point ip ospf dead-interval minimal hello-multipl 4 load-interval 30 carrier-delay msec 0 srr-queue bandwidth share 10 10 60 20 srr-queue bandwidth shape 10 0 0 0 mls qos trust dscp auto qos voip trust interface FastEthernet2/0/1 description Host port with IP Phone switchport access vlan 2 switchport voice vlan 102 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 auto qos voip cisco-phone spanning-tree portfast spanning-tree bpduguard enable</pre>	<pre>interface Vlan2 description Data VLAN for 3750 Data ip address 10.120.2.1 255.255.255.0 ip helper-address 10.121.0.5 no ip redirects ip pim sparse-mode ip igmp snooping fast-leave ! interface Vlan102 description Voice VLAN for 3750-access ip address 10.120.102.1 255.255.255.0 ip helper-address 10.121.0.5 no ip redirects ip pim sparse-mode ip igmp snooping fast-leave ! router ospf 100 router-id 10.120.250.2 log-adjacency-changes area 120 stub no-summary timers throttle spf 10 100 5000 timers throttle lsa all 10 100 5000 timers lsa arrival 80 passive-interface default no passive-interface GigabitEthernet1/0/1 no passive-interface GigabitEthernet3/0/1 network 10.120.0.0 0.0.255.255 area 120</pre>

PIM

Core Layer RP Configuration—1

6k-core Left Anycast-RP configuration	6k-core Right Anycast-RP configuration
<pre>ip multicast-routing ! interface Loopback0 description MSDP PEER INT ip address 10.122.10.1 255.255.255.255 ! interface Loopback1 description ANYCAST RP ADDRESS ip address 10.122.100.1 255.255.255.255 ! interface Loopback2 description Garbage-CAN RP ip address 2.2.2.2 255.255.255.255 ! interface TenGigabitEthernet M/Y ip address 10.122.0.X 255.255.255.252 ip pim sparse-mode ! ip pim rp-address 2.2.2.2 ip pim rp-address 10.122.100.1 GOOD-IPMC override ip pim accept-register list PERMIT-SOURCES ip msdp peer 10.122.10.2 connect-source Loopback0 ip msdp description 10.122.10.2 ANYCAST-PEER-6k-core-right ip msdp originator-id Loopback0</pre>	<pre>ip multicast-routing ! interface Loopback0 description MSDP PEER INT ip address 10.122.10.2 255.255.255.255 ! interface Loopback1 description ANYCAST RP ADDRESS ip address 10.122.100.1 255.255.255.255 ! interface Loopback2 description Garbage-CAN RP ip address 2.2.2.2 255.255.255.255 ! interface TenGigabitEthernet M/Z ip address 10.122.0.X 255.255.255.252 ip pim sparse-mode ! ip pim rp-address 2.2.2.2 ip pim rp-address 10.122.100.1 GOOD-IPMC override ip pim accept-register list PERMIT-SOURCES ip msdp peer 10.122.10.1 connect-source Loopback0 ip msdp description 10.122.10.1 ANYCAST-PEER-6k-core-left ip msdp originator-id Loopback0</pre>

PIM

Distribution and Access Layer

6k-dist-left configuration	4507k-access configuration
<pre>ip multicast-routing ! interface Loopback2 description Garbage-CAN RP ip address 2.2.2.2 255.255.255.255 ! interface Y description GigE to Access/Core ip address 10.122.0.Y 255.255.255.252 ip pim sparse-mode !<snip> ! ip pim rp-address 10.122.100.1 GOOD-IPMC override ip pim rp-address 2.2.2.2 ! ip access-list standard Default permit 10.0.0.0 ip access-list standard GOOD-IPMC permit 224.0.1.39 permit 224.0.1.40 permit 239.192.240.0 0.0.3.255 permit 239.192.248.0 0.0.3.255</pre>	<pre>ip multicast-routing ip igmp snooping vlan 4 immediate-leave ip igmp snooping vlan 104 immediate-leave no ip igmp snooping ! interface VlanX ip address 10.120.X.1 255.255.255.0 ip helper-address 10.121.0.5 no ip redirects ip pim sparse-mode ! ip pim rp-address 10.122.100.1 GOOD-IPMC override ! ip access-list standard Default permit 10.0.0.0 ip access-list standard GOOD-IPMC permit 224.0.1.39 permit 224.0.1.40 permit 239.192.240.0 0.0.3.255 permit 239.192.248.0 0.0.3.255</pre>

PIM

Core Layer RP Configuration—2

6k-core Left Anycast-RP configuration	6k-core Right Anycast-RP configuration
<pre>! Continued from previous slide ! ip access-list standard GOOD-IPMC permit 224.0.1.39 permit 224.0.1.40 permit 239.192.240.0 0.0.3.255 permit 239.192.248.0 0.0.3.255 ! ip access-list extended PERMIT-SOURCES permit ip 10.121.0.0 0.0.255.255 239.192.240.0 0.0.3.255 permit ip 10.121.0.0 0.0.255.255 239.192.248.0 0.0.3.255</pre>	<pre>! Continued from previous slide ! ip access-list standard GOOD-IPMC permit 224.0.1.39 permit 224.0.1.40 permit 239.192.240.0 0.0.3.255 permit 239.192.248.0 0.0.3.255 ! ip access-list extended PERMIT-SOURCES permit ip 10.121.0.0 0.0.255.255 239.192.240.0 0.0.3.255 permit ip 10.121.0.0 0.0.255.255 239.192.248.0 0.0.3.255</pre>