



# Лучшие практики при проектировании/внедрении классической ЛВС/LAN

Что нужно сделать сразу и чего следует избегать в ЛВС?

Денис Коденцев, Старший архитектор EN, CCIE  
март 2021


# О чем пойдет речь?

L2 – топологии, STP, EtherChannel

L3 – топологии, маршрутизация,  
суммаризация и прочие вопросы

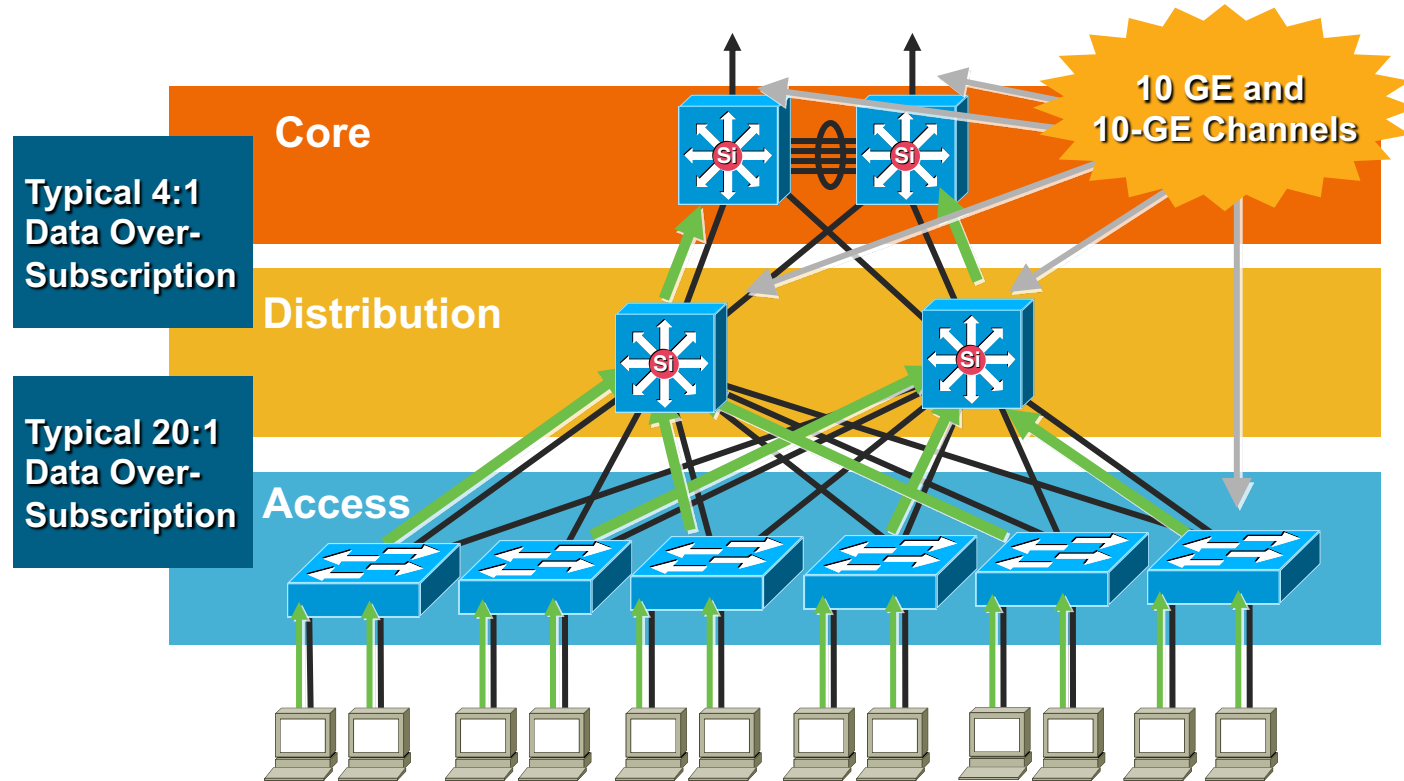
Быстрая сходимость – собирая все вместе

# L2 – топологии, STP и EtherChannel



# EtherChannels


10/100/1000 How Do You Aggregate It?



# Resiliency features for LAN switches

## Global LAN switch configuration

- Rapid PVST+ – improved topology change detection over classic STP Layer 2 loop detection
- BPDUguard default – detect spanning tree BPDUs on portfast-enabled ports for L2 loop prevention
- UDLD – detect and protect against unidirectional links caused by incorrect physical interconnects that can cause spanning tree loops
- Error disable recovery – allows recovery without intervention of automatically disabled ports, post-event
- VTP transparent – ignore VTP updates to avoid accidental outages from unplanned VLAN changes
- Load-interval – reduce time to compute interface load for better visibility to traffic bursts




### Protection across the LAN

```
spanning-tree mode rapid-pvst
spanning-tree portfast bpduguard default
udld enable
errdisable recovery cause all
vtp mode transparent
load-interval 30
```

# Layer 2 Hardening

Spanning Tree Should Behave the Way You Expect



- Place the root where you want it
  - Root primary/secondary macro
- 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
  - Port-security



# EtherChannel member interfaces

## Uplink interface configuration

- Layer 2 EtherChannels are used to interconnect the switch to upstream devices.
- Member interfaces should be on different switches or linecards for resiliency.
- Configure the physical interfaces before configuring the logical portchannel interface.
- Uses LACP for EtherChannel protocol
- Add egress QoS macro for trust inbound traffic and queue outbound (if not using Application Policy or EasyQoS)



```
interface range [type] [port], [type] [port]
 switchport
 channel-protocol lacp
 channel-group 10 mode active
 macro apply EgressQoS
 logging event link-status
 logging event trunk-status
 logging event bundle-status
```



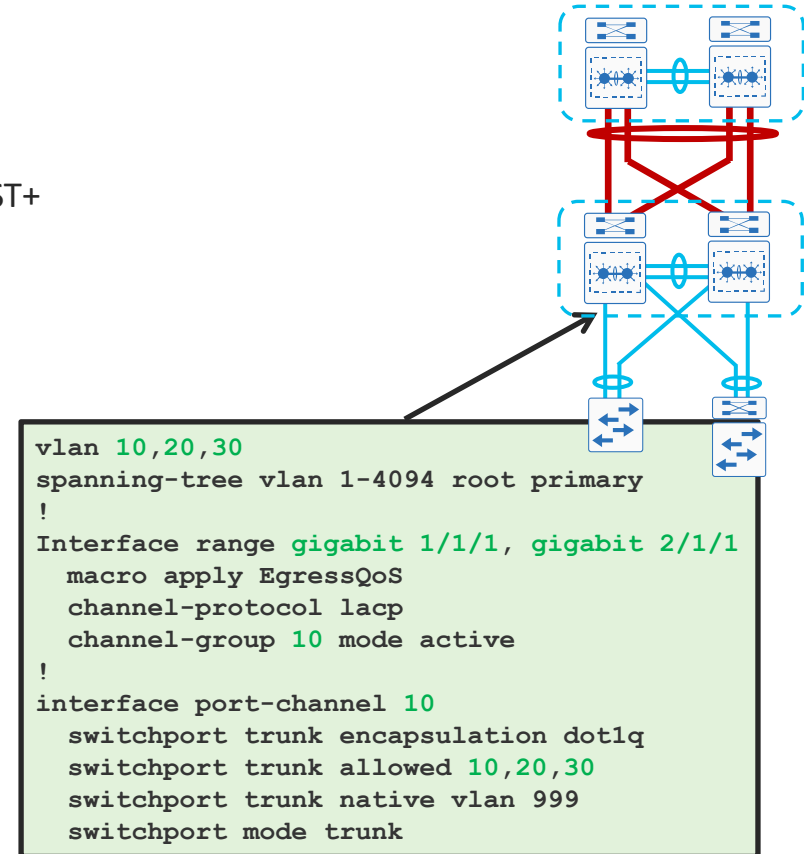
Note: ISR routers do not support LACP. Therefore, when connecting a remote site access switch to an ISR router with an EtherChannel you must configure the switch with mode forced on.

```
interface range [type] [port], [type] [port]
 switchport
 channel-group 10 mode on
 macro apply EgressQoS
```

# Layer 2 connectivity to access layer

## LAN distribution layer

- Configure Layer 2
  - With hub-and-spoke design, no STP loops, still enable RPVST+
  - Configure VLANs servicing access layer
  - Set distribution layer to be STP root for access layer VLANs
- Configure EtherChannel member interfaces
  - Uses LACP for EtherChannel protocol
  - For Layer 2 EtherChannel, configure physical interfaces prior to logical interface
  - Apply egress QoS macro (if not using Application Policy or EasyQoS)
- Configure 802.1Q trunk on EtherChannel logical port (port-channel) interface






# Multichassis EtherChannel load sharing



Reference

- MEC hash algorithm is computed independently by each virtual-switch to perform load share via its local physical ports.
- 8 bits computation on each member link of an MEC is independently done on per virtual-switch node basis.
- Total number of member link bundling in single MEC recommendation remains consistent as described in single chassis EtherChannel section.
- Recommendation to deploy EtherChannel in 2n ratio evenly distributed to each virtual-switch for best load-sharing result.



| Per Switch MEC Flow Distribution Matrix |           |           |           |           |           |           |           |           |
|---|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| Member Links                            | Port1 Bit | Port2 Bit | Port3 Bit | Port4 Bit | Port5 Bit | Port6 Bit | Port7 Bit | Port8 Bit |
| 1                                       | 8         | X         | X         | X         | X         | X         | X         | X         |
| 2                                       | 4         | 4         | X         | X         | X         | X         | X         | X         |
| 3                                       | 3         | 3         | 2         | X         | X         | X         | X         | X         |
| 4                                       | 2         | 2         | 2         | 2         | X         | X         | X         | X         |
| 5                                       | 2         | 2         | 2         | 1         | 1         | X         | X         | X         |
| 6                                       | 2         | 2         | 1         | 1         | 1         | 1         | X         | X         |
| 7                                       | 2         | 1         | 1         | 1         | 1         | 1         | 1         | X         |
| 8                                       | 1         | 1         | 1         | 1         | 1         | 1         | 1         | 1         |

Recommended MEC Bundle link configuration

# Optimize EtherChannel load balancing



Reference

- Load share egress data traffic based on input hash
- Optimal load sharing results with :
  - Bucket-based load-sharing –  
Bundle member-links in power-of-2 (2/4/8)
  - Multiple variation of input for hash (L2 to L4)
- Recommended algorithm \* :
  - Access – Src/Dst IP
  - Dist/Core – Src/Dst IP + Src/Dst L4 Ports
  - Dist – Src/Dst IP

\* May vary based on your network traffic pattern



Core

Default : src-dst-ip vlan

Recommended : src-dst-mixed-ip-port



Dist

Default : src-dst-ip vlan

Recommended : src-dst-mixed-ip-port vlan




Default : src-mac

Recommended : src-dst-ip



Access


L3 – топологии,  
маршрутизация, суммаризация  
и прочие вопросы



# Best Practice - Passive Interfaces for IGP

## Limit IGP Peering Through the Access Layer

- Limit unnecessary peering using passive interface:
  - Four VLANs per wiring closet
  - 12 adjacencies total
  - Memory and CPU requirements increase with no real benefit
  - Creates overhead for IGP



### OSPF Example:

```
Router(config)#router ospf 1
Router(config-router)#passive-
interface Vlan 99

Router(config)#router ospf 1
Router(config-router)#passive-
interface default

Router(config-router)#no passive-
interface Vlan 99
```

### EIGRP Example:

```
Router(config)#router eigrp 1
Router(config-router)#passive-
interface Vlan 99

Router(config)#router eigrp 1
Router(config-router)#passive-
interface default


Router(config-router)#no passive-
interface Vlan 99
```

# Layer 3 connectivity to distribution layer

## LAN core layer

- Links from core layer are Layer 3 links (no SVIs)
- Use MEC to SWV/VSS in distribution layer
- Configure Layer 3 EtherChannel interface
  - When creating L3 EtherChannel, create the logical (port-channel) interface first
- Configure EtherChannel Member Interfaces
  - Configure the physical interfaces to tie to the logical port-channel
- Dual home to WAN or data center to core

```
interface port-channel 20
  no switchport
  ip address [ip address] [mask]
  ip pim sparse-mode
interface range teng1/1/8 , teng2/1/8 , teng1/2/8 , teng2/2/8
  channel-protocol lacp
  channel-group 20 mode active
  macro apply EgressQoS
  no shutdown
```



# Distribution layer IP unicast routing – EIGRP

## LAN distribution layer

EIGRP was chosen for...


simplicity, scalability, and flexibility

- Named mode configuration
- Tie EIGRP router-id to loopback 0 for max stability
- Enable all routed links to be passive by default
- Enable EIGRP for address space
- Each distribution is a stub network

```
router eigrp [NAME]
 address-family ipv4 unicast autonomous-system [AS]
  af-interface default
    passive-interface
  exit-af-interface
  network [network] [inverse mask]
  eigrp router-id [ip address of loopback 0]
  eigrp stub summary
  nsf
exit-address-family
```

Single logical distribution layer design

- Uses stateful switchover (SSO) and non-stop forwarding (NSF)
- SSO provides sub-second failover to redundant supervisor
- NSF maintains packet forwarding while control plane recovers



NSF aware

- Nothing to enable.
- Only need IOS version that supports NSF for EIGRP

NSF capable

- Works on dual supervisor system
- Signals peer of SSO and to delay adjacency timeout
- Once control plane recovers, re-establishes peering

# Distribution layer IP unicast routing – OSPF

## LAN distribution layer

OSPF is available for...

### compatibility


- Tie OSPF router-id to loopback 0 for max stability
- Enable all routed links to be passive by default
- Enable OSPF for address space
- Each distribution is a stub area and ABR

```
router ospf [process]
  router-id [ip address of loopback 0]
  nsf
  area [area number] stub no-summary
  passive-interface default
  network [network] [inv. mask] area [area #]
  network [network] [inverse mask] area 0
```

### Single logical distribution layer design

- Uses stateful switchover (SSO) and non-stop forwarding (NSF)
- SSO provides sub-second failover to redundant supervisor

• NSF maintains packet forwarding while control plane recovers




#### NSF aware


- Nothing to enable.
- Only need IOS version that supports NSF for OSPF

#### NSF capable

- Works on dual supervisor system
- Signals peer of SSO and to delay adjacency timeout
- Once control plane recovers, re-establishes peering

# Provide Alternate Paths

- What happens if  fails?
- No route to the core anymore?
- Allow the traffic to go through the access?
  - Do you want to use your access switches as transit nodes?
  - How do you design for scalability if the access used for transit traffic?
- Install a redundant link to the core
- Best practice: install redundant link to core and utilize L3 link between distribution layer





## Limit EIGRP Queries and OSPF LSA Propagation

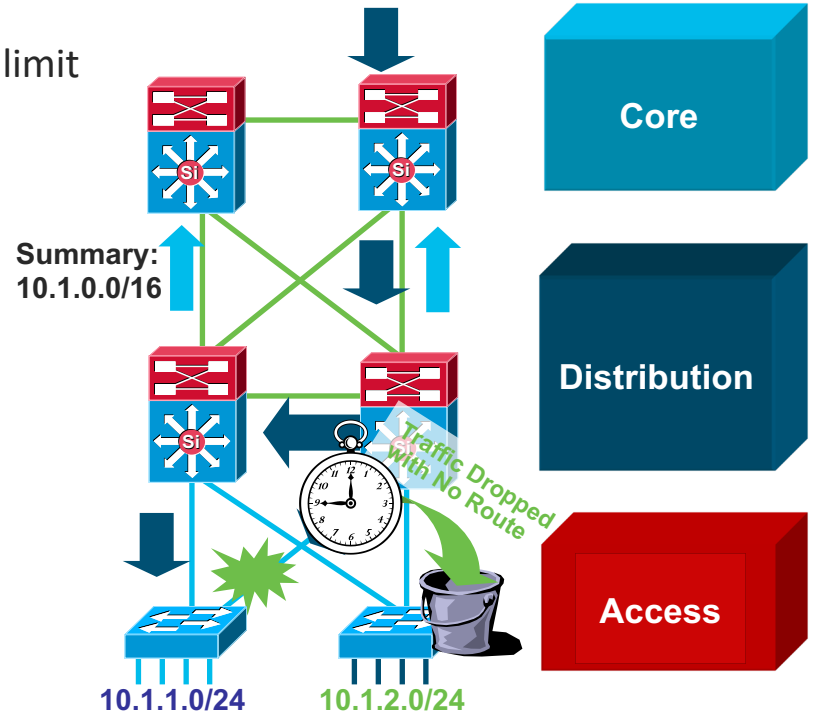
- ```
interface Port-channel1
description to Core#1
ip address 10.122.0.34
255.255.255.252
ip hello-interval eigrp 100 1
ip hold-time eigrp 100 3
ip summary-address eigrp 100
10.1.0.0 255.255.0.0 5
```



# Best Practice - Summarize at the Distribution

## Gotcha—Distribution-to-Distribution Link Required


- Best practice - summaries at the distribution layer to limit EIGRP queries or OSPF LSA propagation
- Gotcha:
  - Upstream: HSRP on left distribution takes over when link fails
  - Return path: old router still advertises summary to core
  - Return traffic is dropped on right distribution switch
- Summarizing requires a link between the distribution switches
- Alternative design: use the access layer for transit



# CEF Load Balancing

## Avoid Under Utilizing Redundant Layer 3 Paths

Redundant Paths Ignored



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

# Synchronize the clock on all devices

## Global LAN switch configuration


- Troubleshooting a network event requires correlation across multiple devices (switches and routers)
- Network devices should be programmed to synchronize time to a local NTP server in the network.
  - allows event log timestamps from multiple devices to be correlated
- Configure console messages, logs, and debug output to provide time stamps

```
ntp server 10.4.48.17
ntp update-calendar
!
clock timezone PST -8
clock summer-time PDT recurring
!
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
```

Update hardware clock on  
Catalyst 6500 and 4500

Set local timezone,  
offset from UTC

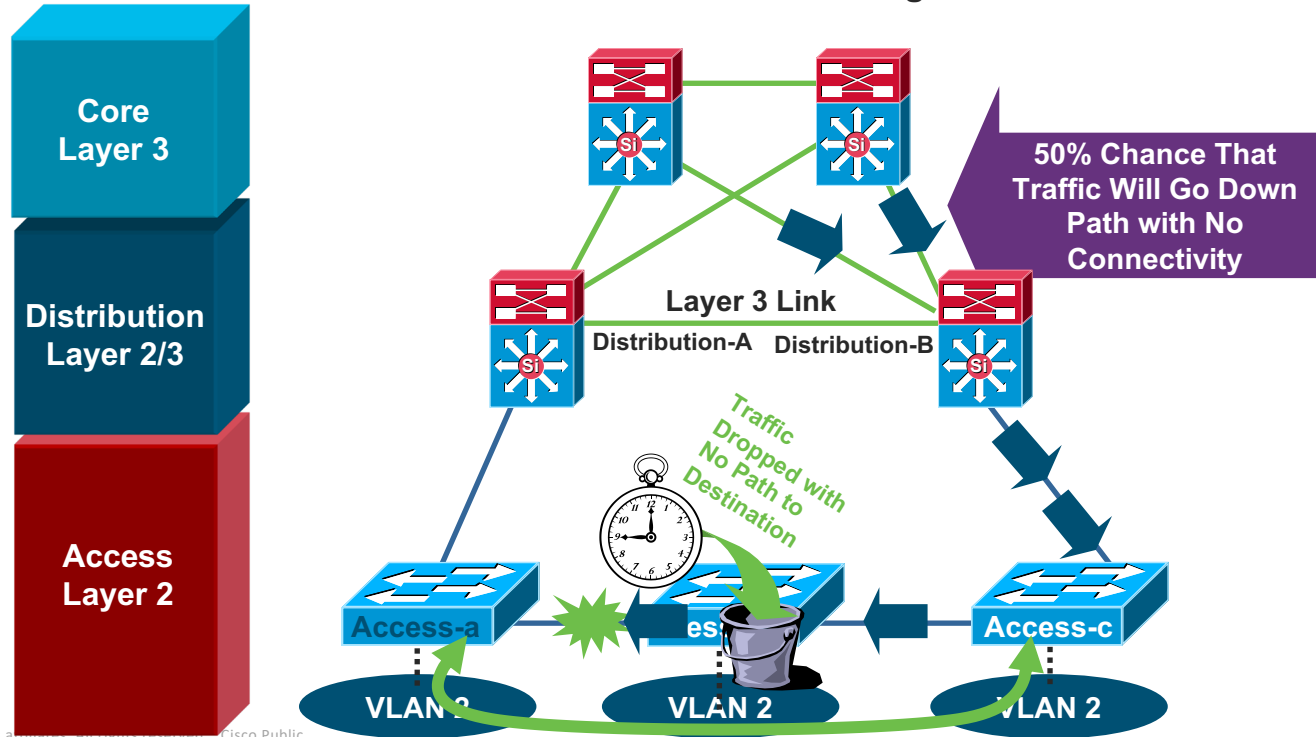
Timestamp output with local  
NTP synchronized time



# Daisy Chaining Access Layer Switches

Avoid Potential Black Holes


Return Path Traffic Has a 50/50 Chance of Being 'Black Holed'




# Daisy Chaining Access Layer Switches

## New Technology Addresses Old Problems

- Stackwise/Stackwise-Plus technology eliminates the concern
  - Loopback links not required
  - No longer forced to have L2 link in distribution
- If you use modular (chassis-based) switches, these problems are not a concern





Быстрая сходимость —  
что будет, если что-то  
пойдет не так?  
И как это улучшить?



# Multilayer campus network design— It is a good solid design, but...

- Utilizes multiple control protocols
  - Spanning tree (802.1w), HSRP / GLBP, EIGRP, OSPF
- Convergence is dependent on multiple factors –
  - FHRP – 900msec to 9 seconds
  - Spanning tree – Up to 50 seconds
- Load balancing –
  - Asymmetric forwarding
  - HSRP / VRRP – per subnet
  - GLBP – per host
- Unicast flooding in looped design
- STP, if it breaks badly, has no inherent mechanism to stop the loop







# Optimizing network convergence

## Failure detection and recovery


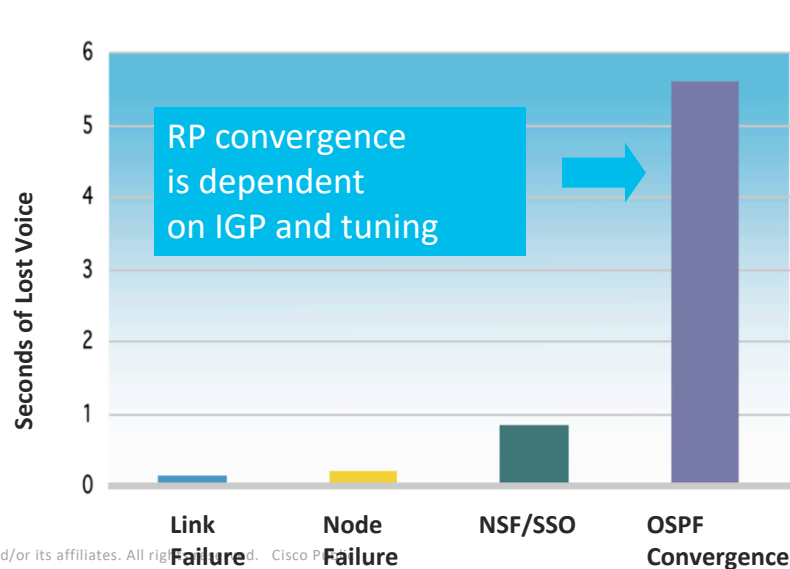
- Optimal high availability network design attempts to leverage 'local' switch fault detection and recovery
- Design should leverage the hardware capabilities of the switches to detect and recover traffic flows based on these 'local' events
- Design principle –  
Hardware failure detection and recovery is both faster and more deterministic
- Design principle –  
Software failure detection mechanisms provide a secondary, not primary, fault detection and recovery mechanism in the optimal design



# Chassis Redundancy at the Core

Depends on topology



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



# Chassis Redundancy at the Distribution

## Recommended


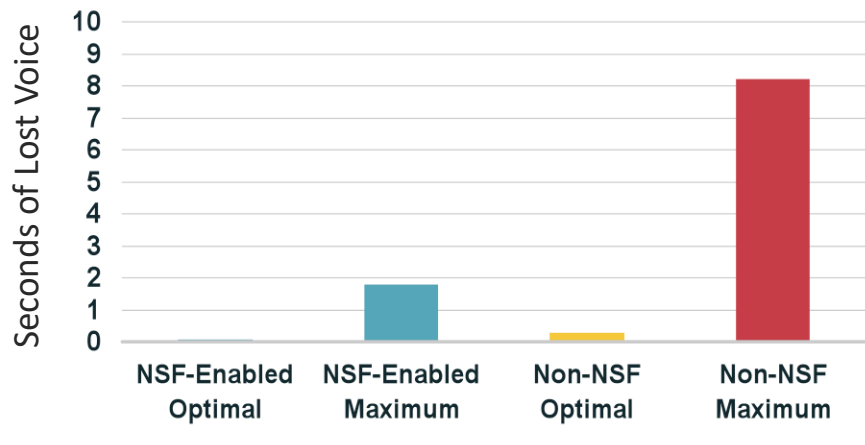
- HSRP doesn't flap on Supervisor SSO switchover
- Reduces the need for sub-second HSRP timers



# Chassis Redundancy at the Access

Recommended for highest availability

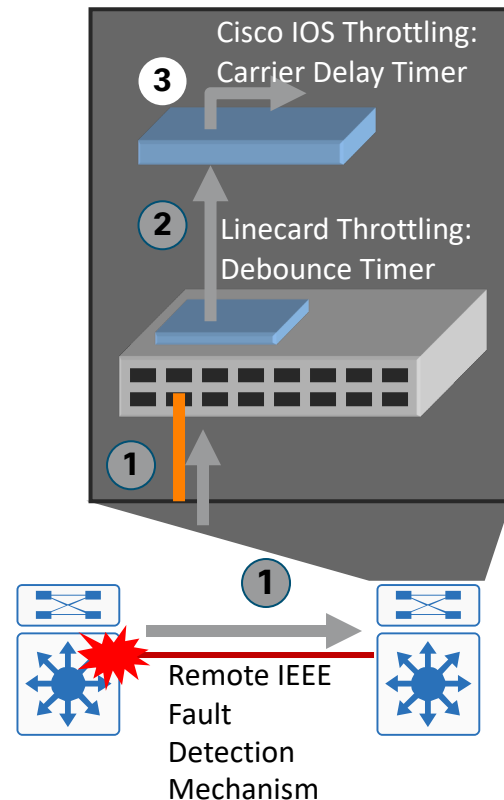
- Access switch is the single point of failure in best practices HA design
- Supervisor failure is most common cause of access switch service outages



# Optimizing network convergence

## Layer 1 link failure fault detection


- Do not disable auto-negotiation on GigE /10GigE ports
- IEEE 802.3z and 802.3ae link negotiation define Remote Fault Indicator & Link Fault Signaling mechanisms
- IOS debounce –
  - GigE/10GigE fiber ports is 10 msec.; copper min. 300 msec.
  - NX-OS debounce – Currently 100 msec. by default
  - All 1G and 10G SFP / SFP+ based interfaces (MM, SM, CX-1) changing to a default of 10 msec.
  - RJ45 based Copper interfaces on NX-OS remains 100 msec.
- Design principle: Understand how hardware choices and tuning impact



# Optimizing network convergence

## Layer 2 software fault detection (e.g. UDLD)

- While 802.3z and 802.3ae link negotiation provide for L1 fault detection, hardware ASIC failures can still occur
- UDLD – L2 based keep-alive mechanism confirms bi-directional L2 connectivity
- Switch ports with UDLD send UDLD protocol packets (at L2) containing:
  - port's own device / port ID
  - neighbor's device / port IDs seen by UDLD on that port
- If port does not see its own device / port ID echoed by incoming UDLD packets, the link is considered unidirectional and is shutdown
- Design principle –
  - Redundant fault detection mechanisms required (SW as a backup to HW as possible)





UDLD Keepalive

# Optimizing Convergence: Trunk Tuning

## Trunk Auto/Desirable Takes Some Time


- DTP negotiation tuning improves link up convergence time
  - IOS(config-if)# switchport mode trunk
  - IOS(config-if)# switchport nonegotiate



# Optimizing L2 Convergence

## PVST+, Rapid PVST+ or MST

- Rapid-PVST+ greatly improves the restoration times for any VLAN that requires a topology convergence due to link UP
- Rapid-PVST+ also greatly improves convergence time over backbone fast for any indirect link failures
- PVST+ (802.1d)
  - Traditional spanning tree implementation
- Rapid PVST+ (802.1w)
  - Scales to large size (~10,000 logical ports)
  - Easy to implement, proven, scales
- MST (802.1s)
  - Permits very large scale STP implementations (~30,000 logical ports)
  - Not as flexible as rapid PVST+





# Optimizing the Layer 2 design


## Complex topologies take longer to converge




Reference

- Time to converge is dependent on the protocol implemented: 802.1D, 802.1s, or 802.1w
- It is also dependent on:
  - Size and shape of the L2 topology (how deep is the tree)
  - Number of VLANs being trunked across each link
  - Number of logical ports in the VLAN on each switch
- Non-congruent topologies take longer to converge.  
Restricting the topology to reduce convergence times
- Prune all unnecessary VLANs from trunk configuration

**400 msec Convergence  
for a Simple Loop**




**900 msec Convergence  
for a More Complex Loop**



# Optimizing the Layer 2 design

## Non-STP-blocking topologies converge fastest

- When STP is not blocking uplinks, recovery of access to distribution link failures is accomplished **based on L2 CAM updates** not on the Spanning Tree protocol recovery
- Time to restore traffic flows is based on:  
Time to detect link failure + Time to purge the HW CAM table and begin to flood the traffic
- No dependence on external events (no need to wait for Spanning Tree convergence)
- Behavior is **deterministic**



- All links forwarding –  
In an environment with all Links active, traffic is restored based on **HW recovery**

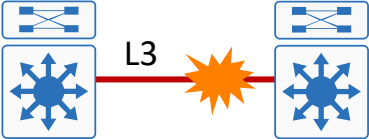
# Optimizing network convergence

## Layer 2 and 3 – Why use routed interfaces?

L3 routed interfaces allow faster convergence than L2 switchport with an associated L3 SVI

~ 8 msec loss


1. Link Down
2. Interface Down
3. Routing Update



```
21:38:37.042 UTC: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet3/1, changed state to down
21:38:37.050 UTC: %LINK-3-UPDOWN: Interface GigabitEthernet3/1, changed state to down
21:38:37.050 UTC: IP-EIGRP(Default-IP-Routing-Table:100): Callback: route_adjust GigabitEthernet3/1
```

~ 200-250 msec. loss

1. Link Down
2. Interface Down
3. Autostate
4. SVI Down
5. Routing Update




```
21:32:47.813 UTC: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet2/1, changed state to down
21:32:47.821 UTC: %LINK-3-UPDOWN: Interface GigabitEthernet2/1, changed state to down
21:32:48.069 UTC: %LINK-3-UPDOWN: Interface Vlan301, changed state to down
21:32:48.069 UTC: IP-EIGRP(Default-IP-Routing-Table:100): Callback: route, adjust Vlan301
```

# Layer 2 access with Layer 3 distribution

## First hop redundancy protocols (FHRP)


- HSRP, GLBP, and VRRP:  
provide a resilient default gateway / first hop address to end stations
- A group of routers act as a single logical router providing first hop router redundancy
- Protect against multiple failures
  - Distribution switch failure
  - Uplink failure
- Default recovery is ~10 Seconds



# Optimizing Convergence: VRRP, HSRP, GLBP

Mean, Max, and Min—Are There Differences?

- VRRP not tested with sub-second timers and all flows go through a common VRRP peer; mean, max, and min are equal
- HSRP has sub-second timers; however all flows go through same HSRP peer so there is no difference between mean, max, and min
- GLBP has sub-second timers and distributes the load amongst the GLBP peers; so 50% of the clients are not affected by an uplink failure



# Even with faster convergence from RPVST+ we still have to wait for FHRP convergence

Subsecond timers improve convergence

## HSRP Config


```
interface Vlan4
ip address 10.120.4.2 255.255.255.0
standby 1 ip 10.120.4.1
standby 1 timers msec 250 msec 750
standby 1 priority 150
standby 1 preempt
standby 1 preempt delay minimum 180
```

## GLBP Config

```
interface Vlan4
ip address 10.120.4.2 255.255.255.0
glbp 1 ip 10.120.4.1
glbp 1 timers msec 250 msec 750
glbp 1 priority 150
glbp 1 preempt
glbp 1 preempt delay minimum 180
```

## VRRP Config

```
interface Vlan4
ip address 10.120.4.1 255.255.255.0
vrrp 1 description Master VRRP
vrrp 1 ip 10.120.4.1
vrrp 1 timers advertise msec 250
vrrp 1 preempt delay minimum 180
```



HSRP is widely used with Its rich feature set

GLBP facilitates uplink load balancing –  
not optimal for L2 looped topology


VRRP for multi-vendor interoperability

HSRP, GLBP and VRRP provide millisecond timers and  
excellent convergence performance

**Critical for VoIP and video recovery in < 1 second**

## HSRP preemption—why it is desirable

- Spanning tree root and HSRP primary are aligned
- When spanning tree root is re-introduced, traffic takes a two-hop path to HSRP active
- **HSRP preemption** allows HSRP to follow the spanning tree topology




Without Preempt Delay, HSRP Can Go Active Before the Switch Is Completely Ready to Forward Traffic – L1 (Linecards), L2 (STP), L3 (IGP Convergence)

# FHRP design considerations

## Preempt delay needs to be longer than boot time

- HSRP is not always aware of the status of the entire switch and network
- Ensure that you provide enough time for the entire system to be up – diagnostics (full or partial), L1 (line cards), L2 (STP), L3 (IGP convergence)
- Tune delay and preempt delay conservatively, as the network is already forwarding data



```
interface Vlan402
. . .
standby delay minimum 60 reload 600
standby 1 ip 10.147.102.1
standby 1 timers msec 250 msec 750
standby 1 priority 110
standby 1 preempt delay minimum 60 reload 600
standby 1 authentication ese
standby 1 name HSRP-Voice
hold-queue 2048 in
```


**standby delay:** Controls time interface needs to be up before HSRP starts.

**preempt delay:** Controls time to wait after HSRP establishes a neighbour relationship. Configure both.




# Multichassis EtherChannel performs better in any network design

- Network recovery mechanic varies in different distribution design –
  - Standalone – protocol and timer dependent
  - StackWise Virtual – hardware dependent
- StackWise Virtual logical distribution system –
  - Single P2P STP Topology
  - Single Layer 3 gateway
  - Single PIM DR system
- Distributed and synchronized forwarding table –MAC address, ARP cache, IGMP
- All links are fully utilized based on Ether-channel load balancing




# FHRP design considerations— asymmetric routing (unicast flooding)

- Alternating HSRP Active between distribution switches can be used for upstream load balancing
- This can cause a problem with unicast flooding
- ARP timer defaults to four hours and CAM timer defaults to five minutes
- ARP entry is valid, but no matching L2 CAM table exists
- In many cases when the HSRP standby needs to forward a frame, it will have to unicast flood the frame since its CAM table is empty




# FHRP design considerations— asymmetric routing (unicast flooding) solutions

- Using 'V' based design with unique voice and data VLANs per access switch, this problem has no user impact
- Don't deploy stacking switches (ie. daisy-chained switches) that depend on spanning tree for managing stack interconnects
- Tune ARP timer to 270 seconds and leave CAM timer to default, unless ARP > 10,000, change CAM timers
- Deploy MultiChassis EtherChannel with StackWise Virtual (SWV), Virtual Switching System (VSS), or Virtual Port Channel (vPC) in the distribution block



CAM timers traditionally default to 5 minutes to allow for MAC addresses (devices) to move in the network. It is safe to increase the CAM timers if the client devices will generate unicast or multicast traffic to refresh the CAM table.

# Выводы




# Layer 2 Distribution Interconnection

## Layer 2 Access—Some VLANs Span Access Layer

- Tune CEF 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/no-negotiate
- Disable Ether Channel unless needed
- RootGuard on downlinks
- LoopGuard on uplinks
- Set port host on access Layer ports:
  - Disable trunking
  - Disable Ether Channel
  - Enable PortFast
- RootGuard or BPDU-Guard
- Use security features


hts reserved. Cisco Public



# Layer 3 Distribution Interconnection


## Layer 2 Access—No VLANs Span Access Layer

- Tune CEF 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/no-negotiate
- Disable Ether Channel unless needed
- Set port host on access layer ports:
  - Disable trunking
  - Disable Ether Channel
  - Enable PortFast
- RootGuard or BPDU-Guard
- Use security features



# Best Practices - Spanning Tree Configuration

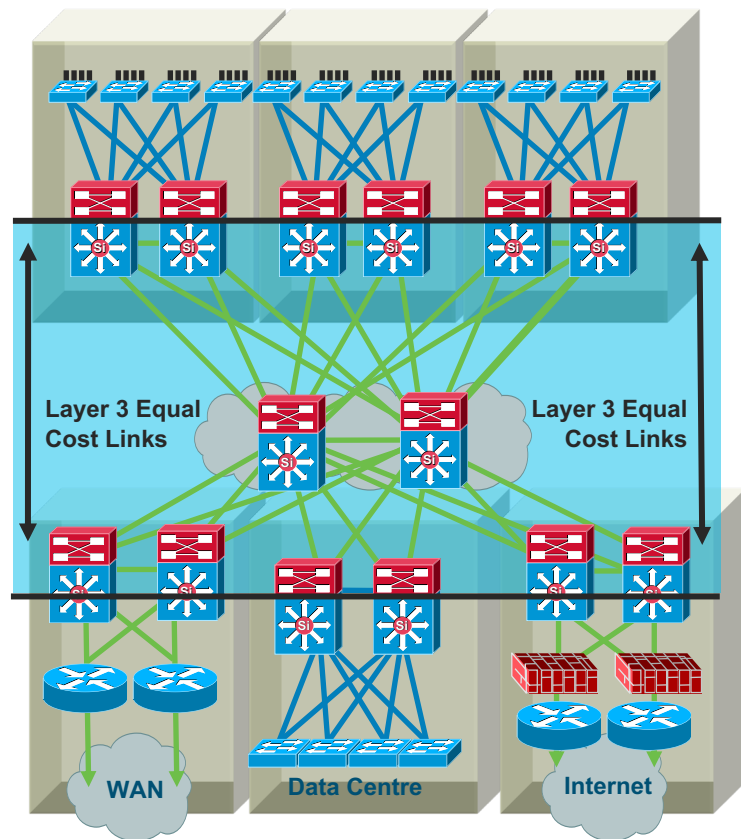
- Only span VLAN across multiple access layer switches when you have to!
- Use rapid PVST+ for best convergence
- More common in the data center
- Required to protect against user side loops
- Required to protect against operational accidents (misconfiguration or hardware failure)
- Take advantage of the spanning tree toolkit



# Best Practices

## Layer 3 Routing Protocols

- Typically deployed in distribution to core, and core-to-core interconnections
- Used to quickly reroute 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)







Спасибо!

[www.cisco.com/go/cvd](http://www.cisco.com/go/cvd)

