Cisco live!

July 10-14, 2016 · Las Vegas, NV

Your Time Is Now

Best practices to deploy high-availability in Wireless LAN Architectures

Patrick Croak, Consulting Systems Engineer – CCIE Wireless #34712

BRKEWN-3014

Cisco live!

“The Wireless network is the projection of my Company brand”

Tech Operation Manager
@Financial Customer

Session Objective

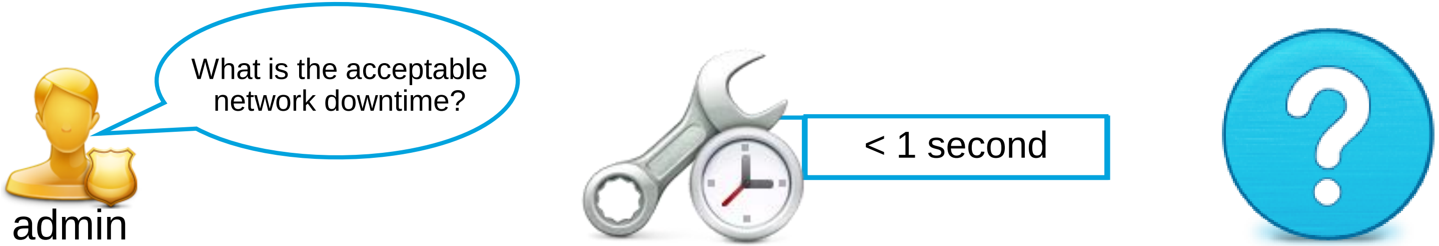

What is the acceptable network downtime?

admin

< 1 second

The goal of this session is to show you how to design and deploy a Highly Available wireless network **to reduce the network downtime**

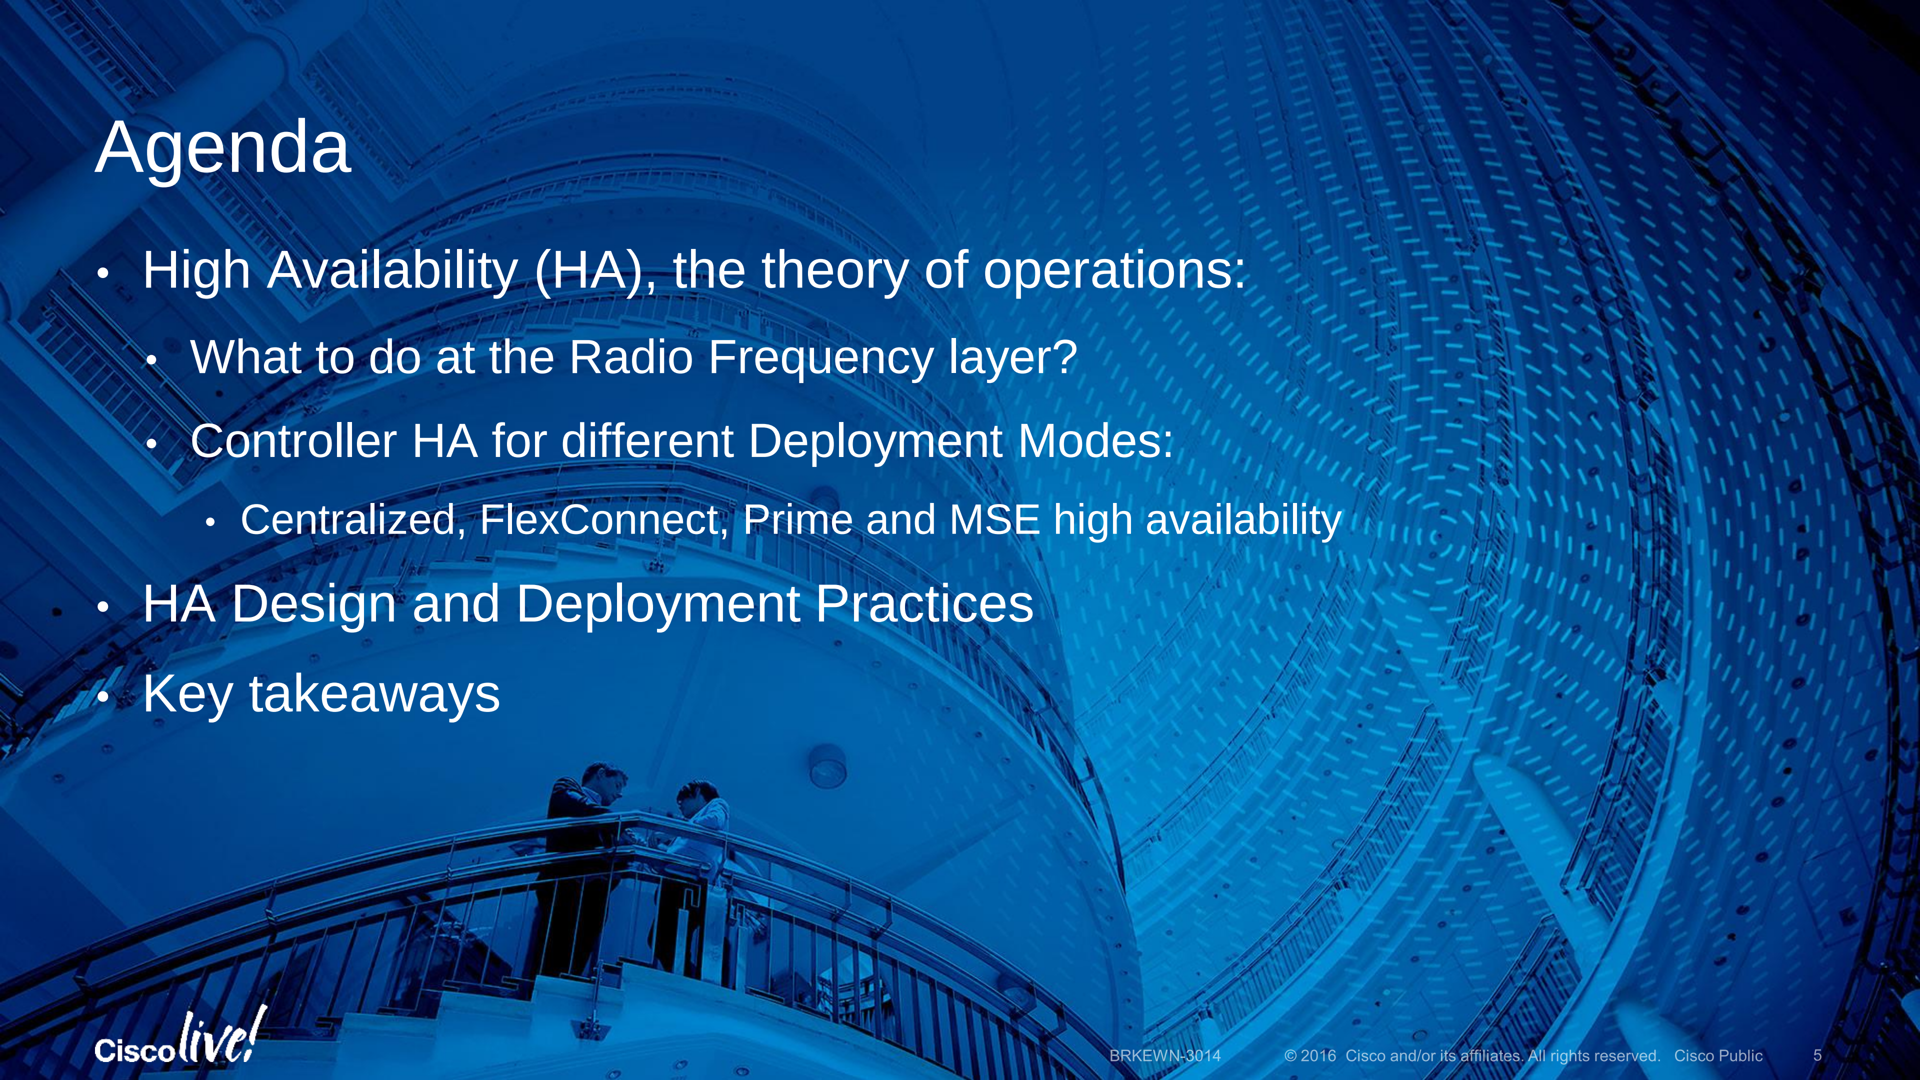

Agenda

- High Availability (HA), the theory of operations:
  - What to do at the Radio Frequency layer?
  - Controller HA for different Deployment Modes:
    - Centralized, FlexConnect, Prime and MSE high availability
- HA Design and Deployment Practices
- Key takeaways

Radio Frequency (RF) considerations

802.11ac is here!!

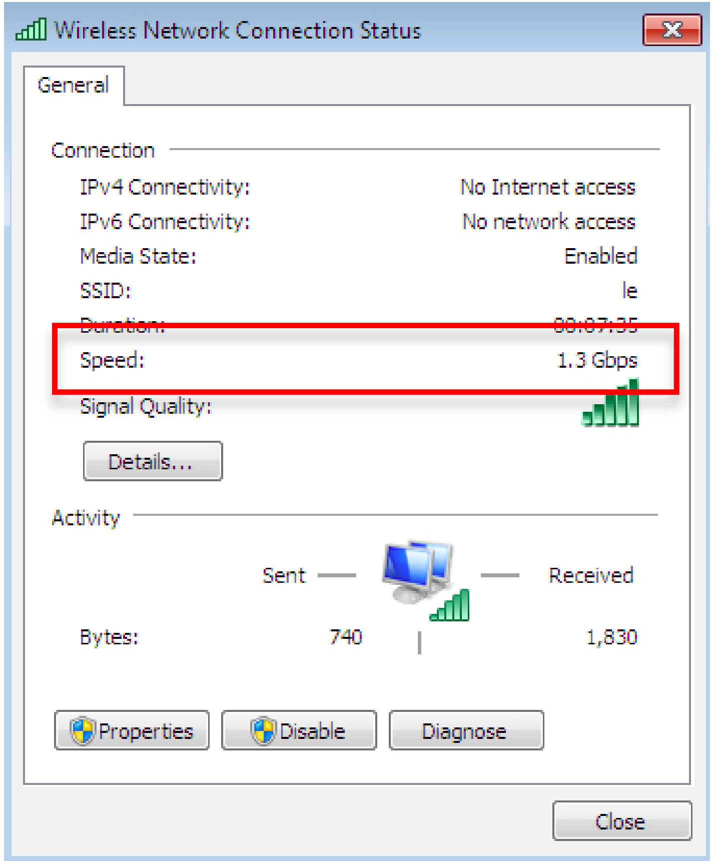

Wireless Network Connection Status

General

Connection

IPv4 Connectivity: No Internet access

IPv6 Connectivity: No network access

Media State: Enabled

SSID: le

Duration: 00:07:35

Speed: 1.3 Gbps

Signal Quality:

Activity

Sent — Received

Bytes: 740 1,830

But it comes with a price

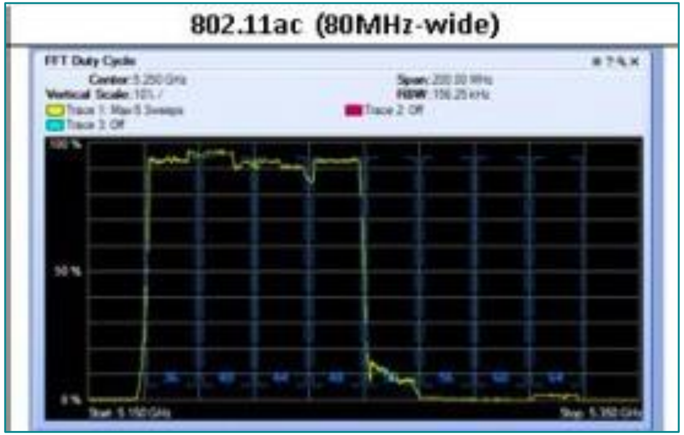

80 MHz channel

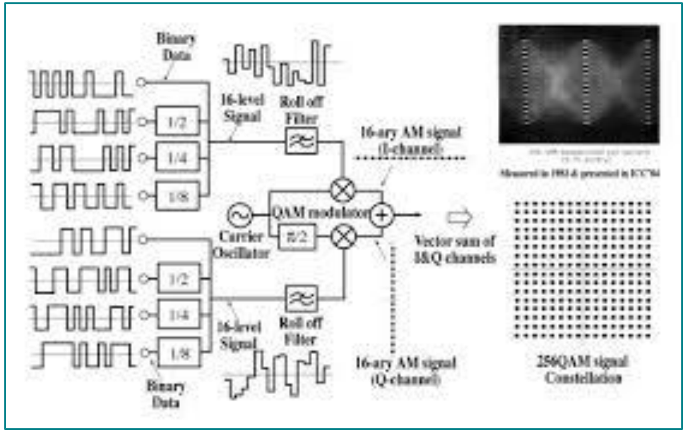
Binary Data
16-level Signal
Roll off Filter
1/2
1/4
1/8
16-ary AM signal (I-channel)
QAM modulator
Carrier Oscillator
π/2
Vector sum of I&Q channels
16-ary AM signal (Q-channel)
256QAM signal Constellation
Measured in 1983 & presented in ICC'84

High Signal at the client for 256QAM

Radio Frequency (RF) High Availability

- RF HA is the ability to build redundancy at the physical layer
- What does it translates to in practice?
  - Creating a pervasive, stable, predictable RF environment (Proper Design, Site Survey, Radio Planning)
  - Dealing with coverage holes if an AP goes down (RF Management)
  - Identifying, Classifying, Mitigating an interference source (Spectrum Intelligence Solution)
  - Improving client (all clients!) received signal (Beamforming)
- BTW…Cisco has differentiating features/functionalities to address all these things

Radio Frequency (RF) High Availability

- Site Survey, site survey….and site survey
  - Use “Active” survey
  - Coverage vs. Capacity
  - Consider Client type (ex. Smartphone vs. Laptop)

My power is half of my brothers MacBook

My antenna gain is 4 times smaller

I try to connect to 5GHz and stay connected until the signal is REALLY bad

and then move to another BSSID if it is REALLY better

Radio Frequency (RF) High Availability

- Site Survey, site survey….and site survey
  - Use “Active” survey
  - Coverage vs. Capacity
  - Consider Client type (ex. Smartphone vs. Laptop)
- AP positioning and antenna choice is Key
  - Use common sense
  - Light source analogy
  - Internal antennas are designed to be mounted on ceiling
  - External antennas: use same antennas on all connectors
- Tools
  - What you use is less important than how you use it
  - Use the same tool to compare results

RF High Availability: Cisco RRM

- What are Radio Resource Manager (RRM)'s objectives?
  - Provide a system wide RF view of the network at the Controller (only Cisco!!)
  - Dynamically balance the network and mitigate changes
  - Manage Spectrum Efficiency so as to provide the optimal throughput under changing conditions
- What's RRM
  - DCA—Dynamic Channel Assignment
  - TPC—Transmit Power Control
  - CHDM—Coverage Hole Detection and Mitigation
- RRM best practices
  - RRM settings to auto for most deployments (High Density is a special case)
  - Design for most radios set at mid power level (lever 3 for example)
  - Use RF Profiles to customize RRM settings per Areas/Groups of APs

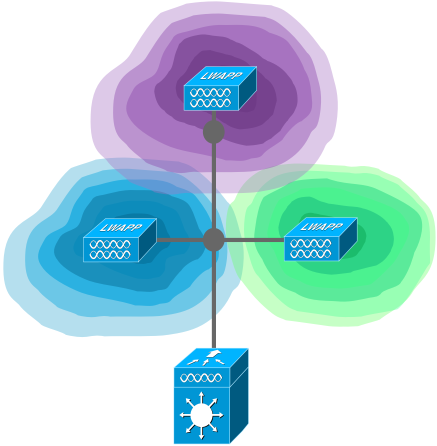

RF High Availability: Cisco RRM

RRM DCA in action

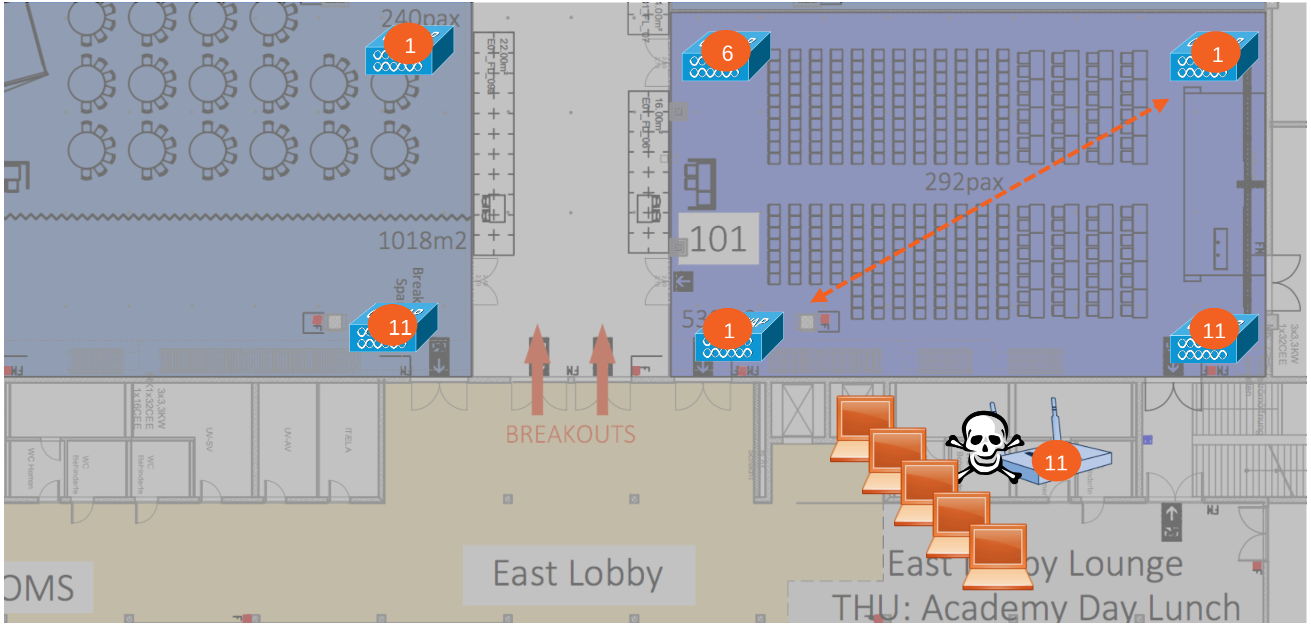

- RRM will determine the optimal channel plan based on AP layout
- A rogue AP is detected on channel 11
- RRM will assess the RF and take a decision in less than 10min
- Channel change is triggered to improve the RF
- Note how the 3 non overlapping channels are still maintained!
- RRM has a RF system view. AP view would be limited and could result in sub-optimal RF plan

RF High Availability: Cisco RRM – RF Profiles

RF profiles = RF Design flexibility

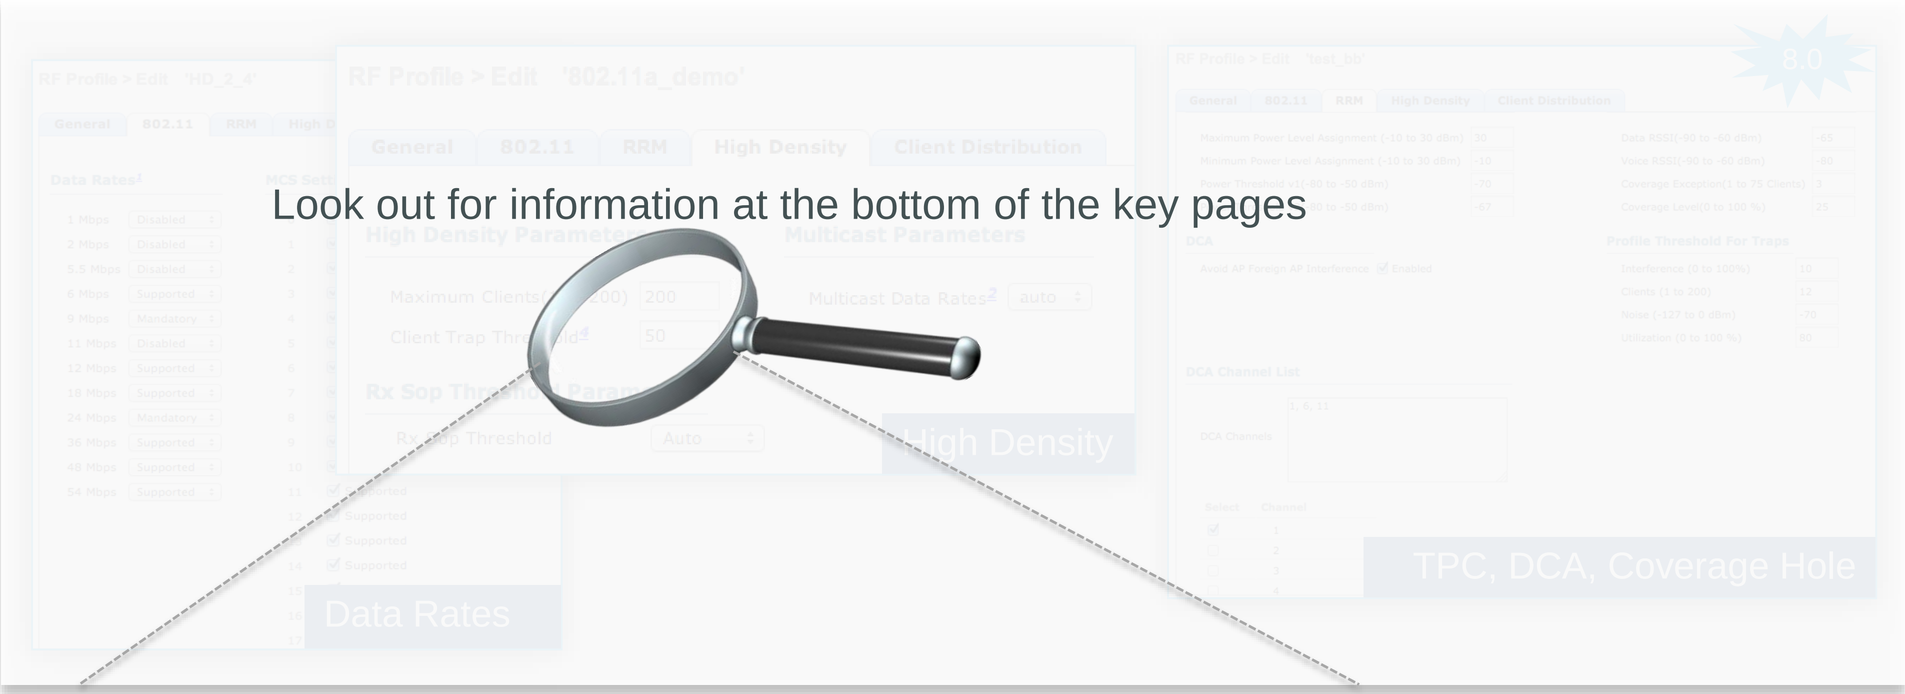

Look out for information at the bottom of the key pages

RF Profile > Edit 'HD_2_4'

General | 802.11 | RRM | High D

Data Rates[1]

1 Mbps	Disabled
2 Mbps	Disabled
5.5 Mbps	Disabled
6 Mbps	Supported
9 Mbps	Mandatory
11 Mbps	Disabled
12 Mbps	Supported
18 Mbps	Supported
24 Mbps	Mandatory
36 Mbps	Supported
48 Mbps	Supported
54 Mbps	Supported

MCS Setti

11	☑ Supported
12	☑ Supported
13	☑ Supported
14	☑ Supported
15	
16	
17	

Data Rates

RF Profile > Edit '802.11a_demo'

General | 802.11 | RRM | High Density | Client Distribution

High Density Parameters

Maximum Clients(1 to 200) 200

Client Trap Threshold[4] 50

Multicast Parameters

Multicast Data Rates[2] auto

Rx Sop Threshold Parameters

Rx Sop Threshold Auto

High Density

8.0

RF Profile > Edit 'test_bb'

General | 802.11 | RRM | High Density | Client Distribution

Maximum Power Level Assignment (-10 to 30 dBm)	30
Minimum Power Level Assignment (-10 to 30 dBm)	-10
Power Threshold v1(-80 to -50 dBm)	-70
Power Threshold v2(-80 to -50 dBm)	-67

Data RSSI(-90 to -60 dBm)	-65
Voice RSSI(-90 to -60 dBm)	-80
Coverage Exception(1 to 75 Clients)	3
Coverage Level(0 to 100 %)	25

DCA

Avoid AP Foreign AP Interference ☑ Enabled

Profile Threshold For Traps

Interference (0 to 100%)	10
Clients (1 to 200)	12
Noise (-127 to 0 dBm)	-70
Utilization (0 to 100 %)	80

DCA Channel List

DCA Channels 1, 6, 11

Select	Channel
☑	1
☐	2
☐	3
☐	4

TPC, DCA, Coverage Hole

For more info: http://www.cisco.com/en/US/tech/tk722/tk809/technologies_tech_note09186a008072c759.shtml

RF High Availability: Cisco CleanAir

- Assess impact of interferences and proactively change channel when needed
- Hardware based Spectrum intelligence solution integrated in Cisco Prime
- Only CleanAir ASIC based solution can reliably detect interference sources:

5MHz scanning resolution

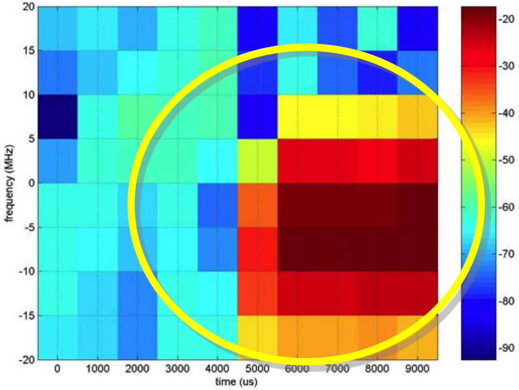

CleanAir
Hardware based Solution

32 times WiFi chip's visibility
Accurate classification
Multiple device recognition

156 kHz scanning resolution

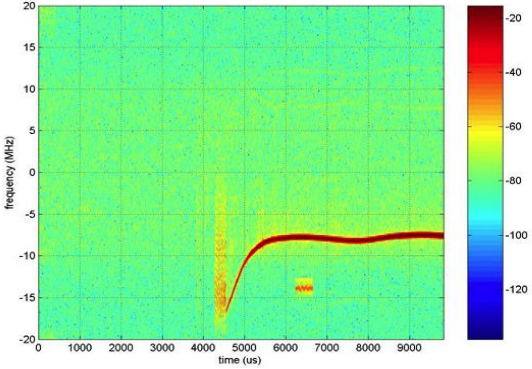

- Best Practice: always turn it on supported APs (all 802.11ac APs are CleanAir capable)

For more info: http://www.cisco.com/en/US/netsol/ns1070

RF High Availability: Cisco ClientLink

- Cisco ClientLink is Beamforming at the chip level:
  - Implemented in hardware, no software component, no performance degradation
- ClientLink creates a better quality RF for all clients (a/g/n/c)
- Do I need a 4x4 AP? Yes, and even more critical with 802.11ac

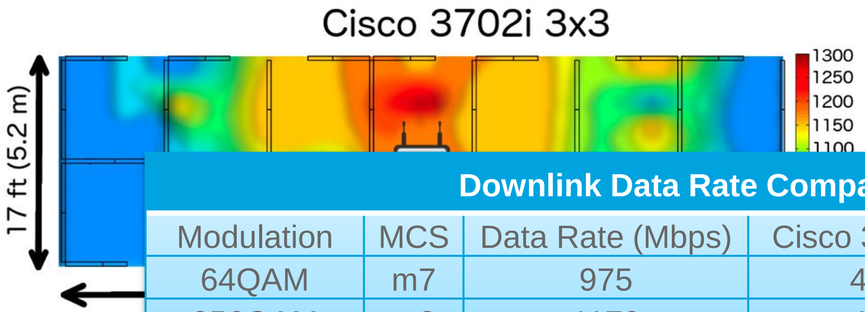

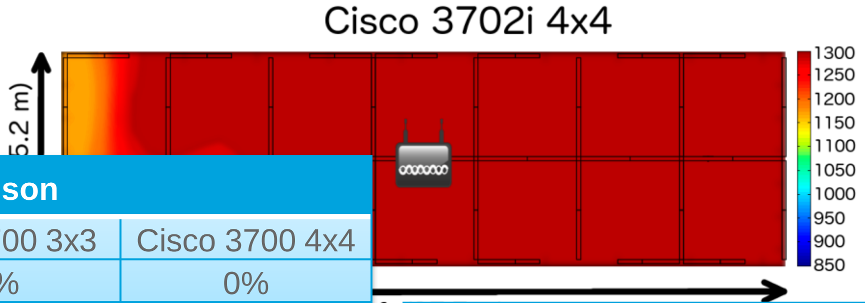

Downlink Data Rate Comparison

Modulation	MCS	Data Rate (Mbps)	Cisco 3700 3x3	Cisco 3700 4x4
64QAM	m7	975	46%	0%
256QAM	m8	1170	49%	15%
256QAM	m9	1300	5%	85%

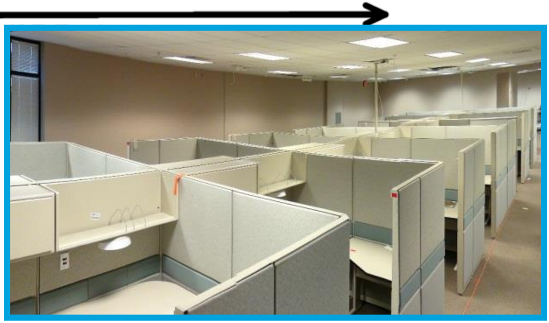

- Best practice: on by default

For more info: http://www.cisco.com/en/US/prod/collateral/wireless/ps5678/ps11983/at_a_glance_c45-691984.pdf

Innovations Only Cisco Delivers

Radio Frequency Excellence for High-Density Environments

Available on new 2800/3800 APs

Zero Impact AVC
Hardware Based Application Visibility and Control without Impact to Performance

Multi-Gigabit Uplinks
Free Up Wireless With Faster Wired Network Offload

Smart Antenna Connector
2nd Antenna Connection
Dual 5 GHz & Band specific Antennas

Turbo Performance
Scales to Support More Devices Running High Bandwidth Apps

Flex Dynamic Frequency Selection
Automatically Adjusts So Not to Interfere With Other Radio Systems

Cisco CleanAir®
Remediates Device Impacting Interference

Flexible Radio Assignment
Adjust Radio Bands to Better Serve the Environment

802.11ac Wave 2 MU-MIMO
Standards based Wave 2 MU-MIMO

Optimized Roaming
Intelligently Connects the Proper Access Point as People Move

Cisco ClientLink
Improves Performance of legacy and 802.11ac Devices

Expandability
Add Functionality Via Module, Smart Antenna Port or **USB Port**

Maximize the Spectrum

Avoiding Excessive Management Traffic

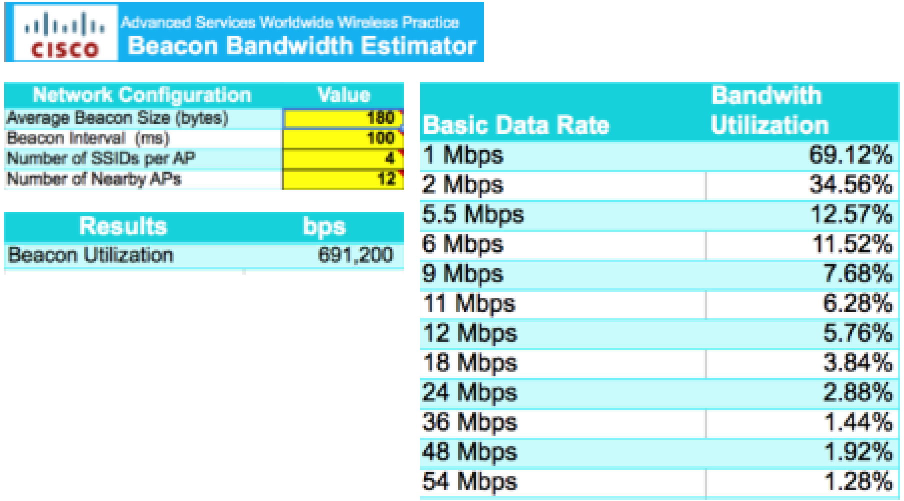

Advanced Services Worldwide Wireless Practice

Beacon Bandwidth Estimator

Network Configuration	Value
Average Beacon Size (bytes)	180
Beacon Interval (ms)	100
Number of SSIDs per AP	4
Number of Nearby APs	12

Results	bps
Beacon Utilization	691,200

Basic Data Rate	Bandwith Utilization
1 Mbps	69.12%
2 Mbps	34.56%
5.5 Mbps	12.57%
6 Mbps	11.52%
9 Mbps	7.68%
11 Mbps	6.28%
12 Mbps	5.76%
18 Mbps	3.84%
24 Mbps	2.88%
36 Mbps	1.44%
48 Mbps	1.92%
54 Mbps	1.28%

- Always aim for 1 SSID
  - More SSID’s = Worse Performance
- Why?
  - Each SSID requires a separate Beacon
  - Each SSID will beacon at the minimum mandatory data rate
- Each broadcast SSID will respond to null probe requests
  - **Exponential** amounts of airtime wasted

Maximize the Spectrum

PHY Rate Tuning: Why PHY Rates Matter

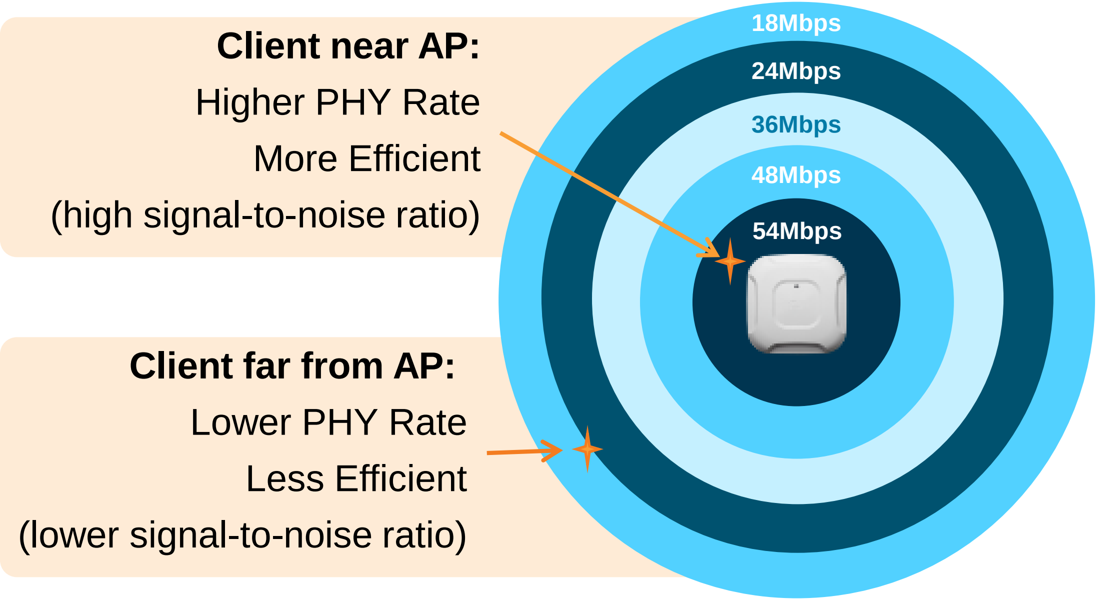

Client near AP:

Higher PHY Rate

More Efficient

(high signal-to-noise ratio)

Client far from AP:

Lower PHY Rate

Less Efficient

(lower signal-to-noise ratio)

- How fast can we talk?
  - Signal (RSSI) and Noise are key factors
- As client moves further from AP or as noise worsens, client rate-shifts downward
- Lower rate, more airtime consumed
- 802.11ac Wave 2 example ~15’

Maximize the Spectrum

PHY Rate Tuning: How-To Basics

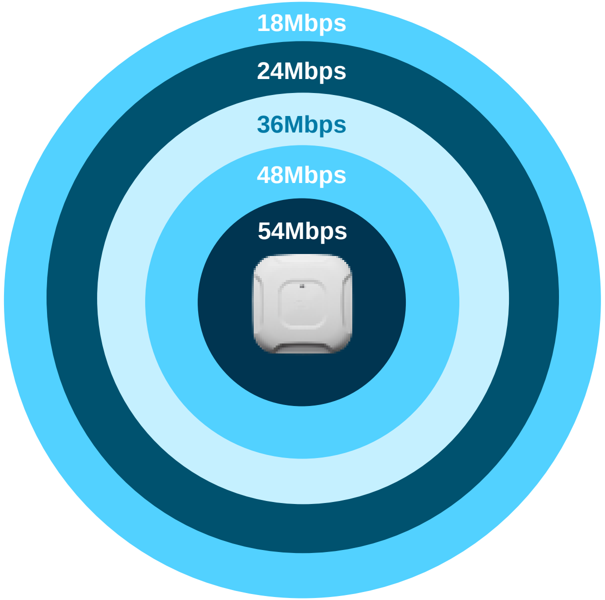

- Position AP’s and antennas to allow elimination of low rates (i.e., <18mbps)
- Eliminate 802.11b rates
- Avoid disabling MCS rates
  - Disabling MCS rates, especially 0-7, can cause significant client issues

Remember the 3 Key RF Relationships!

Client-Induced Interference: What is it?

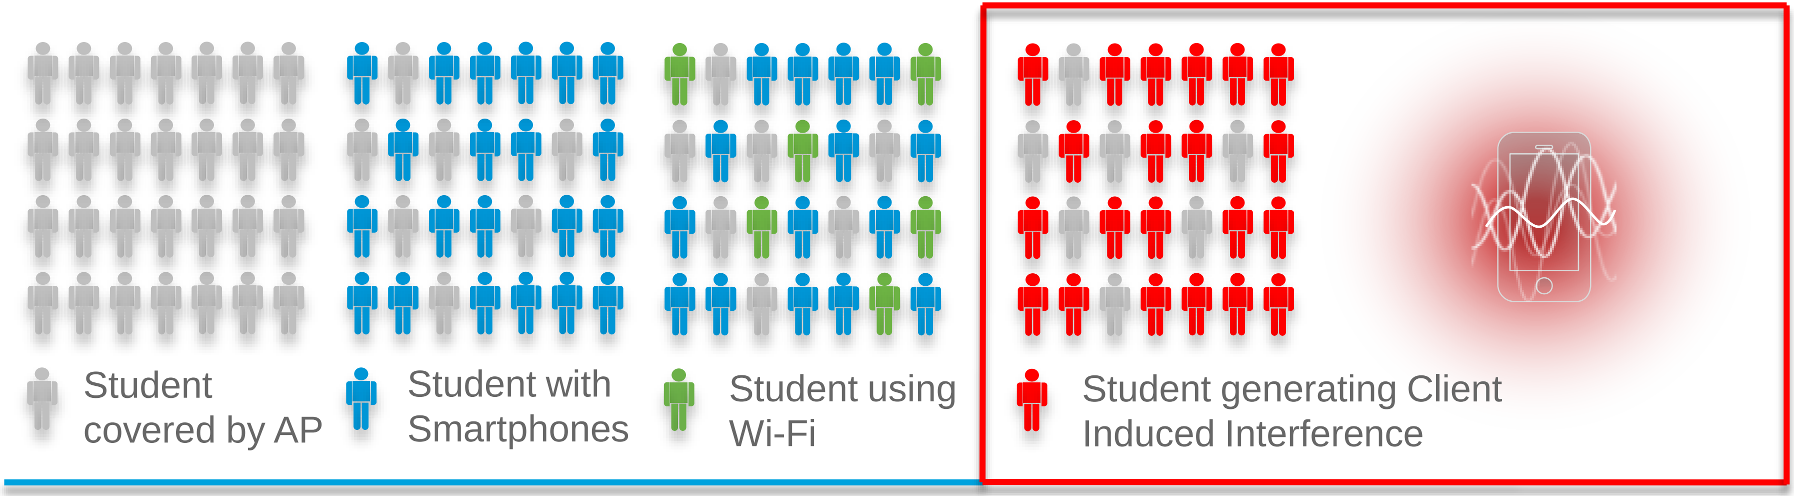

Student covered by AP

Student with Smartphones

Student using Wi-Fi

Student generating Client Induced Interference

Common Assumptions

- 75% of Students will have a Smartphone
- 30% of Smartphone users will utilize Wi-Fi

- **But what is everyone else doing?**

Client-Induced Interference

What does it look like and how can we mitigate?

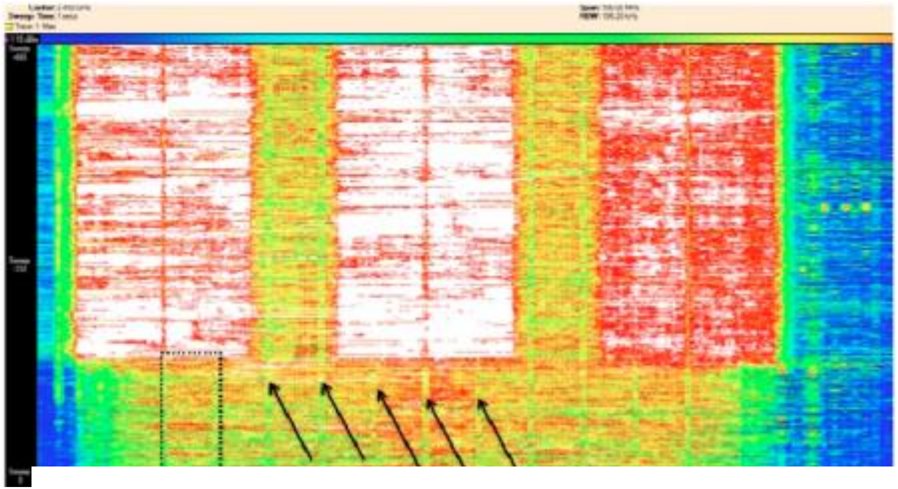

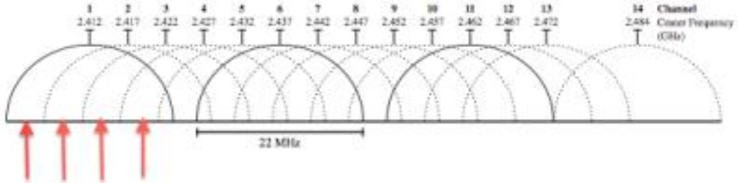
1 2 3 4 5 6 7 8 9 10 11 12 13 14 Channel
2.412 2.417 2.422 2.427 2.432 2.437 2.442 2.447 2.452 2.457 2.462 2.467 2.472 2.484 Center Frequency (GHz)
22 MHz

- Client-induced interference: especially damaging on 2.4GHz but also impacts 5GHz via ACI (Adjacent Channel Interference)
- Probe requests sent on *all* channels
  - Many frames on overlapping channels, **driving noise floor to be higher/worse**
- Getting these devices on your network can help
  - Probe frequency diminishes significantly on an associated device

Maximizing the Spectrum

Ease-of-Use & Client Induced Interference

- Ask yourself - how difficult is it to get on your WiFi network?
- Ease-of-use directly impacts airtime efficiency
- Low take rate = lots of probe request noise (1mb, max power, all channels)
  - Results in **Client Induced Interference**
- Design for seamless end-user experience
  - Captive portals for T&C: necessary?
- **A device on the network is <u>far</u> less damaging than a device off the network!**

Reforming 5 GHz to Optimize for 802.11ac

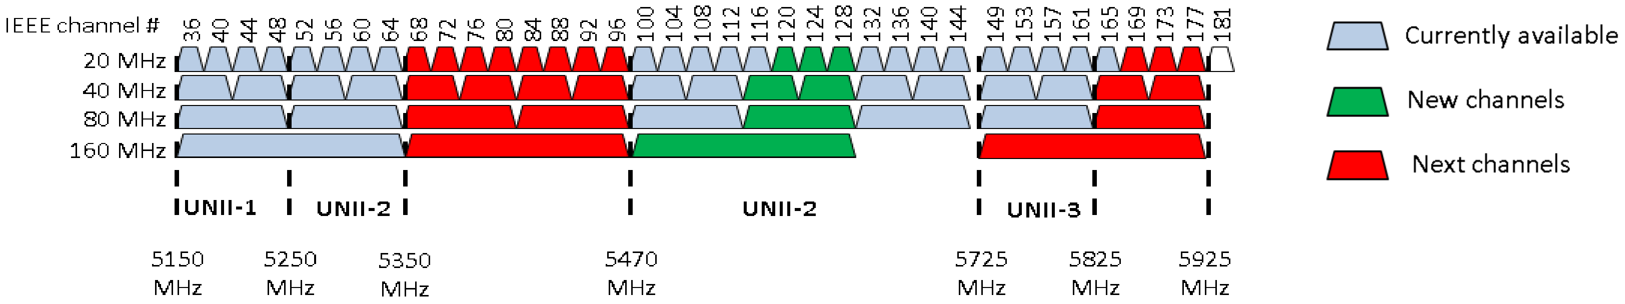

- More non-overlapping channels enabling better 802.11ac experience
- 6x 80 MHz channels (5 in Canada and Europe)
- 2x 160 MHz channels (1 in Canada)
- Additional 5GHz spectrum liberalization (5.35-5.47 GHz and 5.85-5.925 GHz) allows:

Future 5GHz Opportunity

Channel Bandwidth (MHz)	No. of Non-overlapping Channels
20	37
40	18
80	9
160	4

RRM’s new Flexible Radio Assignment (FRA)

- Manage the Flexible Radio Hardware
  - Determine Coverage Overlap Factor (COF) at 2.4 GHz
  - Evaluate Radios as potentially Redundant
  - Determine best role for Flexible Radio
  - Assign
- Radio role determination and assignment is Automatic If radio’s FRA Auto and FRA is enabled.
- FRA calculates COF for Manual assigned radios and Administrator can make Role choices

FRA – COF, Coverage Overlap Factor

- 2.4 GHz Radios that are members of the "Same" AP Group will be calculated together
- Coverage Overlap is the percentage (%) of a given cell that is covered by other AP's at -65 dBm or greater
- All AP models considered in the coverage calculation
- Neighbors above -60 dBm will be used for coverage
- Only 2800/3800 can be marked as Redundant

Cisco 2700 Guest
3800-6
3800-5
3800-4
3800-3
3800-2
3800-1
4.9 ft / 3.0dBm
Aruba
-89
-65
-37
-70
dBm
>= 0

FRA – Assignment Priority

1

5GHz Serving

2.4GHz Serving

Coverage too dense – Mark Redundant

2

5GHz Serving

5GHz Serving

DCA will determine suitability, and
If Unsuitable – then Monitor

e
ps
nts

3

5GHz Serving

Wireless Security Monitor

- Secure Network from Non-Wi-Fi Interference, wIPS

Wireless monitoring of 2.4 and 5 GHz

- Scan both 2.4GHz and 5GHz for security threats

For more information on FRA

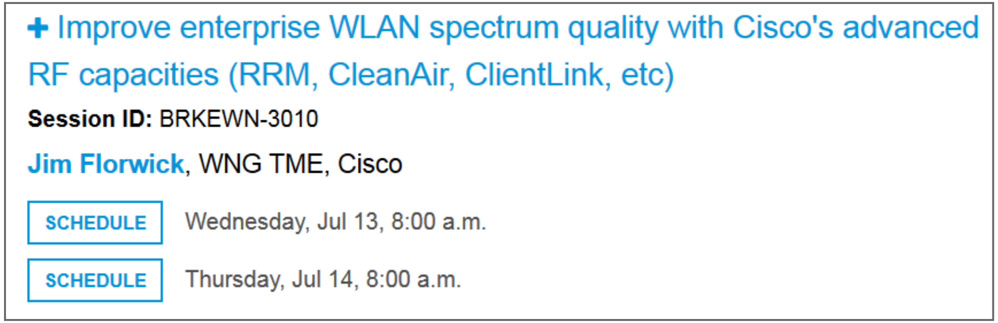

+ Improve enterprise WLAN spectrum quality with Cisco's advanced RF capacities (RRM, CleanAir, ClientLink, etc)

Session ID: BRKEWN-3010

Jim Florwick, WNG TME, Cisco

Wednesday, Jul 13, 8:00 a.m.

Thursday, Jul 14, 8:00 a.m.

- Radio Resource Management White Paper

http://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-3/b_RRM_White_Paper.html

RF sessions you don't want to miss…

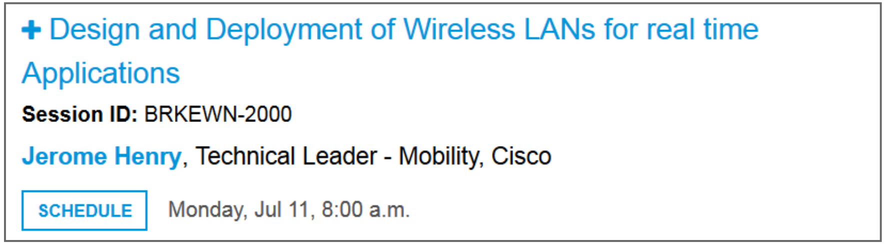

+ Design and Deployment of Wireless LANs for real time Applications

Session ID: BRKEWN-2000

Jerome Henry, Technical Leader - Mobility, Cisco

SCHEDULE Monday, Jul 11, 8:00 a.m.

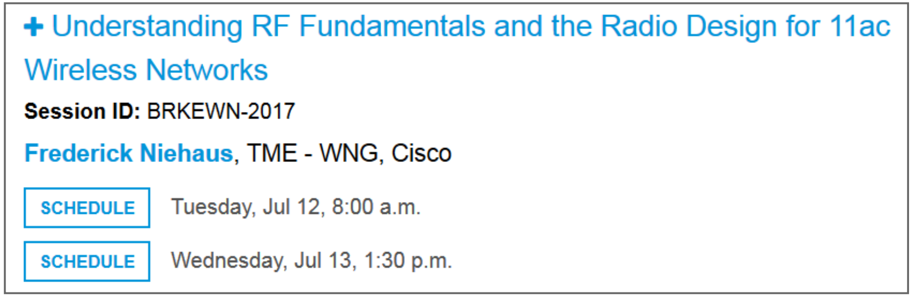

+ Understanding RF Fundamentals and the Radio Design for 11ac Wireless Networks

Session ID: BRKEWN-2017

Frederick Niehaus, TME - WNG, Cisco

SCHEDULE Tuesday, Jul 12, 8:00 a.m.

SCHEDULE Wednesday, Jul 13, 1:30 p.m.

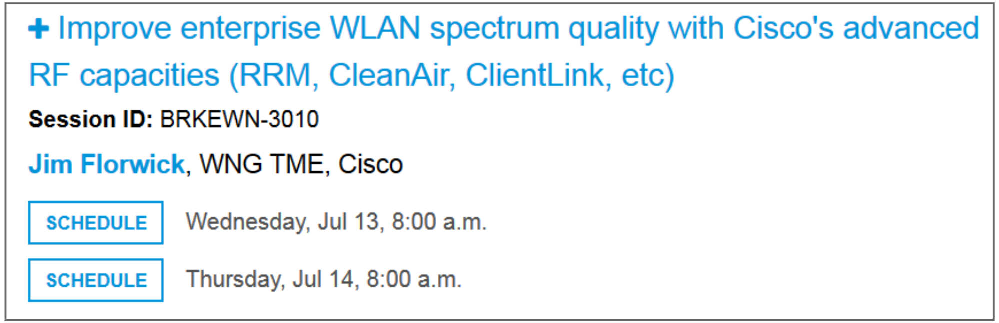

+ Improve enterprise WLAN spectrum quality with Cisco's advanced RF capacities (RRM, CleanAir, ClientLink, etc)

Session ID: BRKEWN-3010

Jim Florwick, WNG TME, Cisco

SCHEDULE Wednesday, Jul 13, 8:00 a.m.

SCHEDULE Thursday, Jul 14, 8:00 a.m.

Wireless Controller HA

Wireless Controller Deployment modes

	Cisco Cloud based	FlexConnect	Centralized	Converged Access
	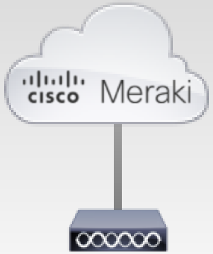	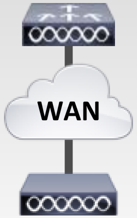	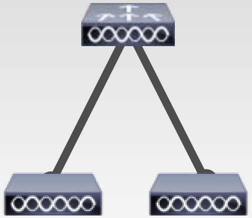	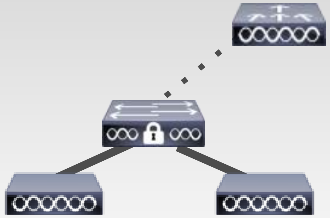
	Traffic Distributed at AP	Traffic Distributed at AP	Traffic Centralized at WLC	Traffic Distributed at Switch
Target Positioning	**Branch and Campus**	**Branch**	**Campus**	**Branch and Campus**
Purchase Decision	Public Cloud	Wireless only	Wireless only	Wired and Wireless
High Availability	• Multiple Geo distributed DCs • Data over at least 3 DCs • Distributed packet processing	• Full RF HA • Client SSO when Local Switching	• Most complete solution	• Exploits HA in IOS switches • Equivalent to AP SSO
Key Considerations	• Ease of Management, scalability, cloud based	• Branch with WAN BW and latency requirements	• Full features	• 3650/3850 at the access layer

Centralized Mode HA

Network Uptime

	Requirements	Benefits
Client SSO	Minimum release: 7.5 5500, WiSM2, 7500, 8500 series L2 connection between boxes Same HW and software 1:1 box redundancy	Active Client State is synched AP state is synched No Application downtime HA-SKU available
AP SSO (SSID stateful switchover)	Release: 7.3 and 7.4 5500, WiSM2, 7500, 8500 series Direct physical connection Same HW and SW 1:1 box redundancy	AP state is synched No SSID downtime HA-SKU available (> 7.4)
N+1 Redundancy (Deterministic/Stateless HA, a.k.a.: primary/secondary/tertiary)	Each Controller has to be configured separately	Available on all controllers Crosses L3 boundaries Flexible: 1:1, N:1, N:N HA-SKU available (> 7.4)

WLC redundancy with Mobility Group

Why not recommended?

- Mobility Group allows controllers to peer with each other to support Seamless and Fast roaming across controller boundaries
  - Support for up to 24 WLCs in the same Mobility Group
- Best Practice is to keep Mobility Group small and limited to the areas where seamless mobility can happen
- APs learn about all the WLCs in a Mobility Group at join time
- NOT Recommended for HA...why?

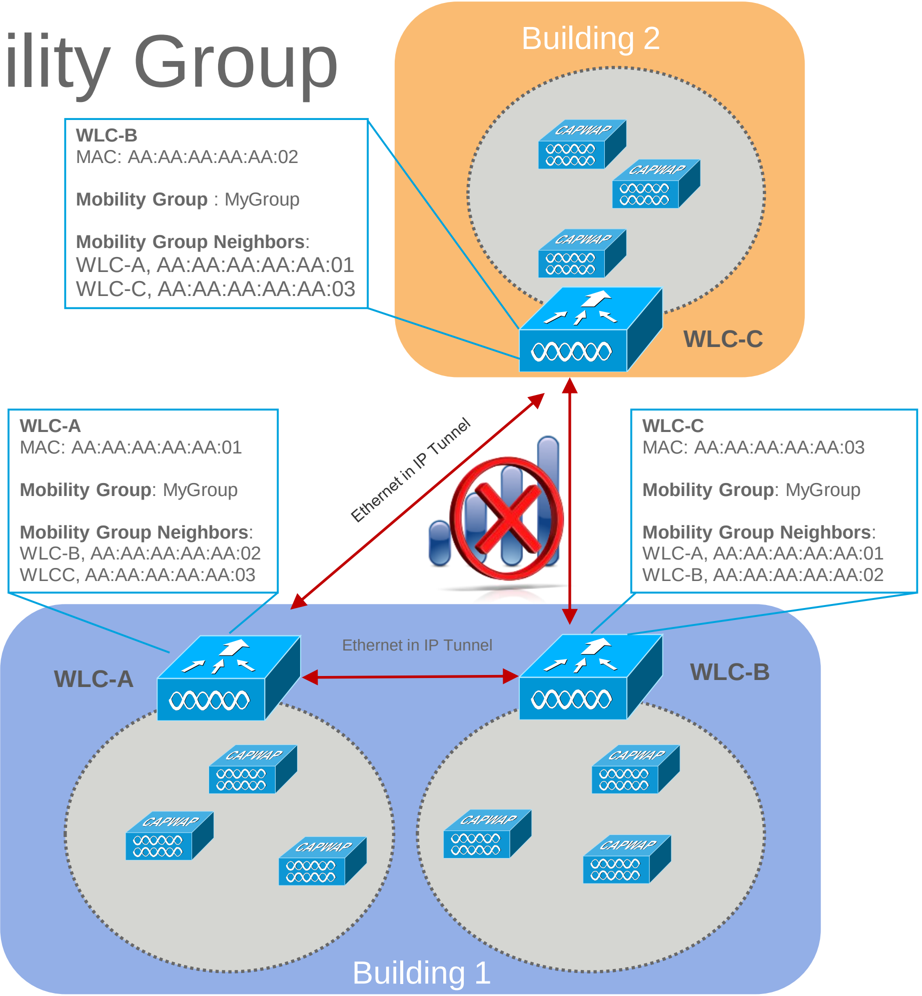

WLC-B
MAC: AA:AA:AA:AA:AA:02

Mobility Group : MyGroup

Mobility Group Neighbors:
WLC-A, AA:AA:AA:AA:AA:01
WLC-C, AA:AA:AA:AA:AA:03

WLC-A
MAC: AA:AA:AA:AA:AA:01

Mobility Group: MyGroup

Mobility Group Neighbors:
WLC-B, AA:AA:AA:AA:AA:02
WLCC, AA:AA:AA:AA:AA:03

WLC-C
MAC: AA:AA:AA:AA:AA:03

Mobility Group: MyGroup

Mobility Group Neighbors:
WLC-A, AA:AA:AA:AA:AA:01
WLC-B, AA:AA:AA:AA:AA:02

WLC redundancy with Mobility Group

Why not recommended?

When relying only on Mobility Group information:

- AP only learns available Controllers at JOIN time
- AP joins the least loaded WLC
- This could lead to “Salt and Pepper” deployment:
    - Same floor AP on different WLCs
    - More inter-controller roaming
    - Harder to troubleshoot
- For High Availability:
    - No concept of backup controller list
    - Upon loosing the registered controller, the AP has to start from scratch the whole Discovery process to all members of Mobility Group

“Salt and Pepper” Deployment

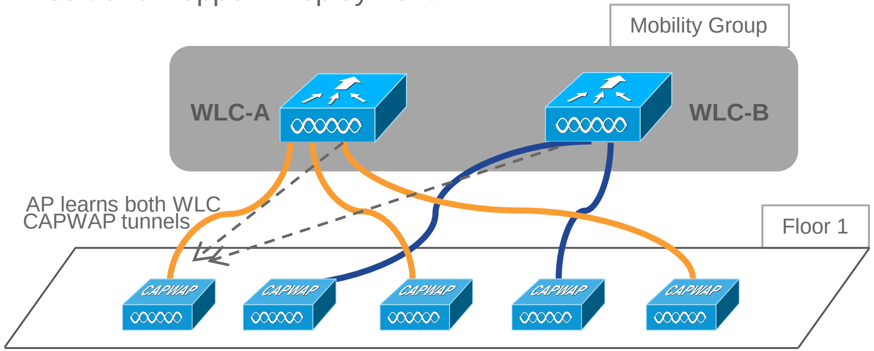

AP state machine

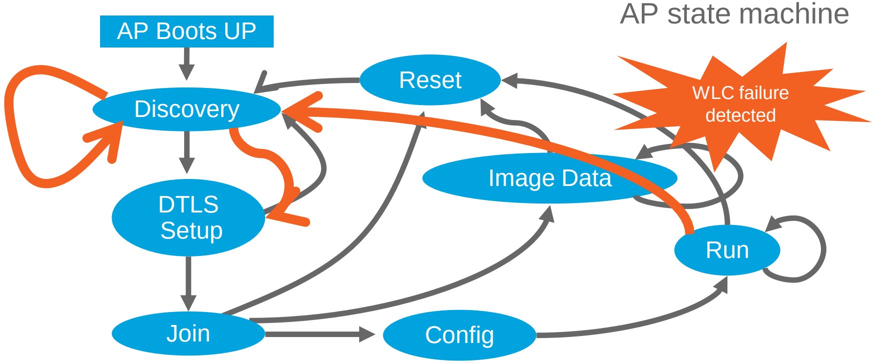

WLC redundancy with Mobility Group

Why not recommended

When relying only on Mobility Group information:

- AP only learns available Controllers at JOIN time
- AP joins the least loaded WLC
- This could lead to "Salt and Pepper" deployment:
  - Same floor AP on different WLCs
  - More inter-controller roaming
  - Harder to troubleshoot
- For High Availability:
  - No concept of backup controller list
  - Upon loosing the registered controller, the AP has to start from scratch the whole Discovery process to all members of Mobility Group
  - Failover takes more time and it's not deterministic: you don't know where the AP will end up

"Salt and Pepper" Deployment

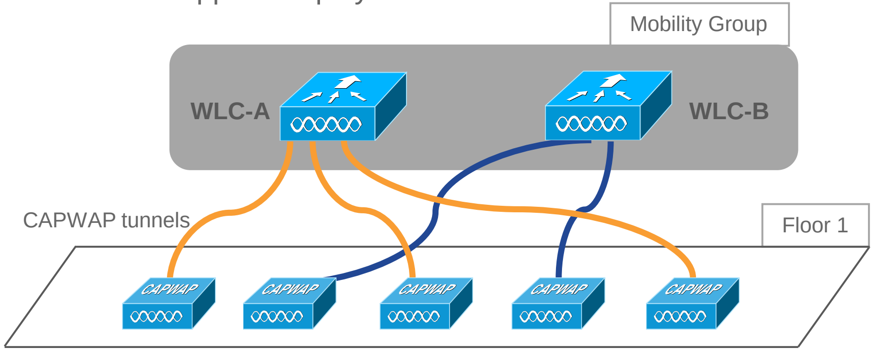

Not deterministic

N+1 Redundancy

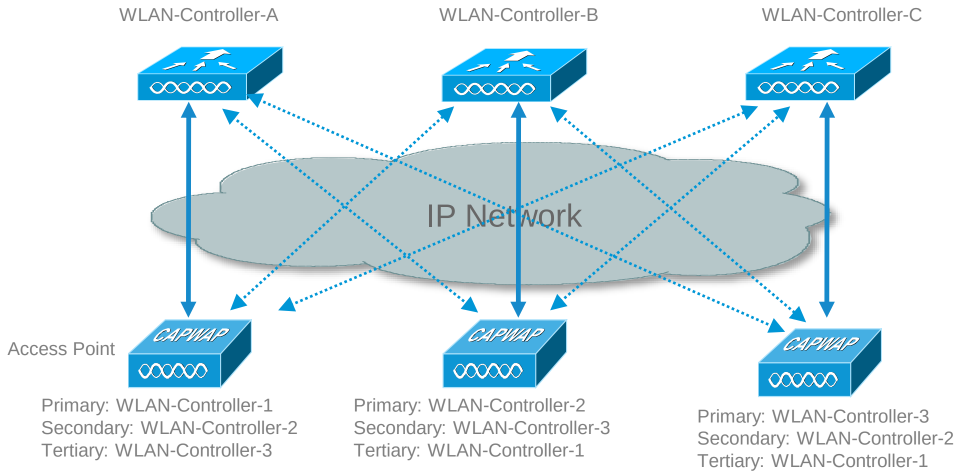

Primary: WLAN-Controller-1
Secondary: WLAN-Controller-2
Tertiary: WLAN-Controller-3

Primary: WLAN-Controller-2
Secondary: WLAN-Controller-3
Tertiary: WLAN-Controller-1

Primary: WLAN-Controller-3
Secondary: WLAN-Controller-2
Tertiary: WLAN-Controller-1

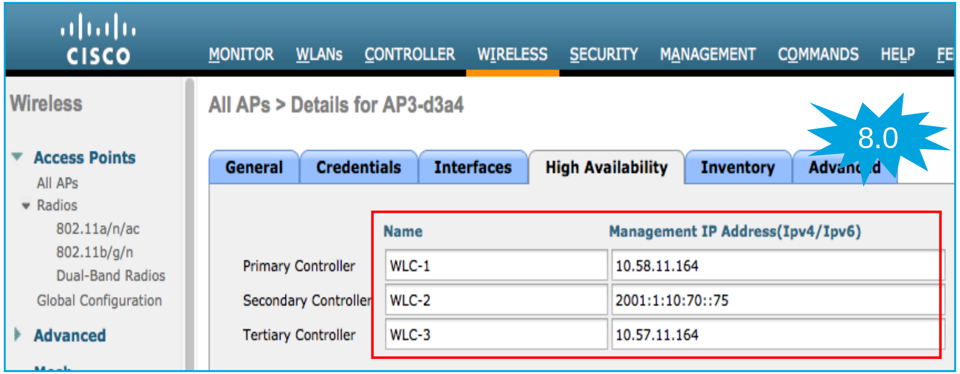

All APs > Details for AP3-d3a4

	Name	Management IP Address(Ipv4/Ipv6)
Primary Controller	WLC-1	10.58.11.164
Secondary Controller	WLC-2	2001:1:10:70::75
Tertiary Controller	WLC-3	10.57.11.164

- Administrator statically assigns APs a primary, secondary, and/or tertiary controller
  - Assigned from controller interface (per AP) or Prime Infrastructure (template-based)
  - You need to specify Name and IP if WLCs are not in the same Mobility Group
- **Pros:**
  - Predictability: easier operational management
  - Support for L3 network between WLCs
  - Flexible redundancy design options:1:1, N:1, N:N:1
  - WLCs can be of different HW and SW (*)
  - “Fallback” option in the case of failover
  - Can overload APs on controllers (using AP priority)
- **Cons:**
  - Stateless redundancy
  - More upfront planning and configuration

(*) AP will need to upgrade/downgrade code upon joining

N+1 Redundancy

Global backup Controllers

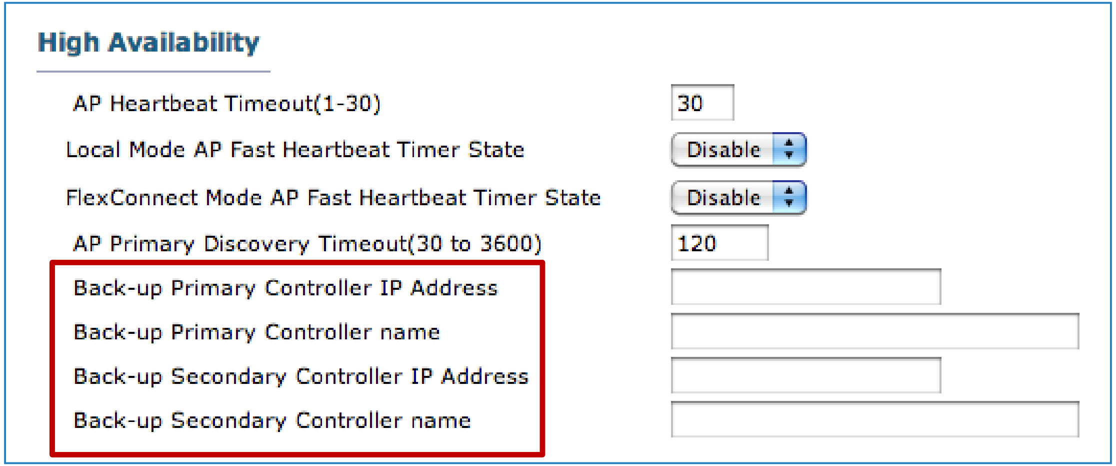

High Availability

AP Heartbeat Timeout(1-30)	30
Local Mode AP Fast Heartbeat Timer State	Disable
FlexConnect Mode AP Fast Heartbeat Timer State	Disable
AP Primary Discovery Timeout(30 to 3600)	120
Back-up Primary Controller IP Address	
Back-up Primary Controller name	
Back-up Secondary Controller IP Address	
Back-up Secondary Controller name	

- Backup controllers configured for all APs under Wireless > High Availability
- Used if there are no primary/secondary/tertiary WLCs configured on the AP
- The backup controllers are added to the primary discovery response message to the AP

N+1 Redundancy

AP Failover mechanism

- When configured with Primary and backup Controllers:
  - AP uses heartbeats to validate current WLC connectivity
  - Upon loosing a heartbeat to the Primary, AP sends 5 consecutives heartbeats every 3 second (default)
    - Configurable to minimum of 3 keepalive every 2 sec
  - If no reply, AP starts the join process to the first backup WLC candidate:
    - Backup is the first alive WLC in this order: primary, secondary, tertiary, global primary, global secondary.
  - With N+1 Failover, AP goes back to discovery state just to make sure the backup WLC is UP and then immediately starts the JOIN process
  - With N+1, AP periodically checks for Primary to come back online and falls back to it (AP fallback can be disabled)

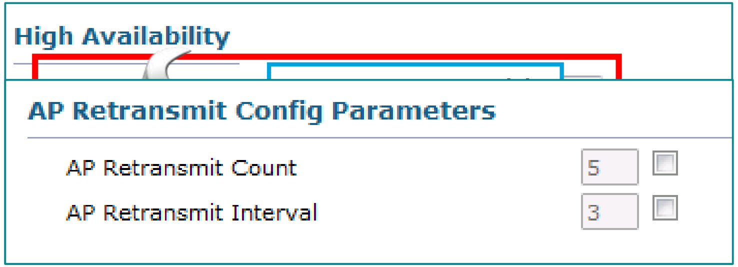

High Availability

AP Retransmit Config Parameters

AP Retransmit Count	5	□
AP Retransmit Interval	3	□

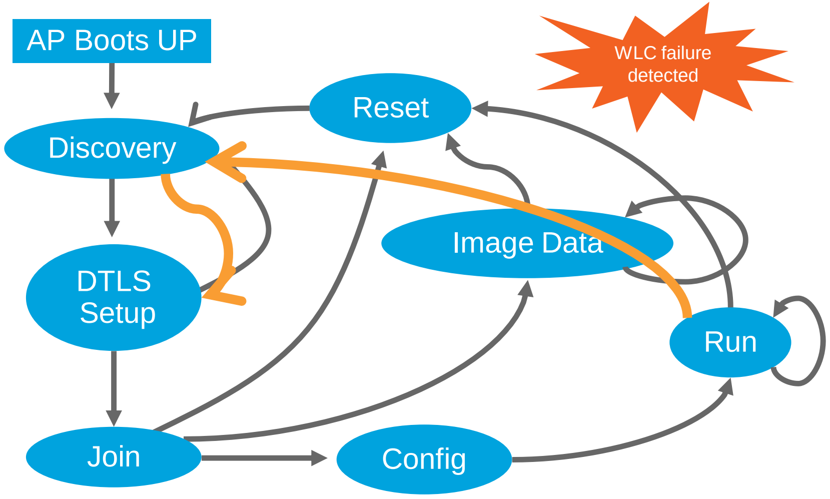

(*) With Fast Heartbeat and minimum values for keepalive

N+1 Redundancy

AP Fast Heartbeat

< 30-45 sec (*)

- Fast Heartbeats lower the amount of time it takes to detect Primary controller failure
- How Fast Heartbeat works
  - AP sends these packets, by default every 1 sec
  - When the fast heartbeat timer expires, the AP sends a 3 fast echo requests to the WLC for 3 times (configurable)

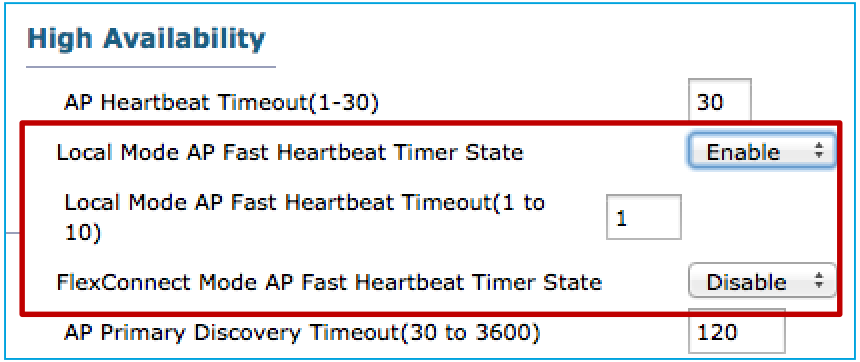

High Availability

AP Heartbeat Timeout(1-30)	30
Local Mode AP Fast Heartbeat Timer State	Enable
Local Mode AP Fast Heartbeat Timeout(1 to 10)	1
FlexConnect Mode AP Fast Heartbeat Timer State	Disable
AP Primary Discovery Timeout(30 to 3600)	120

  - If no response primary is considered dead and the AP selects an available controller from its “backup controller” list in the order of primary, secondary, tertiary, primary backup controller, and secondary backup controller.
- Fast Heartbeat only supported for Local and Flex mode

N+1 Redundancy

AP Primary Discovery Request Timer

- The access point periodically sends primary discovery requests to the Primary WLC to know when it is back online. Default is 120 sec.
- If AP Fallback is enabled (default), the AP automatically joins back the Primary controller

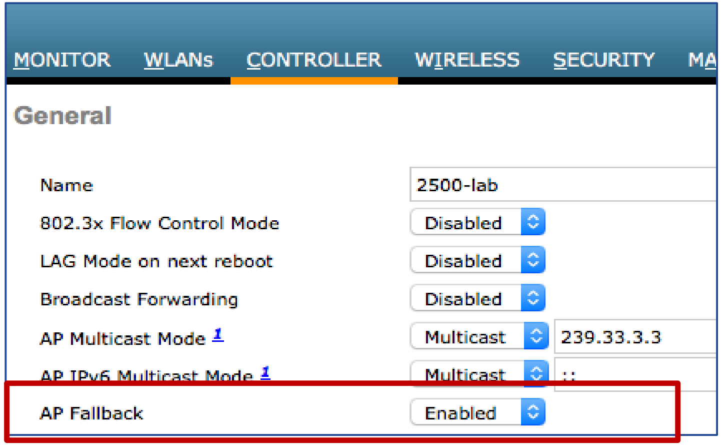

General

Name	2500-lab
802.3x Flow Control Mode	Disabled
LAG Mode on next reboot	Disabled
Broadcast Forwarding	Disabled
AP Multicast Mode [1]	Multicast 239.33.3.3
AP IPv6 Multicast Mode [1]	Multicast ::
AP Fallback	Enabled

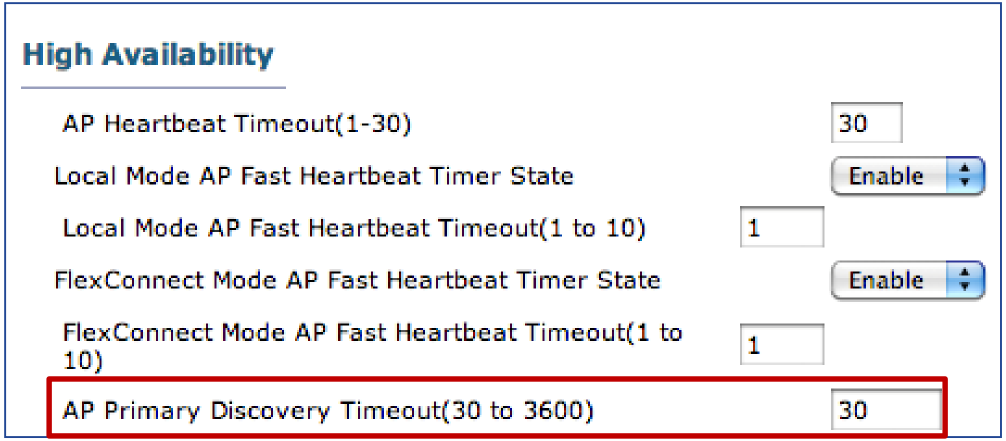

High Availability

AP Heartbeat Timeout(1-30)	30
Local Mode AP Fast Heartbeat Timer State	Enable
Local Mode AP Fast Heartbeat Timeout(1 to 10)	1
FlexConnect Mode AP Fast Heartbeat Timer State	Enable
FlexConnect Mode AP Fast Heartbeat Timeout(1 to 10)	1
AP Primary Discovery Timeout(30 to 3600)	30

N+1 Redundancy

AP Failover Priority

- Assign priorities to APs: Critical, High, Medium, Low
- Critical priority APs get precedence over all other APs when joining a controller
- In a failover situation, a higher priority AP will be allowed to join ahead of all other APs
- If backup controller doesn't have enough licenses (ex. multiple Primary WLCs fail), existing lower priority APs will be dropped to accommodate higher priority APs

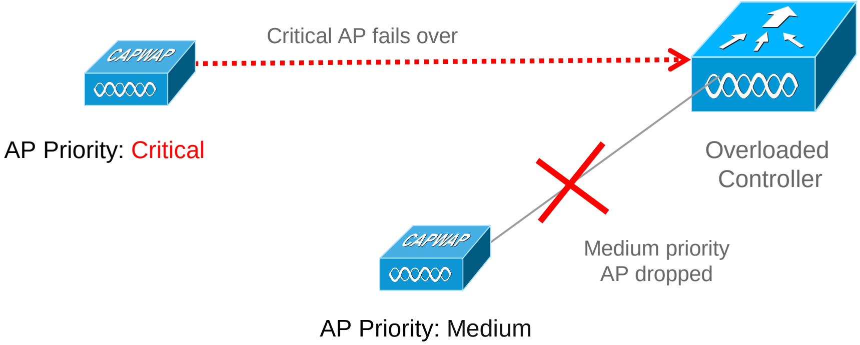
Critical AP fails over

AP Priority: Critical

Overloaded Controller

Medium priority AP dropped

AP Priority: Medium

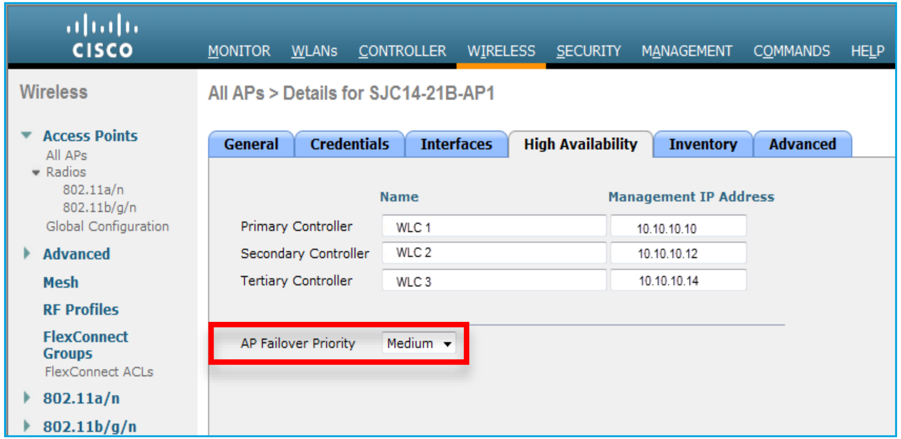

All APs > Details for SJC14-21B-AP1

	Name	Management IP Address
Primary Controller	WLC 1	10.10.10.10
Secondary Controller	WLC 2	10.10.10.12
Tertiary Controller	WLC 3	10.10.10.14

AP Failover Priority: Medium

N+1 Redundancy

Controller HA SKU

- The HA-SKU was introduced in 7.4 for 5508, WiSM2, Flex7500, 8510 and in 7.5 for 2504
  - It provides the support for the maximum number of APs on the specific hardware platform
  - It needs to be configured as you would with the secondary controller (no auto synch with Primary).

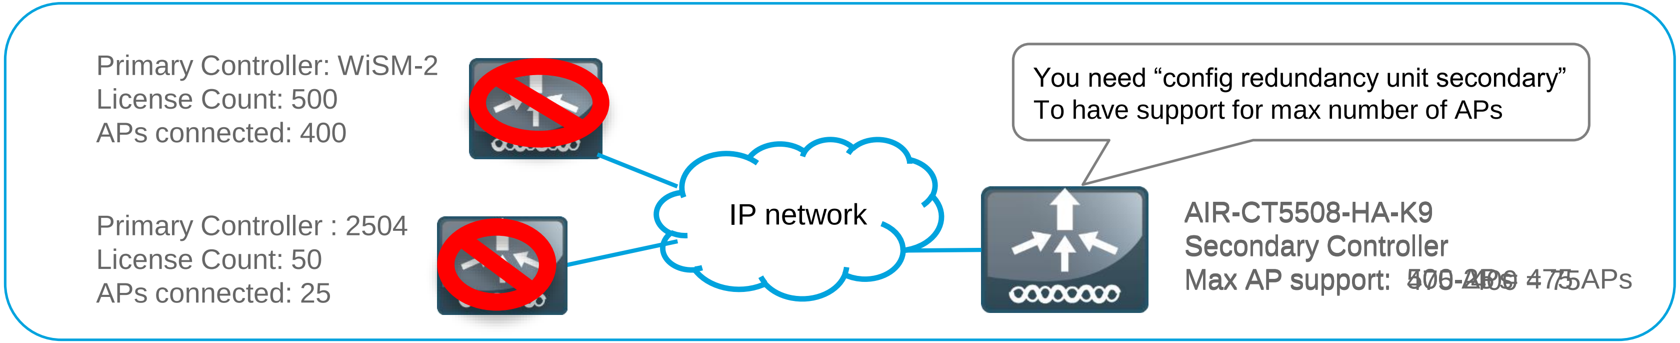

Primary Controller: WiSM-2
License Count: 500
APs connected: 400

Primary Controller : 2504
License Count: 50
APs connected: 25

IP network

You need "config redundancy unit secondary"
To have support for max number of APs

AIR-CT5508-HA-K9
Secondary Controller
Max AP support: 500-25 = 475 APs / 475-400 = 75 APs

- Other important information:
  - For 5508 (2504) you need a minimum of 50 (5) PERMANENT licenses to convert it into HA-SKU
  - From 7.6 you can convert HA-SKU to Primary and use it as Active controller *(you'd need to add licenses, of course*)
  - In 8.0 no more nagging message on the console after 90 days from first AP joining
  - New 5520 and 8540 Controllers do not have an HA SKU, use the zero AP SKU instead

N+1 Redundancy

Typical Design

< 30-45 sec (*)

- Most common Design is N+1 with Redundant WLC in a geographically separate location
- Can provide 30-45 sec of downtime when use faster heartbeat to detect failure
- Use AP priority in case of over subscription of redundant WLC
- Use HA SKU for the backup Controller
  - available for 5508, 7500, 8500 since 7.4 and for 2500 from release 7.5

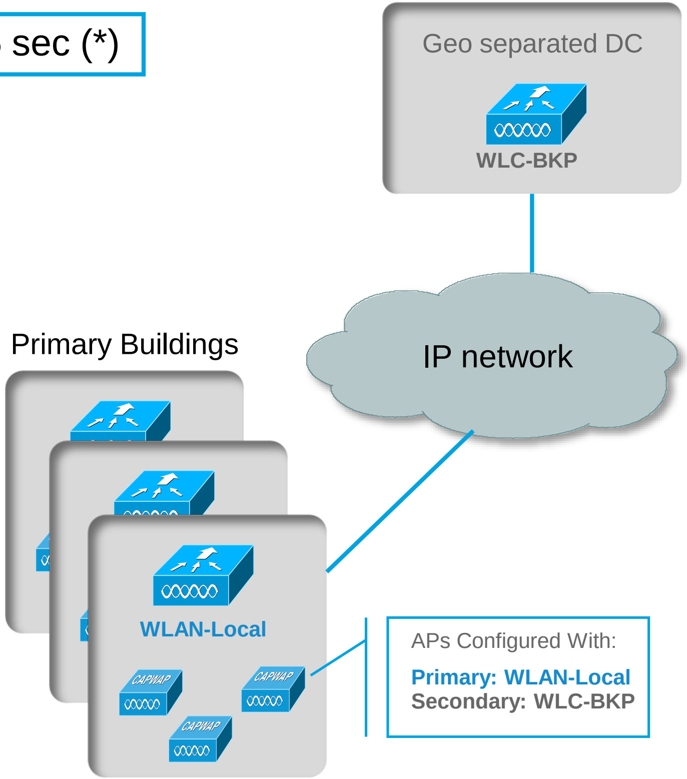

APs Configured With:
Primary: WLAN-Local
Secondary: WLC-BKP

For more info: http://www.cisco.com/en/US/docs/wireless/technology/hi_avail/N1_HA_Overview.html or http://www.cisco.com/en/US/prod/collateral/wireless/ps6302/ps8322/ps10315/qa_c67-714540.html

Wireless Controller HA

Centralized Mode – Stateful Switch Over (SSO)

Stateful Switchover (SSO)

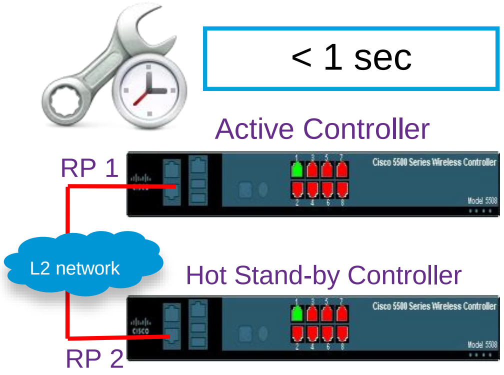

- True Box to Box High Availability i.e. 1:1
  - One WLC in Active state and second WLC in Hot Standby state
  - Secondary continuously monitors the health of Active WLC via dedicated link
- Configuration on Active is synched to Standby WLC
  - This happens at startup and incrementally at each configuration change on the Active
- What else is synched between Active and Standby?
  - AP CAPWAP state in 7.3 and 7.4: APs will not restart upon failover, SSID stays UP – AP SSO
  - Client in “RUN”/active state in 7.5: client will not disconnect – Client SSO
- Downtime during failover reduced is greatly reduced:
  - **2 - 100 msec** for a box failover (Active WLC crashes, system hangs, manual reset or forced switch-over)
  - **350-500 msec** in the case of power failure on the Active WLC (no direct command for switchover is possible)
  - **Few seconds** in the case of network failover (gateway not reachable)

Stateful Switchover (SSO)

What’s the impact on client applications?

	Ping	May loose one ping
	VoIP Call	Voice call stays up
	MS Lync	No session drop
	Citrix VDI	No impact

video: https://www.youtube.com/watch?v=If5F7eZkC3w

Stateful Switchover (SSO)

Failover sequence

1. Redundancy role negotiation and config sync
2. APs associates with Active controller
3. Client associates with Active through AP
4. Active failure: notify peer / or missing keep alive
5. Standby WLC sends out GARP
6. Standby becomes Active:

   AP DB and Client DB (7.5) is already synced with standby controller

   AP CAPWAP tunnel session intact

   Client session intact, client does not re-associate*

Effective downtime for the client is:
Detection time + Switchover time
+ (client association if AP SSO)

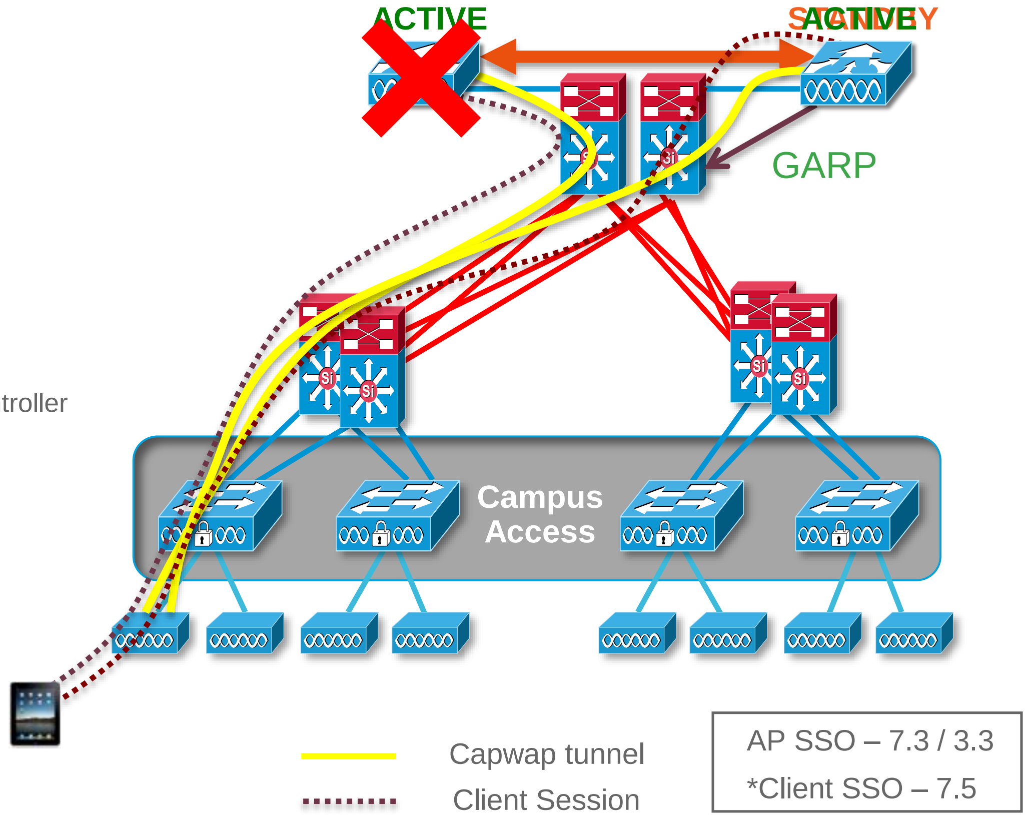

AP SSO – 7.3 / 3.3

*Client SSO – 7.5

Stateful Switchover (SSO)

Pairing the boxes

- HA Pairing is possible only between the same type of hardware and software versions
- 5500/7500/8500 have dedicated Redundancy Ports
  - Direct connection supported in 7.3 and 7.4
  - L2 connection supported in 7.5 and above
- WiSM-2 has dedicated Redundancy VLAN
  - Redundancy VLAN should be a non-routable VLAN
  - WISM-2 can be deployed in single chassis OR multiple chassis
    - WISM-2 in multiple chassis needs to use VSS (7.3, 7.4)
    - WISM-2 in multiple chassis can be L2 connected in 7.5 and above
- Requirements for L2 connection: RTT Latency: < 80 ms; Bandwidth: > 60 Mbps; MTU: 1500

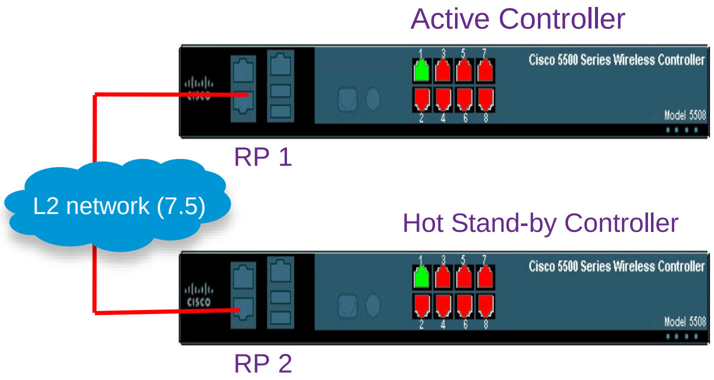

Stateful Switch Over (SSO)

Redundancy Management Interface

- Redundancy Management Interface (RMI)
  - To check gateway reachability sending ICMP packets every 1 sec
  - Peer reachability once the Active does not respond to Keepalive on the Redundant Port
  - Notification to standby in event of box failure or manual reset
  - Communication with Syslog, NTP, TFTP server for uploading configurations
  - Must be in same subnet as Management Interface. **From 8.0 the Management VLAN needs to be tagged**

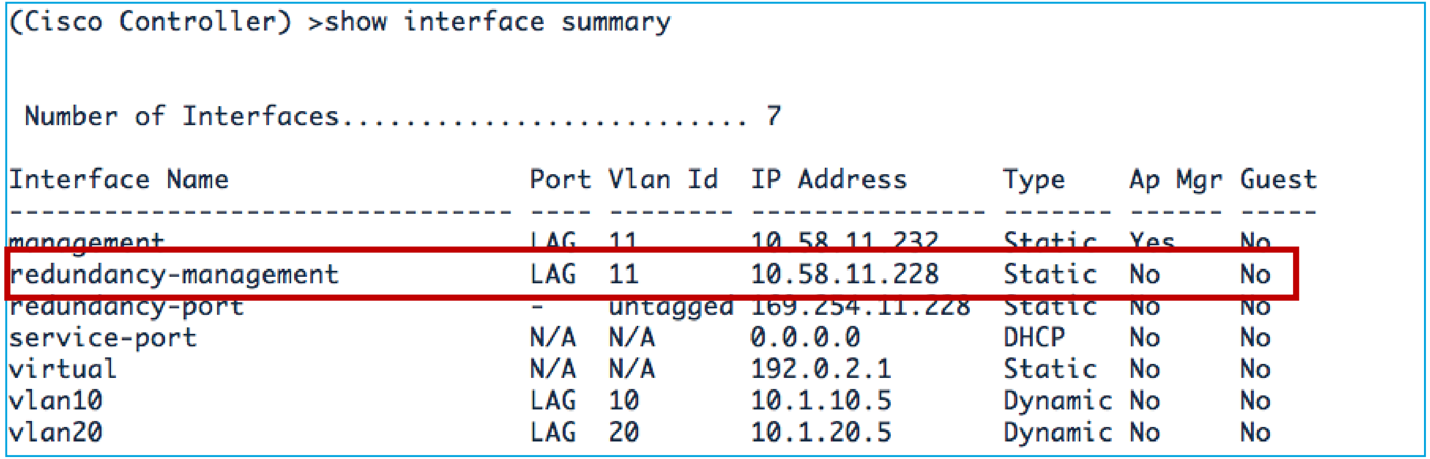

```
(Cisco Controller) >show interface summary


 Number of Interfaces.......................... 7

Interface Name                   Port Vlan Id  IP Address      Type    Ap Mgr Guest
-------------------------------- ---- -------- --------------- ------- ------ -----
management                       LAG  11       10.58.11.232    Static  Yes    No
redundancy-management            LAG  11       10.58.11.228    Static  No     No
redundancy-port                  -    untagged 169.254.11.228  Static  No     No
service-port                     N/A  N/A      0.0.0.0         DHCP    No     No
virtual                          N/A  N/A      192.0.2.1       Static  No     No
vlan10                           LAG  10       10.1.10.5       Dynamic No     No
vlan20                           LAG  20       10.1.20.5       Dynamic No     No
```

Stateful Switchover (SSO)

Redundancy Port

- Redundancy Port (RP):
  - Active/Standby role negotiation
  - Configuration synch from Active to Standby (bulk and incremental configuration)
  - Peer reachability sending UDP keep alive messages every 100 msec
  - Notification to standby in event of box failure
  - Time synch with peer, if NTP not available
  - Auto generated IP Address where last 2 octets are picked from the last 2 octets of RMI

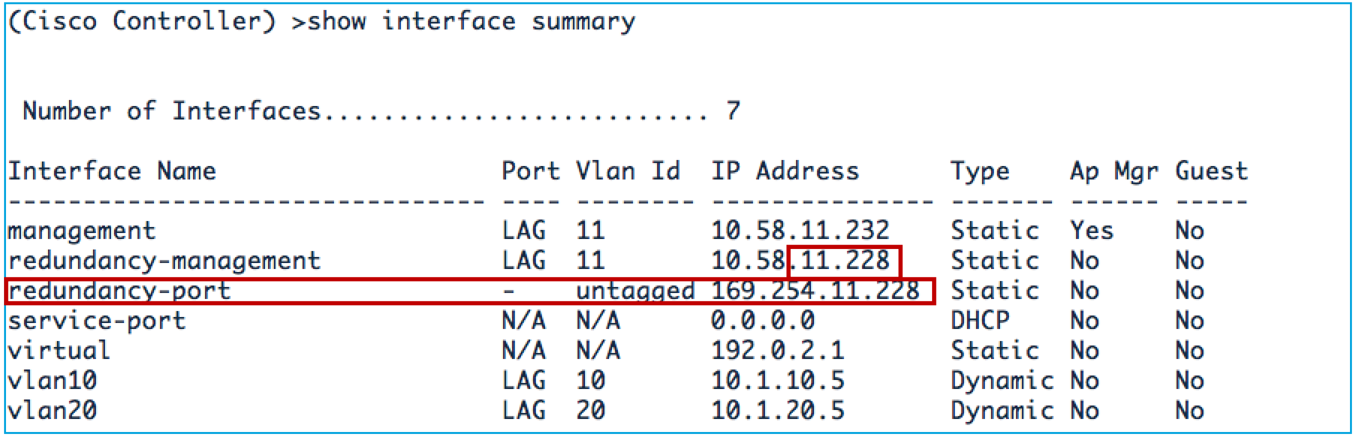

```
(Cisco Controller) >show interface summary


 Number of Interfaces.......................... 7

Interface Name                   Port Vlan Id  IP Address      Type    Ap Mgr Guest
-------------------------------- ---- -------- --------------- ------- ------ -----
management                       LAG  11       10.58.11.232    Static  Yes    No
redundancy-management            LAG  11       10.58.11.228    Static  No     No
redundancy-port                  -    untagged 169.254.11.228  Static  No     No
service-port                     N/A  N/A      0.0.0.0         DHCP    No     No
virtual                          N/A  N/A      192.0.2.1       Static  No     No
vlan10                           LAG  10       10.1.10.5       Dynamic No     No
vlan20                           LAG  20       10.1.20.5       Dynamic No     No
```

Stateful Switchover (SSO)

For Your Reference

Configuration

- Management interfaces on both WLCs must be on the same subnet
- Mandatory Configuration for HA setup:
  - Redundant Management IP Address
  - Peer Redundant Management IP Address
  - Redundancy Mode set to SSO enable (7.3 and 7.4 would show AP SSO)
  - Primary/Secondary Configuration – Required if peer WLC's UDI is not HA SKU
  - The Primary HA must have valid AP licenses
  - Unit can be secondary if it has at least 50 AP (5508) permanent licenses (no restrictions for other WLCs)

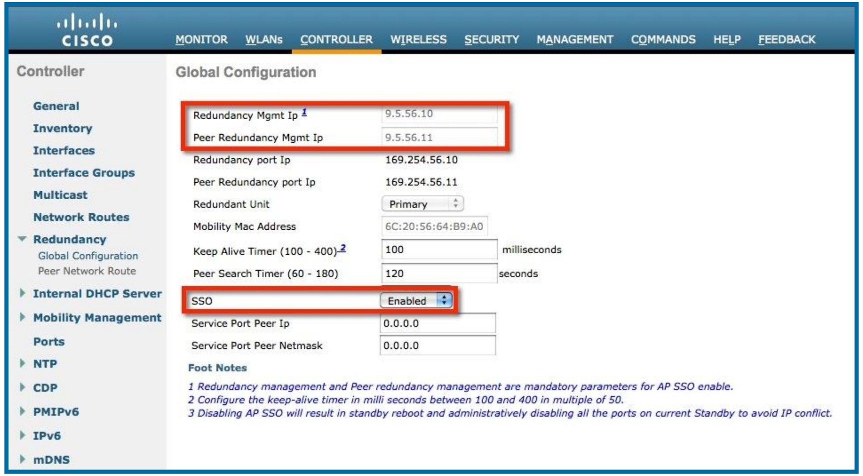

Global Configuration

Redundancy Mgmt Ip [1]	9.5.56.10	
Peer Redundancy Mgmt Ip	9.5.56.11	
Redundancy port Ip	169.254.56.10	
Peer Redundancy port Ip	169.254.56.11	
Redundant Unit	Primary	
Mobility Mac Address	6C:20:56:64:B9:A0	
Keep Alive Timer (100 - 400)[2]	100	milliseconds
Peer Search Timer (60 - 180)	120	seconds
SSO	Enabled	
Service Port Peer Ip	0.0.0.0	
Service Port Peer Netmask	0.0.0.0	

Foot Notes

1 Redundancy management and Peer redundancy management are mandatory parameters for AP SSO enable.
2 Configure the keep-alive timer in milli seconds between 100 and 400 in multiple of 50.
3 Disabling AP SSO will result in standby reboot and administratively disabling all the ports on current Standby to avoid IP conflict.

Optional Configuration:

- Service Port Peer IP
- Mobility MAC Address
- Keep Alive and Peer Search Timer

Stateful Switchover (SSO)

For Your Reference

Connectivity to the boxes

- Once SSO is enabled:
  - Connect to Standby WLC using console or SSH to Service Port and RMI
  - TFTP, NTP and Syslog traffic use the RMI interface on the Standby WLC
  - Telnet / SSH / SNMP / Web Access is not available on Management and Dynamic interface on Standby WLC
  - There is no SNMP or GUI access on the service port for both WLCs in the HA setup
- When SSO is disabled:
  - Configuration done on Active is pushed to Standby; after rebooting all the ports will come up on Active and will be disabled on Standby
  - This is to avoid network conflicts because the two WLCs have the same configuration

Stateful Switchover (SSO)

For Your Reference

Maintenance Mode

- Standby transitions to Maintenance Mode if:
  - Gateway not reachable via RMI Interface
  - Software mismatch
  - WLC with HA SKU has never discovered its peer
  - Redundant Port is down

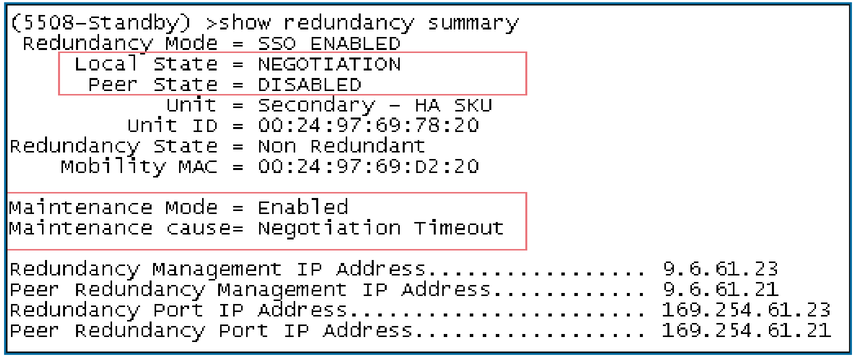

```
(5508-Standby) >show redundancy summary
 Redundancy Mode = SSO ENABLED
     Local State = NEGOTIATION
      Peer State = DISABLED
            Unit = Secondary - HA SKU
         Unit ID = 00:24:97:69:78:20
Redundancy State = Non Redundant
    Mobility MAC = 00:24:97:69:D2:20

Maintenance Mode = Enabled
Maintenance cause= Negotiation Timeout

Redundancy Management IP Address................. 9.6.61.23
Peer Redundancy Management IP Address............ 9.6.61.21
Redundancy Port IP Address....................... 169.254.61.23
Peer Redundancy Port IP Address.................. 169.254.61.21
```

- Important info:
  - When one of the conditions above is met, Standby reboots and goes to Maintenance mode
    - From 8.0 it will go directly without reboot
  - In Maintenance mode same rules to connect to standby box apply (console or Service port)
  - WLC should be rebooted to bring it out of Maintenance Mode
    - From 7.6 it will reboot automatically when pbs are fixed

Stateful Switchover (SSO)

What is synched/not synced between Active and Standby

All Configuration
Access Point data base
Client in “RUN” state (7.5)
Sleeping Client data base (8.0)
OEAP clients (8.0)
Internal DHCP (8.0)
Static CAP config and stats (8.0)

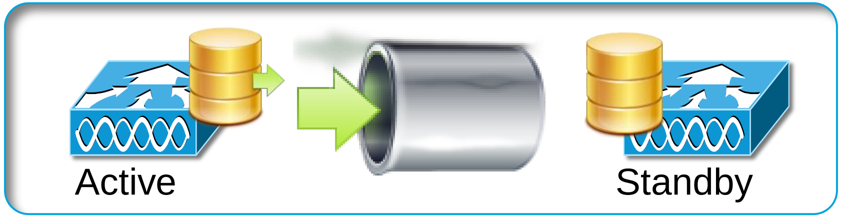

Not synched as of 8.2

Client Statistics
PMIPv6 info
NBAR (AVC) info
SIP static CAC info
WGB and associated clients
Passive clients DB
Location DB
Rogue AP and client info
Device and root certificates

Stateful Switchover (SSO)

Other important things to keep in mind..

- There is no preemption in Controller SSO:
  - when the failed Active WLC comes back online it will joining as Hot Standby
- Recommendations:
  - In Service Software Upgrade (ISSU) is not supported: plan for down time when upgrading software
  - Physical connection between Redundant Ports should be done first before HA configuration
  - Keepalive and Peer Discovery timers should be left at default values for better performance
- SSO and MESH APs:
  - only RAP are supported from 7.5, for MAPs the state is not synched
  - Use N+1 redundancy for a mesh based network

Stateful Switchover (SSO)

Changes introduced from release 8.0

- Gateway (GW) reachability changes:

GW reachability	7.3/7.4	7.5	8.0
protocol	ICMP	ICMP	ICMP
# of keepalives	3	12	6
ARP check	n.a.	n.a	Yes

- Release 8 introduced IPv6 support for the wireless infrastructure:
  - SSO (AP and Client) is supported with IPv6
  - Redundancy Management/Redundancy port interface supports only IPv4 addresses.

Stateful Switchover (SSO)

Changes introduced from release 8.0

- Peer Redundancy Management interface (RMI) reachability check:

RMI reachability	Before 8.0	8.0
protocol	ICMP	UDP
interval	1 sec	1 sec

- IEEE 802.1Q tag for Management VLAN: starting 8.0, Management and RMI interfaces are highly RECOMMENDED to be tagged
  - If upgrading from a previous release with untagged interface, the controller will show a warning message “ Untagged configuration is not recommended”

For Your Reference

Stateful Switchover (SSO)

Changes introduced from release 8.0

- Peer configuration:
  - new range for Keep Alive and Peer search timers
  - new Keep Alive Retries parameter

Before 8.0

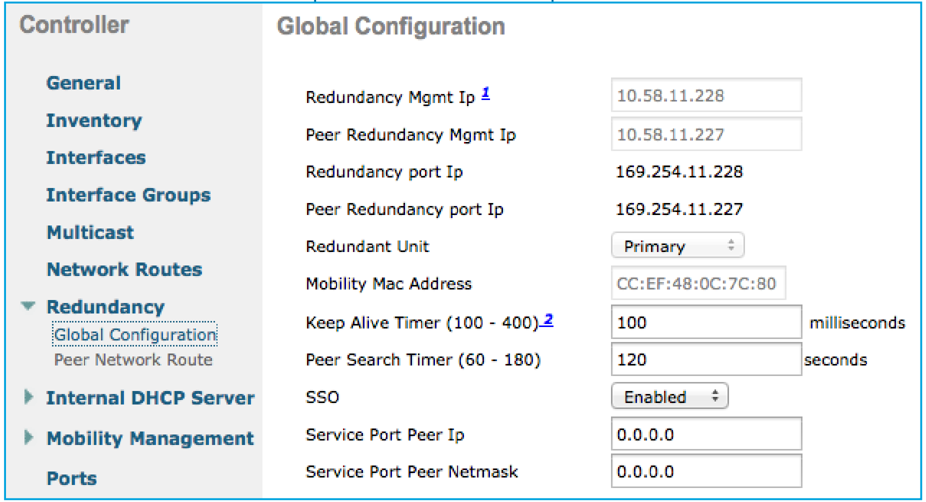

Global Configuration

Redundancy Mgmt Ip [1]	10.58.11.228	
Peer Redundancy Mgmt Ip	10.58.11.227	
Redundancy port Ip	169.254.11.228	
Peer Redundancy port Ip	169.254.11.227	
Redundant Unit	Primary	
Mobility Mac Address	CC:EF:48:0C:7C:80	
Keep Alive Timer (100 - 400) [2]	100	milliseconds
Peer Search Timer (60 - 180)	120	seconds
SSO	Enabled	
Service Port Peer Ip	0.0.0.0	
Service Port Peer Netmask	0.0.0.0	

From 8.0

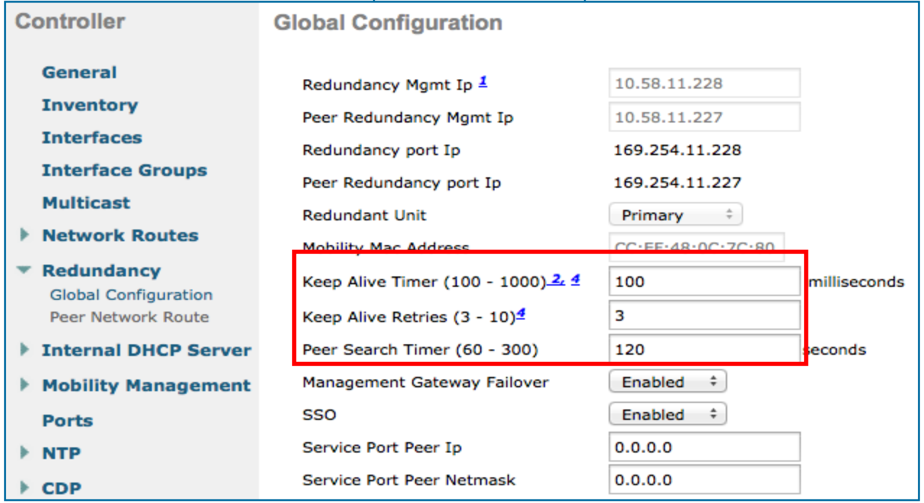

Global Configuration

Redundancy Mgmt Ip [1]	10.58.11.228	
Peer Redundancy Mgmt Ip	10.58.11.227	
Redundancy port Ip	169.254.11.228	
Peer Redundancy port Ip	169.254.11.227	
Redundant Unit	Primary	
Mobility Mac Address	CC:EF:48:0C:7C:80	
Keep Alive Timer (100 - 1000) [2, 4]	100	milliseconds
Keep Alive Retries (3 - 10) [4]	3	
Peer Search Timer (60 - 300)	120	seconds
Management Gateway Failover	Enabled	
SSO	Enabled	
Service Port Peer Ip	0.0.0.0	
Service Port Peer Netmask	0.0.0.0	

Stateful Switchover (SSO)

Changes introduced from release 8.1: Fast Restart

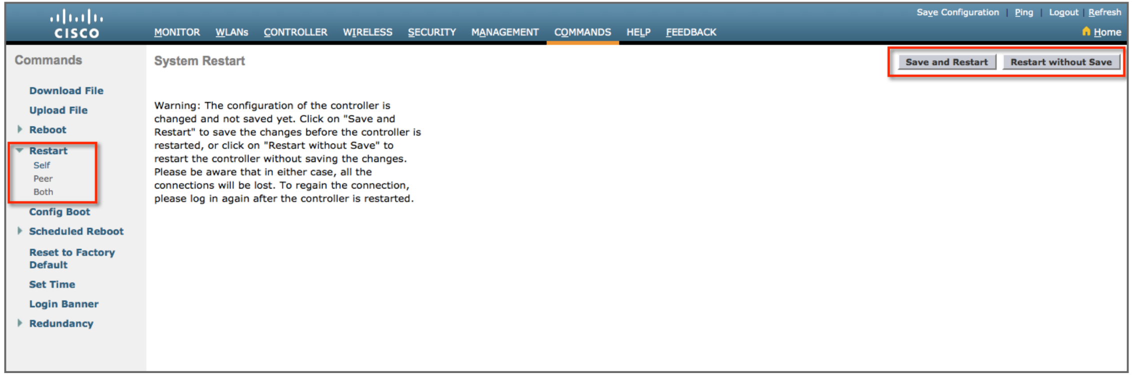

Warning: The configuration of the controller is changed and not saved yet. Click on "Save and Restart" to save the changes before the controller is restarted, or click on "Restart without Save" to restart the controller without saving the changes. Please be aware that in either case, all the connections will be lost. To regain the connection, please log in again after the controller is restarted.

- Process restart to reduce network and service downtime
- Supported on Cisco WLC 7510, 8510, 5520 8540 and vWLC
- CLI Command “`restart`”

When useful:

- ✓ LAG Configuration change
- ✓ Mobility Mode change
- ✓ Web-auth certificate installation
- ✓ Clear Configuration
- ✓ Post Configuration Wizard
- ✓ Transfer Download of configuration

8.1 release: Fast Restart

System Downtime with ‘reset system’

7510/8510

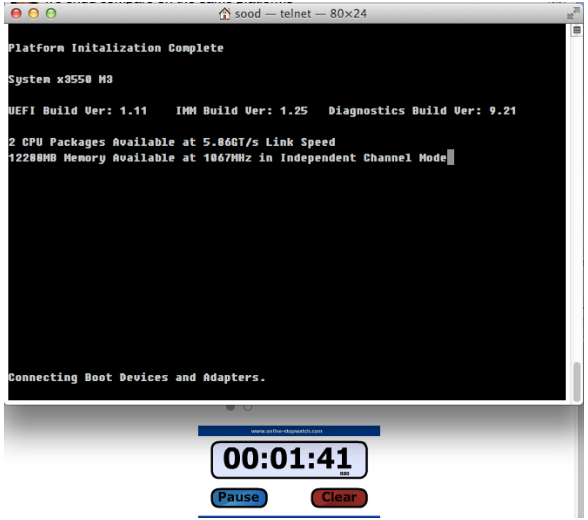

```
Platform Initalization Complete

System x3550 M3

UEFI Build Ver: 1.11    IMM Build Ver: 1.25    Diagnostics Build Ver: 9.21

2 CPU Packages Available at 5.86GT/s Link Speed
12288MB Memory Available at 1067MHz in Independent Channel Mode









Connecting Boot Devices and Adapters.
```

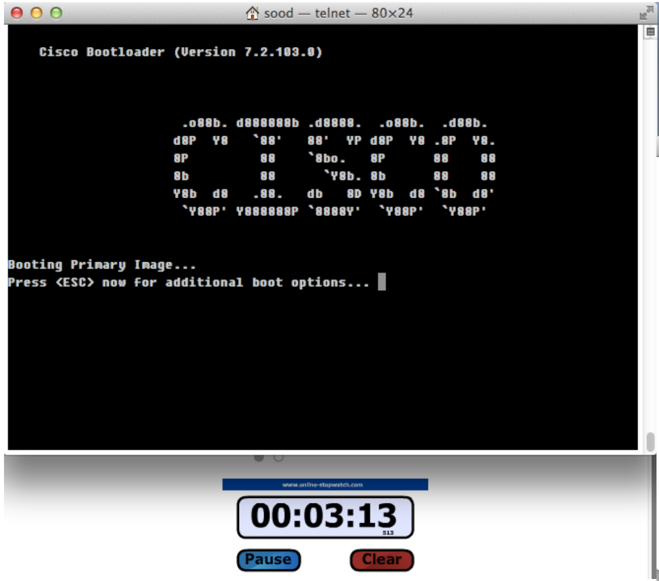

```
   Cisco Bootloader (Version 7.2.103.0)



              .o88b. d888888b .d8888.  .o88b.  .d88b.
             d8P  Y8   `88'   88'  YP d8P  Y8 .8P  Y8.
             8P         88    `8bo.   8P      88    88
             8b         88      `Y8b. 8b      88    88
             Y8b  d8   .88.   db   8D Y8b  d8 `8b  d8'
              `Y88P' Y888888P `8888Y'  `Y88P'  `Y88P'



Booting Primary Image...
Press <ESC> now for additional boot options...
```

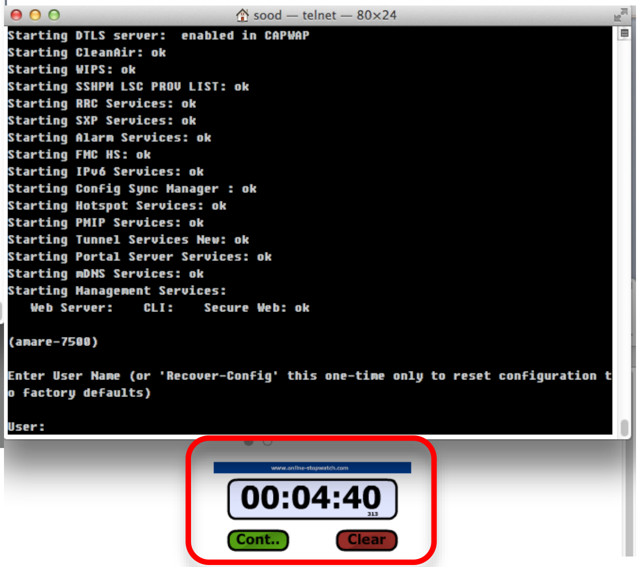

```
Starting DTLS server:  enabled in CAPWAP
Starting CleanAir: ok
Starting WIPS: ok
Starting SSHPM LSC PROV LIST: ok
Starting RRC Services: ok
Starting SXP Services: ok
Starting Alarm Services: ok
Starting FMC HS: ok
Starting IPv6 Services: ok
Starting Config Sync Manager : ok
Starting Hotspot Services: ok
Starting PMIP Services: ok
Starting Tunnel Services New: ok
Starting Portal Server Services: ok
Starting mDNS Services: ok
Starting Management Services:
   Web Server:    CLI:    Secure Web: ok

(amare-7500)

Enter User Name (or 'Recover-Config' this one-time only to reset configuration t
o factory defaults)

User:
```

Cisco live!

New in 8.1 release - Fast Restart

System Downtime with Fast ‘restart’

7510/8510

73% Faster

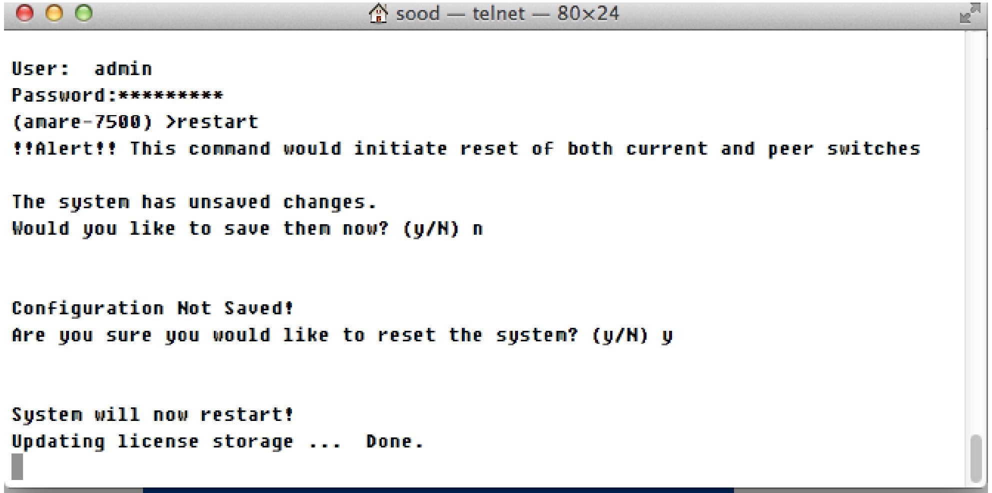

```
User:  admin
Password:*********
(amare-7500) >restart
!!Alert!! This command would initiate reset of both current and peer switches

The system has unsaved changes.
Would you like to save them now? (y/N) n


Configuration Not Saved!
Are you sure you would like to reset the system? (y/N) y


System will now restart!
Updating license storage ...  Done.
```

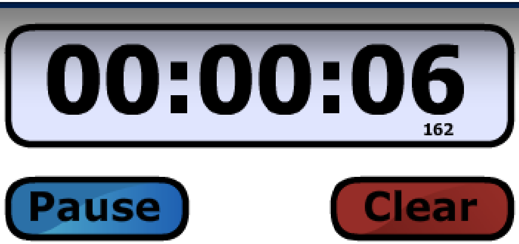

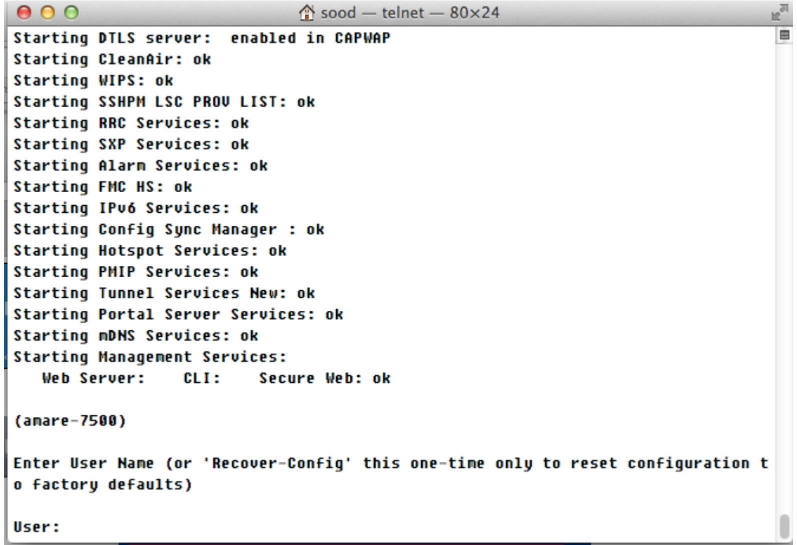

```
Starting DTLS server:  enabled in CAPWAP
Starting CleanAir: ok
Starting WIPS: ok
Starting SSHPM LSC PROV LIST: ok
Starting RRC Services: ok
Starting SXP Services: ok
Starting Alarm Services: ok
Starting FMC HS: ok
Starting IPv6 Services: ok
Starting Config Sync Manager : ok
Starting Hotspot Services: ok
Starting PMIP Services: ok
Starting Tunnel Services New: ok
Starting Portal Server Services: ok
Starting mDNS Services: ok
Starting Management Services:
   Web Server:    CLI:    Secure Web: ok

(amare-7500)

Enter User Name (or 'Recover-Config' this one-time only to reset configuration t
o factory defaults)

User:
```

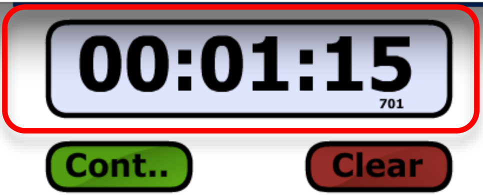

Cisco live!

New in 8.1 release - HA Standby Monitoring

For Your Reference

HA-SKU Trap, Events and Logging

1 WLC Turns Hot Standby - Trap

2 Bulk Sync Completion - Trap

3 Standby Reboot - Trap

4 Peer System, CPU, Memory details on Active GUI and CLI

5 Admin Login on Standby RMI - Syslog

Stateful Switchover (SSO)

Licensing

- You need valid licenses on the Active for HA to work (permanent or evaluation)
- As Standby you can use HA-SKU or a “converted” existing WLC
- To convert any existing WLC to a Standby WLC:
    - Use the “`config redundancy unit secondary`” command in the CLI or GUI equivalent.
    - **Restriction**: on the 5508 a minimum of 50 AP **Permanent** licenses are needed.
- What happens to licenses when you create a HA pair? Example with HA-SKU:
    - The device with HA-SKU becomes Standby first time it pairs up
    - AP-count licenses will be pushed from Active to Standby
    - On event of Active failure, HA-SKU will let APs join and start a 90-day count-down.
    - After 90-days, HA-SKU WLC starts nagging messages but won’t disconnect connected APs
    - New WLC joins as Standby and timer is reset if the new WLC has a number of licenses >= to the failed one.

Stateful Switchover (SSO)

Adding licenses to a SSO pair

- The licenses are added to the ACTIVE controller
  - If using the HA-SKU make sure that the ACTIVE is the Primary controller
- No need to break the pair and/or reboot. The HOT-STANDBY inherits the new added licenses
  - From 8.1, a reboot is recommended for the 5508 and WISM2 (not needed for 8510/8540/5520)
- Let’s see the actual steps:

Active

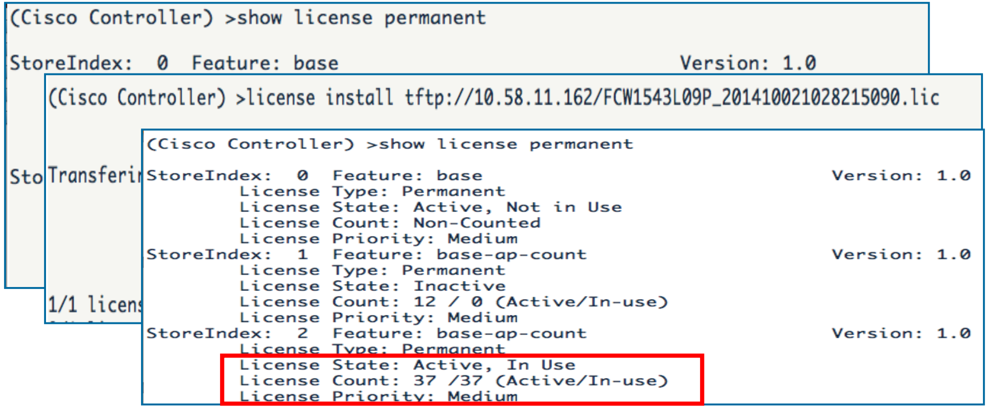

```
(Cisco Controller) >show license permanent

StoreIndex:  0  Feature: base                              Version: 1.0



Sto
```

```
(Cisco Controller) >license install tftp://10.58.11.162/FCW1543L09P_201410021028215090.lic


Transferir



1/1 licens
```

```
(Cisco Controller) >show license permanent

StoreIndex:  0  Feature: base                              Version: 1.0
        License Type: Permanent
        License State: Active, Not in Use
        License Count: Non-Counted
        License Priority: Medium
StoreIndex:  1  Feature: base-ap-count                     Version: 1.0
        License Type: Permanent
        License State: Inactive
        License Count: 12 / 0 (Active/In-use)
        License Priority: Medium
StoreIndex:  2  Feature: base-ap-count                     Version: 1.0
        License Type: Permanent
        License State: Active, In Use
        License Count: 37 /37 (Active/In-use)
        License Priority: Medium
```

Standby

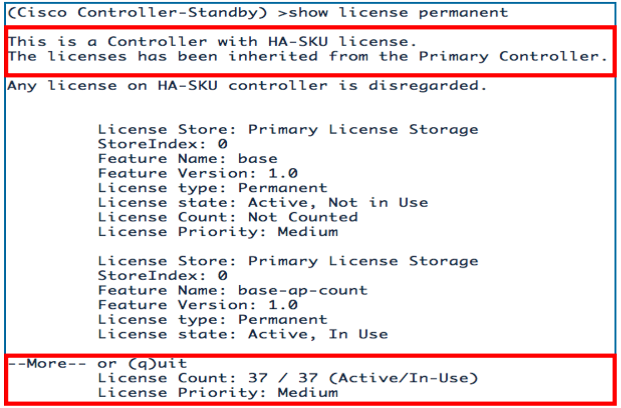

```
(Cisco Controller-Standby) >show license permanent

This is a Controller with HA-SKU license.
The licenses has been inherited from the Primary Controller.

Any license on HA-SKU controller is disregarded.


        License Store: Primary License Storage
        StoreIndex: 0
        Feature Name: base
        Feature Version: 1.0
        License type: Permanent
        License state: Active, Not in Use
        License Count: Not Counted
        License Priority: Medium

        License Store: Primary License Storage
        StoreIndex: 0
        Feature Name: base-ap-count
        Feature Version: 1.0
        License type: Permanent
        License state: Active, In Use

--More-- or (q)uit
        License Count: 37 / 37 (Active/In-Use)
        License Priority: Medium
```

Wireless Controller HA

FlexConnect Mode

FlexConnect quick recap....

- Control plane, two modes of operation:
  - Connected (when WLC is reachable)
  - Standalone (when WLC is not reachable)
- Data Plane can be:
  - Centralized (split MAC architecture) switching
  - Local (local MAC architecture) switching
- Traffic Switching mode is configured per AP and per SSID
  - From 7.3 split tunneling is supported on a WLAN basis
- FlexConnect Group:
  - Defines the Key caching domain for Fast Roaming, allows backup Radius scenarios
- From 8.0 Flex + Mesh mode supported
- WAN recommendations:
  - Minimum bandwidth 12.8 kbps per AP
  - Round trip latency no greater than 300 ms for data deployments and 100 ms for data + voice deployments

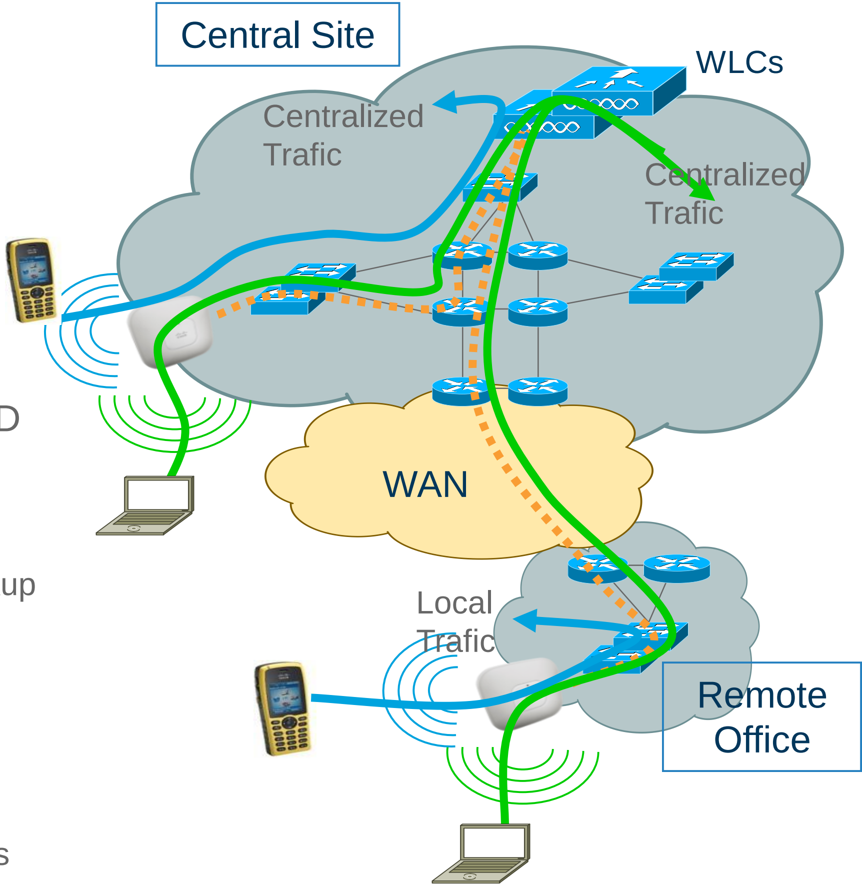

FlexConnect HA

	Limitations	Benefits
FlexConnect Local Switching	L2 roaming Flex Groups for AAA Local Auth. Fault Tolerance: Identical configuration on N+1 controllers	Upon WLC failure AP stays up and clients are not disconnected Equivalent to Client SSO AAA survivability available
FlexConnect Central Switching	Same as Centralized mode	Same as Centralized mode

For more info: http://www.cisco.com/en/US/products/ps11635/products_tech_note09186a0080b7f141.shtml

FlexConnect

WAN Failure (or single central WLC failure)

- HA considerations:
  - Disconnection for centrally switched SSIDs clients
  - No impact for connected clients on locally switched SSIDs
- Fast roaming allowed within FlexConnect group for already connected clients
- What about new clients?
  - Static keys are locally stored in FlexConnect AP: new clients can join if authentication is PSK
  - Can design for AAA survivability (see next slides)
- Lost features
  - RRM, CleanAir, WIDS, Location, other AP modes
  - Web authentication, NAC

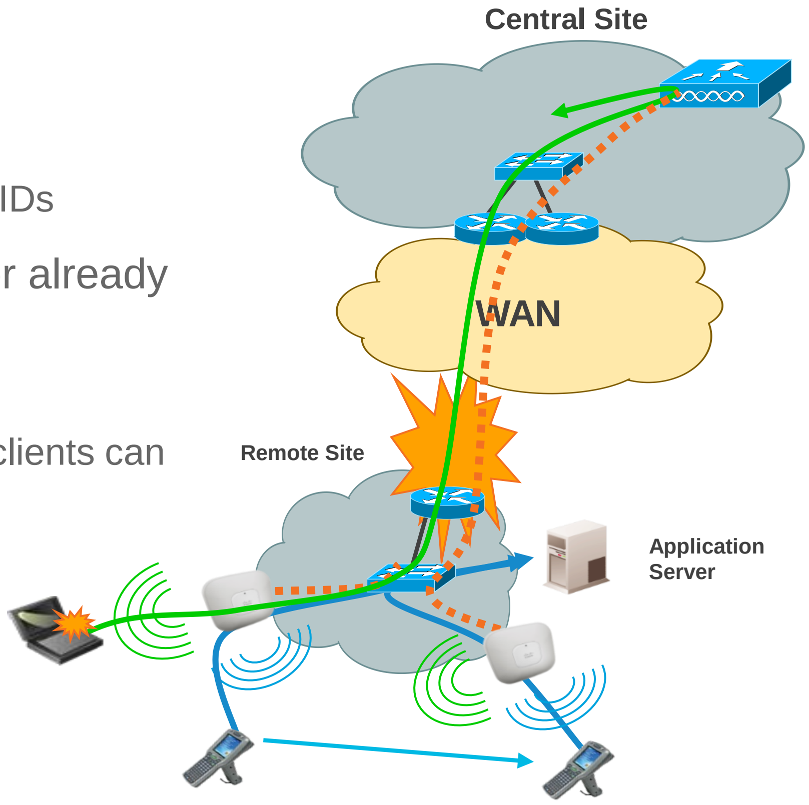

Centrally switched traffic

Locally switched traffic

FlexConnect

WLC failure with Deterministic N+1 HA

- HA considerations:
  - Disconnection for centrally switched SSIDs clients
  - No impact for connected clients on locally switched SSIDs
- FlexConnect AP transitions to Standalone and then to Connected when joins the Secondary
- When in Standalone mode, Fast roaming is allowed within the FlexConnect Group
- Fault Tolerant: upon re-syncing with Secondary, client sessions for local traffic are not impacted, provided that the configuration on the WLCs are identical

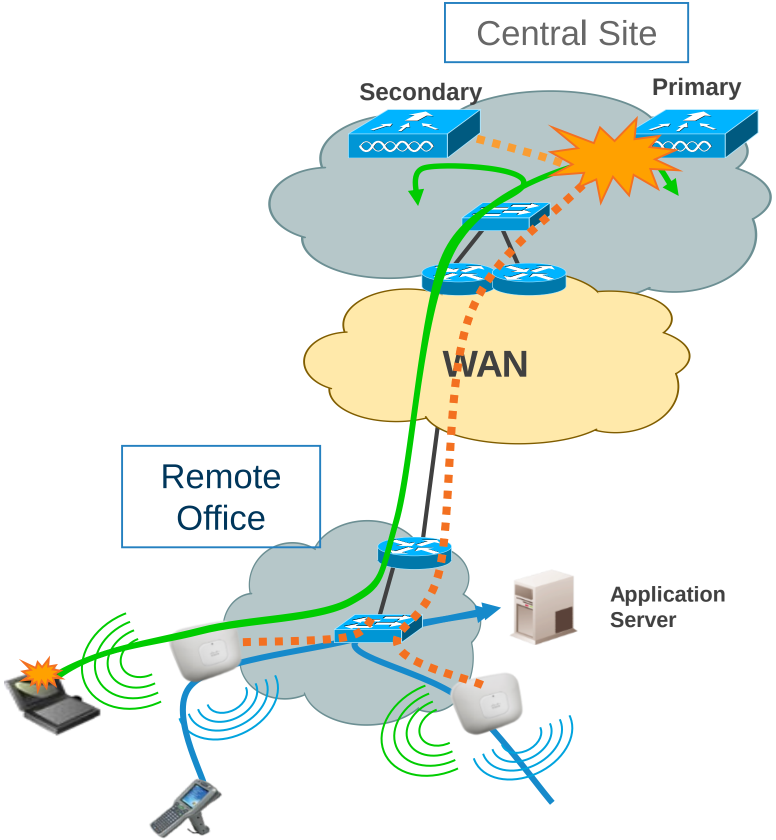

🟩 Centrally switched traffic

🟦 Locally switched traffic

FlexConnect

WLC failure scenario with SSO

- HA considerations:
  - No impact for locally switched SSIDs
  - Disconnection of centrally switched SSIDs clients with AP SSO
  - No/minimal impact for centrally switched client with Client SSO (7.5 and above)
- FlexConnect AP will NOT transition to Standalone because SSO kicks in
- AP will continue to be in Connected mode with the Standby (now Active) WLC

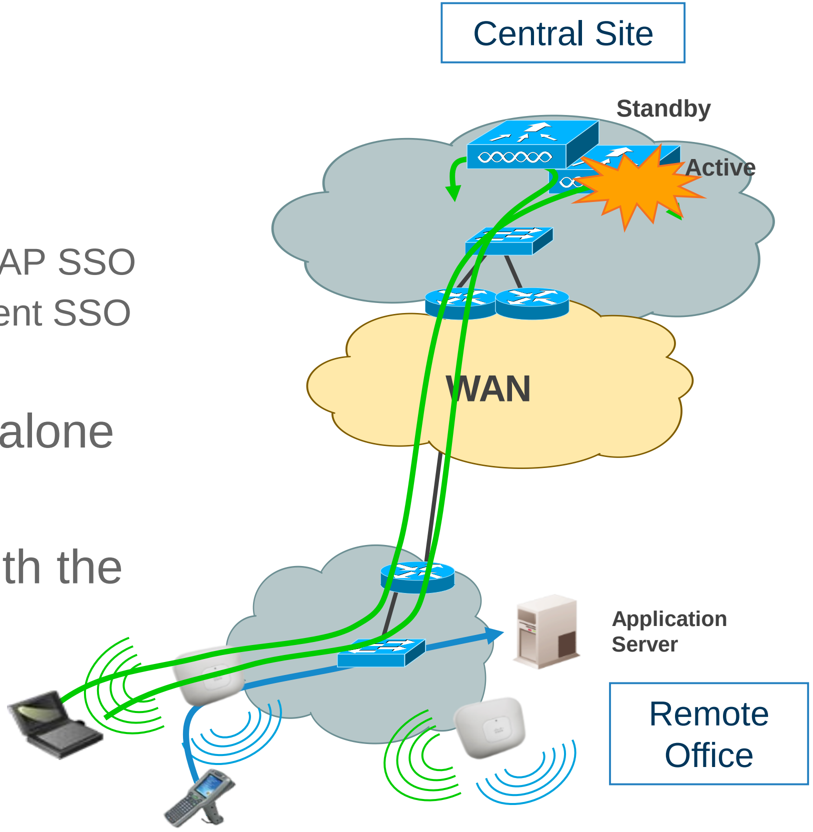

🟩 Centrally switched traffic

🟦 Locally switched traffic

FlexConnect AAA Survivability

AAA Server Backup

- By default authentication is done centrally in connected mode
- Backup AAA servers are configured at FlexConnect Group level
- When WLC/WAN fails, AP goes in Standalone mode
- In Standalone mode, the AP can be configured to authenticate new clients with backup RADIUS defined locally at the AP
- Upon WAN/WLC failure:
  - Existing connected clients stay connected
  - New clients are authenticated to the locally defined AAA

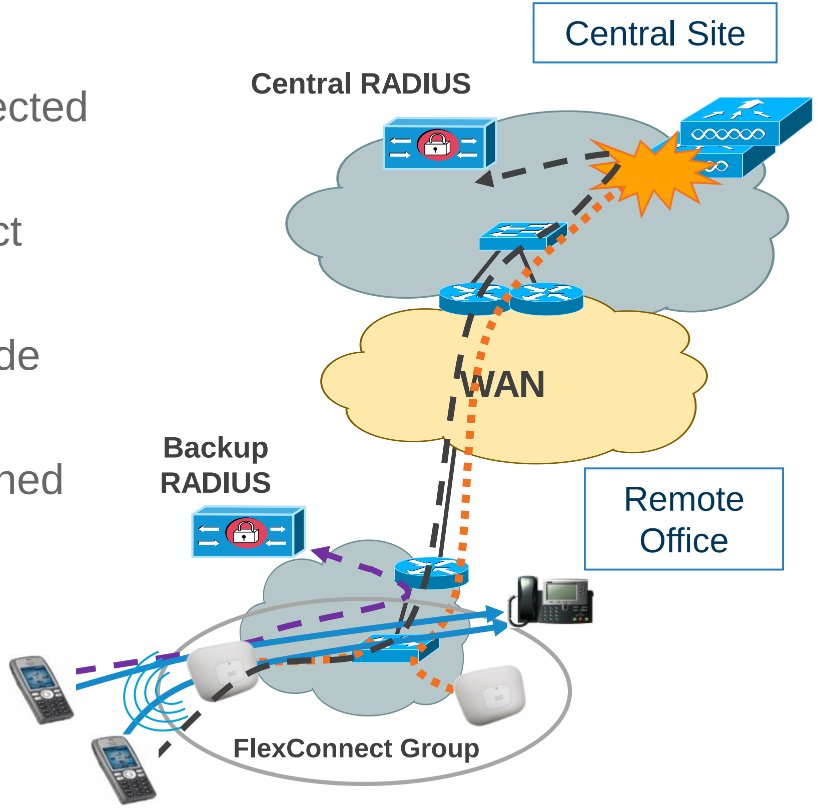

– – – – – Central authentication traffic

– – – – – Local authentication traffic

FlexConnect AAA Survivability

AAA Server Backup Configuration

For Your Reference

- Define primary and secondary local backup RADIUS server under FlexConnect Group configuration

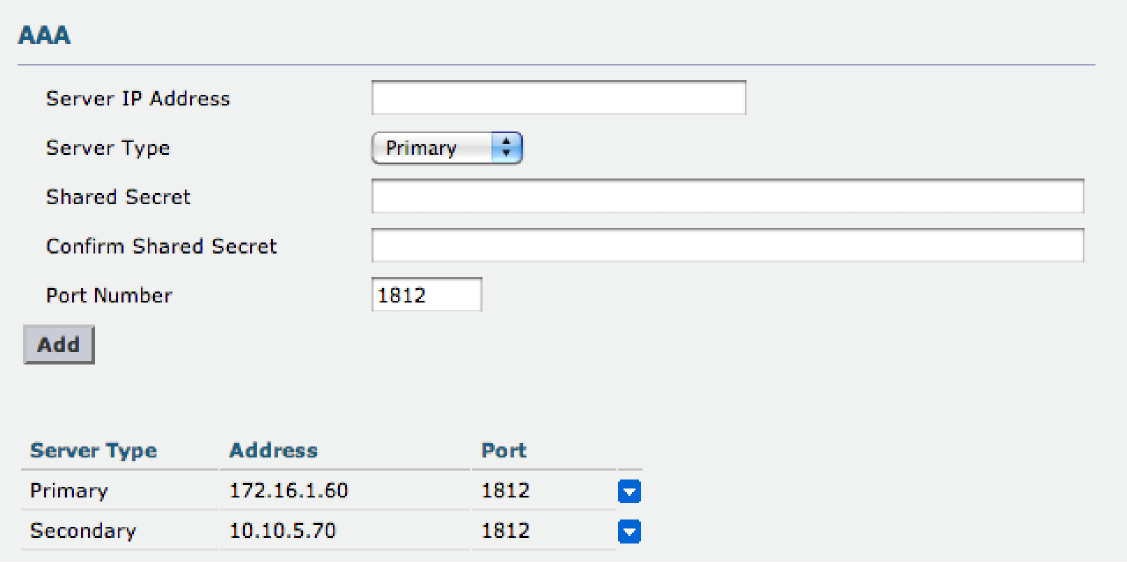

AAA

Server IP Address	
Server Type	Primary
Shared Secret	
Confirm Shared Secret	
Port Number	1812

Server Type	Address	Port
Primary	172.16.1.60	1812
Secondary	10.10.5.70	1812

FlexConnect AAA Survivability

FlexConnect Local Auth

- By default FlexConnect AP authenticates clients through central controller when in Connected mode
- This feature allows AP to act as an Authenticator even in Connected mode
- AAA servers are defined at the FlexGroup level
- Useful HA scenarios:
  - Independent branch: AAA is local at the branch, no AAA traffic goes through WAN
  - WLC goes down but WAN is up. Local users are authenticated from AP to Central site AAA

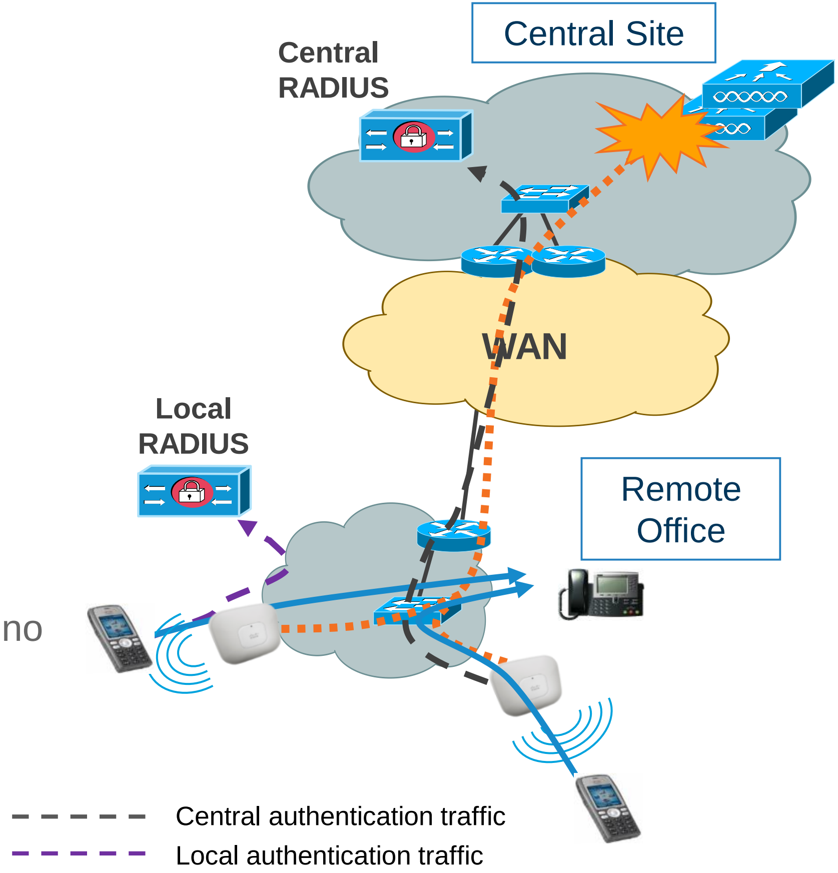

- - - - - Central authentication traffic

- - - - - Local authentication traffic

FlexConnect AAA Survivability

FlexConnect Local Auth: configuration

For Your Reference

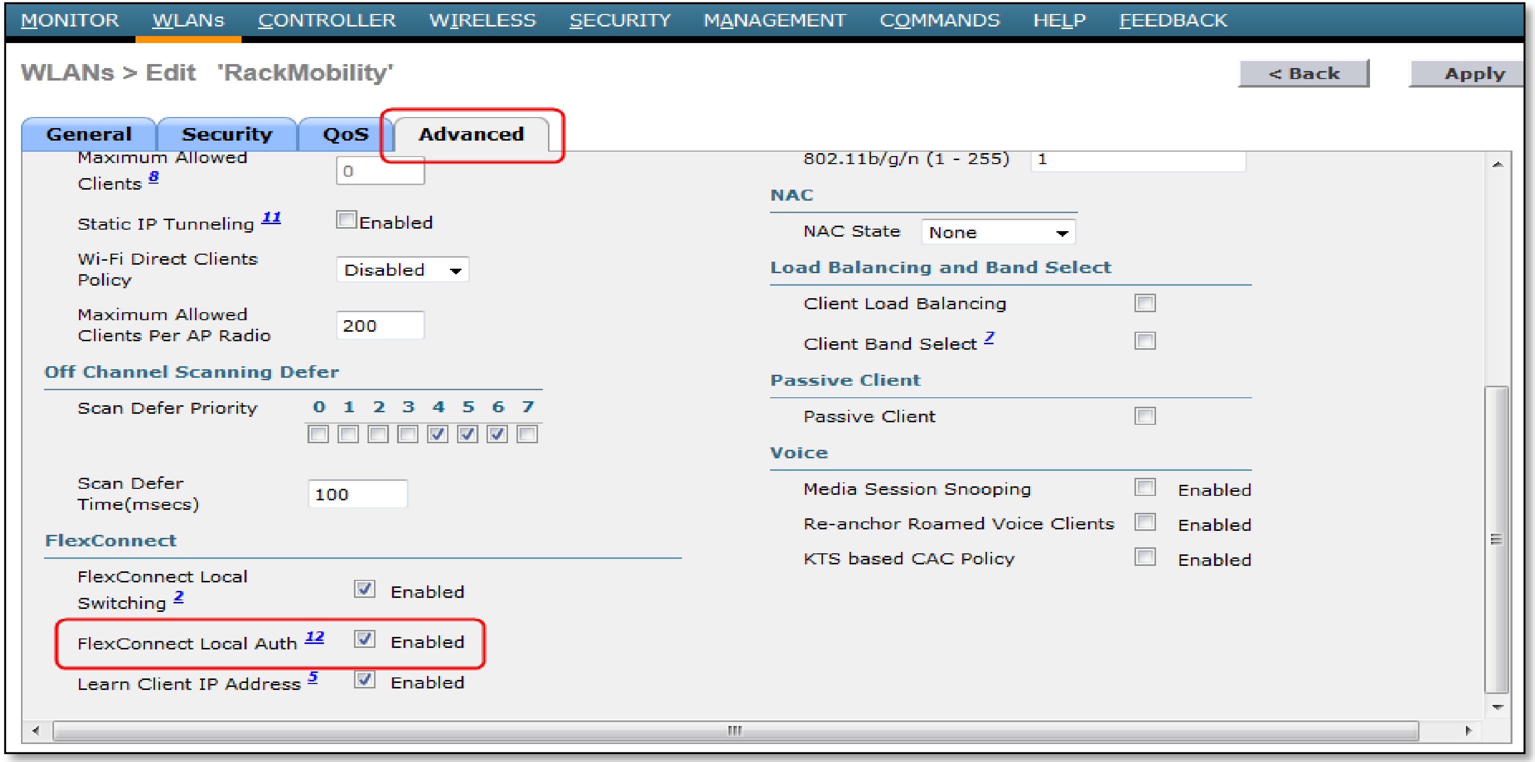

WLANs > Edit 'RackMobility'

Maximum Allowed Clients [8]: 0

Static IP Tunneling [11]: □ Enabled

Wi-Fi Direct Clients Policy: Disabled

Maximum Allowed Clients Per AP Radio: 200

Off Channel Scanning Defer

Scan Defer Priority: 0 □ 1 □ 2 □ 3 □ 4 ☑ 5 ☑ 6 ☑ 7 □

Scan Defer Time(msecs): 100

FlexConnect

FlexConnect Local Switching [2]: ☑ Enabled

FlexConnect Local Auth [12]: ☑ Enabled

Learn Client IP Address [5]: ☑ Enabled

802.11b/g/n (1 - 255): 1

NAC

NAC State: None

Load Balancing and Band Select

Client Load Balancing: □

Client Band Select [7]: □

Passive Client

Passive Client: □

Voice

Media Session Snooping: □ Enabled

Re-anchor Roamed Voice Clients: □ Enabled

KTS based CAC Policy: □ Enabled

FlexConnect AAA Survivability

AAA Server on AP

- By default authentication is done centrally in connected mode
- When WLC/WAN fails AP goes in Standalone mode
- In Standalone, the AP can act as a AAA server
- EAP-FAST, LEAP, PEAP*, EAP-TLS* and a max of 100 clients supported
- Upon WAN/WLC failure:
  - Existing connected clients stay connected
  - New clients are authenticated to the locally defined AAA

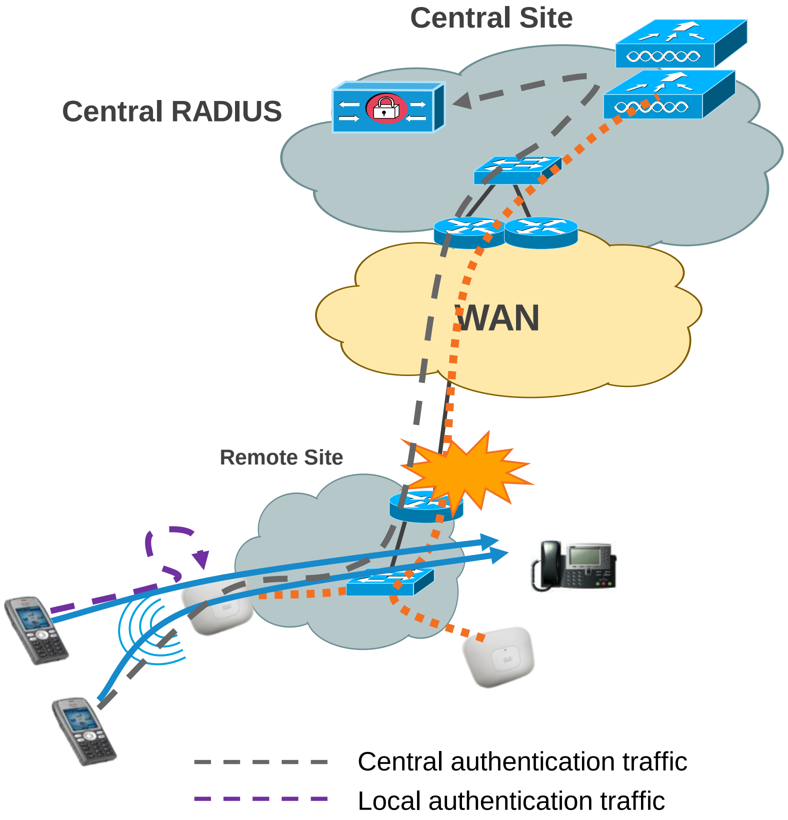
Central Site
Central RADIUS
WAN
Remote Site
- - - - - Central authentication traffic
- - - - - Local authentication traffic

* 7.5 Code and above

FlexConnect

AAA server on AP - Configuration

For Your Reference

- Check “Enable AP Local Auth” under the FlexConnect Group “General” tab
- Under the “Local Authentication” tab:
  - Define EAP parameters (LEAP, EAP-FAST, PEAP, EAP-TLS)
  - Define users (max 100) and passwords

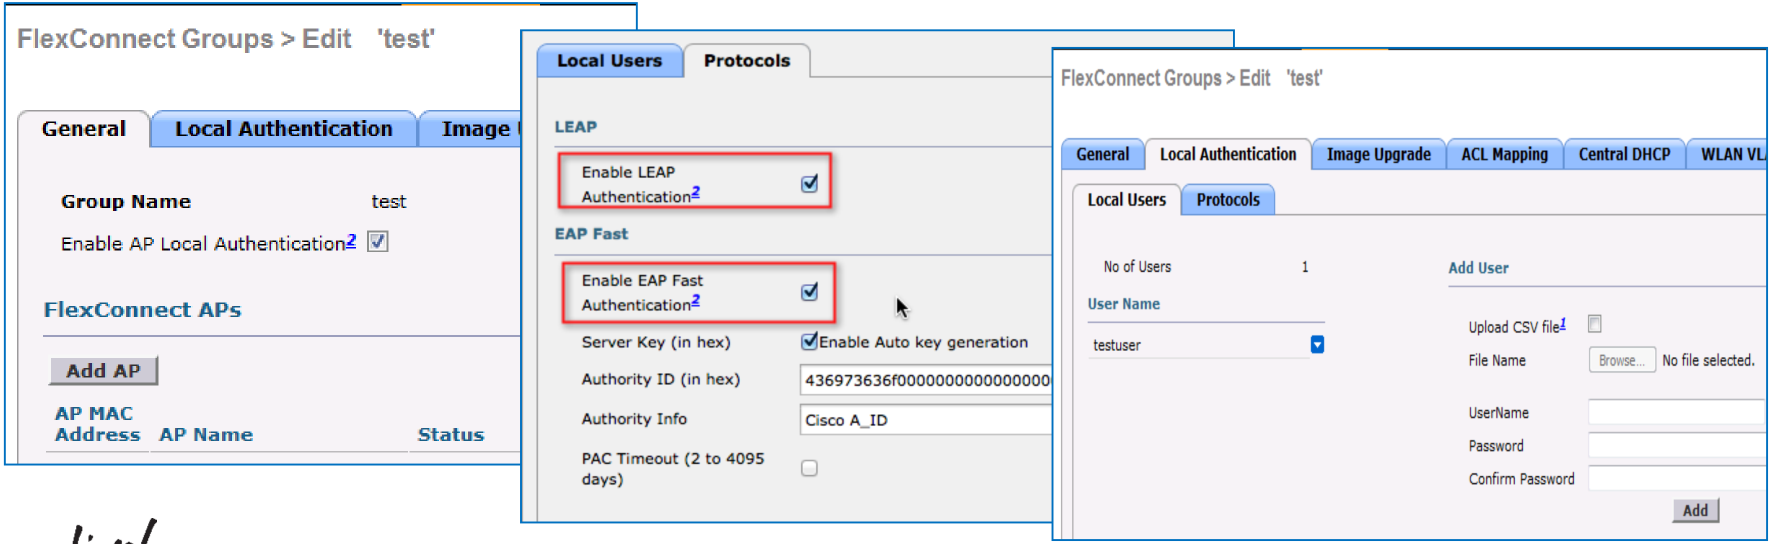

FlexConnect Groups > Edit 'test'

Group Name test

Enable AP Local Authentication[2] ☑

FlexConnect APs

AP MAC Address	AP Name	Status

LEAP

Enable LEAP Authentication[2] ☑

EAP Fast

Enable EAP Fast Authentication[2] ☑

Server Key (in hex) ☑ Enable Auto key generation

Authority ID (in hex) 436973636f0000000000000000

Authority Info Cisco A_ID

PAC Timeout (2 to 4095 days) □

FlexConnect Groups > Edit 'test'

No of Users 1

User Name

testuser

Add User

Upload CSV file[1] □

File Name Browse... No file selected.

UserName

Password

Confirm Password

Management and Mobility Services HA

Prime and MSE HA

	Requirements	Benefits
Prime HA	▪ Active / Standby (1:1) mode ▪ Same software & hardware ▪ Minimum failover time is 15 s ▪ PI 2.2 supports Virtual IP (VIP) ▪ HA SKU from PI 2.0 and later	▪ No database loss upon failover ▪ Failover Automatic or Manual ▪ Failback is always manual ▪ No AP licenses on Secondary ▪ Supported across WAN
MSE HA	▪ Active / Standby (1:1) mode ▪ Same software and hardware ▪ Same subnet only (no WAN) ▪ Same software version ▪ Release 8.0 recommended	▪ HA for all Services supported ▪ Failover times < 1 min ▪ No HA licenses needed ▪ Services licenses on Primary ▪ Failover Automatic or Manual

Prime Infrastructure HA

How it works

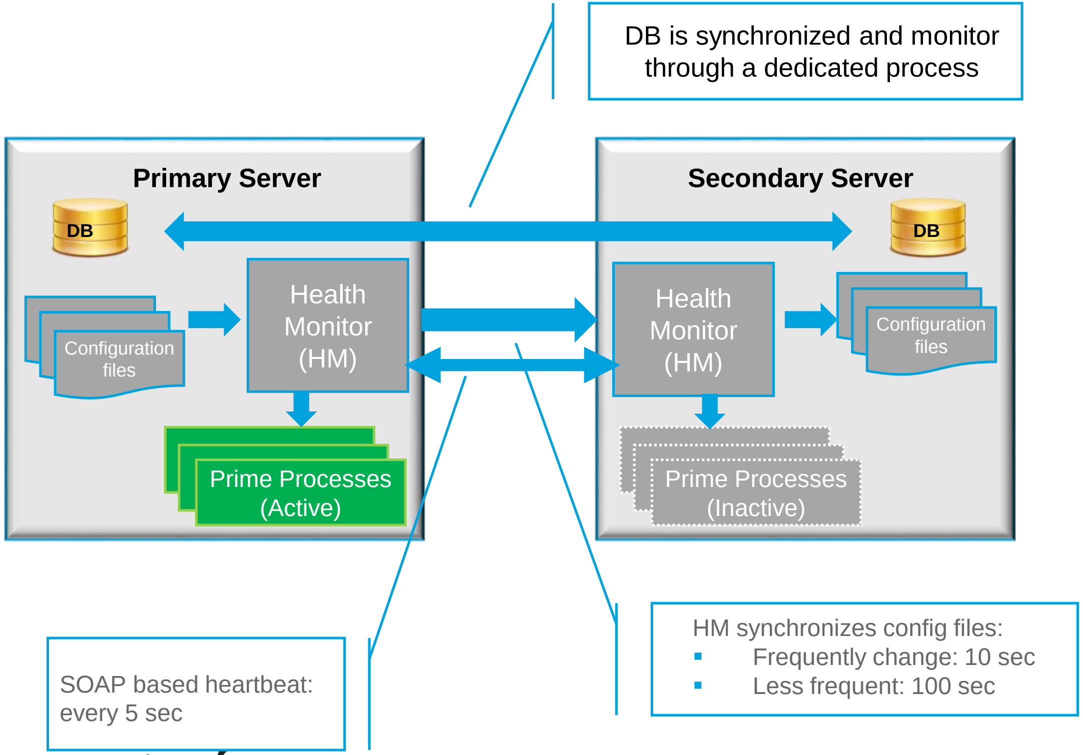

DB is synchronized and monitor through a dedicated process

SOAP based heartbeat: every 5 sec

HM synchronizes config files:
- Frequently change: 10 sec
- Less frequent: 100 sec

- The Health Monitor (HM) is the primary component for HA operation of the system:
  - Synchronizes configuration related to HA
  - Synchronizes the database
  - Exchanges heartbeat messages
  - Checks the available disk space on both servers
  - Triggers failover
  - Connect to https://<IP>:8082 of Primary/Secondary
- All configuration is done on the Primary
  - Secondary needs only authentication key at setup
  - Pi 2.2 introduces Virtual IP for same subnet deployments to simply configuration on monitored devices
  - Manual Failover is recommended
- Prime HA is supported in 3 scenarios:
  - Same LAN: Virtual IP can be used
  - Campus: usually different subnets
  - Remote: across WAN

MSE HA

How it works

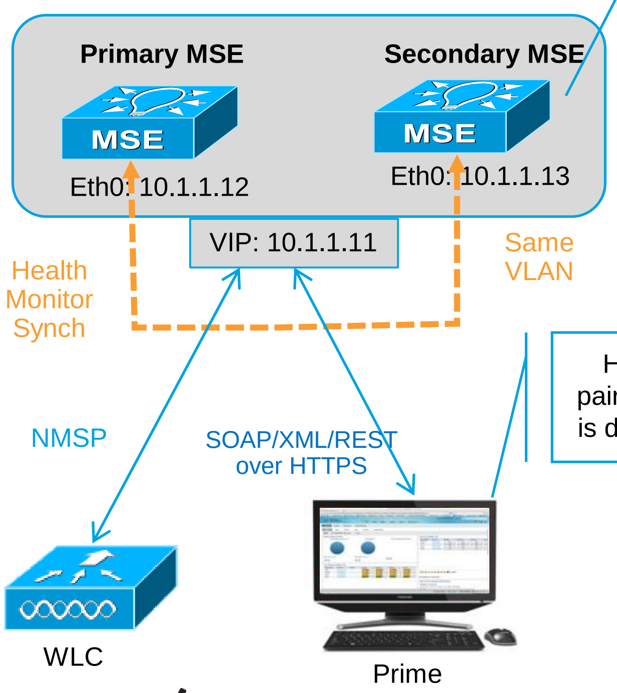

Role configuration is done through the setup MSE script

HA Configuration, pairing and Monitoring is done through Prime

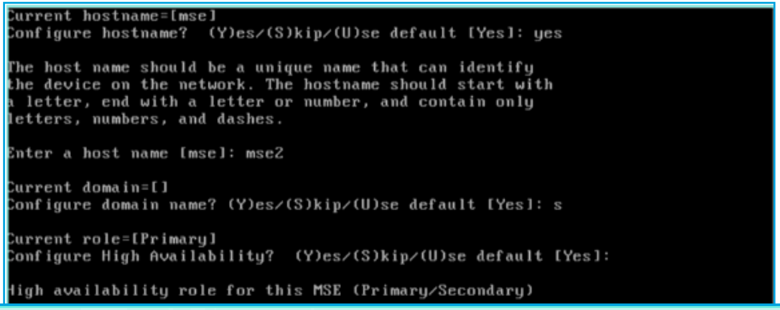

```
Current hostname=[mse]
Configure hostname?  (Y)es/(S)kip/(U)se default [Yes]: yes

The host name should be a unique name that can identify
the device on the network. The hostname should start with
a letter, end with a letter or number, and contain only
letters, numbers, and dashes.

Enter a host name [mse]: mse2

Current domain=[]
Configure domain name? (Y)es/(S)kip/(U)se default [Yes]: s

Current role=[Primary]
Configure High Availability?  (Y)es/(S)kip/(U)se default [Yes]:

High availability role for this MSE (Primary/Secondary)
```

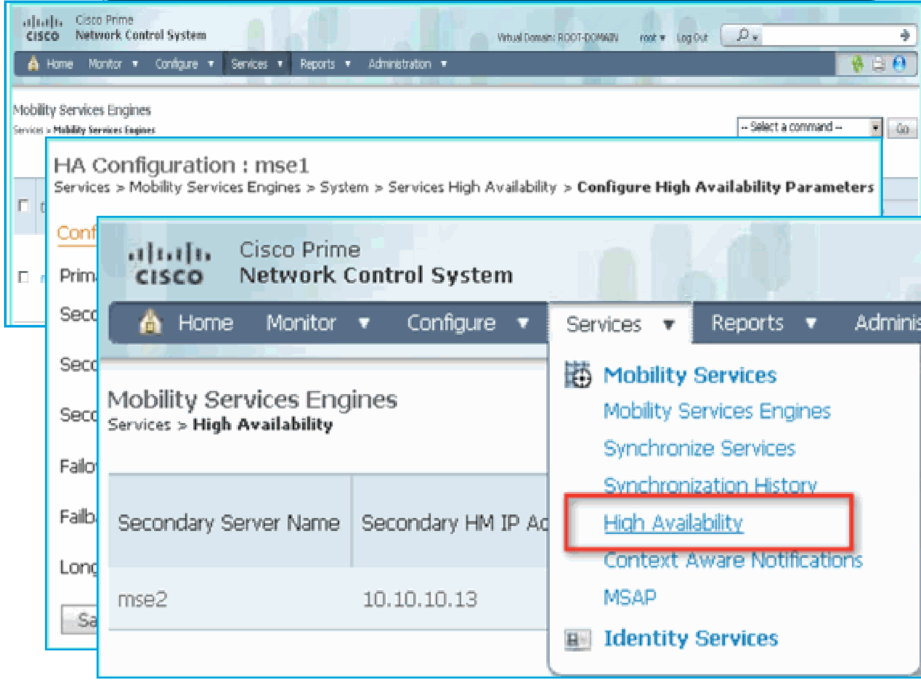

Mobility Services Engines

Services > Mobility Services Engines

HA Configuration : mse1

Services > Mobility Services Engines > System > Services High Availability > **Configure High Availability Parameters**

Mobility Services Engines

Services > **High Availability**

Secondary Server Name	Secondary HM IP Ac
mse2	10.10.10.13

HA Design and Deployment Practices

HA Design and Deployment Practices

Connecting an AP to the wired network

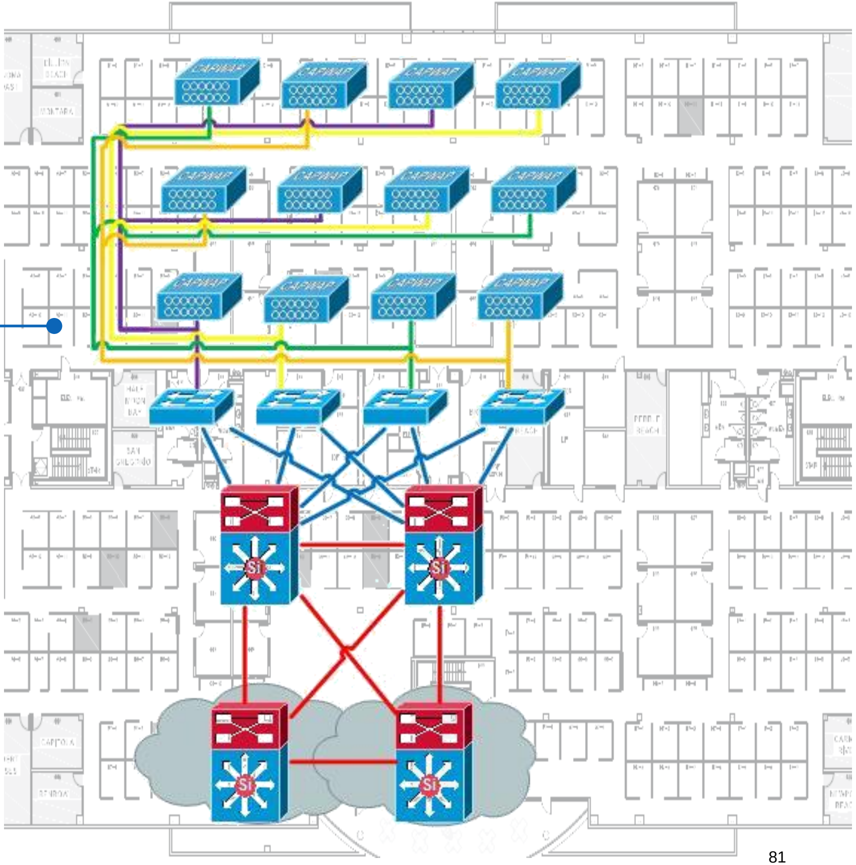

Recommendations:

- Create redundancy throughout the access layer by homing APs to different switches
- If the AP is in Local mode, configure the port as access with SPT PortFast, BPDU guard, etc.
- If the AP is in Flex mode and Local Switching, configure the port as trunk and allow only the VLANs you need

HA Design and Deployment Practices

Connecting a Controller to the wired network: options

1) To a single Modular Switch or StackWise

- Use Trunk EtherChannel(EC)/LAG
- Trunk only the required VLANs to the Controller
- 2/4/8 ports in a bundle to optimize load sharing
- Spread ports across Line Cards/Stack members

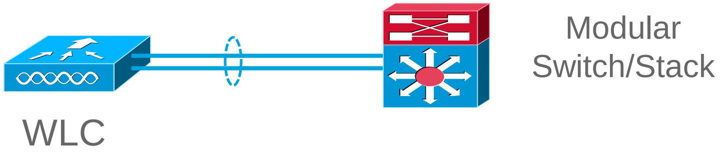

2) To a VSS pair

- Same as Option 1
- Spread ports across VSS members

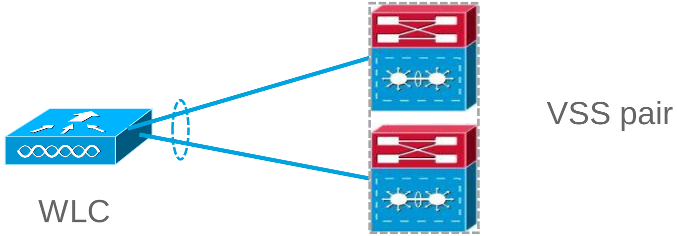

Connecting a Controller to the wired network

Single AireOS Controllers (2504/5508/7500/8500/Wism2)

Option 1: to single Modular Switch or StackWise

- Identical configuration on WLC and switch side (EC mode, trunk mode, allowed VLANs, native VLAN, etc.)
- EC mode: only mode "ON" supported; no LACP, PAgP
- EC load-balancing: no restriction for 5508/2500/7500/8500
  - Recommended to include L3 and L4 port for better hash results
- EC load-balancing for WISM2:
  - Need to set the EC load balancing method on the switch to "src-dest-IP". Use CLI "port-channel load-balance src_dest_ip"
- Note: no STP supported on AireOS Controllers. Do not disable it on switch side. Use "switchport portfast trunk"

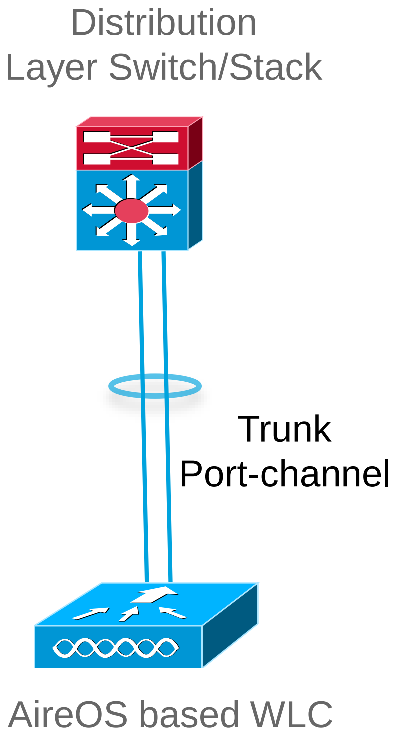

Connecting a Controller to the wired network

Single AireOS Controllers (2504/5508/7500/8500/Wism2)

Option 1: to single Modular Switch or StackWise

- Identical config
  allowed VLANs
- EC mode: only
- EC load-balanc
  - Recomme
  - On the swi
- EC load-balanc
  - Need to se
    “port-chan
  - For Cataly
    vlan” (command supported in 12.2(33)SXH6 and 12.2(33)SXI3 and above)
- Note: no STP supported on AireOS Controllers. Do not disable it on switch side. Use “switchport portfast trunk”

```
port-channel load-balance src-dst-mixed-ip-port
!
interface GigabitEthernet1/0/1
 description to_WLC-1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,11,20,30,40
 switchport mode trunk
 channel-group 1 mode on
 switchport portfast trunk
```

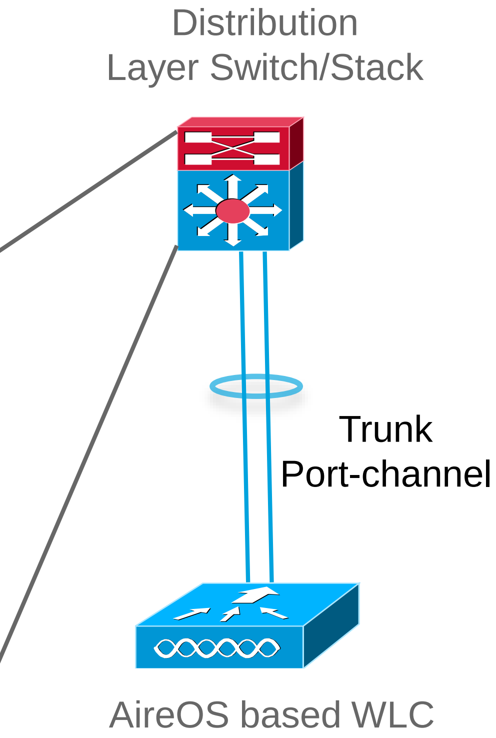

Cisco live!

Connecting a Controller to the wired network

Option 2: to a VSS pair

- Single LAG to the VSS pair
- Spread ports across VSS pair
- In case of failure of Primary switch traffic continues to flow through Secondary switch in the VSS pair
- Same recommendations given for Option 1 also apply

Recommended Network Design

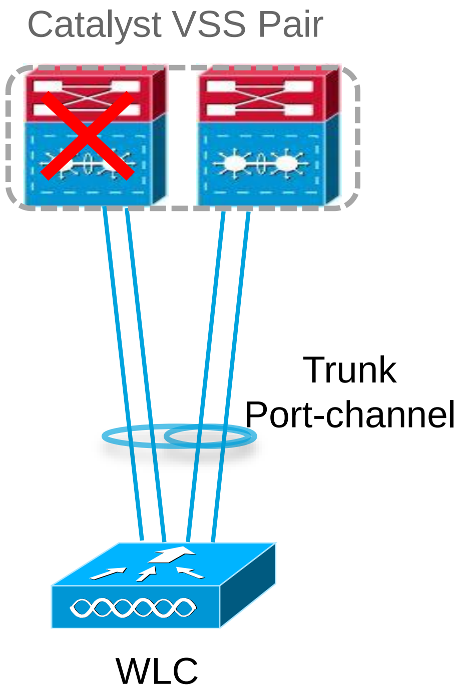

Design & Deployment Practice

Connecting a Controller HA pair

HA Design and Deployment Practices

Connecting AireOS HA Pair to the wired network

Option 1: to single Modular Switch or StackWise

- The HA pair of AireOS WLCs should be considered as separated WLCs with the same exact configuration
- Ports on both WLCs are UP but only the ones on the Active WLC are forwarding data traffic
- On WLC side: use same physical ports are connected to the network, for ex.: port 1-4 on WLC1 and port 1-4 on WLC2
- On switch side the configuration has to be the same. If using LAG, for example, two Port-channel should be used with the same configuration (same mode, same VLANs, same native, etc.)
- General recommendations for Option 1 AireOS WLC also apply

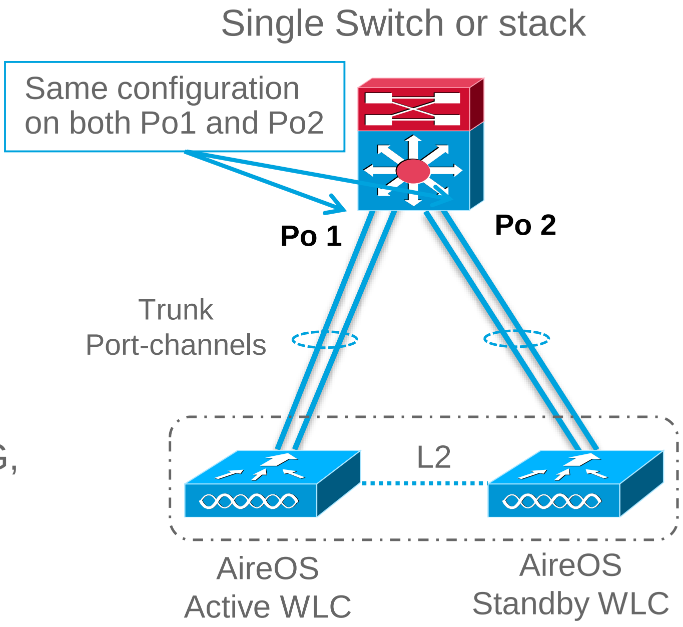

HA Design and Deployment Practices

Connecting AireOS HA Pair to the wired network

Recommended Network Design

Option 2: to VSS pair

- Use EC from each WLC to Distribution VSS
- Spread the links in each EC among the two physical switches: this will prevent a WLC switchover upon a failure of one of the VSS switch
- Same considerations for connecting to a single Distribution switch apply
- General recommendations for Option 1 AireOS WLC also apply

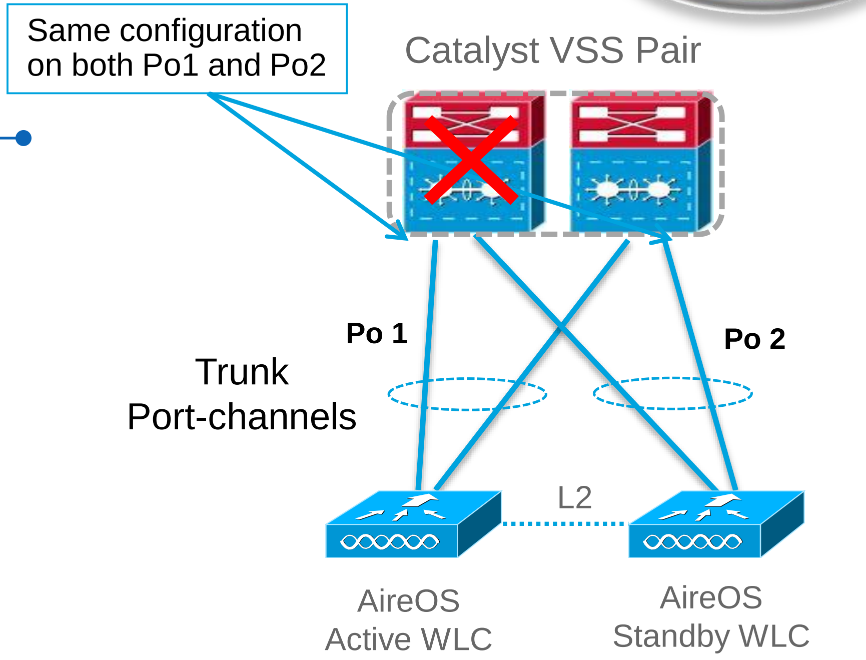

HA Design and Deployment Practices

Connecting AireOS HA Pair to the wired network

Option 3: to Pair of Distribution switches

- Use ECs to connect to Distribution switches
- Same exact configuration on both Dist. switches
- Use same physical ports on the WLCs
- Layer 2 between the distribution switches for the Wireless VLANs
- Use STP on the Distribution switches

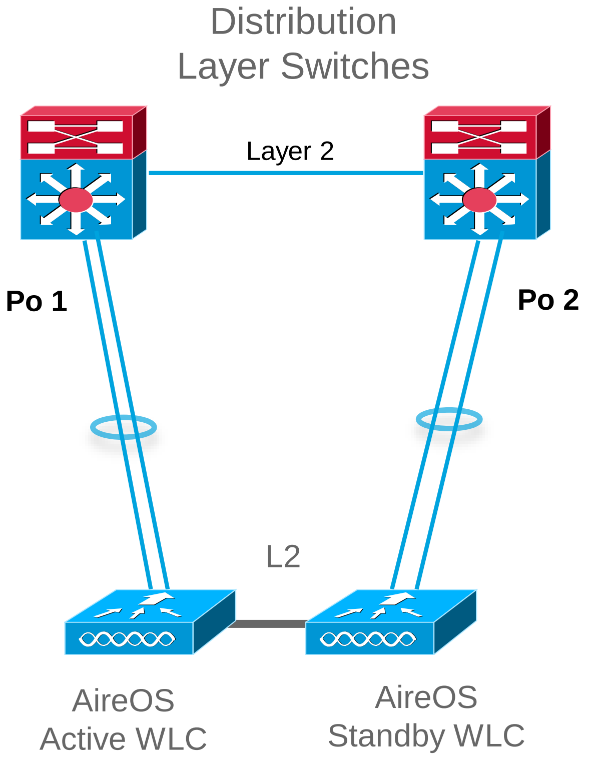

HA Design and Deployment Practices

Controller considerations

Recommended Network Design

- VSS is the recommended Design choice as it provides:
  - Redundancy at distribution layer
  - Efficient use of all links with Multi-Link EtherChannel
  - Fast convergence, no spanning tree
- How many WLC ports do I need to connect?
  - Multiple interfaces for redundancy
  - Consider the wireless over-subscription (80:1 is considered normal)
- Choose the right model of switch to connect to:
  - Some controllers have only 10GE interfaces (8510, 7510)
  - Consider TCAM scalability for the number of client MACs
    - Sup2T and Nexus 7000 supports 128k MAC addresses
    - 3850 supports up to 32k MAC addresses

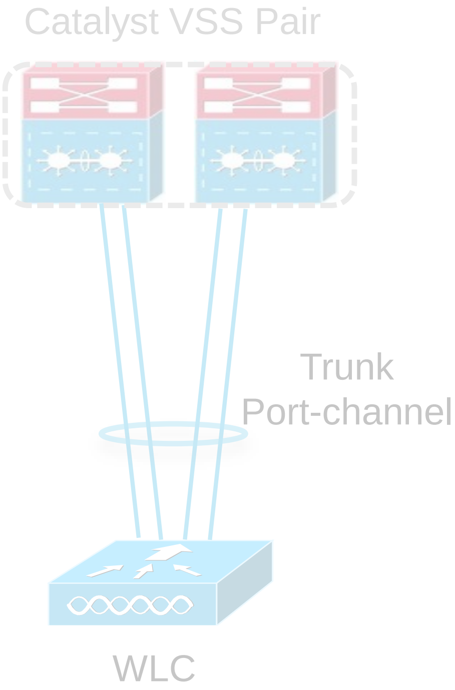

HA Design and Deployment Practices

Campus

HA Design and Deployment Practices

Campus

- What is the acceptable downtime for your business applications?
  - Are 30 sec to few minutes ok? Go with N+1 to have more deployment flexibility
  - No downtime? Go with AireOS Stateful Switchover
- SSO: what is the downtime to upgrade a HA pair and how to minimize it?

HA Design and Deployment Practices

Upgrading an SSO Pair - standard procedure

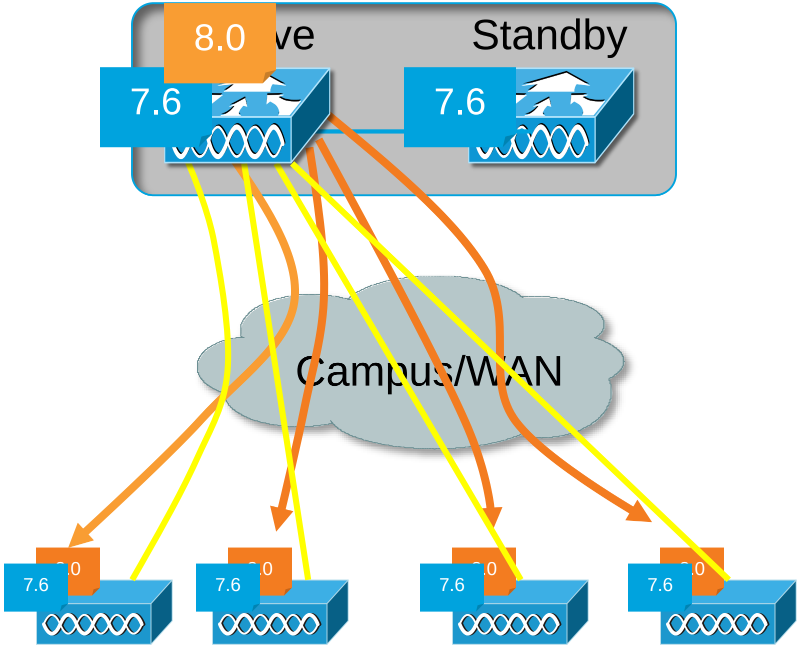

1. Download the new code on Active

2. Code transferred to Standby:

   Do NOT reboot at this time!

3. **Pre-download** software on APs

CAPWAP tunnel ▬

HA Design and Deployment Practices

Upgrading an SSO Pair - standard procedure

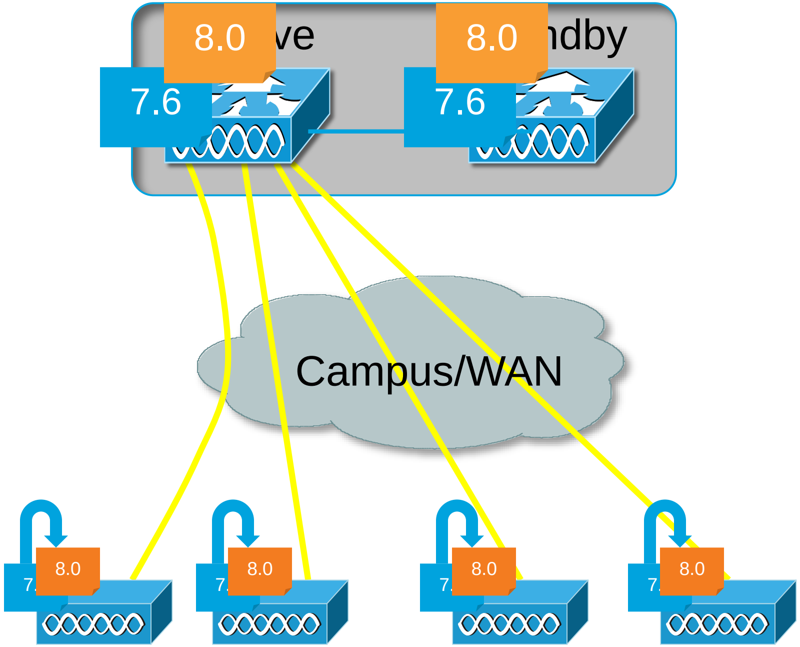

1. Download the new code on Active
2. Code transferred to Standby
3. Pre-download software on APs
4. Swap the images on APs
5. Reboot the HA pair
   - APs will reboot and join when Active is UP

Total Network Downtime:
Time for HA pair to reboot + the APs to join

5min:12sec with fully loaded 5508
(500 APs/7000 clients)

HA Design and Deployment Practices

Upgrading an SSO Pair – Efficient procedure

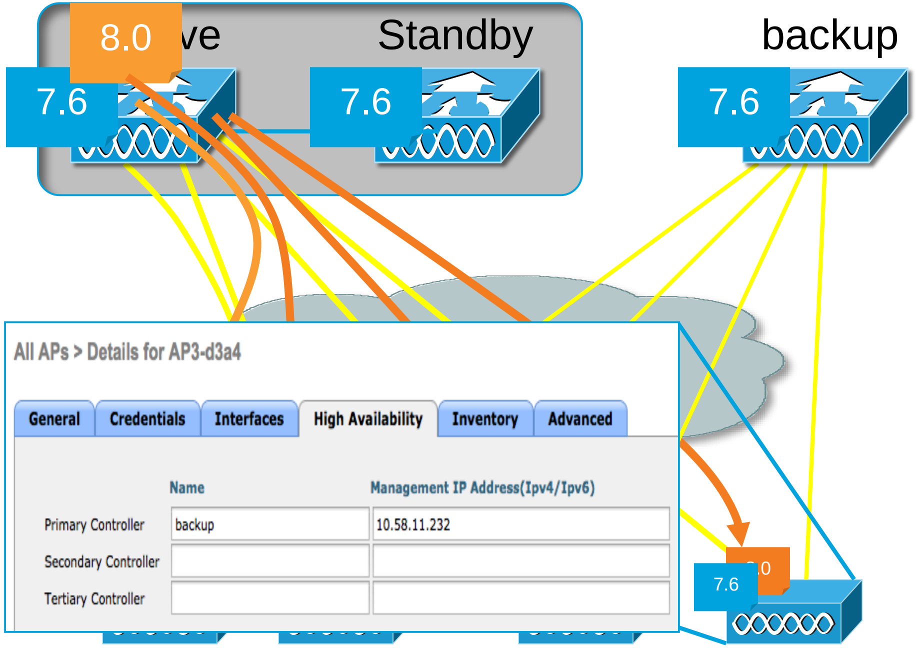

All APs > Details for AP3-d3a4

General | Credentials | Interfaces | **High Availability** | Inventory | Advanced

	Name	Management IP Address(Ipv4/Ipv6)
Primary Controller	backup	10.58.11.232
Secondary Controller		
Tertiary Controller		

1. Download the new code on Active

2. Code transferred to Standby

   Do NOT reboot at this time!

3. **Pre-download** software on APs

4. Configure APs to join the backup controller
   - This can be done per group of APs/Areas
   - This can be automated using Prime

5. The APs join the backup WLC (no reboot)
   - This takes less than 30sec
   - Downtime can be isolated per area

CAPWAP tunnel ▬

HA Design and Deployment Practices

Upgrading an SSO Pair – Efficient procedure

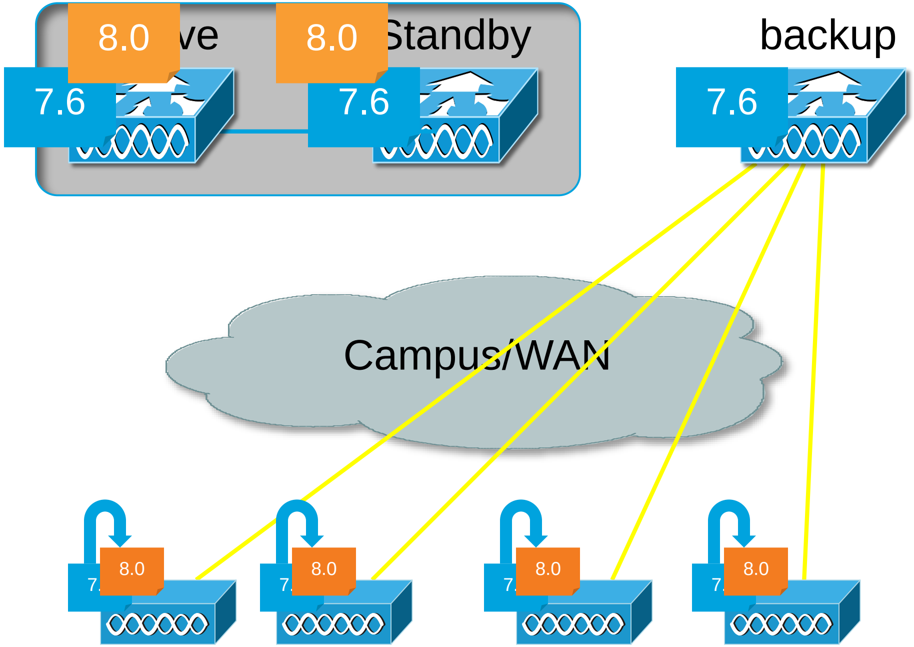

1. Download the new code on Active
2. Code transferred to Standby
3. **Pre-download** software on APs
4. Configure APs to join the backup controller
5. The APs join the backup WLC (no reboot)
6. Swap the images on Aps

Do this for all the APs in your network

HA Design and Deployment Practices

Upgrading an SSO Pair – Efficient procedure

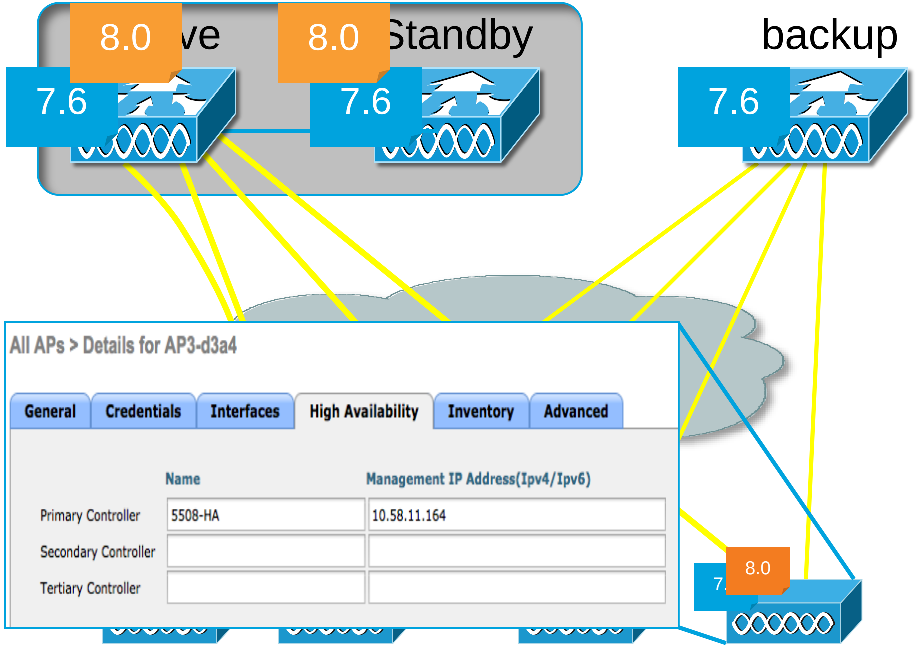

All APs > Details for AP3-d3a4

	Name	Management IP Address(Ipv4/Ipv6)
Primary Controller	5508-HA	10.58.11.164
Secondary Controller		
Tertiary Controller		

7. Reboot the HA pair
8. Configure the APs to join the HA pair
   - This can be done per group of APs/Areas
   - This can be automated via Prime
9. APs will join the Active WLC and reboot because of new code:

Network Downtime:
Time for the APs to move to Active, reboot and join back: **3min**

Main Advantage: downtime is per Area

HA Design and Deployment Practices

Campus

- What is the acceptable downtime for your business applications?
  - No downtime? Go with AireOS Stateful Switchover
  - Are 30 sec to few minutes ok? Go with N+1 to have more deployment flexibility
- SSO: what is the downtime to upgrade a HA pair and how to minimize it?
- Would like to deploy SSO across a L3 network, what are the implications?

HA Design and Deployment Practices

SSO across L3 domain, what are the implications?

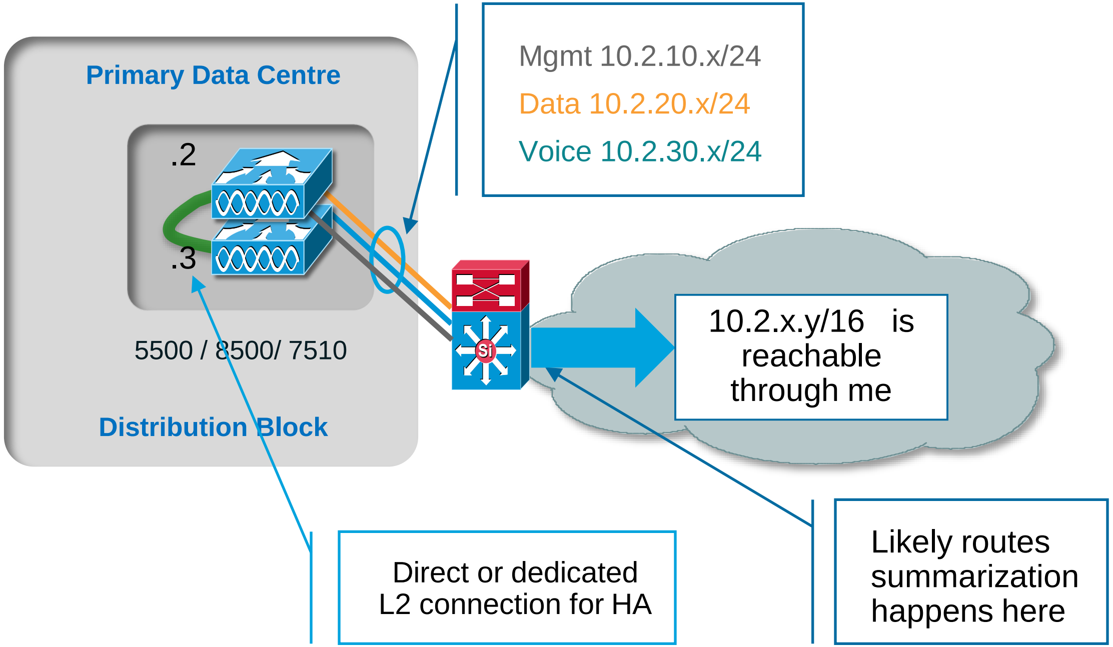

Mgmt 10.2.10.x/24

Data 10.2.20.x/24

Voice 10.2.30.x/24

10.2.x.y/16 is reachable through me

Direct or dedicated L2 connection for HA

Likely routes summarization happens here

HA Design and Deployment Practices

SSO across L3 domain, what are the implications?

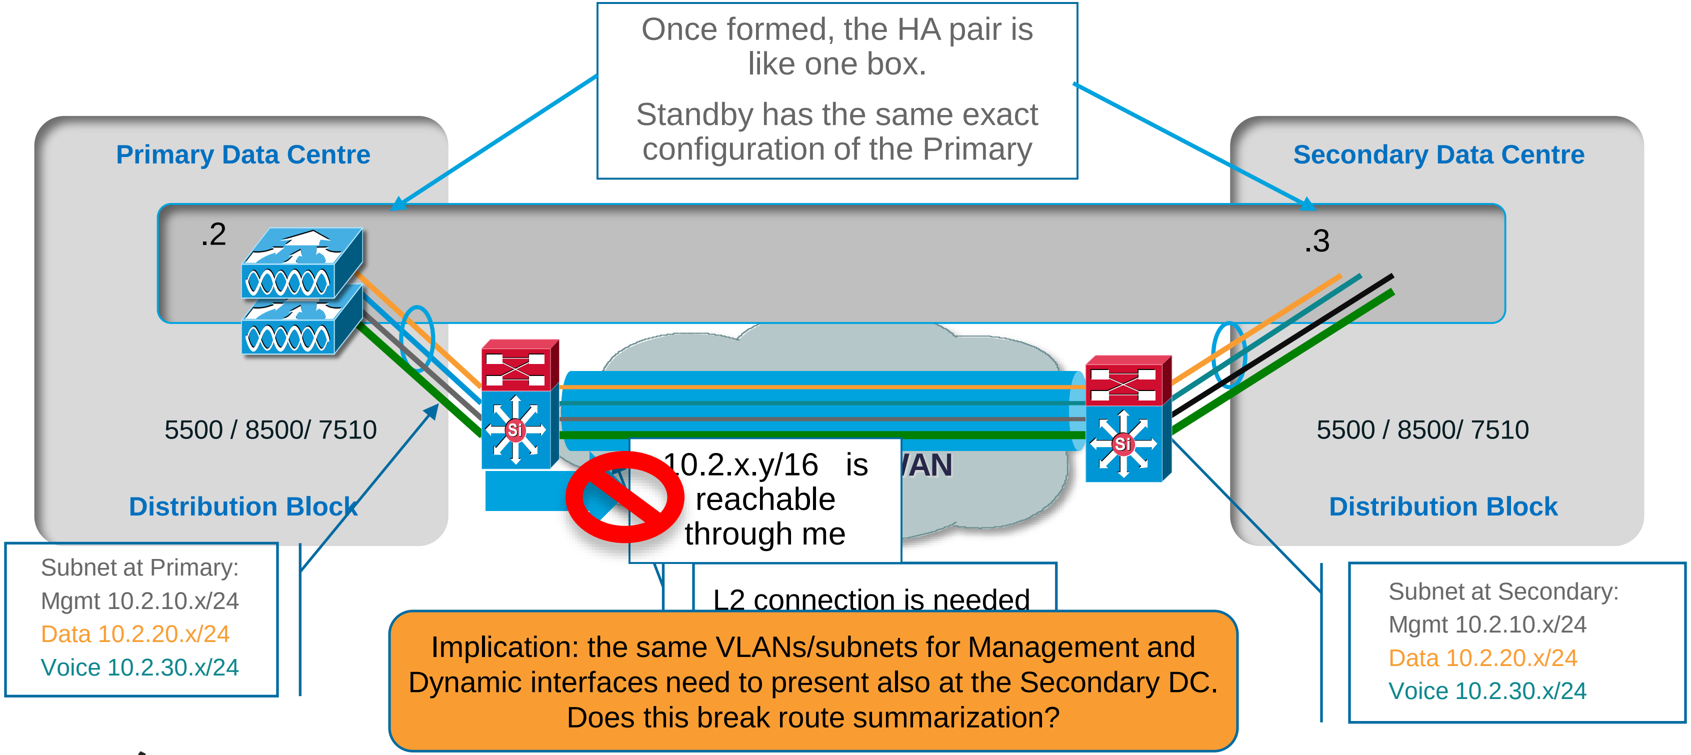

Once formed, the HA pair is like one box.

Standby has the same exact configuration of the Primary

10.2.x.y/16 is reachable through me

L2 connection is needed

Subnet at Primary:
Mgmt 10.2.10.x/24
Data 10.2.20.x/24
Voice 10.2.30.x/24

Subnet at Secondary:
Mgmt 10.2.10.x/24
Data 10.2.20.x/24
Voice 10.2.30.x/24

Implication: the same VLANs/subnets for Management and Dynamic interfaces need to present also at the Secondary DC. Does this break route summarization?

HA Design and Deployment Practices

Campus

- What is the acceptable downtime for your business applications?
  - No downtime? Go with AireOS Stateful Switchover
  - Are 30 sec to few minutes ok? Go with N+1 to have more deployment flexibility
- What is the downtime to upgrade a HA pair and how to minimize it?
- Would like to deploy SSO across a L3 network, what are the implications?
- What is the recommended HA deployment in a multi site Campus?
  - Use Hybrid (SSO and N+1) HA deployment
    - Use SSO in the main site (Primary WLC)
    - Use Secondary/Tertiary in redundancy sites with HA-SKU
  - For max resiliency use SSO in all sites

HA Design & Deployment Practices

HA Design in a multi site Campus

Combine SSO with N+1

- SSO pair can act as the Primary Controller and be deployed with single Secondary and Tertiary WLC
- Network downtime:
  - No network downtime for single controller failure in the Primary DC
  - On failure of both Active and Standby WLC, APs will fall back to secondary and further to configured tertiary controller
- Recommendations:
  - Make sure that AP Fallback is enabled
  - Use AP Failover priority in case of oversubscription of the backup WLC
  - Useful to reduce downtime for SSO pair software upgrade

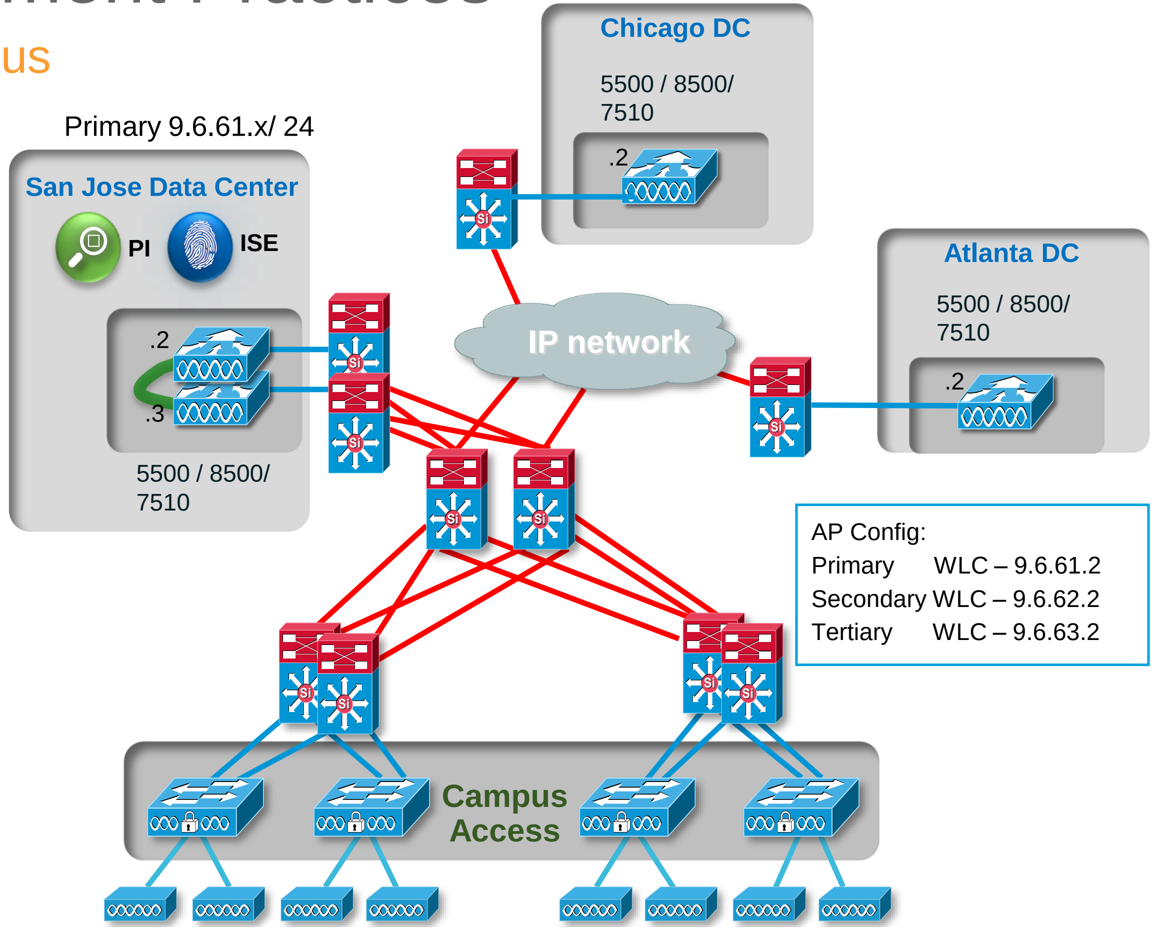

AP Config:

Primary	WLC – 9.6.61.2
Secondary	WLC – 9.6.62.2
Tertiary	WLC – 9.6.63.2

HA Design & Deployment Practices

HA Design in a multi site Campus

SSO everywhere!

- Each site can be its own separated SSO architecture
- Full site redundancy by assigning primary, secondary, tertiary to the APs.
- Max level of High Availability: no network downtime upon controller failure within any site

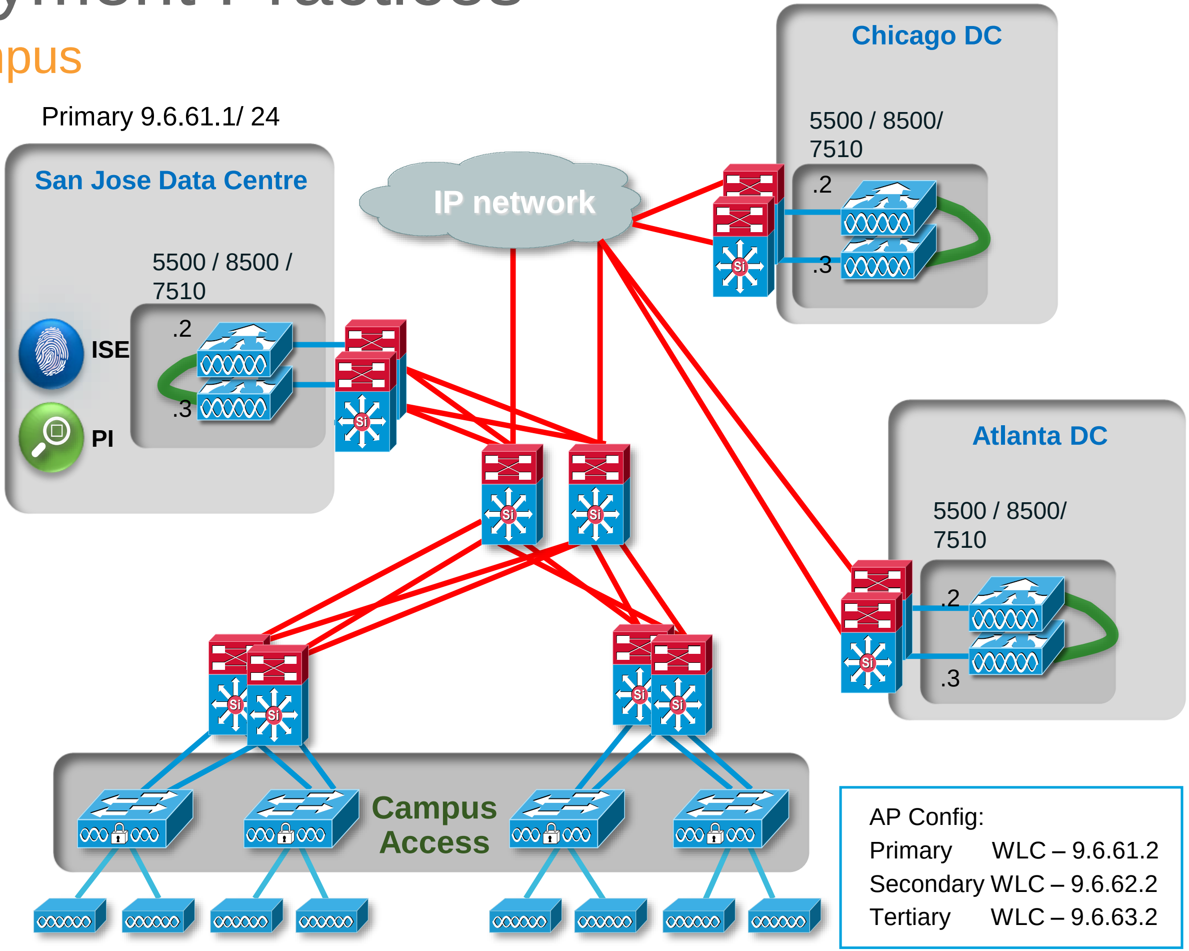

AP Config:

Primary	WLC – 9.6.61.2
Secondary	WLC – 9.6.62.2
Tertiary	WLC – 9.6.63.2

HA Design and Deployment Practices

Campus Guest Access

- How can I make the Guest Access highly available?
- Customer design requirements:
  - Redundancy at the Anchor level controller
  - Two DC sites, A and B, with direct access to Internet
    - Guest traffic needs to go out from site A (Primary)
    - If there is a failure at site A, traffic should go out at site B (Secondary)

HA Design & Deployment Practices

Guest Access HA – Round Robin Option

- Add a second Anchor Controller in the DMZ
- A Foreign controller load balances guest traffic across the Anchor controllers with same priority configured on the WLAN.

Advantage:

- Add a basic type of redundancy
- Guest session capacity is the sum of the capacity of each controller used as Anchor
  - ex. 14k users for 5508

Disadvantage:

- Doesn’t fully meet the requirements of customer in terms of traffic handling
- No SSO and no deterministic redundancy

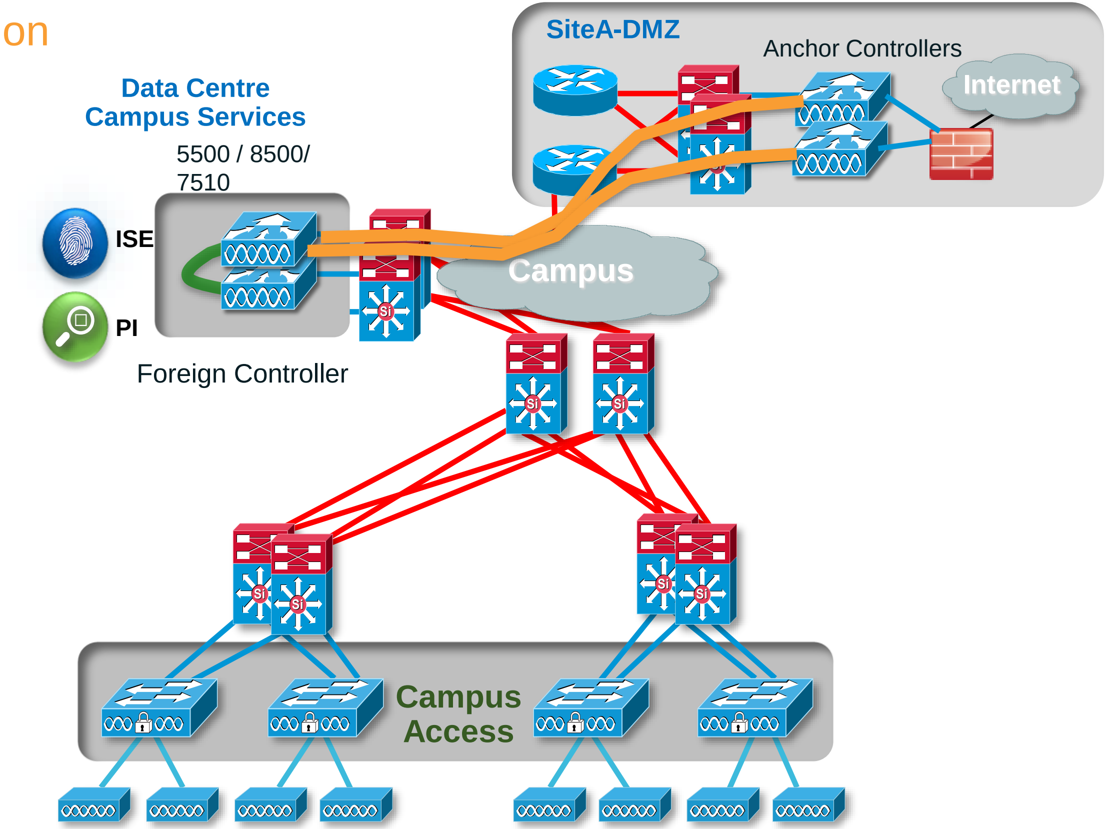

HA Design & Deployment Practices

Guest Access HA – SSO + Anchor Priority

- Create an SSO pair at the anchor
  - HA-SKU can be used as Anchor standby
  - Use same software (AireOS or IOS) on Foreign and Anchor pair
- Add an anchor at site B and use priority to define which is the Primary (AireOS 8.1)

Advantage:

- No guest client disconnection upon anchor WLC failover (AireOS 7.5 and above)
- Met customer requirements: traffic goes out from site A unless there is a failure

Disadvantage:

- Guest client sessions at site A limited to capability of one anchor controller
  - Example: 7k clients on 5508

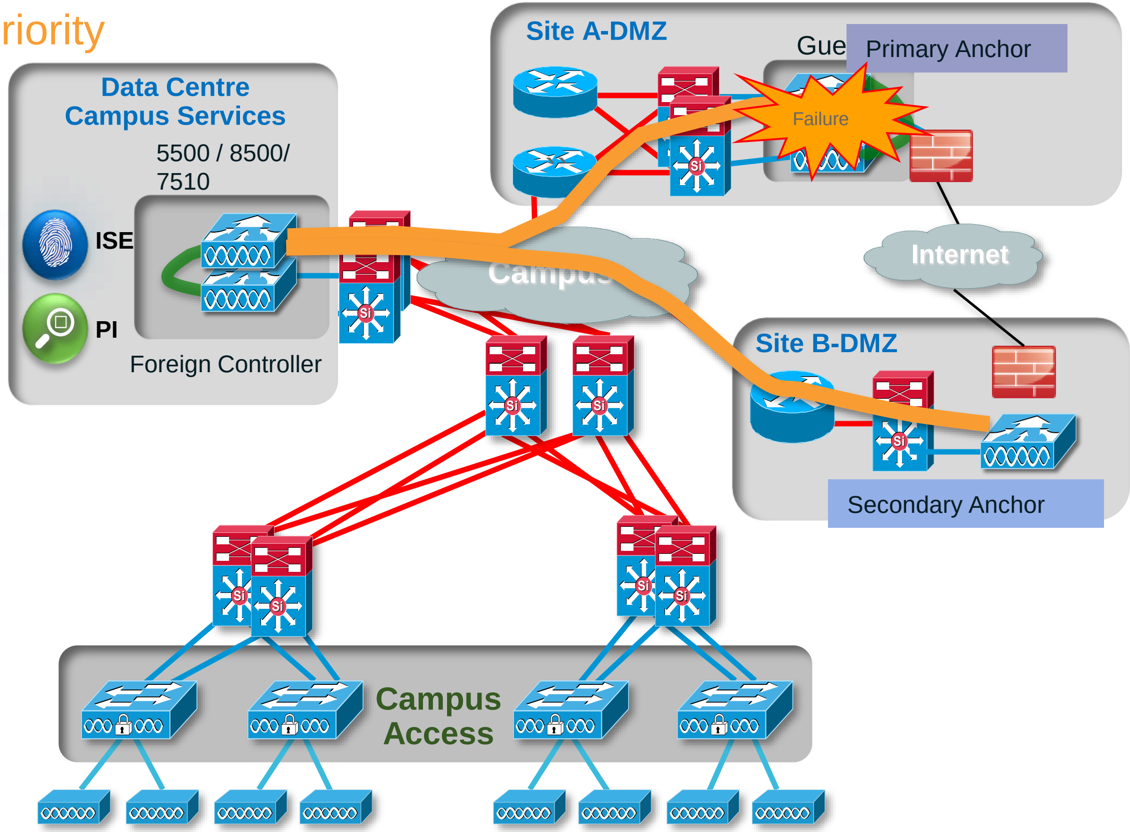

HA Design & Deployment Practices

Guest Access HA – SSO Option (before 8.1)

- SSO is supported on the anchor controllers
- HA-SKU can be used as Anchor
- Use same software (AireOS or IOS) on Foreign and Anchor pair

Advantage:

- Using client SSO (7.5 or above) no guest client disconnection upon Anchor WLC failover
- Geo Separated Anchor with SSO to determine Primary and Secondary Guest exit to internet

Disadvantage:

- Guest client sessions limited to capability of one Anchor controller
  - Example: 7k clients on 5508

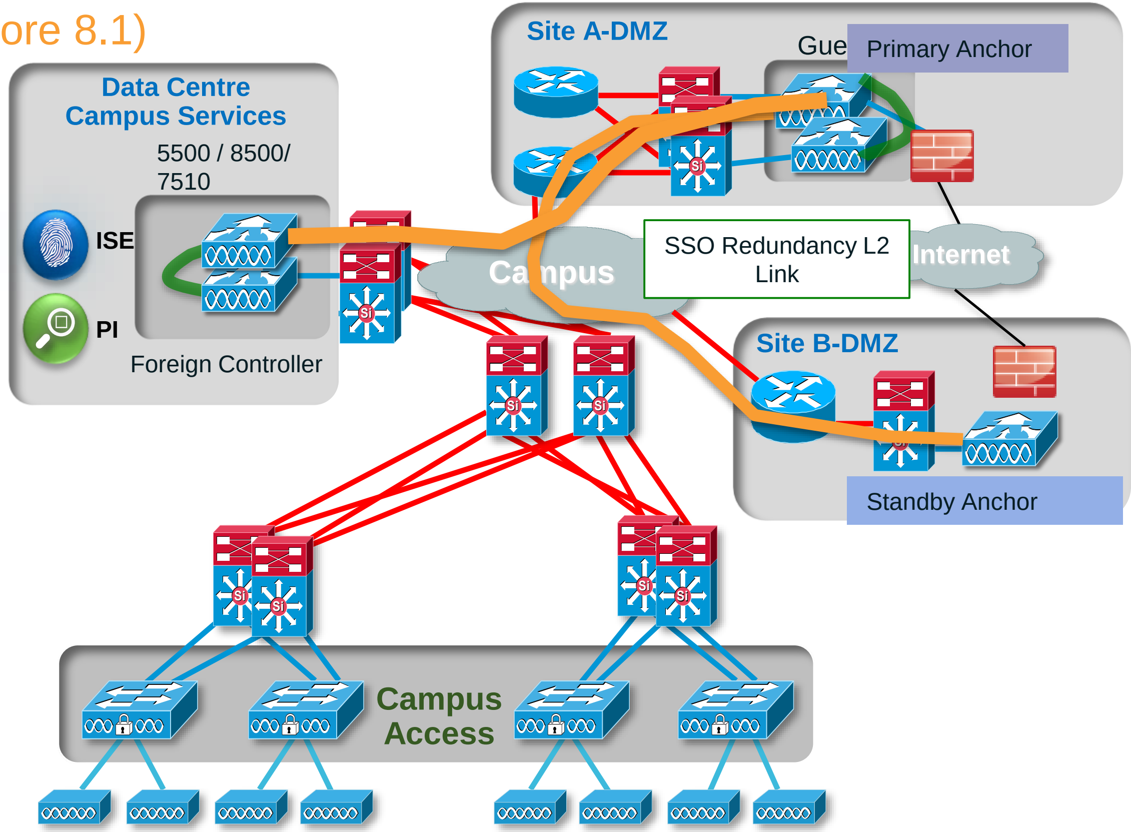

HA Design and Deployment Practice

Branch

HA Design and Deployment Practices

Branch: some key Design questions

- General considerations:

Local Controller

- Specific per branch configuration
- Independency from WAN quality
- Reduced configuration on switches
- Full feature support
- L3 roaming supported

VS.

FlexConnect

- Branch cookie cutter configuration
- Single pane of Mgmt. & Troubleshooting
- Reduced branch footprint
- Cheapest AP controller licenses
- Built-in resiliency
- Perfect fit for centralized IT Team

- HA considerations:
  - Is the branch independent from the Central site from an operation prospective?
    - What is the traffic flow of your application? Are the APP servers centrally located?
    - Is there a local Internet breakout?
    - How do you authenticate new users if WAN/Controller is down? Where is the AAA server located?
  - FlexConnect is inherently designed for HA and offers:
    - Data plane resiliency upon Central WLC failure or WAN outage
    - AAA survivability options

HA Design and Deployment Practices

Branch Redundancy: Centralized Controller & Flex (local switching)

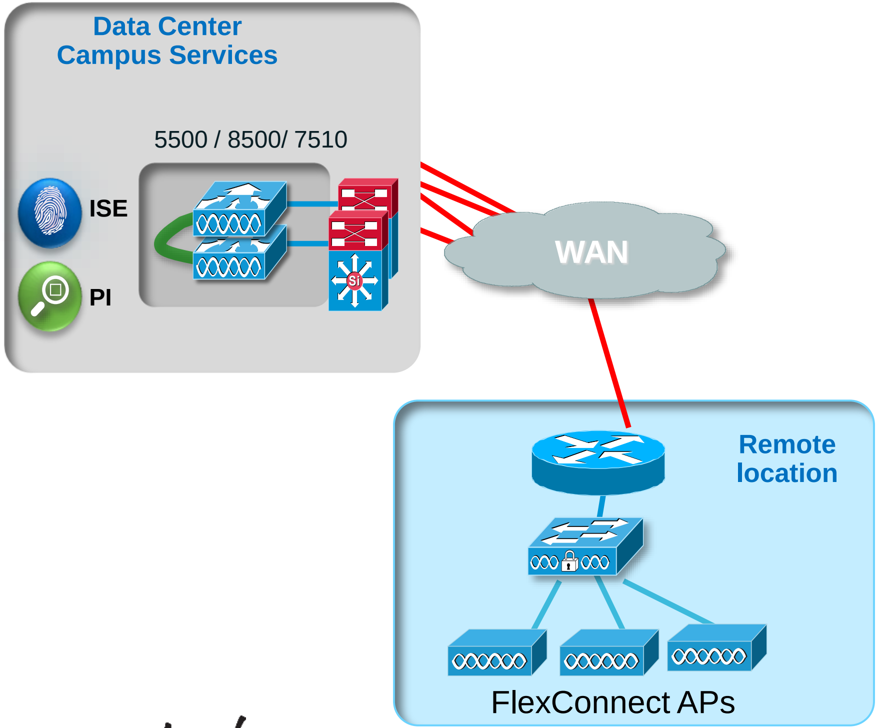

HA considerations:

- if WAN fails, Flex APs allow a level of redundancy:
  - Local Data path stays UP
  - Control plane features go down: RRM, CleanAir, WebAuth, etc.
- WLC SSO at central site provides Control plane survivability

Design considerations:

- WAN requirements:
  - General: 12kbps per AP, 300 ms RTT, 500B MTU
  - More info here: http://tiny.cc/FlexDG
- APs are in Flex Mode = less features and functionalities compare to Local Mode. Key features missing:
  - No L3 roaming, No Bonjour Gateway
- Flex Groups have AP count limit
  - 25 APs for 2500/5508, 100 APs for 7500/85x0/5520
- Switchport as Trunk if SSID/VLAN separation needed

HA Design and Deployment Practices

Branch Redundancy: Local Controller, Flex local switching & Central backup Controller

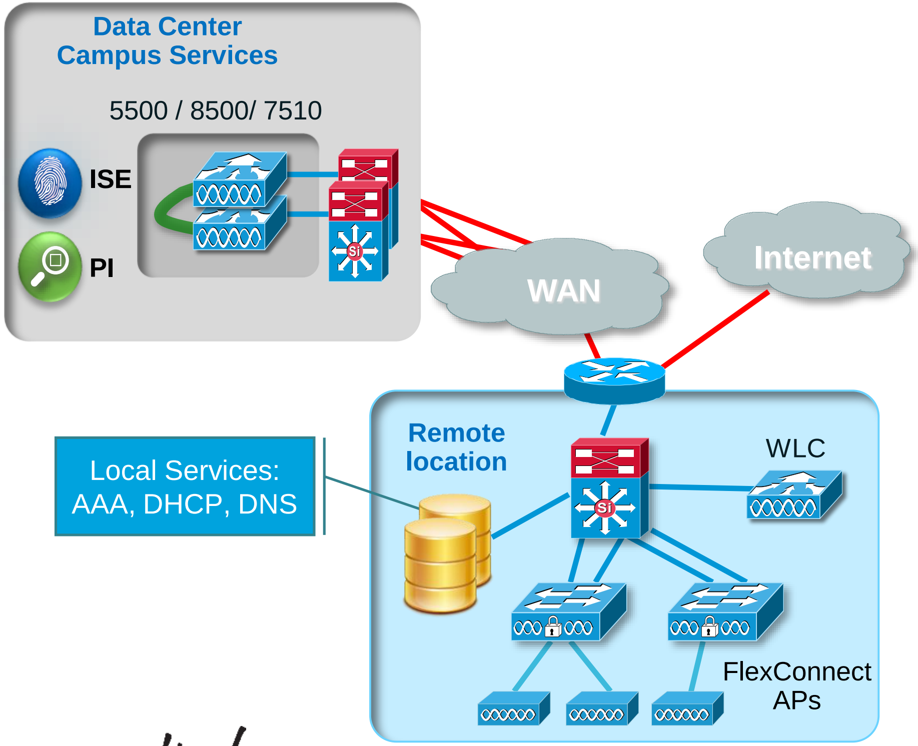

High Availability considerations:

- Local Controller for managing the APs and for providing **Control plane survivability** in the event of a WAN failure (RRM, CleanAir, WebAuth, etc.)
- Why AP in Flex? So that if the local controller fails, the APs can failover to the central controller but traffic still remains local

Design considerations:

- AP in Flex mode = less features and functionality compare to Local Mode. Key features missing:
  - No L3 roaming, No Bonjour GW
- If using Flex Groups be aware of the AP count limit (25 APs for 2500/5508, 100 APs for 7500 / 8500)
- Switchport as Trunk if SSID/VLAN separation needed
- For large branch is recommended to have DHCP, DNS and AAA services running locally for better reliability

Key takeaways

Cisco live!

Key Takeaways

- High Availability for Wireless is a multi level approach, starting from Level 1 (RF)
- You have different solutions to chose based on the downtime that is acceptable for your business application
- Cisco Controller SSO eliminates the network downtime upon a controller failure

Key Takeaways

Network Uptime

	Requirements	Benefits	Downtime
Stateful Switchover (SSO)	Minimum release: 7.5 5500, WiSM2, 7500, 8500 series L2 connection Same HW and software 1:1 box redundancy	Active Client State is synched AP state is synched No Application downtime HA-SKU available	Predictable < 1 sec
N+1 Redundancy	Each Controller has to be configured separately	Available on all controllers Crosses L3 boundaries Flexible: 1:1, N:1, N:N HA-SKU available (> 7.4)	Predictable <30 sec
Mobility Group	Each Controller has to be configured separately	Available on all controllers Crosses L3 boundaries No specific HA configuration	Unpredictable

Complete Your Online Session Evaluation

- Give us your feedback to be entered into a Daily Survey Drawing. A daily winner will receive a $750 Amazon gift card.
- Complete your session surveys through the Cisco Live mobile app or from the Session Catalog on CiscoLive.com/us.

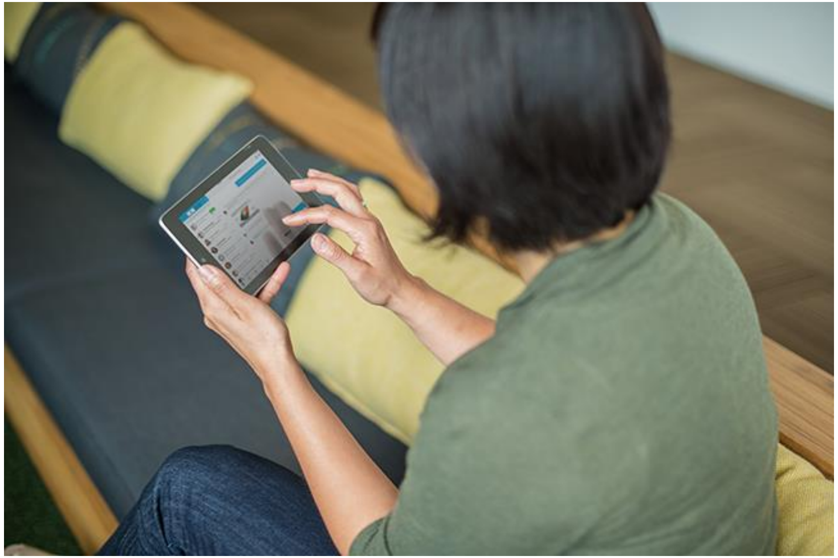

Don’t forget: Cisco Live sessions will be available for viewing on-demand after the event at CiscoLive.com/Online

Continue Your Education

- Demos in the Cisco campus
- Walk-in Self-Paced Labs
- Lunch & Learn
- Meet the Engineer 1:1 meetings
- Related sessions

Please join us for the Service Provider Innovation Talk featuring:

Yvette Kanouff | Senior Vice President and General Manager, SP Business

Joe Cozzolino | Senior Vice President, Cisco Services

Thursday, July 14th, 2016

11:30 am - 12:30 pm, In the Oceanside A room

What to expect from this innovation talk

- Insights on market trends and forecasts
- Preview of key technologies and capabilities
- Innovative demonstrations of the latest and greatest products
- Better understanding of how Cisco can help you succeed

Register to attend the session live now or watch the broadcast on cisco.com

Wireless Cisco Education Offerings

Course	Description	Cisco Certification
• Designing Cisco Wireless Enterprise Networks • Deploying Cisco Wireless Enterprise Networks • Troubleshooting Cisco Wireless Enterprise Networks • Securing Cisco Wireless Enterprise Networks	Professional level instructor led trainings to prepare candidates to conduct site surveys, implement, configure and support APs and controllers in converged Enterprise networks. Focused on 802.11 and related technologies to design, deploy, troubleshoot as well as secure Wireless infrastructure. Course also provide details around Cisco mobility services Engine, Prime Infrastructure and wireless security.	CCNP® Wireless Version 3.0 (Available March 22nd, 2016)
Implementing Cisco Unified Wireless Network Essential	Prepares candidates to design, install, configure, monitor and conduct basic troubleshooting tasks of a Cisco WLAN in Enterprise installations.	CCNA® Wireless (Available Now)
Deploying Basic Cisco Wireless LANs (WDBWL)	Understanding of the Cisco Unified Wireless Networking for enterprise deployment scenarios. In this course, you will learn the basics of how to install, configure, operate, and maintain a wireless network, both as an add-on to an existing wireless LAN (WLAN) and as a new Cisco Unified Wireless Networking solution.	1.2
Deploying Advanced Cisco Wireless LANs (WDAWL)	The WDAWL advanced course is designed with the goal of providing learners with the knowledge and skills to successfully plan, install, configure, troubleshoot, monitor, and maintain advanced Cisco wireless LAN solutions such as QoS, “salt and pepper” mobility, high density deployments, and outdoor mesh deployments in an enterprise customer environment.	1.2
Deploying Cisco Connected Mobile Experiences (WCMX)	WCMX will prepare professionals to use the Cisco Unified Wireless Network to configure, administer, manage, troubleshoot, and optimize utilization of mobile content while gaining meaningful client analytics.	2.0

For more details, please visit: http://learningnetwork.cisco.com
Questions? Visit the Learning@Cisco Booth or contact ask-edu-pm-dcv@cisco.com

Thank you

cisco

Cisco live!

July 10-14, 2016 · Las Vegas, NV