



---

# **CONVERGED ACCESS – WIRED / WIRELESS SYSTEM ARCHITECTURE, DESIGN, AND OPERATION**

**JUNE 2013**

# Contents

Introduction.....	10
User Trends.....	10
Wireless Evolution.....	10
High Availability.....	11
Mobile Devices.....	11
Unified and Converged Access.....	11
Converged Access Platform Overview.....	12
Cisco Catalyst 3850.....	12
Cisco WLC 5760.....	13
Wireless Deployment Options.....	14
Use Cases for Cisco Converged Access.....	15
System Architecture, Roaming, and High Availability.....	16
Existing Architecture.....	16
Tunnels.....	16
Control.....	17
Mobility Groups and Domains.....	18
PoP and PoA.....	19
Layer 2 Roaming.....	19
Layer 3 Roaming.....	21
Roaming with Mobility Anchors.....	22
Traffic Flow.....	23
Control Plane and Data Plane.....	23
Next Generation Cisco Converged Access Architecture.....	24
Mobility Architecture.....	25
Physical Entities.....	26
Logical Entities.....	27
Scalability.....	27
Traffic Flows and Roaming.....	28
Basic Branch Roaming.....	29
Intra-SPG Branch Roaming.....	29
Intra-SPG Roaming – Details.....	30

Intra-SPG Branch Roaming with a Separate MC .....	35
Roaming Across SPGs Under Control of the Same MC .....	37
Inter-SPG Roaming Under Control of the Same MC – Details .....	39
Roaming Across SPGs Under Control of Different MCs .....	41
Roaming with Mobility Anchors .....	42
Layer 2 or Non-Sticky Roaming – Optional .....	43
Scaling with Catalyst 3850 based MCs .....	44
High Availability .....	45
MC Redundancy Information .....	51
Multicast, Deployment Options, and Migration .....	53
Multicast for Converged Access .....	53
Multicast Process Details .....	53
Multicast Process Flows .....	55
Multicast with IPv6 .....	57
Design Options .....	58
Small Branch .....	58
Small/Medium Branch .....	60
Large Branch .....	61
Small Campus .....	62
Campus .....	63
Large Campus .....	64
IP Addressing .....	64
Option 1. Separate Wired and Wireless VLANs Per Wiring Closet .....	64
Option 2. Merged Wired and Wireless VLANs Per Wiring Closet .....	65
Option 3. Merged Wired and Wireless VLANs Per Wiring Closet with Spanned Wireless VLAN .....	66
Deployment and Implementation Examples .....	67
Basic Deployment Example .....	68
Small Campus Deployment Example .....	70
Large Campus Deployment Example .....	71
RRM and CleanAir .....	73
RRM .....	73
CleanAir .....	75

Migration.....	76
Quality of Service .....	79
QoS Elements.....	79
Existing QoS deployments and Challenges .....	80
Existing Campus QoS Architecture .....	81
QoS Policies .....	82
QoS with Converged Access .....	83
QoS Behavior .....	84
QoS Design Options and Policies .....	85
Management.....	87
Cisco Prime Infrastructure .....	87
Features .....	87
Lifecycle Management.....	87
Platform and Technology Support .....	88
Simplifying Deployment.....	88
Improved Day Zero Experience – Plug and Play.....	88
Design and Deploy .....	91
Single Pane Glass Visibility.....	101
Security and Guest Access.....	103
Integrated Policies.....	103
Policy Definition and Application .....	104
Session-Aware Networking .....	104
Policy Enforcement.....	104
Per-Session VLAN Assignment .....	105
Service Templates .....	105
ACLs .....	105
Device Enrollment and Provisioning .....	105
Catalyst Integrated Security Features .....	106
Security Features Summary .....	106
Guest Access.....	107
Deployment Examples .....	107
Guest Access Configuration .....	109



Release Compatibility for Guest Access.....	112
Troubleshooting .....	113
Catalyst 3850 Hardware Architecture.....	114
Catalyst 3850 Platform Overview.....	114
Network Modules.....	114
Power Modules .....	115
Stacking and Stacking Cables on the Catalyst 3850 .....	116
IOS-XE on the Catalyst 3850 .....	116
Catalyst 3850 Platform Architecture .....	117
ASIC Details .....	118
Packet Walks.....	118
CAPWAP Frames .....	123
Traffic Walks .....	124
Catalyst 3850 Stacking Architecture .....	127
Catalyst 3850 High Availability .....	130
HA Redundancy .....	130
Stacking vs. Catalyst 6500 .....	130
Catalyst 3850 Software HA Processes.....	130
Stack Discovery and Formation.....	130
Stack Addition and Deletion.....	131
Performance and Scalability .....	132
UADP Performance .....	132
TCAM and ACL Sizing .....	132
FNF Scale and Performance .....	133
QoS Scalability .....	133
Scalability Comparison.....	134
Migration.....	134
Conclusion .....	138

## List of Figures

Figure 1. Main Campus Deployment .....	13
Figure 2. Deployment Options.....	14
Figure 3. Catalyst 3850 Use Cases .....	15
Figure 4. Existing Architecture .....	16
Figure 5. Existing Architecture – Control Functions .....	18
Figure 6. MD with MGs.....	18
Figure 7. Existing Architecture – PoP and PoA .....	19
Figure 8. Existing Architecture – Layer 2 Roaming .....	20
Figure 9. Existing Architecture – Layer 3 Roaming .....	21
Figure 10. Existing Architecture – Mobility Anchors .....	22
Figure 11. Existing Architecture – Traffic Flow .....	23
Figure 12. Converged Access Architecture – Deployment Overview .....	24
Figure 13. Converged Access – Mobility Architecture .....	26
Figure 14. Mobility Packet Format .....	28
Figure 15. Roaming Within an SPG .....	29
Figure 16. Intra-SPG Roaming Example – Before Roaming .....	31
Figure 17. Intra-SPG Roaming Example – Before Roaming with DHCP Snooping .....	31
Figure 18. Intra-SPG Roaming Example - Roaming .....	32
Figure 19. Intra-SPG Roaming Example – Details.....	33
Figure 20. Intra-SPG Roaming Example – Additional Details.....	34
Figure 21. Intra-SPG Roaming Example – Additional Details.....	35
Figure 22. Intra-SPG Configuration with Separate MC .....	35
Figure 23. Roaming in an Intra-SPG Configuration with Separate MC.....	37
Figure 24. Roaming Across SPGs Under Control of the Same MC .....	38
Figure 25. Inter-SPG Branch Roaming with an MC.....	39
Figure 26. Inter-SPG Branch Roaming with an MC – Data Plane .....	40
Figure 27. Intra-SPG Branch Roaming with an MC – Roamed and Non-Roamed Clients.....	41
Figure 28. Roaming Across SPGs Under Control of the Different MCs.....	41
Figure 29. Roaming with a Mobility Anchor.....	42
Figure 30. Roaming Across SPGs Under Control of Different MCs – Layer 2 Roaming .....	44
Figure 31. State Behavior for Local and Roamed Users .....	45

Figure 32. State Behavior for Local and Roamed Users – View from the MC .....	46
Figure 33. State Behavior for Local and Roamed Users – View from the MAs .....	46
Figure 34. MC Failure – Effect on MC Sub-Domain and Anchor Connections .....	47
Figure 35. MC Failure – Effect on Local (Non-Roamed) Users .....	47
Figure 36. MC Failure – Effect on Previously Roamed Clients .....	48
Figure 37. MC Failure – Effect on Previously Roamed Clients .....	49
Figure 38. MC Failure – Effect on Intra-MC Roams Following MC Down .....	49
Figure 39. MC Failure – Effect on Inter-MC Roams Following MC Down .....	50
Figure 40. MC Redundancy .....	51
Figure 41. HA with Catalyst 3850 Based MCs.....	52
Figure 42. Multicast Process Examples .....	54
Figure 43. Multicast Forwarding .....	55
Figure 44. IGMP Leave .....	56
Figure 45. Broadcast Forwarding .....	56
Figure 46. Multicast with Video Streams .....	57
Figure 47. Small Branch Deployment .....	59
Figure 48. Small/Medium Deployment .....	60
Figure 49. Large Branch Deployment .....	61
Figure 50. Small Campus Deployment.....	62
Figure 51. Small Campus with Multiple MGs.....	63
Figure 52. Campus Deployment .....	63
Figure 53. Large Campus Deployment.....	64
Figure 54. IP Addressing – Option 1 .....	65
Figure 55. IP Addressing – Option 2 .....	66
Figure 56. IP Addressing – Option 3 .....	67
Figure 57. Branch Deployment Example.....	68
Figure 58. Small Campus Deployment Example.....	70
Figure 59. Large Campus Deployment Example.....	71
Figure 60. Overview of RRM and CleanAir.....	73
Figure 61. Wired and Wireless Considerations for QoS.....	79
Figure 62. Existing Wired QoS Environment.....	80
Figure 63. Wired QoS Environment .....	81

Figure 64. Policies Overlay in Existing QoS Environment .....	82
Figure 65. QoS Policies – Wired to Wireless .....	83
Figure 66. QoS Policies – Wireless to Wired .....	84
Figure 67. Example of Design Options to Support QoS.....	85
Figure 68. Cisco PI .....	87
Figure 69. Existing Day 0 Experience Versus Converged Access Day 0 Experience .....	88
Figure 70. Design and Deploy Workflow .....	99
Figure 71. Client Listing Interface .....	102
Figure 72. Users, Devices, and Policies.....	103
Figure 73. Guest Access – Small to Mid-Size Branch.....	108
Figure 74. Guest Access – Small to Mid-Size Branch with Guest Anchor .....	108
Figure 75. Guest Access – Large Campus .....	109
Figure 76. Guest Access Example for Large Campus .....	111
Figure 77. Constraints .....	113
Figure 78. Catalyst 3850 Platform .....	114
Figure 79. Catalyst 3850 Network Modules .....	115
Figure 80. Catalyst 3850 Power Modules.....	115
Figure 81. Catalyst 3850 Stacking Cables .....	116
Figure 82. Catalyst 3850 Hardware Architecture .....	117
Figure 83. WS-C3850-48 Layout .....	117
Figure 84. Packet Walk for Local Switching.....	119
Figure 85. Packet Walk for Remote Switching – Ingress .....	120
Figure 86. Packet Walk for Remote Switching – Egress.....	120
Figure 87. Packet Walk for Multicast Local Replication .....	121
Figure 88. Packet Walk for Multicast Across Stacks – Ingress .....	121
Figure 89. Packet Walk for Multicast Across Stacks - Egress .....	122
Figure 90. Packet Walk for Recirculation.....	123
Figure 91. CAPWAP Frames – Wireless to Wired.....	123
Figure 92. CAPWAP Frames – Wired to Wireless.....	124
Figure 93. Wired to Wireless Feature Flow - Simple .....	125
Figure 94. Wireless to Wired Feature Flow - Simple .....	125
Figure 95. Wired to Wireless Feature Flow – Frame and Encrypt .....	126

Figure 96. Stack Ring .....	127
Figure 97. Unicast Packet Path .....	128
Figure 98. Unicast Packet Path – Spatial Reuse .....	128
Figure 99. Multicast Packet Path on the Stack Ring.....	129
Figure 100. Stacking Ring Healing.....	129
Figure 101. TCAM and ACL Sizing.....	132
Figure 102. TCAM and ACL Sizing.....	133
Figure 103. Example Deployment Prior to Migration .....	135
Figure 104. Beginning Migration to Converged Access .....	135
Figure 105. Continued Migration to Converged Access .....	136
Figure 106. Migration to Converged Access .....	137

# Introduction

The need for pervasive connectivity continues to grow as mobile devices proliferate along with user expectations of high quality access to systems and applications from wherever they happen to be. IT organizations are facing increased pressure to provide greater flexibility in devices clients use to access IT services, and they are looking for effective solutions.

## User Trends

Several major trends have combined to increase the pressure on IT organizations to deliver an uncompromised user experience anywhere and on any device.

**Bring your own device (BYOD).** Users are increasingly relying on their own mobile devices (laptops, tablets, smart phones) for network access, whether at work, home, or in public places. Throughout the enterprise, secure access must be provided for contractors and guests in addition to employees. Regardless of the access type, users expect a customized and rich user experience.

**Mobility.** Users want to retain wireless connectivity as they move within a location or from one location to another. Enterprise employees need to connect to resources at home, on-the-go in hotels or airports, and in different areas of the enterprise, including branch offices and campus offices. They expect seamless roaming with no degradation in client performance. As they increasingly rely on the cloud for access to services and data, they don't want to lose connections to cloud services as they roam.

**Video.** Video on demand and video streaming to mobile devices has become an essential part of the user experience. Video conferencing is no longer restricted to specially configured enterprise locations, but has now moved to the desktop, and even to mobile devices. The bandwidth requirements of video mean the performance of the underlying network is a major concern. As video is intolerant of loss and jitter, Quality of Service (QoS) across wired and wireless networks is a must.

## Wireless Evolution

Wireless has evolved dramatically over the past 15 years. Initially, customer connectivity was characterized as best effort and casual. Deployments tended to be confined to conference rooms or small office environments. First generation products enabled this type of coverage, operating in a hotspot mode. As enterprises looked to move into more pervasive wireless deployments, the need to manage the network centrally became increasingly apparent. The introduction of controller-based architectures allowed IT to manage the wireless network as a system and made it easier to cover the entire enterprise appropriately for wireless.

Next, the introduction of 802.11n allowed the wireless network to respond to the demand for more capacity, driven by media-rich applications. 802.11n provided the throughput that many leading adopters of wireless were looking for. Yet, while 802.11n solved the capacity problem, many companies still struggled with how to make the network more reliable and easier to operate. The need to deliver a truly mission-critical wireless network has driven the development of network infrastructure that has the intelligence to be self-healing and self optimizing.

The growth has been dramatic. From 1997 until today, data rates have grown from 2 Mbps to 11 Mbps to 54 Mbps, up to the latest rates of 450 Mbps. Moreover, 802.11ac Wave-1 products are now available with up to 1300 Mbps of wireless data rates provided within a single access point radio. Future 802.11ac Wave-2 products will be able to reach up to 3500 Mbps wireless data rates, allowing for up to 50 high definition rate videos at 720p for 802.11ac Wave-1, and 1080p for 802.11ac Wave-2 devices.

## High Availability

With many more devices connecting through the network and greater reliance on such devices for essential work, high availability (HA), resiliency, and visibility of service impacting incidents have become critical to the running of many businesses. In just the past few years, HA for Wi-Fi has become a reality from an infrastructure perspective, in terms of the time it takes to roam for voice and video, and for clients to reconnect after an outage without losing an IP address or connection state. Industries such as transportation, healthcare, and mining are demanding management and visibility with as much as 99.999% uptime.

## Mobile Devices

Mobile devices have become ubiquitous, and it is now common for people to use multiple mobile devices. That trend is expected to grow, with some devices providing general functions and others used for specific applications.

## Unified and Converged Access

Based on “One Policy, One Management, One Network,” the Cisco Unified Access Network solution delivers an integrated, simplified, and intelligent network platform that enables IT to spend less time running the network and more time collaborating and innovating with stakeholders to differentiate and transform the business.

The Cisco Unified Access solution encompasses the following:

- **Converged Access.** One physical infrastructure for wired and wireless connectivity, which increases business agility, simplicity, and scale, and delivers significant operational efficiencies. The Cisco Catalyst 3850 is a Converged Access switch platform with integrated wireless controller functionality, and is the foundation of the unified wired and wireless network.
- **Consistent network-wide intelligence and operations.** One common set of network capabilities and context-aware intelligence for policy, visibility, analytics, and granular QoS across the entire wired-wireless infrastructure, providing operational simplicity and a consistent user experience. The Cisco Converged Access solution is based on a common ASIC design and a common operating system for wired and wireless to further enhance feature consistency.
- **Integration into Cisco Open Network Environment (ONE).** The industry’s first common set of interfaces across wired and wireless that support a programmable data plane using Open Network Environment Platform Kit (OnePK) to enhance business agility.

Converged Access has important benefits to support the demands of pervasive connectivity:

- **Single platform for wired and wireless.** Connectivity is all based on IOS with the same administration point for both the wired and wireless aspects of connectivity, and a single code release for enhancements and new features.
- **Network-wide visibility.** Wired and wireless visibility into the network is available at every hop. This means faster troubleshooting and trending of network traffic.
- **Consistent security and QoS.** Hierarchical bandwidth management and distributed policy enforcement can start taking effect at the first point of ingress into the network, for both wired and wireless traffic.

- Maximum resiliency. The Catalyst 3850 solution supports HA with a stack using StackPower, StackWise redundancy, and the active-standby relationship of the master switch in the stack. This level of resiliency now benefits both wired and wireless traffic.
- Scaling with distributed wired and wireless data plane. With support for up to 480G stack bandwidth, up to 40G wireless switching capacity per Catalyst 3850 switch, and efficient multicast traffic handling, the Cisco Converged Access solution is fully 802.11ac-ready.

This paper focuses on the Converged access aspect of Cisco Unified Access and how this approach can effectively meet the ever increasing demands for connectivity and performance in wired and wireless networks.

## Converged Access Platform Overview

Cisco Unified Access structure is built on three pillars: one policy, one management, and one network. *One policy* is provided through the Identity Services Engine (ISE), which supports BYOD policy management, device profiling and posture, and guest access portals. *One management* is provided through Cisco Prime 2.0, which offers complete wired and wireless management, a user and device-centric view of the network, and intuitive troubleshooting workflows. *One network* is provided through Converged Access, which combines wireless networks into a single unified infrastructure

The Cisco solution for converged access within the Unified Access framework is based on two major platforms: the Cisco Catalyst 3850, and the Cisco WLC 5760.

### Cisco Catalyst 3850

The Cisco Catalyst 3850 is a single wired/wireless platform with a common IOS implementation, the same administrative access for both wired and wireless aspects of connectivity, and a single code release for improvements and new features. It incorporates a wireless LAN controller (WLC) capability natively in IOS to support functions such as Layer 2 and Layer 3 fast roaming, CleanAir, VideoStream, Radio Resource Management (RRM), and Dynamic ARP inspection. DHCP snooping protections enhance wireless security beyond what is available in wireless-only type of deployments, while wireless Switch Peer Groups (SPGs) provide faster roaming, along with support for latency-sensitive applications and mobility for IPv4 and IPv6 clients.

On the wired side, features include StackPower, advanced identity services, network visibility and control, granular QoS, and scripting capabilities. Wired and wireless features are built on the innovative Unified Access Data Plane (UADP) Flexparser ASIC technology, which eliminates the need to replace hardware each time a new feature is made available.

The Catalyst 3850 itself is a powerful switch with up to 480 Gbps of stacking bandwidth and up to 4x10 Gbps modular uplinks. Each stack can support to four stack members and up to 2000 clients per stack. It offers full Power over Ethernet + (PoE+) support, and with StackPower, the ability to load balance and load share power to the different switches in the stack, which alleviates the need to calculate the exact power requirement for each switch in the stack. Field replaceable fans and power supplies augment the ease-of-maintenance and HA of the entire solution.

Using the flexible power offered by the UADP ASIC, a CAPWAP tunnel from a directly-connected Cisco 802.11n-capable AP can be terminated directly in hardware on the ingress Catalyst 3850 switchport, enabling great scalability in wireless performance while also enabling end-to-end traffic visibility for QoS and security. PoE is available for APs, IP Phones, security cameras, and other devices.



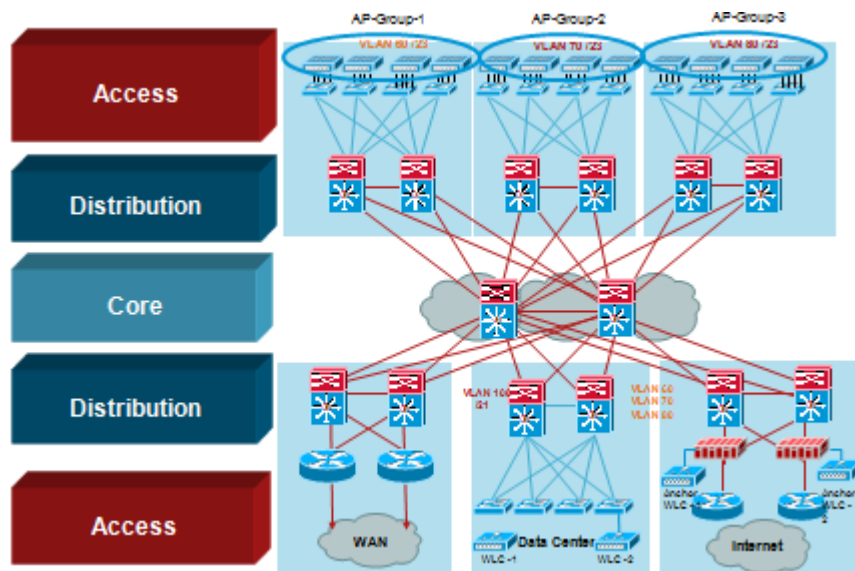
A stack of Catalyst 3850s supports up to fifty directly connected APs, which can cover up to five floors of a building 60 by 30 meters or 45 by 40 meters. The wireless SPGs allow faster roaming between different floors in the building and different buildings in the campus. SPGs thus provide faster switch stack roaming for latency sensitive applications.

## Cisco WLC 5760

The Cisco WLC 5760 is a next generation WLC that is fully 802.11ac optimized. It supports up to 60 Gbps of wireless bandwidth on six 10 Gbps SFP ports, which also support link aggregation to multiple uplink switches. As a centralized wireless controller, it supports up to 1000 APs. As the first IOS based wireless controller, the WLC 5760 includes IOS features such as granular QoS and Flexible NetFlow. Traffic from up to 12,000 concurrent clients can pass through the WLC 5760 at any time. As with the Catalyst 3850, field replaceable fans and power supplies improve maintenance and HA.

For campus deployments, a three-tier architecture includes access, distribution, and core tiers, with routed access or Layer 2 trunks between access and distribution. There is no difference in the placement of the WLC 5760 compared to current WLC 5508 or WiSM-2 integrated controllers used with the Catalyst 6500 switch platform. Placing the WLC 5760 in the services building block, mobility building block, or data center allows distribution of APs by groups and use of mobility groups (MGs) to enhance the size of the virtual cell in which clients can move. Clients can move from AP to AP without losing their credentials using Cisco Centralized Key Management (CCKM) or Opportunistic Key Caching (OKC)/ Proactive Key Caching (PKC).

**Figure 1. Main Campus Deployment**







## Wireless Deployment Options

The following figure compares different wireless deployment mode options based on the type of network.

- Autonomous mode is appropriate for standalone APs or small cost-effective deployments. Support is provided for roaming (Layer 2 only), basic RRM capability, and security features.
- FlexConnect mode allows break-out of traffic at the AP, while control is handled by a centralized controller. This approach has been used to scale performance in distributed networks. A centralized controller helps in deploying configurations and providing local network access in the branches. As with autonomous APs, this option does not support Layer 3 roaming, and it has some bandwidth limitations on the WAN between the controller at a central location and APs that are deployed in the branch office.
- Centralized mode is a common option for campuses and allows for Layer 3 roaming, RRM, and advanced security, including rogue detection and mitigation. It simplifies the operations of the wireless network with centralized control and traffic visibility at the controller.
- Finally, Converged Access mode uses the Catalyst 3850 at the access layer. The switch with integrated wireless controller functionality is a single enforcement point for QoS and policy at the access layer and provides traffic visibility through advanced functionality such as Flexible Netflow. Converged Access is optimized for performance of 802.11ac, as wireless traffic is terminated as soon as the traffic enters the wired network. For deployments over 250 APs, a centralized controller is needed to coordinate Catalyst 3850s.

**Figure 2. Deployment Options**

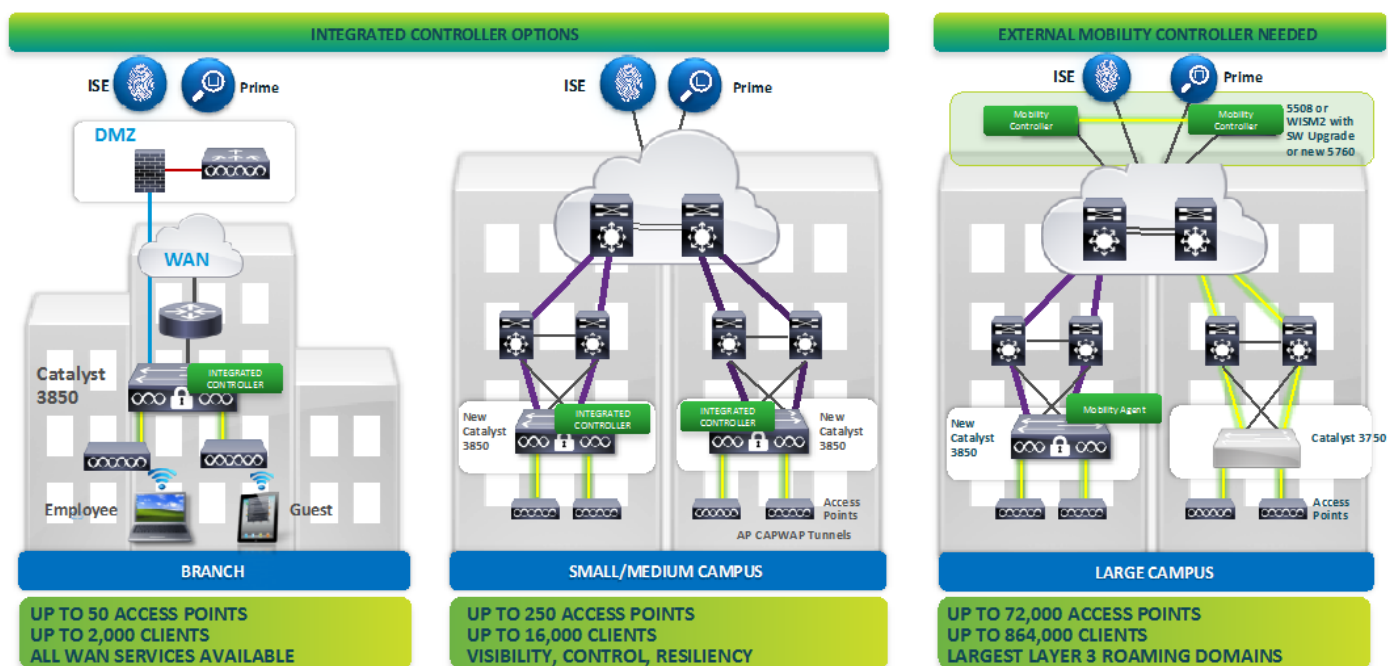
	Autonomous	FlexConnect	Centralized	Converged Access
	 Standalone APs	 Traffic Distributed at AP	 Traffic Centralized at Controller	 Traffic Distributed at Switch
Target Positioning	Small Wireless Network	Branch	Campus	Branch and Campus
Purchase Decision	Wireless only	Wireless only	Wireless only	Wired and Wireless
Benefits	<ul style="list-style-type: none"> <li>• Simple and cost-effective for small networks</li> </ul>	<ul style="list-style-type: none"> <li>• Highly scalable for large number of remote branches</li> <li>• Simple wireless operations with DC hosted controller</li> </ul>	<ul style="list-style-type: none"> <li>• Simplified operations with centralized control for Wireless</li> <li>• Wireless Traffic visibility at the controller</li> </ul>	<ul style="list-style-type: none"> <li>• Wired and Wireless common operations</li> <li>• One Enforcement Point</li> <li>• One OS (IOS)</li> <li>• Traffic visibility at every network layer</li> <li>• Performance optimized for 11ac</li> </ul>
Key Considerations	<ul style="list-style-type: none"> <li>• Limited RRM, no Rogue detection</li> </ul>	<ul style="list-style-type: none"> <li>• L2 roaming only</li> <li>• WAN BW and latency requirements</li> </ul>	<ul style="list-style-type: none"> <li>• System throughput</li> </ul>	<ul style="list-style-type: none"> <li>• Catalyst 3850 in the access layer</li> </ul>

## Use Cases for Cisco Converged Access

The following figure illustrates typical use cases for the Cisco Converged Access:

- **Small to medium-sized branch.** Up to 50 APs and 2000 clients in the network. The integrated controller inside the Catalyst 3850 handles a single stack of switches, or a group of stacks, allowing traffic visibility and control, without the need to deploy a discrete controller at the site.
- **Larger branch or small to medium-sized campus.** Multiple stacks of Catalyst 3850s allow for scalability up to 250 APs and 16,000 clients. This type of deployment supports visibility, resiliency, and control, still without the necessity to deploy a discrete controller as part of the solution.
- **Large campus.** Multiple stacks of switches are deployed inside the network, and one or more central Mobility Controllers (MCs) in the form of discrete controllers are deployed in the data center. This arrangement supports APs that are connected directly to Catalyst 3850s, as well as APs that are not connected to any Converged Access switches which are handled by the discrete controller or controllers. This option allows for tremendous scalability of large campuses with up to 72,000 AP and 800+ thousand clients into the network.

Figure 3. Catalyst 3850 Use Cases



# System Architecture, Roaming, and High Availability

The changes encompassed by Converged Access are being driven by the need for scale and performance relating to the number of devices, speed of devices coming onto the network, and the demand for a unified experience. The goal is to expand the rich experience of functionality on wired networks into the wireless realm while maintaining scale and performance standards.

To fully understand the enhancements and benefits of Converged Access, it is important to first review the existing Unified Access architecture and how it manages the overlaid wired/wireless network.

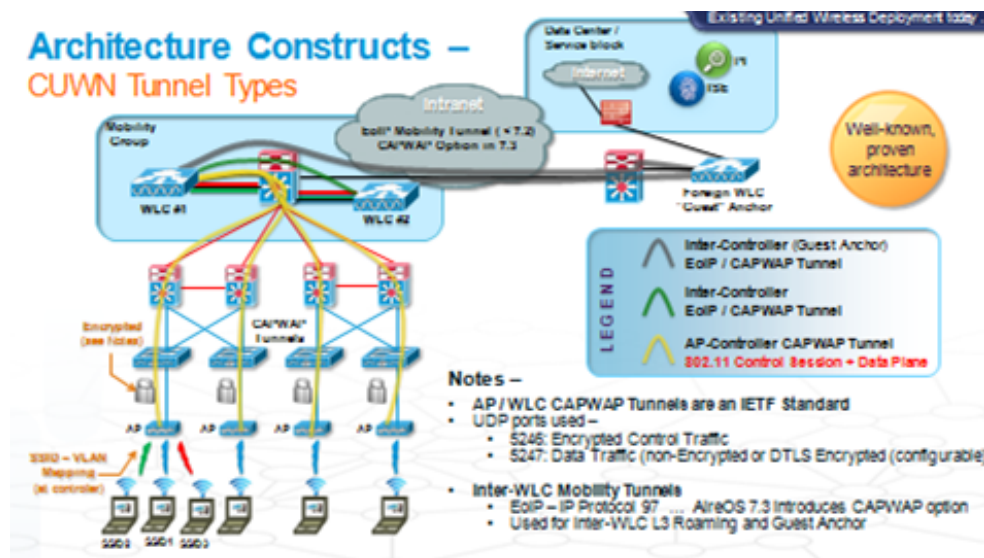
## Existing Architecture

The existing Cisco Unified Wireless Network (CUWN) architecture incorporates previous changes that were driven around a functional split for control and provisioning of wireless APs (CAPWAP). System-wide functionality moved to the controller level, where a controller terminates the wireless data plane through CAPWAP tunnels that return from the APs. The controller also terminates the control plane for the wireless network and handles functions such as channel and power assignment for APs, rogue detection, and wireless intrusion prevention. The controller approach frees the AP to do deal with AP-specific time sensitive functions such as frame queuing and packet prioritization, encryption and beacons and probes response. This approach is called a *split MAC architecture*, in which the MAC layer is split so that part of it runs on the controller and part of it runs on the AP. In this existing architecture, the wireless network is built as an overlay of tunnels on top of the wired network.

## Tunnels

To understand the architecture, consider a typical campus network, as shown the following figure. The network includes APs and CAPWAP tunnels running back from the APs up to the controller.

Figure 4. Existing Architecture



The yellow lines show the CAPWAP tunnels for the control and data planes, and the blue lines show the VLAN/SSID mappings. The WLAN controllers are grouped together in an MG that supports Ethernet over IP (EoIP), or CAPWAP tunnels, which are automatically added as part of the MG. The green lines show these inter-controller tunnels with EoIP or CAPWAP.

Gray lines show CAPWAP tunnels running from the WLCs up to a foreign WLC or a guest anchor controller. These are typically EoIP or CAPWAP tunnels running from the individual WLAN controllers up to a guest anchor controller. The guest traffic typically drops to a firewall and is forwarded to the Internet. The entire deployment is managed by ISE and Cisco Prime Infrastructure (PI). The bottom of the figure shows multiple SSIDs (SSID2 and SSID3), with an SSID to VLAN mapping. This mapping could be static (1:1 mapping of SSIDs to associated wired-wide VLANs), or dynamic based on RADIUS, where users within the same SSID can be mapped to different wired-side VLANs based on their Authentication, Authorization, and Accounting (AAA) attributes.

The following subsections highlight the evolution from this well-known, well-understood, and widely used deployment mode to Converged Access. Because Converged Access has evolved from Cisco's overall wireless architecture, understanding the existing CUWN deployment model provides a base for understanding Converged Access. Many of the architectural elements are identical, but they now operate in a more widely-distributed Converged Access deployment model instead of the centralized CUWN model.

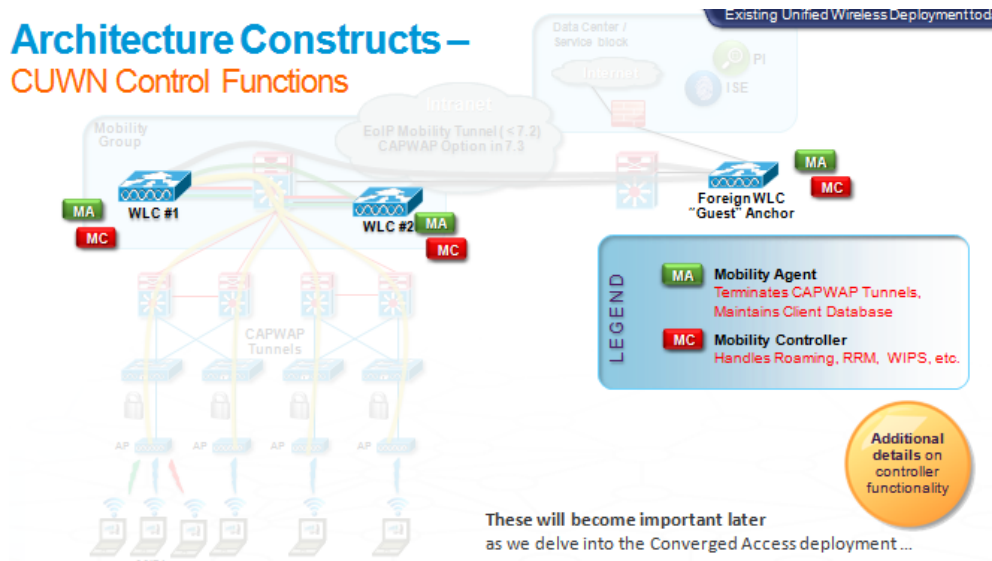
## **Control**

The controller supports two key functions. The *mobility agent* (MA) function terminates the CAPWAP tunnels in the network and maintains the client database. This is a major responsibility, with potentially thousands or tens of thousands of clients in a network associating, de-associating, and roaming.

The *mobility controller* (MC) is responsible for system-wide functions such as roaming, RRM, and wireless intrusion prevention. In each of the following diagrams, the MA function is shown in green and the MC function is shown in red.

The MC and MA functions are already part of the existing CUWN architecture. They are reintroduced here because they play a central role in the distributed Converged Access architecture, and where they are hosted and placed affects the operation and capabilities of the specific Converged Access solution.

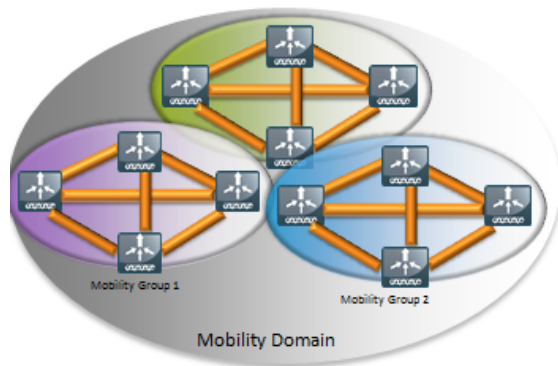
Figure 5. Existing Architecture – Control Functions



### Mobility Groups and Domains

A *mobility group* (MG) is a set of WLCs that are assigned the same MG name. An MG can include up to 24 controllers that share a full mesh of tunnels and provides seamless mobility and fast roaming. A *mobility domain* (MD) is a group of controllers configured on each WLC that specifies members in different MGs. It is a cross-population of up to three MGs for a total of up to 72 controllers and defines the seamless MD.

Figure 6. MD with MGs



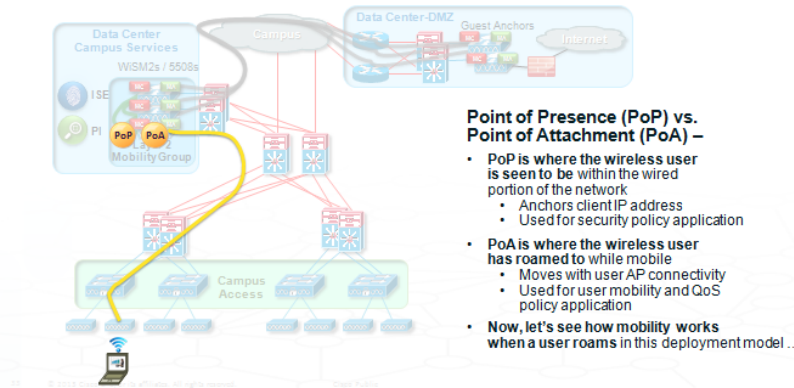
## PoP and PoA

The following figure shows user traffic coming up through CAPWAP tunnels at the bottom of the figure to the group of controllers in the data center block. The user traffic enters through an AP that is terminated at the bottom controller. The controllers form a Layer 2 MG with a common set of VLANs on the back end, in this example.

Figure 7. Existing Architecture – PoP and PoA

### Architecture Constructs –

Point of Presence (PoP), Point of Attachment (PoA)



The *point of presence* (PoP) is where the wireless user is seen to be within the wired portion of the network. The user might be roaming anywhere in the wireless network, but from the wired network perspective (for example with respect to network ARP for the user), the PoP is the .1Q trunk that leads to the controller where the user's initial AP is hosted. The PoP anchors the client IP address so the wireless client can maintain the same IP address while moving around the network. It is also used for its security policy application, for example, for application of access control lists (ACLs).

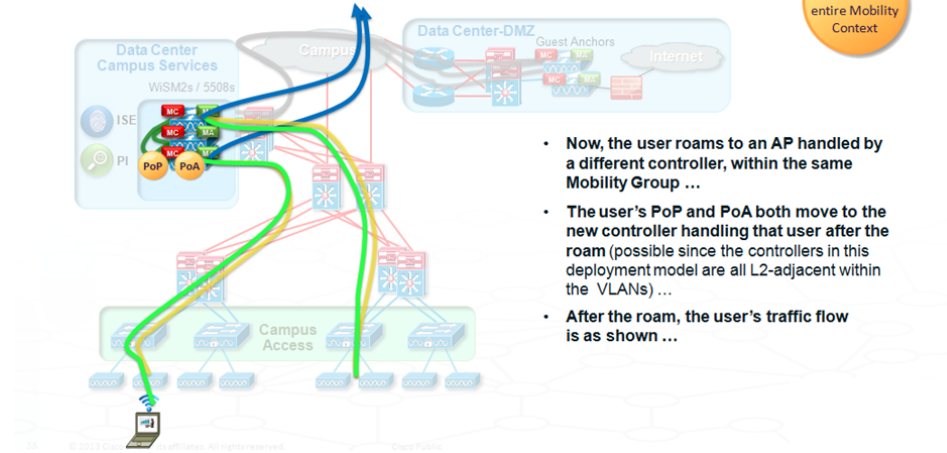
The *point of attachment* (PoA) is where the CAPWAP tunnel from the AP is terminated. It is where the wireless user traffic first is first seen in the network. The PoA moves with the user roaming as the user associates to a different AP. The PoA is used for user mobility and QoS.

## Layer 2 Roaming

To understand roaming, consider the three controllers in the following figure. The user's PoP and PoA are initially at the bottom controller in this example, as the AP that the user happens to be associated to is hosted on that controller. The bright green line on the left shows the user's traffic moving through a CAPWAP tunnel to the bottom controller. The lower blue line with an arrow on it shows that the user's traffic is decapsulated and is moving into the larger Ethernet / IP network.

Figure 8. Existing Architecture – Layer 2 Roaming

## Architecture Constructs – Layer 2 Roaming (Campus Deployment)



Assume that the user now roams over to the AP shown on the right. The traffic moves up the CAPWAP tunnel to the new AP that the user is now associated to, which happens to be hosted by the top controller in the figure. As part of the roam process, the user's PoA and PoP have moved from the bottom controller to the top controller. All of the controllers are grouped together into a common MG, and communicate the user's mobility event. .

User traffic flow now comes up to the top controller, is decapsulated, and moves out to the network. The PoA and PoP were both able to move because all controllers in the MG share a common set of back end Layer 2 VLANs in this example. The user's identity and IP address are valid on all of the controllers. Because it is a Layer 2 MG, all of that information is retained.

As the user roams, mobility messages are exchanged, and the user's entire client context moves from the original controller to the roamed-to controller.



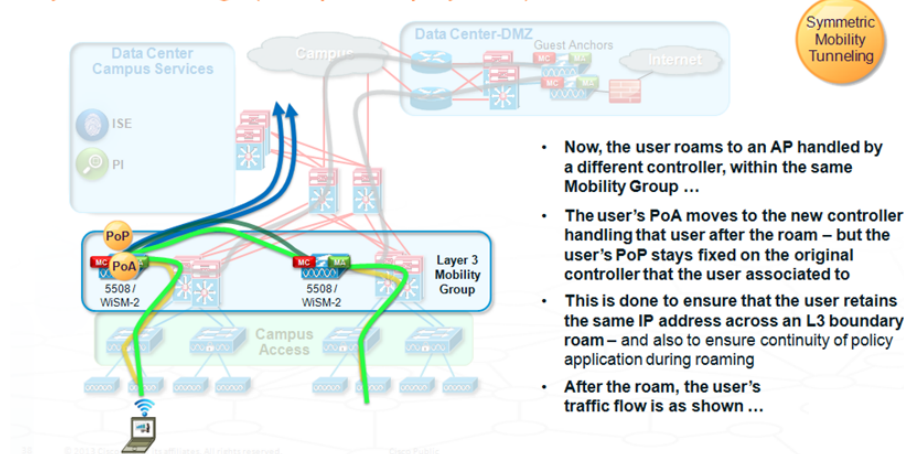
## Layer 3 Roaming

Layer 3 roaming uses a Layer 3 MG. In the following example, two controllers are plugged into two different sets of distribution layer switches, and they are separated by a routed core. The red lines (routed links) in the following figure separate the two controllers. Because the controllers are Layer 3-separated, they cannot share a common set of back end VLANs.

Figure 9. Existing Architecture – Layer 3 Roaming

### Architecture Constructs –

#### Layer 3 Roaming (Campus Deployment)



In this example, the PoP and PoA for the user are initially hosted on the left-hand controller. The user's traffic flows through this controller and out to the network.

Assume now that the user roams from the left-hand controller to the right-hand controller, or from the left-hand AP to the right-hand AP (handled by the left-hand controller and the right-hand controller, respectively). Only the user's PoA moves to the new roamed-to controller. The user's PoP stays on the controller where the user originally associated because the PoP handles the user's IP address. Because the controllers are Layer 3-separated, the user's IP address, and identity are not valid on the right-hand controller. They are valid only on the left-hand side controller because the subnets where the user was assigned apply only to the left-hand side of that network. Because the user's IP address is retained while roaming (along with Layer 3 policies), only the user's PoA is moved with this type of roam – the PoP stays fixed.

The user traffic is then re-tunneled from the PoA controller back to the PoP controller, because all of the controllers involved here have been grouped together in a common MG. This facilitates system-wide roaming via the auto-built full-mesh group of EoIP / CAPWAP tunnels formed between all of the controllers that are grouped together in a common MG. From the perspective of the wired network, the wireless user has not moved – yet the user's wireless connectivity and movement has been accommodated.

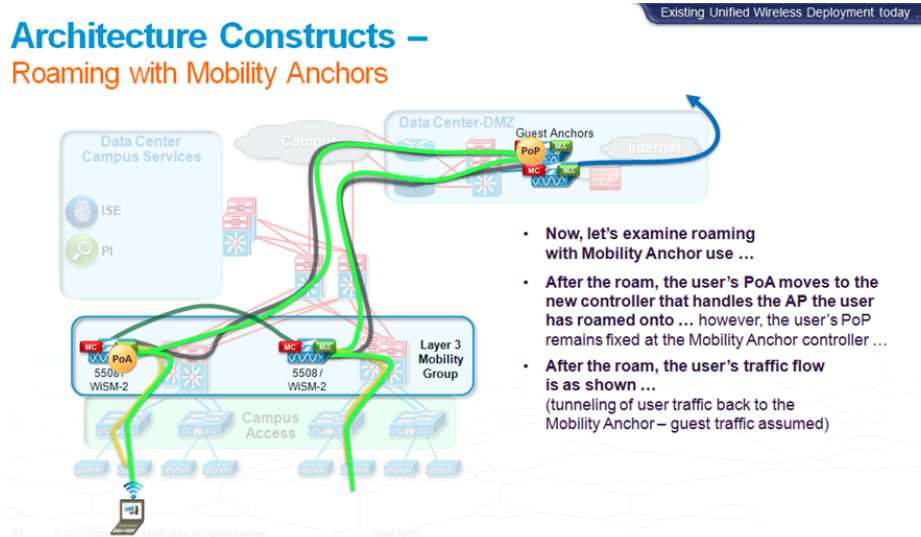
By default, every roaming event in Converged Access appears as this type of Layer 3 roaming event. It is called *symmetric mobility tunneling* because user traffic follows the same path downstream and upstream. A mobility message is exchanged, and the entire database entry is copied (and not transferred as in a Layer 2 roam), as different portions of it are used on different controllers. The left hand controller, WLC 1, is still the anchor for this session. All the user's traffic goes up to WLC 2 and is tunneled back to WLC 1 and out to network. There is no need for the

user to change IP address. The terms *foreign* and *anchor* are also often used to describe the roles within the wireless network for this type of roam. On a Layer 3 roam, the foreign controller is the one that owns the user's PoA, and the anchor controller is the one that own the user's PoP.

## Roaming with Mobility Anchors

Mobility anchors are often used for guest traffic. In the following figure, a user from the guest SSID associates to the left-hand controller, but the user traffic flows up the grey tunnels to the guest anchor controller. That is where the traffic is decapsulated and moved out to through a firewall to the Internet.

Figure 10. Existing Architecture – Mobility Anchors



The guest user's PoP is on the guest anchor controller, because that is where the guest user is seen from the perspective of the wired network. The PoA is on the controller the user has associated to. The AP and the controller the user roamed to handles the user's PoA, but the user's PoP is always up at the guest anchor.

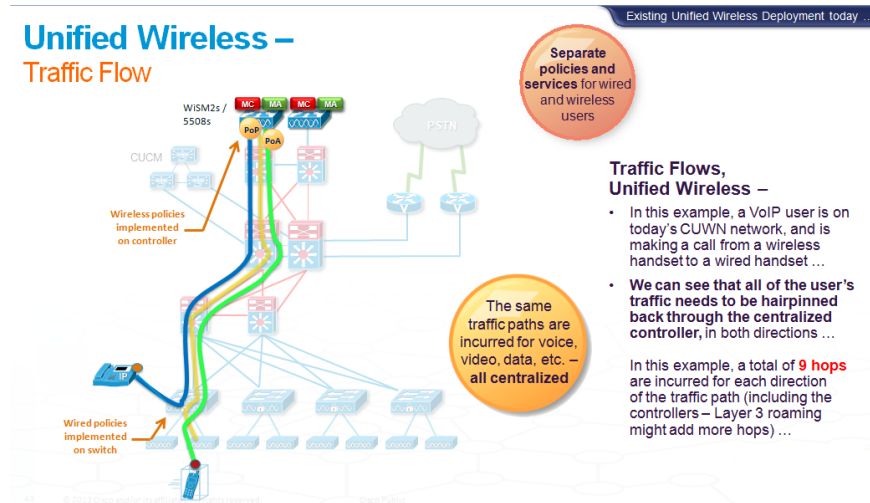
The guest user maintains the same IP address throughout the roaming environment. Before the guest user roams, the user's traffic comes into the PoA controller. It is re-tunneled up to the guest anchor controller and out to the network. This is the desired behavior – the guest user should be dropped off at the Internet edge up at the DMZ.

When a guest user roams, the user's PoA moves to the newly roamed-to AP and associated roamed-to controller. The traffic is then once again tunneled up to the guest anchor and out to the network. From perspective of the wired portion of the network, the user has not moved; however, from a wireless point of view, a mobility event has occurred.

## Traffic Flow

Consider a wired phone plugged into a switch at the bottom of the following figure and a wireless handset that is associated through an AP on the same switch. The PoA and PoP for the wireless handset user are at the controller because the wireless user's traffic flows through the CAPWAP tunnel from the AP up to the controller. At the controller, the user traffic is decapsulated and moved back down, for example to a wired IP phone. All user traffic flows over this path through the switches and controllers as shown. If there are Layer 3 roaming events, there might be additional hops and additional load on the controller.

Figure 11. Existing Architecture – Traffic Flow



All of the voice, video, and data traffic for wireless users is centralized on the controllers. That implies a potential scalability limit with 802.11ac, which requires scaling massive amounts of bandwidth and numbers of wireless clients to the network. Moreover, there are separate policies and services for wired and wireless users. Wired policies, such as ACL and QoS mechanisms reside on the switches; however, the wireless users gets a different set of ACLs and QoS policies applied on the controllers.

The approach described here has worked well to deliver wireless service over the past years; however, to support future options for scaling, services, and performance, the approach needs to be enhanced with more powerful and scalable controllers and new deployment options offered as part of the Converged Access architecture.

## Control Plane and Data Plane

As a last observation before turning to the Converged Access architecture, consider the behavior of the existing architecture with regard to the control plane and data plane. When a user associates and roams in a network, a mobile announcement in the control plane is sent to all of the controllers in the network so that mobile handoff events can occur. This is supported by the full mesh of CAPWAP tunnels that are automatically set up within the MG when the MG is formed.

From a data plane perspective, the wireless network is entirely centralized. Wireless traffic is overlaid on top of the wired network. All traffic is sent back through the local controller in local mode, using CAPWAP encapsulation. The encapsulated user traffic isn't visible as it crosses the wired network, as it is inside a CAPWAP tunnel. For example, Netflow data for this end-user traffic in a CUWN centralized deployment model is not available until the traffic emerges from the CAPWAP tunnel at the centralized controller.

## Next Generation Cisco Converged Access Architecture

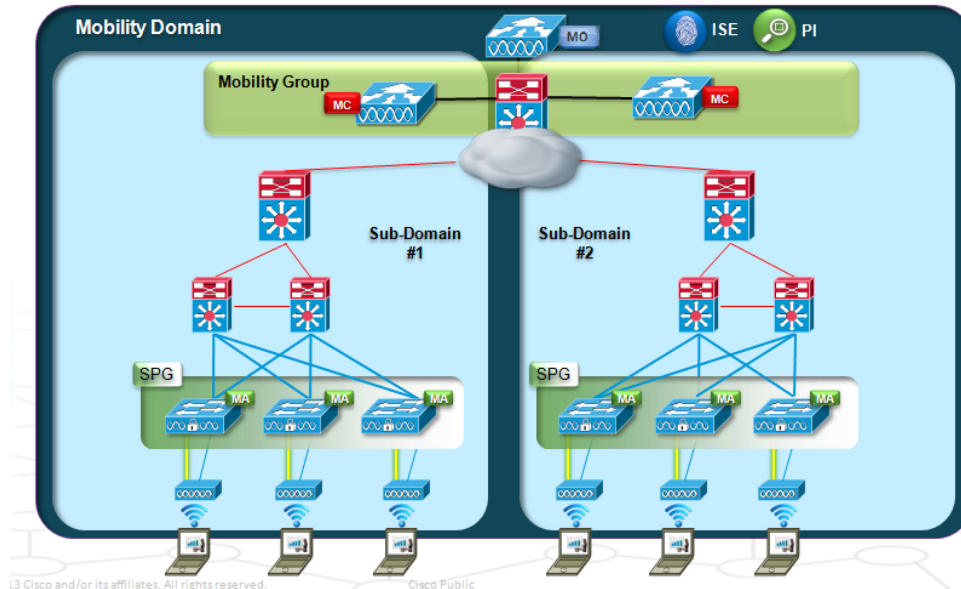
The Converged Access approach extends the existing CUWN architecture to provide additional deployment options for increased scale and performance. Converged Access optimizes the architecture for greater scale, greater wired/wireless traffic visibility, and integrated wired/wireless common services, through a distributed deployment model in which the Catalyst 3850, acting as the ingress access switch, terminates the CAPWAP tunnel from the AP, and thereby becomes a full partner in the mobility roaming domain.

In the Converged Access approach, a controller operating as an MC provides system-wide control plane functionality in the network. The controller can be a software-upgraded existing WLC 5508 or WiSM-2 controller, or it can be a newer WLC 5760. To address scale and performance and offer a unified wired and wireless experience, data plane termination functionality (CAPWAP termination capability) moves into the Catalyst 3850 in the Converged Access deployment mode. The Catalyst 3850 operates as an MA (and can optionally also operate as an MC for smaller deployments).

At a high level, Converged Access takes the capability that runs on a controller, splits it into control plane functions that run on a centralized controller, and provides the option to terminate data plane functionality (CAPWAP) on the Catalyst 3850.

To see how Converged Access works, consider a typical campus network with wireless users coming in the bottom through APs, as shown in the following figure. The CAPWAP tunnels, shown as yellow lines coming from the APs, do not go back to a central controller. Instead they terminate on the access switch, a Catalyst 3850. The access switch acts as an MA, and it also handles roaming. It is a full partner in the wireless MD.

**Figure 12. Converged Access Architecture – Deployment Overview**



The switch peer group (SPG), new in Converged Access, localizes roaming within the environment by introducing a layer of hierarchy into MG design and deployment. An SPG is typically built around a building, or floors within a building – areas where users commonly roam. All the switches within the SPG form a full mesh of CAPWAP tunnels with each other. As with an MG, the full mesh of CAPWAP tunnels is built automatically within the SPG,

at the time that the SPG is formed. By introducing a sub-hierarchy within a larger MG infrastructure, SPGs enhance scalability and performance.

There is always a need in any wireless deployment for an MC, which provides the control plane functionality for system-wide capabilities such as inter-controller roaming and RRM. With Converged Access, the MC function can operate on discrete controllers, such as software-upgraded WLC 5508 or WiSM-2, or on new WLC 5760s added to the network. For smaller deployments, the MC function can also operate on the Catalyst 3850 itself.

In either case, each MC is in charge of one or more downstream SPGs, depending on how the network is scaled. All of the MAs point up to the MC as the control plane. As with the existing architecture, MCs can be grouped together in an MG. There is also an option to build a larger MD.

Very large deployments can optionally use a *Mobility Oracle* (MO), which is an existing controller with MO functionality enabled. The operation of an MO is somewhat analogous to a route reflector in BGP. In a BGP deployment, instead of having a full mesh of BGP peering relationships among all BGP speakers, the BGP speakers can be pointed at a route reflector, which simplifies the deployment and allows it to more readily scale up without introducing undue complexity. Similarly, if a Converged Access deployment has many MCs, instead of fully meshing them all within a common MG, inter-controller messaging exchange can be optimized by pointing all of the controllers at an MO. The MO serves as a “single source of truth” for roaming within the MD. An MO is needed in only the most large-scale deployments, where many controllers operating as MCs are in use. As in the existing CUWN architecture, ISE and PI serve as the authentication and management points of the whole network.

To summarize, Converged Access still uses controllers, but some of the controller functions, such as the MA function, are moved to the switches, and a new construct, the SPG, localizes and optimizes roaming at the switches. An MO can be optionally added for the largest scale deployments.

## **Mobility Architecture**

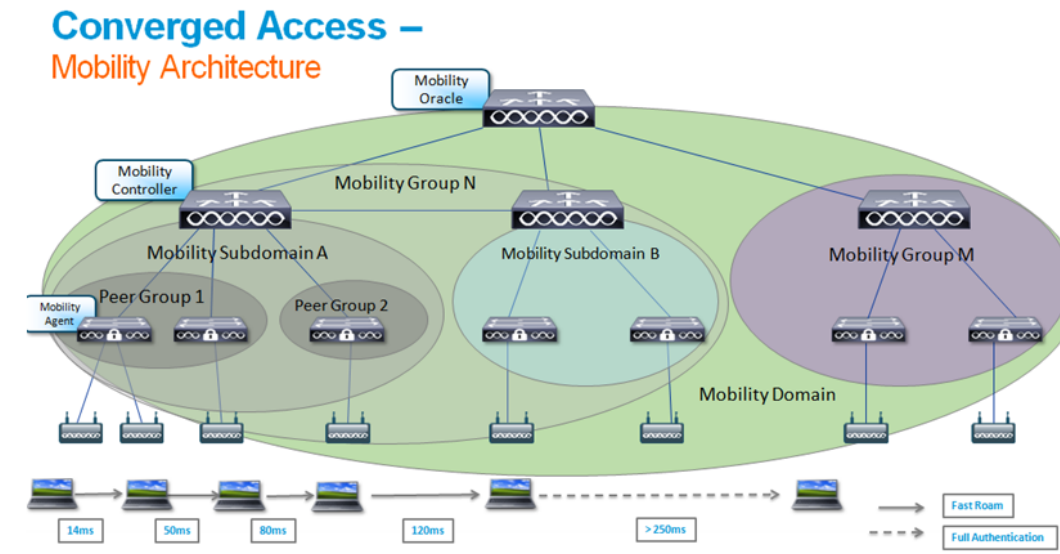
Now for a deeper look at the Converged Access mobility architecture. Within a single Catalyst 3850 switch or switch stack, roaming times are typically less than 20 ms (shown in the following figure as 14 ms). This is the time required to roam from AP to AP on the same switch.

Roaming to another switch in the same SPG typically increases the roaming time to 50 ms because messaging exchange is required between the switches. The following figure shows a separate MC (for smaller deployments, the controller could be running as an MC function on the switch). This leads to the concept of a *mobility subdomain*, which operates as a peer group of multiple SPGs (the figure shows it handling two SPGs). With mobility subdomains, multiple roaming scenarios can exist.

One scenario involves the user roaming from one Catalyst 3850 to another, with those two switches located in the same SPG. In this case, the target roaming time is 50 ms, which is optimized due to the smaller size and local proximity of the switches within the common SPG.

Another scenario involves a user who is roaming from a switch that is in peer group number one, for example, on one set of floors within a building, to another peer group on a different set of floors within the building. The target roam time in this case is somewhat longer (approximately 80 ms), as messaging is required between the roamed-from switch, the roamed-to switch, and the MC. However, 80 ms is still a rapid roam time, and works well for many deployments.

Figure 13. Converged Access – Mobility Architecture



Extending the model, a larger MG can be built out between MCs. In the previous figure, roaming from subdomain A to subdomain B takes approximately 120 ms because messaging now needs to move between multiple switches and multiple controllers to handle the roaming event. These times are still rapid, and are considered to be fast roaming events because the user doesn't have to go through a full reauthentication (the security keys are cached and distributed at the MG level).

If two completely different MGs are set up under the control of an MO and the user roams across these two separate MGs, user-level re-authentication takes place, which can increase the roaming time to more than 250 milliseconds. The exact roaming time depends on needed interactions with, and the performance of, the AAA server in this case.

### Physical Entities

The Converged Access architecture involves physical entities and logical entities. The physical entities are items that have a definite physical presence (as opposed to being only logical entities). Physical entities within the Converged Access deployment mode include the MA, MC, and MO.

- The MA terminates the CAPWAP tunnel from the directly attached AP and maintains the client database for the user. One MA function runs across the stack and provides an interface to the MC. It is the first level within the physical hierarchy.
- The MC manages mobility in and across sub-domains. As with the MA, the MC can run on a switch for smaller deployments, whereas a discrete controller platform is recommended for larger deployments. An MC is always required in a wireless deployment. System-wide functions such as RRM and roaming are handled by the MC.

The Catalyst 3850 usually runs the MA functionality and can also potentially operate as an MC for smaller deployments. If the deployment is fewer than 50 APs or 2,000 clients, the MA and MC can both run together on the Catalyst 3850 stack. In these scenarios one Catalyst 3850 can serve as the combined MA/MC for multiple other switches that act only as MAs.

- The optional MO allows for more scalable mobility across a very large domain.

## Logical Entities

Logical entities are virtual groupings that are built within the network, and include MGs, MDs, and SPGs.

- An MG is the first entity used to build a wireless deployment. It is a group of controllers operating in the network to allow for system-wide functions such as fast roaming, and RRM.
- MDs build up to a larger scale as a grouping of MCs, with mobility lists to allow very large scale deployments.
- The SPG is designed to localize roaming within the distribution block in the network, typically within a building or floors within a building.

The next sections provide a more detailed discussion of how roaming events actually work in a Converged Access deployment.

## Scalability

The following table is a condensed summary of scaling information for the Converged Access solution.

**Table 1. Scaling for Converged Access**

Scalability	Catalyst 3850 as MC	WLC 5760	WLC 5508	WiSM-2
Max number of MCs in an MD	8	72	72	<b>72</b>
Max number of MCs in an MG	8	24	24	<b>24</b>
Max number of MAs in a Sub-domain (per MC)	16	350	350	<b>350</b>
Max number of SPGs in a Mobility Sub-Domain (per MC)	8	24	24	<b>24</b>
Max number of MAs in a SPG	16	64	64	<b>64</b>
Max number of WLANs	<b>64</b>	<b>512</b>	<b>512</b>	<b>512</b>

The first Catalyst 3850 as MC column examines the scalability of the Converged Access solution if the Catalyst 3850 hosts the MC (as well as MA) function. The first cell in this column shows the maximum number of MCs supported in a MD. The maximum number of switch stacks operating as MCs in a common mobility roaming domain is 8, but each of those switch stacks can handle a downstream switch stack operating as an MA (itself plus up to 15 others – third row in the Catalyst 3850 as MC column).

As shown in the table, fairly large network deployments can be built around the Catalyst 3850 operating as an MC. Using one or more discrete controllers as MCs within a deployment allows for even greater scalability, and has the added advantage that these scaling limits can be realized with no future change in design.

While the switch-as-MC deployment model for Converged Access works well for smaller deployments, mid-sized to larger deployments are typically better served by employing one or more discrete controllers as MCs, with Catalyst 3850s deployed as MAs. Both methods provide the Converged Access advantages of greater traffic growth, greater traffic visibility, and increased traffic control through converged wired / wireless policies. They do so at varying levels of CapEx expenditure and future growth potential, which can readily be matched to the specifics of any individual network to “right-size” the deployment for the needs at hand.



## Traffic Flows and Roaming

To see how traffic flows and roaming work in the Converged Access solution, consider user traffic coming into a Catalyst 3850 operating as the PoA switch (where the CAPWAP tunnel is terminated). The user's traffic is decapsulated at the PoA and moved out to the network. Services such as Netflow and QoS can both be applied to traffic at this point.

When users first associate into the network, there is a pairwise master key push and a mobile announce for the users as they enter the wireless environment. The user information is sent out so that both the MC and MO can learn about it, in addition to the MA. Both the PoP and PoA are initially located at this initial ingress switch, ARP messages are sent to the network to announce the user, and the user's traffic can flow.


The format selected for the mobility and guest tunnels is CAPWAP with 802.3 encapsulated inside. CUWN deployments need to support this format as well for hybrid deployments (CUWN support for new mobility packet format was added in the AireOS 7.3.112.0 release on the existing WLC 5508 and WiSM-2 controllers).

The following figure shows the packet format for mobility and UDP ports.

UDP ports:

- Control plane 16666 (always DTLS encrypted)
- Data plane 16667 (not encrypted and cannot be encrypted at FCS)
- MO listens on UDP port 16668

**Figure 14. Mobility Packet Format**



The next sections discuss the how roaming works at different levels within the Converged Access architecture, beginning with simple branch roaming within a single switch and moving on to more complex roaming examples within and between SPGs. In a well-designed network, as the types of roaming grow in complexity, there is a decreasing probability that many of these types of roam events will occur. If SPGs are set up around buildings or floors within a building, the more complex types of roaming happen only when users move between buildings or between SPGs, which is generally less frequent than users moving around within a building or floor.

In addition, it is important to preface the following roaming discussion by noting that some events that might be thought of as roams might not in fact generate a roaming event. For example, closing up a laptop or tablet and then moving within a building typically drops the device off the network after a short period (immediately, in the case of the laptop). Unless the device wakes up quickly and re-joins the network, moving the device in this state is not a roaming event. Instead the device might simply re-join the network as a new device as it comes back online. While this is not a roam, it can represent a major portion of actual user/device behavior. True roams typically involve devices that are always on when moving around the network. Thus, when reviewing the following discussion of roaming, it is important to place into context as to whether the users on the network are actually incurring roams, device moves, or a combination of the two.



## Basic Branch Roaming

Layer 2 roaming can optionally occur when the switches involved share a common set of VLANs on the back end, whereas Layer 3 roaming must occur when they do not share a common set of VLANs. By default, all roaming events in Converged Access are handled as Layer 3 events. It is possible to handle switch roaming at Layer 2, but this is not the default behavior.

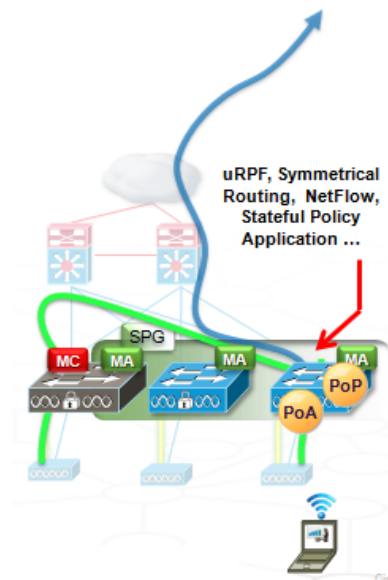
Consider the simplest possible case for roaming – within an individual switch stack. The user associates into the network through an AP on the switch, and has a PoP and PoA at that location. Traffic flow is handled at the switch without the need to involve another controller. This is highly optimized – providing excellent scalability, comprehensive traffic visibility and control, and common policy definition and enforcement for both wired and wireless traffic flows at the access switch – all of the benefits that accrue to a Converged Access deployment.

Now assume that the user roams to another AP on the same switch, for example, within a floor in a building. The PoA moves to a new switch port on the same switch. The PoP doesn't move, the user's IP address doesn't change, and the roam and the resulting traffic flow are fast, simple, and highly optimized.

## Intra-SPG Branch Roaming

Now consider how roaming operates with a SPG. The following figure shows an SPG where one of the switches acts as both the MA and MC (possible because this is a relatively small deployment).

**Figure 15. Roaming Within an SPG**



Assume that a user associates through the right hand side switch, and the user's PoA and PoP is located on that switch. Traffic enters, is terminated on the switch, is decapsulated, and moves out to network for simple, highly optimized traffic flow.

Now assume that the user roams over to a switch on the left side within the SPG. (In the figure, that switch is also the MC, but that is not relevant for the roaming event.) Because the user's IP address came from the right side switch, it might only be valid on that switch, depending on the wired infrastructure. To maintain the IP address, the

traffic is re-tunneled back to the right hand side switch (using the full mesh of CAPWAP tunnels that was auto-generated at the formation of common SPG that both of these switches are part of), and then moves out to the network. This behavior is similar to a Layer 3 roaming event in a CUWN network.

The SPG has a full mesh of CAPWAP tunnels that is built between all the switches, so the traffic is simply re-tunneled back to the origin switch, and then it is moved out to the network. There are several benefits to this approach, including the ability to do unicast RPF check, symmetrical routing, NetFlow, and stateful policy application. All of these are handled at PoP, which has not moved.

Whether or not users have roamed across a Layer 3 boundary, their traffic is always (by default) taken back their PoP for policy application. It is possible to modify the configuration to allow for a Layer 2 roam if the two switches are Layer 2-adjacent – however, this is not the default behavior, and it has implications which are discussed later in this document.

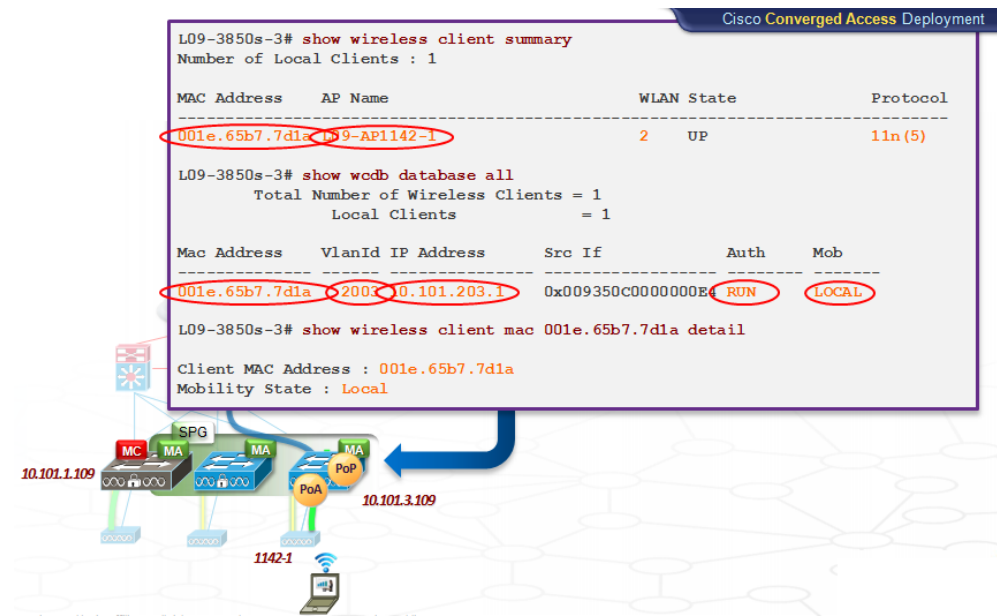
Messaging works as follows for intra-SPG roaming. On client roams, the roamed-to (foreign) switch does a Mobile Announce after associating the client, and the switch the user has roamed from (the anchor) does a handoff message over to the foreign switch where the user has roamed to. The handoff complete message goes up to the MC and is flooded out to all the other switches within the common SPG, so that all devices know the user's roaming destination. The handoff complete messages that come back from the MC are kept to a minimum. By minimizing the number of messages, roaming within an SPG can work much faster.

The following section describes intra-SPG roaming in greater detail. It can be skipped if this level of additional detail on the specifics of inter-switch roaming within an SPG is not needed.

### **Intra-SPG Roaming – Details**

This section examines the details of intra-SPG roaming. In the following figure, the user associates to the AP labeled as AP1142-1. The user's PoA and PoP are on the switch to which the AP they have associated through is connected. The **show wireless client summary** command on this switch shows that the traffic from the user in this example has a MAC address that ends in 7d1a, is associated to AP 1142-1 on VLAN 2003, and has IP address 10.101.203.1. The user is fully authorized and authenticated into the network, and the mobility state is local, which indicates that the PoA and PoP are in the same place (the user hasn't roamed yet).

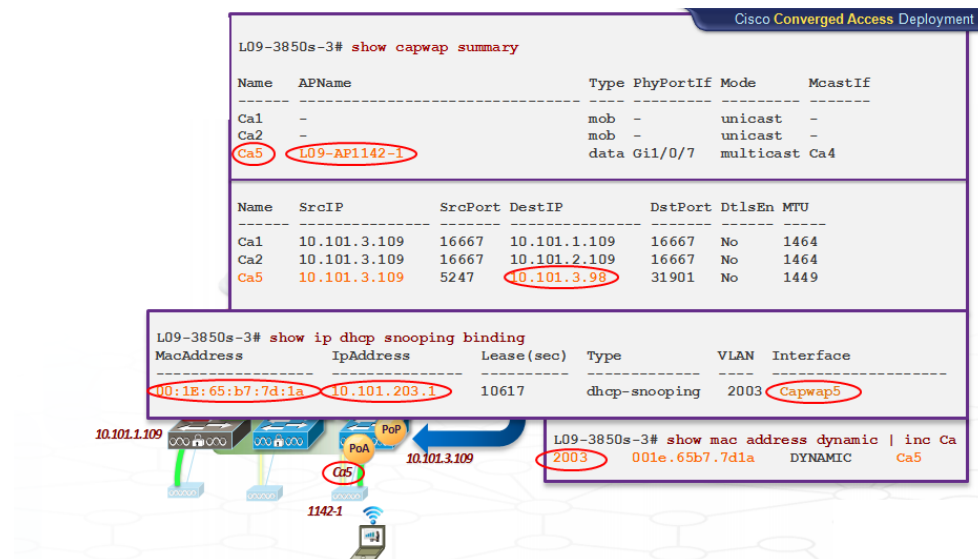
Figure 16. Intra-SPG Roaming Example – Before Roaming



DHCP snooping is normally an optional function for a wired network, but in the Converged Access deployment it is required to have DHCP snooping turned on. DHCP snooping is how the network keeps track of wireless users as they roam. The following figure shows the DHCP snooping binding table, with the user with IP address 10.101.203.1 (MAC address ending in 7d1a) reachable through the Capwap5 interface. The Capwap5 interface is the tunnel from the switch to the AP in this example, because that is where the user is at this time. To verify that, the **show capwap summary** command indicates that Capwap5 leads to 1142-1, the IP address of the AP is 10.101.3.98, and the switch IP address is 10.101.3.109.

The output of the **show mac address dynamic** command shows that the user is on VLAN 2003 and has not yet roamed.

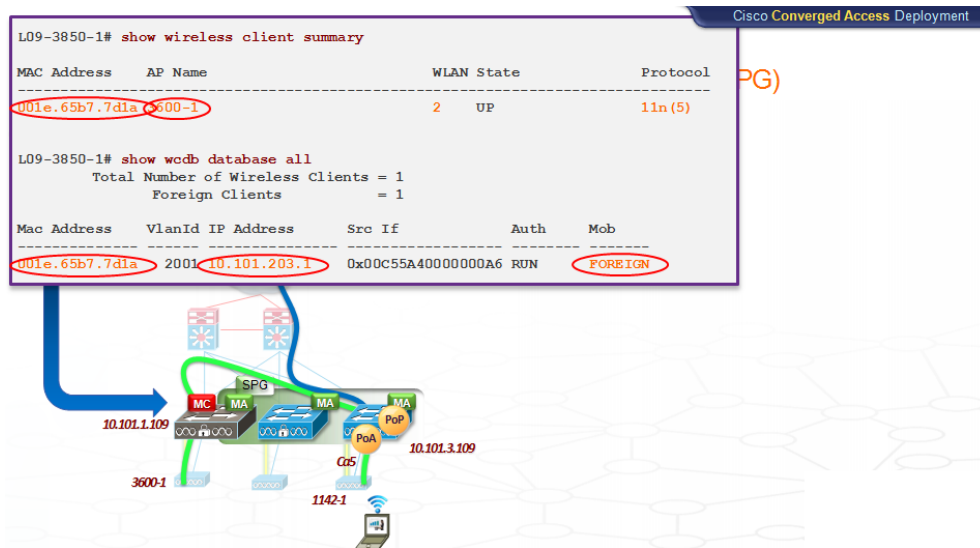
Figure 17. Intra-SPG Roaming Example – Before Roaming with DHCP Snooping



Now assume that the user roams to an AP labeled 3600-1, which is managed by switch 10.101.1.109, a switch in the same SPG as the 10.101.3.109 switch. (In the example, the 10.101.1.109 switch is the MC for this SPG, but that is not relevant for this particular roam event.) The user's PoA has moved to the roamed-to (foreign) switch, but the PoP has stayed back on the origin (anchor) switch.

The **show wireless client summary** command shows the roamed-to switch with the 7d1a user reachable through AP 3600-1. The 7d1a user has the same IP address, but the mobility state is foreign on the roamed-to switch because the user has roamed. A foreign mobility state means that the switch from which the command is executed owns the PoA, but does not own the PoP, which remained on the origin (roamed-from, or anchor) switch.

**Figure 18. Intra-SPG Roaming Example - Roaming**



To continue this example, the user's traffic comes into the PoA switch and is re-tunneled across the SPG back to the PoP switch, and then is decapsulated and moves out to the network. The behavior is that of a Layer 3 roam.

The **show wireless client** for user MAC address 7d1a on the roamed-to switch shows that this switch is in the foreign state for this user, indicating that it owns this user's PoA. The mobility anchor IP address is also displayed, indicating where the user started, namely on the 10.101.3.109 switch.

The **show mac address dynamic** command on this roamed-to switch shows that the 7d1a user is reachable through Ca3, the Capwap3 tunnel. Capwap 3 is the CAPWAP tunnel down to the 3600-1 AP, because that is where the user has roamed to. The **show capwap summary** command shows that Capwap3 leads to 3600-1 with IP address 10.101.1.98.

The **show mac address** command output shows that the VLAN show for this user is VLAN 4095. However, VLAN 4095 is not a real VLAN; it is used as a flag, or a trigger, in the MAC address table. It indicates that traffic coming from this user is not to be switched out locally, as if it were coming from wired port on a switch and switched out locally to the network, but instead that the traffic from this roamed user needs special handling. The special handling involved is that when the traffic from this roamed user comes in on the Capwap3 tunnel from the attached AP, it is to be re-tunneled over the network back to the 10.101.3.109 switch, and from there the user's traffic is to be decapsulated and moved out to the network.

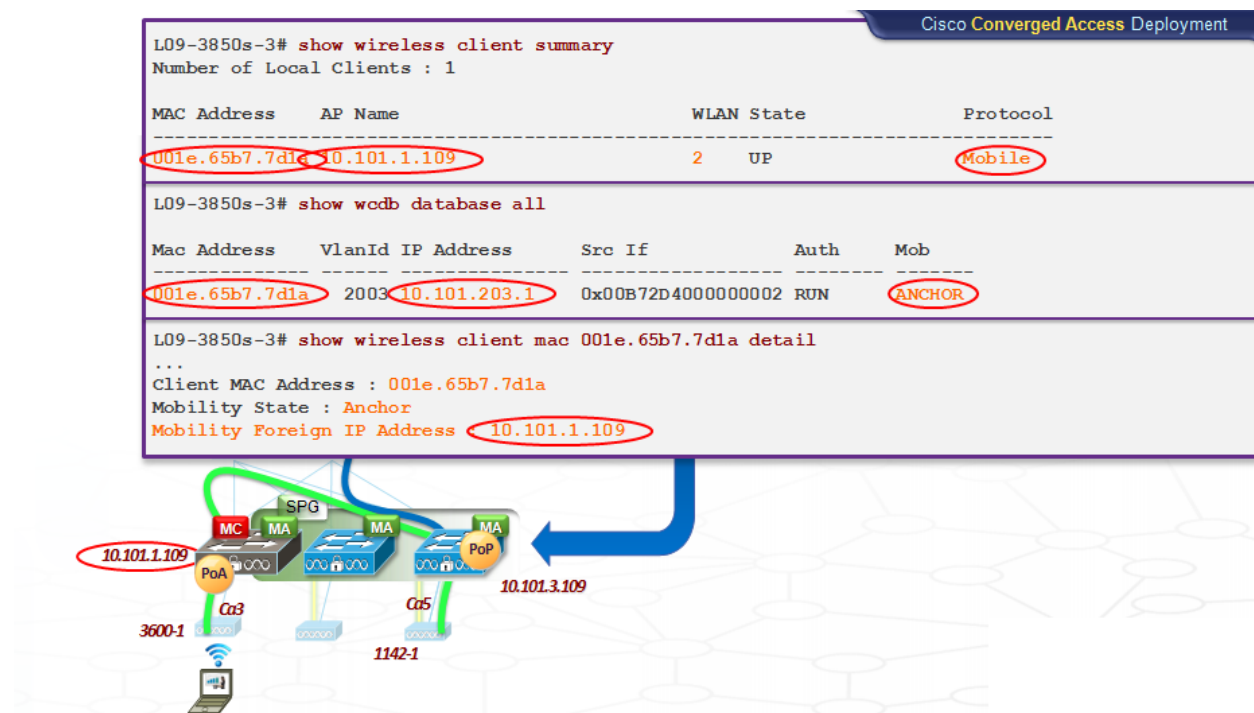
Figure 19. Intra-SPG Roaming Example – Details



Detailed information about this roamed user is also available on the roamed-from switch. The 7d1a user is still known on the roamed-from switch, and the user state is shown as mobile. The **show wireless client summary** command indicates where the user has roamed to. The AP the user is shown as being connected to is actually the 10.101.1.109 switch that the user has roamed to. The roamed-from (anchor) switch knows the roamed-to (foreign) switch and its IP address, which it learned during the roaming process and the associated mobility messaging that was exchanged between the switches when this user roamed.

The **show wcdb database all** command shows the user and IP address, which was obtained from the IP address pool available at the roamed-from (anchor) switch, while the **show wireless client mac address detail** command shows information about the switch the user has roamed to (the foreign switch).

Figure 20. Intra-SPG Roaming Example – Additional Details

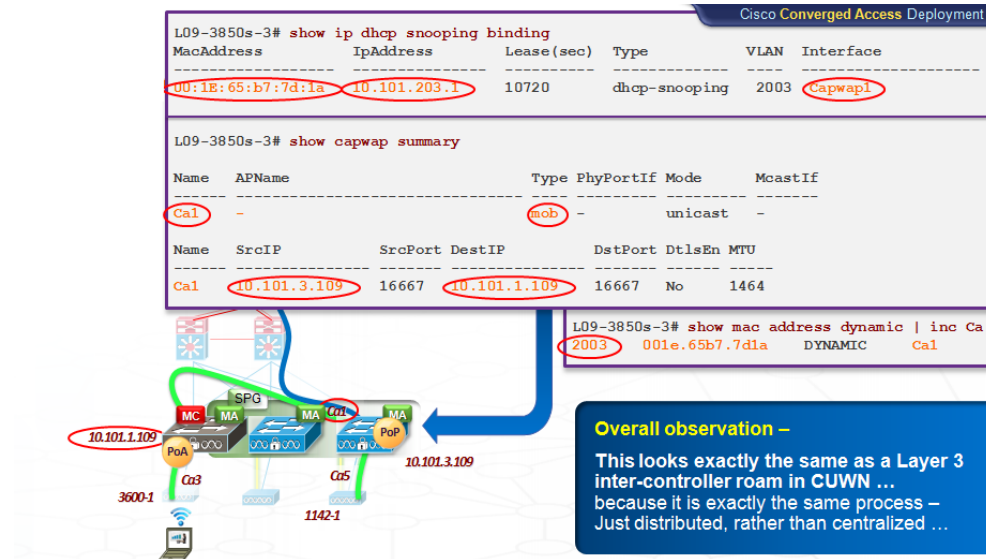


Additional information about the user can also be reviewed via DHCP snooping. The **show ip dhcp snooping binding** command on the roamed-from switch shows that the user is still known here as reachable through the Capwap1 tunnel. Capwap 1 is a mobility tunnel that was built automatically when the SPG was formed. The **show capwap summary** command shows that the user is reachable over the Capwap1 mobility tunnel from the 10.101.3.109 switch where the user started (anchor) to the 10.101.1.109 roamed-to (foreign) switch.

The **show mac address** command on the PoP switch shows that the user traffic is still switched out locally on VLAN 2003, as it was prior to roaming. The behavior looks exactly like a Layer 3 inter-controller roam in CUWN, because it is the same procedure. The primary difference is that the access switch, which previously would have just passed through the CAPWAP tunnel, is now terminating the CAPWAP tunnel, so the switch is a full partner in the MD, and has all the level of visibility and control associated with acting as an MA for the user's traffic.

Because the CAPWAP tunnel is terminated locally, it is possible to scale to a much greater extent than with a centralized CUWN approach. A Converged Access deployment also provides better traffic visibility, and the benefits of applying a common set of services and policies to both wired and wireless traffic.

**Figure 21. Intra-SPG Roaming Example – Additional Details**




### Intra-SPG Branch Roaming with a Separate MC

Moving beyond the simple intra-SPG roaming discussed in the previous section, consider intra-SPG roaming, where the switches serve only the MA function and a separate controller is used for the MC. Roaming is still within the SPG, but the SPG in this example now consists of switches that are operating as MAs only.

The following figure shows this type of configuration with a WLC 5760 operating as the controller.

**Figure 22. Intra-SPG Configuration with Separate MC**



The purple lines in the figure shows the CAPWAP tunnels operating as control sessions (which can also potentially carry user data for some roaming scenario types) that run from the MAs back up to the MC. As with the previous intra-SPG roaming example, assume that the user has roamed from the left hand to the right hand switch. The PoA moves to the roamed-to MA, while the PoP stays on the roamed-from MA (on the roamed-from switch).

The user's traffic comes into the PoA switch and moves back directly to its PoP switch and out to the network. It does not go across those connections up to the MC. The MC is involved here from a control plane perspective, but from a data plane perspective all the traffic stays localized below the distribution layer. From a data plane and scalability point of view, therefore, the MC is not involved in this intra-SPG roam (which also tends to be the most common type of roam in most deployments, given the recommendation that SPGs should be constructed around buildings, floors within a building, or other areas where users may commonly be expected to roam).

Now consider traffic flows in this type of deployment. The Converged Access approach differs from the earlier CUWN approach, because the CAPWAP tunnels from the attached APs are now terminated directly on the Catalyst 3850. Because the CAPWAP tunnels are now terminating on the switches, the MA function has moved from the discrete controllers, where it previously lived, down onto the switches themselves. The PoP and PoA are located on the switch itself, whereas previously they lived on the controller in a CUWN deployment (because all the traffic was centralized through the controller in that deployment mode). This makes the Catalyst 3850, operating as an MA, a full partner within the overall MD for roaming and other wireless functionality.

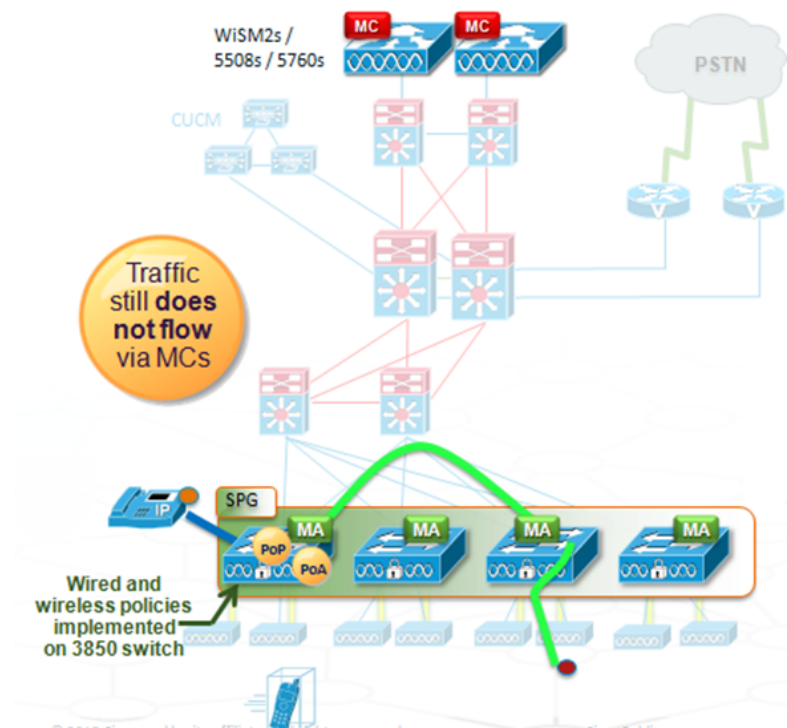
In this example, the user traffic flows into the switch, where it is decapsulated and moved out to the wired phone that is attached to the switch. The traffic is highly optimized and does not need to move from the switch. Where previously there were nine switch hops in this same deployment with CUWN in this example, the traffic now flows within that same switch, with full visibility to, and control of, that traffic from both a wired and wireless perspective.

Because the traffic does not need to flow through separate MCs, performance increases, and converged policies can be applied for wired and wireless in the same place and in a similar fashion on the Catalyst 3850. The traffic has converged to the same place in the network. This avoidance of traffic flow using the MC also eliminates the MC as a bottleneck in the deployment. This is an important consideration as wireless traffic scales up in the future with 802.11ac deployments and the ever-increasing number of devices associating into the wireless network.

The following figure shows what happens when the user roams from the first switch to the third switch in this example, with both switches in a common SPG. The PoA for the user involved moves, while the PoP for this user stays fixed. As in the previous roaming example, the traffic is tunneled back from the user's PoA (foreign) switch, to the user's PoP (anchor) switch. Again the traffic flow is highly optimized, as it needs to cross only three switch hops in this example, as opposed to the nine switch hops in the example given previously. There is still full visibility to, and control of, the traffic back on the PoP switch, and wired and wireless policies can be applied in the same place (at this anchor switch for the user). Because it is still an intra-SPG roam, traffic doesn't need to flow through the MCs, and thus performance and scalability are increased, and policies are converged at the same place in the network.



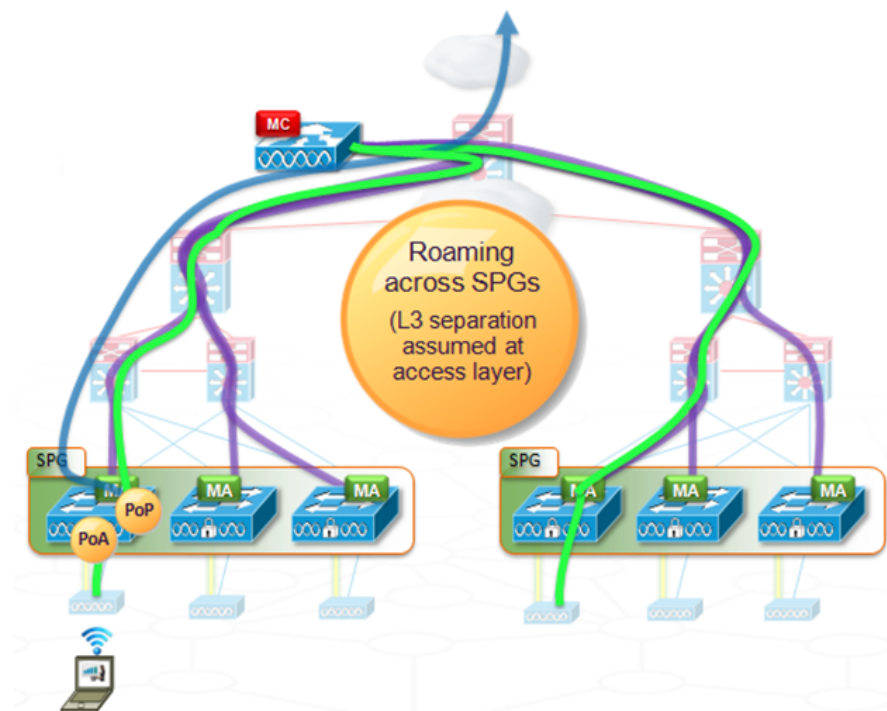
**Figure 23. Roaming in an Intra-SPG Configuration with Separate MC**



### **Roaming Across SPGs Under Control of the Same MC**

The following figure shows different SPGs that are both under control of the same MC. In this example, the user roams from the left hand MA to the right hand MA. In this case, the two MAs are located within separate SPGs (in the previous example, both MAs were within the same SPG). Here also, both SPGs are under the control of the same MC.

**Figure 24. Roaming Across SPGs Under Control of the Same MC**



Recall that each SPG is a full mesh of CAPWAP tunnels formed automatically between all the MAs within the SPG. These CAPWAP tunnels between the MAs in the same SPG are built automatically; however, there is no direct connection (or meshing of tunnels) between the left side and right side SPGs in the deployment. The MC connects the two separate SPGs within this deployment example, thus allowing for seamless roaming and wireless services across both SPGs.

As in the previous examples, the purple lines show the CAPWAP tunnels that are built between all the switches and their MC. Those are mostly used for control plane purposes, but in this case where the user is moving from one SPG to another, the same CAPWAP tunnels can also carry data during inter-SPG roaming.

The user traffic comes into the roamed-to switch where it terminates and is re-tunneled back up to the MC. At this point, the MC performs a mobility tunnel endpoint (MTE) function, where it terminates the CAPWAP traffic and re-tunnels it back down to the PoP switch. This is where it is decapsulated and moved out to the network. From the wired network perspective, even on a more complex roam like this, the wireless user has remained fixed, and has not moved. The roamed wireless user retains the same IP addresses, and the same set of policies, as the distributed wireless network handles the roaming event.

Examining the roaming event in more detail, when the user roams from SPG A to SPG B, a mobility announcement is made from SPG B out to the MC and back down to the anchor switch. A handoff happens directly from the anchor switch to the foreign switch as an optimization of the control plane messaging. When the handoff is complete, a notification is sent to all the other MAs that are part of the SPG that the user has roamed into. A station left message is flooded into the origin SPG, and acknowledged from the MC to complete the roam. While this roaming scenario is certainly possible, it is less likely than an intra-SPG roam in most deployments. Recall that some device roams are actually moves – and device movement between two SPGs (for example, between two buildings) might be a device move rather than a roam.

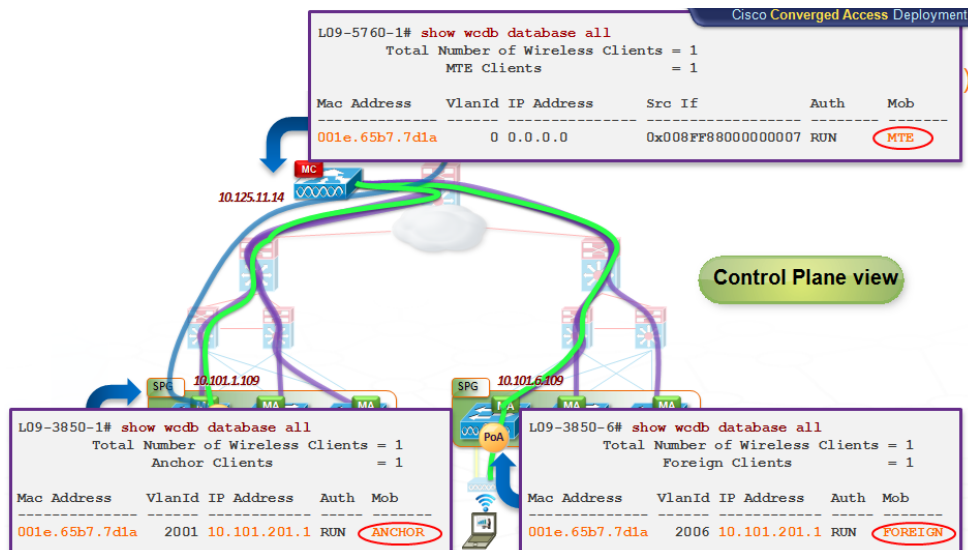
The following section describes roaming across SPGs under control of a single MC in greater detail. It can be skipped if this level of additional detail concerning inter-SPG roaming is not required.

### Inter-SPG Roaming Under Control of the Same MC – Details

This section examines the details of inter-SPG roaming from a control and data plane perspective. The roamed user started from the left side (switch 10.101.1.109) in this example, and has roamed to the right side (switch 10.101.6.109). In this case, both of the switches involved are located in different SPGs, under the control of a single MC. The controller acting as the MC for these two SPGs has the IP address 10.125.11.14.

The switch the user roamed to is in the foreign state for this user, because that switch owns the user’s PoA, but not the user’s PoP. The switch the user roamed from is in the anchor state for this user, because that switch owns the user’s PoP. However, these two switches are located in different SPGs, and thus have no direct CAPWAP peering relationship to each other (as they do to switches located within their own SPG, with its full mesh of CAPWAP inter-switch tunnels). The MC on top, which has a CAPWAP connection to all of the switches in both SPGs, completes the paths for the roamed user, and owns the MTE state for this user. MTE is essentially a functional state within the MC that “glues together” the two data plane paths for this user, from the switch in the roamed-to SPG and the switch in the roamed-from SPG. This MTE function for this user operates at the MC, because the MC is the common network element that both switches in SPGs share. This is an example of the normally control-plane-oriented CAPWAP tunnels from MA to MC also carrying user data for this particular type of inter-SPG roam event (purple lines in the figure).

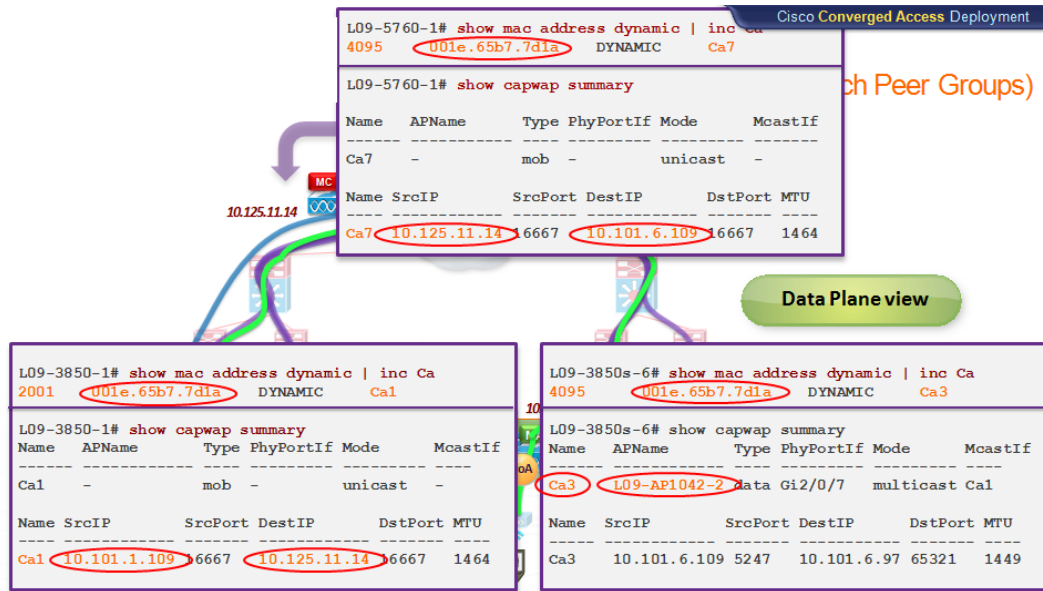
Figure 25. Inter-SPG Branch Roaming with an MC



From a data plane perspective, traffic traveling from the network towards the user whose MAC address ends in 7d1a moves from the 10.101.1.109 MA switch (anchor), to the 10.125.11.14 controller acting as the MC (MTE). If a packet comes into the anchor switch and needs to go to the roamed-to user, it follows this path. The MC output shows traffic flowing from the MC acting as the MTE (10.125.11.14), down to the roamed-to (foreign) switch (10.101.6.109). This roamed-to switch shows that traffic from the 7d1a user then flows down the Capwap3 tunnel to the AP named AP1042-1 where the user is now associated, after which the user’s traffic is then sent out via RF to

the roaming wireless client. Traffic coming from the same roaming wireless user flows symmetrically, back to the anchor switch, where the user's policies and network control capabilities are maintained.

**Figure 26. Inter-SPG Branch Roaming with an MC – Data Plane**

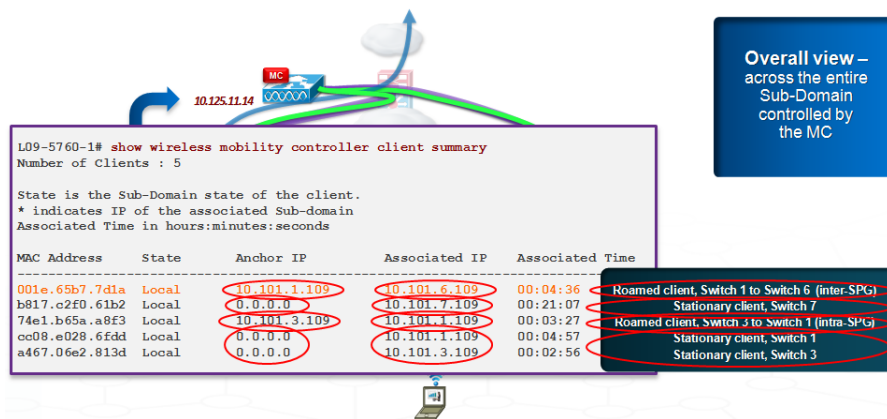


Although a view into all of the roaming events within the network can be obtained by visiting each MA and MC within the network, it is not necessary to do so to view the roaming state of users and devices. In other words, it isn't necessary to go to all of the MAs in the network to know what users exist and where they may have roamed to. The **show wireless mobility controller client summary** command on the MC suffices to show all of the roamed and non-roamed users across the entire mobility sub-domain, or sub-domains, that are under that MC's control.

The output from this command at that MC in this example shows multiple clients in various mobility states. The first is the 7d1a client that started out on the 10.101.1.109 switch and roamed to the 10.101.6.109 switch three or four minutes previously. In this example topology, this is an inter-SPG roam. There is also another roamed client, with MAC address ending in a8f3, which started on the 10.101.3.109 switch and roamed to the 10.101.1.109 switch, which is an intra-SPG roam.

Finally, users are also shown on the 10.101.7.109, 10.101.1.109, and 10.101.3.109 switches. The anchor IP for these users is shown as all zeros, which indicates that these are stationary clients. These clients have associated into the network, and are operational, but haven't yet roamed.

Figure 27. Intra-SPG Branch Roaming with an MC – Roamed and Non-Roamed Clients

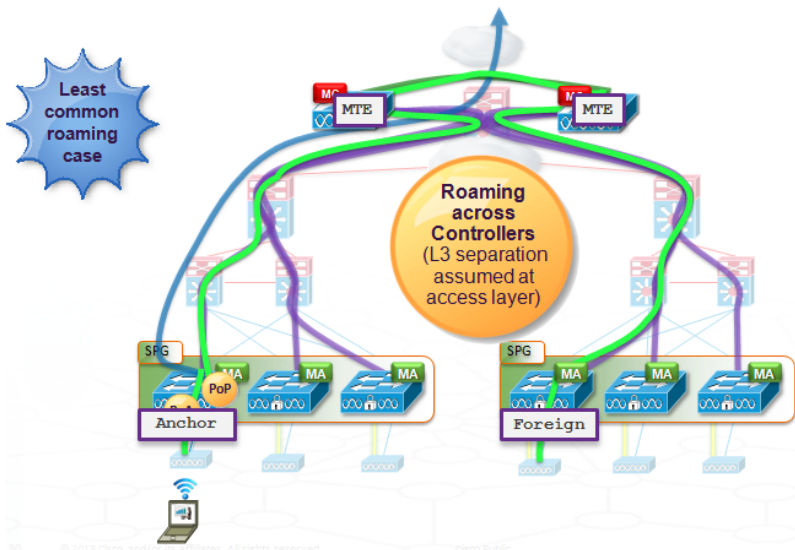


### Roaming Across SPGs Under Control of Different MCs

To further examine possible (but decreasingly less-likely) roam cases, the following figure shows different SPGs that are both under control of different MCs within a common MG. This is type of roaming is possible across the network, but is likely the least common example that would be seen in most deployments.

In the figure, the user roams from the left hand MA to the right hand MA. Both MCs at the top of the figure operate as MTEs and the anchor switch handles the PoP. Traffic flow is similar to roaming across SPGs under control of the same MC, except that in this case, traffic moves between the two controllers (each handling its own sets of MAs and SPGs, and grouped together into a common MG), back down to the anchor switch, and out to the network.

Figure 28. Roaming Across SPGs Under Control of the Different MCs



When the user roams, a mobile announcement is made to the MC. If there is an MO, the mobile announcement is made to the MO (which knows where all control devices are in the wireless network). The MO then makes an announcement back to the roamed-from controller. There is a direct handoff from the anchor switch to the foreign switch in this case, which again serves to speed up and simplify the roaming process as much as possible. When the mobility handoff is complete, a mobility complete message is sent up to the MC, which provides a station left

message. Handoff notifications are sent throughout network, and it is now known to all control devices that the user has roamed.

## Roaming with Mobility Anchors

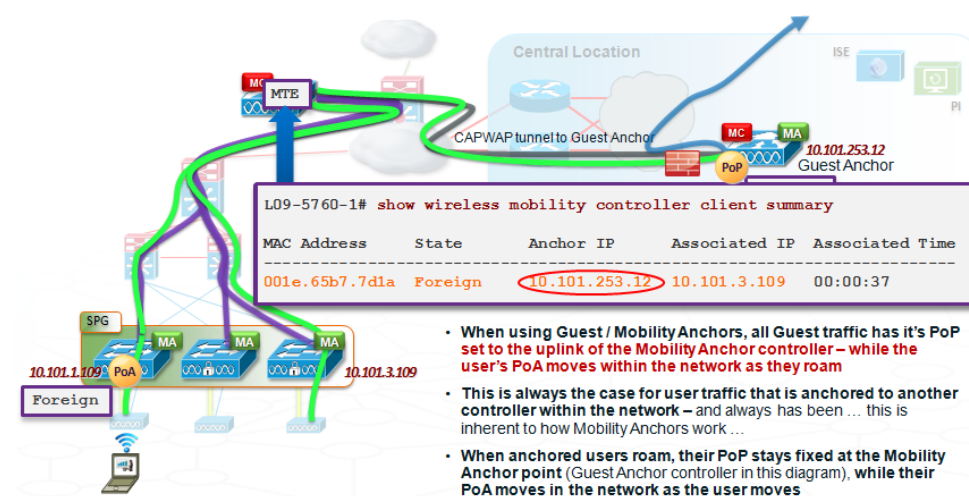
With an understanding of the operation of MAs, MCs, and PoA / PoP, and how roaming using these elements operates within the network, consider in detail how mobility anchors work and how roaming is performed with mobility anchors.

Mobility anchors (for example, for guest traffic) place the PoA on the controller or switch the user is currently associated to, but always place the PoP up at the mobility anchor location. The mobility anchor's location is where the user's traffic is ultimately de-capsulated to and moved out into the wired network (for example, in the DMZ area of the network for guests). In the example in the following figure, the PoA is on the left MA switch on the bottom. The traffic in this example is traffic from a guest user (typically associated into a guest SSID), out to a guest anchor somewhere out in the network (most likely in the DMZ area of the network, for Internet access for the guest user).

In this example, the guest user's anchor switch, acting as an MA, is associated with a discrete controller acting as the MC for the switch's SPG. The discrete controller handling the SPG is thus intermediate in the traffic path between the switch with the associated guest user (acting as MA), and the controller acting as the mobility anchor up in the DMZ. As such, the discrete controller acting as the MC in this example has this guest user in the MTE state. The **show wireless mobility controller client summary** command on this controller shows that the guest user's state is foreign on the 10.101.3.109 switch (which is where the guest user's PoA is), but the guest user is still anchored at the guest anchor controller in the DMZ (which is where the user's PoP is). By simply placing the PoP for guest users at the guest anchor controller in the DMZ, all of the guest users within the network can have their traffic centrally hosted and controlled at this point. Catalyst 3850 switches can serve as the origination point for guest anchor tunnel traffic, but cannot serve as the termination point for guest anchor tunnels (this function must be performed via a discrete controller – up in the DMZ area of the network in this example).

When the guest user roams, only the guest user's PoA moves to the new roamed-to (foreign) switch – the guest user's PoP remains located up at the guest anchor controller in the DMZ. With this understanding of the function and operation of PoA, PoP, MA, and MC, it is straightforward to understand how roaming for guest users is accommodated within the wireless and wired network infrastructure.

**Figure 29. Roaming with a Mobility Anchor**



## Layer 2 or Non-Sticky Roaming – Optional

As mentioned previously, by default all roams within Converged Access operate as Layer 3 roams. This has several important benefits. First, using Layer 3 roaming decreases roaming times, as only the user's data plane connectivity needs to move on a roam (the IP address and associated security policy stays fixed at the PoP switch). Second, Layer 3 roaming is deterministic, because no outside device with a possibly unpredictable response time (such as a AAA server) needs to be involved in the roam. Finally, the use of Layer 3 roaming makes no assumptions about the underlying network topology – for example, it can work across Layer 2 boundaries, across Layer 3 boundaries, or a combination of the two.

Nevertheless, a Layer 2 roaming approach might be considered as desirable in some deployment scenarios. The following figure shows how optional Layer 2 roaming works within an SPG. The SPG is deployed underneath a single distribution layer where a wired-side VLAN is deployed end-to-end. For example, the SPG might be beneath a Virtual Switching System (VSS) type of deployment at the distribution layer, where spanning such a wired-side VLAN end-to-end might be done without incurring undesirable Layer 2 loops within the wired network infrastructure.

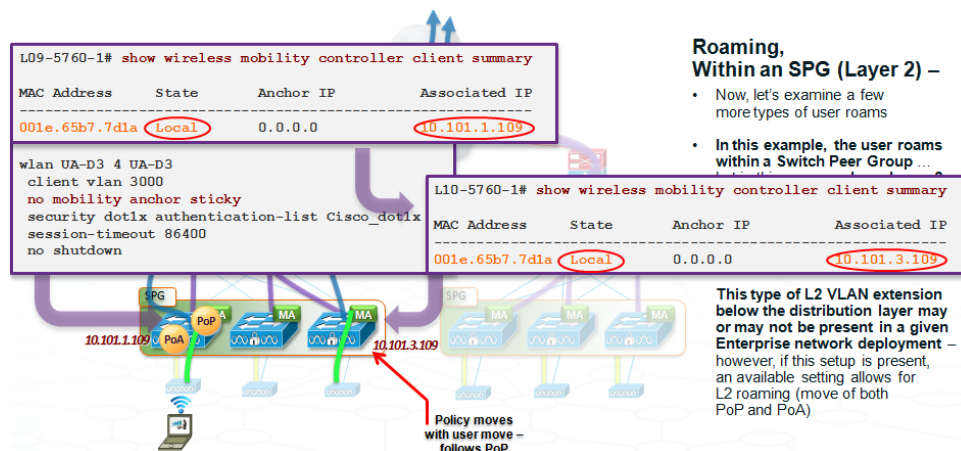
To enable Layer 2 roaming for a WLAN, one can enable the **no mobility anchor** sticky option under the particular SSID involved. With this configuration option in place, when a user roams, both the PoP and the PoA for that user are moved from the roamed-from switch to the roamed-to switch. However, it is important to note that for this type of roam to be successful, the user's IP address must continue to be valid on roamed-to switch. Accordingly, the VLANs within the wired network must be stretched between the access switches involved, with this group of access switches located below a common distribution layer. The wired network topology for a given network deployment might or might not permit this stretching of VLANs between the access switches involved. This example assumes that it does.

When the user roams, because the PoP and PoA have moved, it is not necessary to re-tunnel back to the anchor switch –there is no concept of anchor and foreign switch for a user in this case, because the user's entire mobility state is simply moved between the two switches. However, because the user's policy is enforced at the PoP switch, and the PoP has moved in this case, the policy must move to the new roamed-to switch as well. The policy must be retrieved by this roamed-to switch from the AAA server, and the roamed-to switch must then reapply this policy locally. Because of the level of additional interaction involved, this optional Layer 2 roaming configuration takes additional time - and the amount of time required is less deterministic than with the default Layer 3 roaming behavior. For this reason, all roams in Converged Access are set up as Layer 3 roams by default. The default of Layer 3 roaming behaviour in Converged Access also supports all access layer network topologies (Layer 2 and Layer 3). A best practice is to retain this default Layer 3 Converged Access roaming configuration and behavior, unless there is a well-defined reason to convert the configuration for Layer 2 roaming support.

Because the user's PoP moved in this Layer 2 roam as illustrated, the user starts as local on the 10.101.1.109 switch, and after the Layer 2 roam, the user becomes local on the 10.101.3.109 switch. As noted previously, there is no concept of foreign and anchor, because the user's entire mobility context has moved.



**Figure 30. Roaming Across SPGs Under Control of Different MCs – Layer 2 Roaming**



### Scaling with Catalyst 3850 based MCs

The previous sections described how that the Catalyst 3850 can be used flexibly as an MA, MC, or both. A single Catalyst 3850 MC can support an SPG, which can have up to 16 x MAs (stacks) maximum, and can support up to 50 APs and 2,000 clients across the entire SPG that is formed with a Catalyst 3850 as that SPG's MC. This is a fixed limit on the capacity of an SPG, if the Catalyst 3850 is used as the MC for that SPG.

MGs allow the construction of larger roaming domains that still use the Catalyst 3850 as the MC. Multiple Catalyst 3850s operating as MCs and each handling its own SPG can be grouped together into a common MG. An MG constructed in this way from Catalyst 3850s can support up to 250 APs total and 16,000 clients supported. The limitation of 250 APs instead of 400 is driven by control plane scaling limits on the Catalyst 3850.

With a CUWN deployment, licensing is per MC – not pooled across MCs. Inter-SPG roaming, RRM, guest access, and other MC capabilities and functions are coordinated across the MCs in the same MG. In this example as shown, it is important to note that the guest access tunnel originates only from the Catalyst 3850 which is acting as the MC, up to the guest anchor controller (even though the guest anchor configuration definition is put into place on all of the MA switches within the SPG). This origination of the guest anchor tunnel from the MC only is done for scalability, due to the fact that a given guest anchor controller can handle only up to 71 tunnels, and that capacity could be easily exhausted if the guest anchor tunnel were originated from each MA. As noted previously, a Catalyst 3850 switch can originate the guest anchor tunnel, but not terminate it (this function must be performed via a discrete controller).

The advantage to using the Catalyst 3850 as an MC is in realizing CapEx cost savings, because it is not necessary to add discrete controllers, and licensing is done at the switch level. For example, it is a cost-effective way to include an MC in a branch environment. The downside is increased operational complexity in the roaming environment, as even a relatively midsized domain it can be necessary to maintain multiple controllers and multiple MCs.

To summarize, using Catalyst 3850s as MCs can be an efficient and cost-effective solution for smaller branch-type deployments. For midsized designs, whether this is still appropriate depends on the size of the environment and how much it is likely to grow over time. In larger environments, however, WLC 5508s, WiSM-2s, or WLC 5760 discrete controllers should be deployed as MCs, as they provide for increased AP and user scale, and they are easier (as a smaller group of devices) to manage, with the Catalyst 3850s deployed in this case as MAs only.




This type of deployment still reaps the benefits that accrue to a distributed wireless deployment (greater total traffic scalability, increased traffic visibility, and integrated wired/ wireless network policies and services), while allowing greater scalability in terms of users and APs via the use of a discrete controller, or group of controllers, as part of the entire integrated wired/ wireless solution.

## High Availability

The design of high availability (HA) in the network is closely tied to the state held in the network for roamed and non-roamed users.

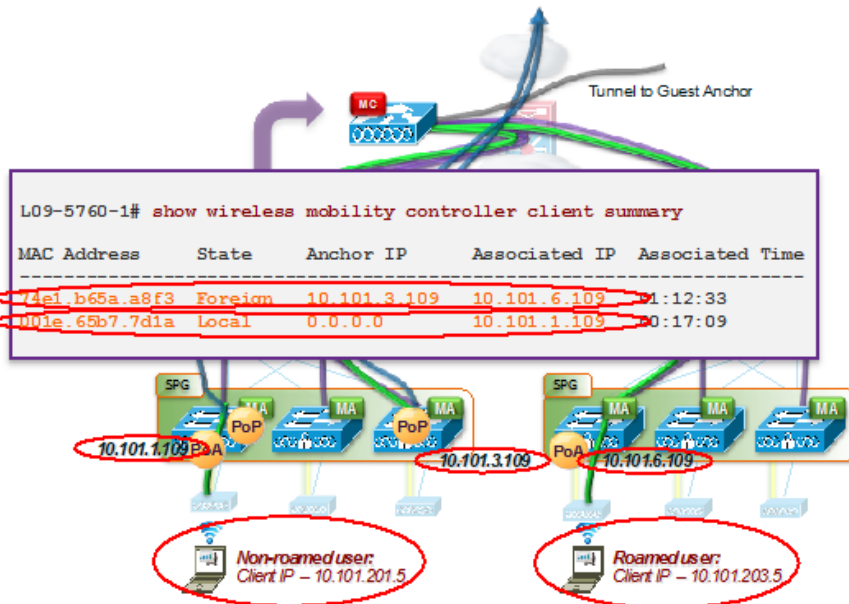
When Layer 3 (sticky) roaming occurs, the information about a user might be held in two or even three places in the network (on MAs and MCs). The following figure uses a discrete controller such as a WLC 5760, WLC 5508, or WiSM-2 as an MC. The figure shows two users, one local (non-roaming), and the other roamed across SPGs under control of the same MC. The roamed user's client IP address is associated with the IP address pool on the right-hand switch in the left-side SPG (where the user originally associated).

**Figure 31. State Behavior for Local and Roamed Users**



The following figures details for the example. The **show wireless mobility controller client summary** command on the MC shows that the a8f3 user is held there. The user is in the foreign state because they started off at switch 10.101.3.109 and roamed to switch 10.101.6.109. The other user has not roamed and is in the local state.

**Figure 32. State Behavior for Local and Roamed Users – View from the MC**

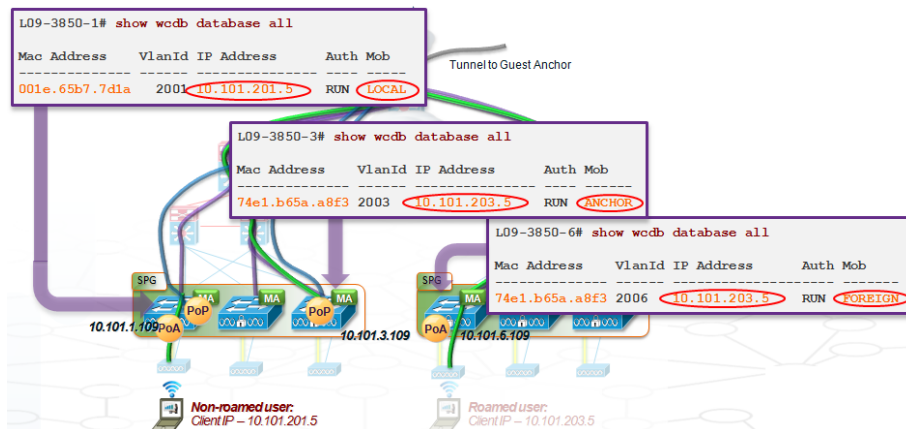


**Roamed and Local users, High Availability Considerations –**

- Here, we can see the state of the network for the roamed and non-roamed clients, as reflected at the MC itself (shows a snapshot of the traffic flows within the Mobility Sub-Domain the MC controls) ...


The following figure shows the MA view. The 7d1a user is the local user who has not yet roamed, while the a8f3 user is foreign on the roamed-to switch, and anchored back on the origin switch.

**Figure 33. State Behavior for Local and Roamed Users – View from the MAs**



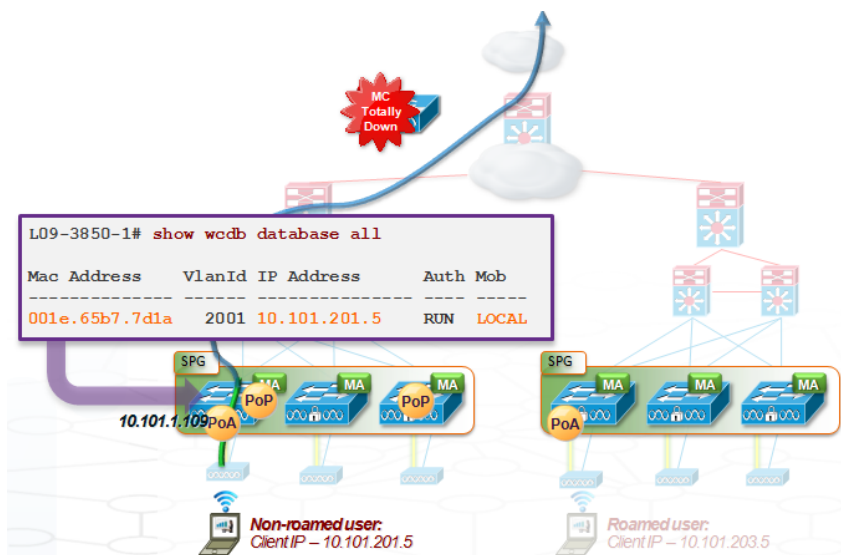
Now, assume that an MC goes down due to a power outage or system reboot. All of the CAPWAP tunnels from the MC to the MAs (purple in the figure) go down because the MC is down. But the tunnels to the APs, and the tunnels within and between the switches in a given SPG, all stay up, because they don't depend on the MC for continued operation. All of the tunnels are pre-formed and continue to function.

**Figure 34. MC Failure – Effect on MC Sub-Domain and Anchor Connections**



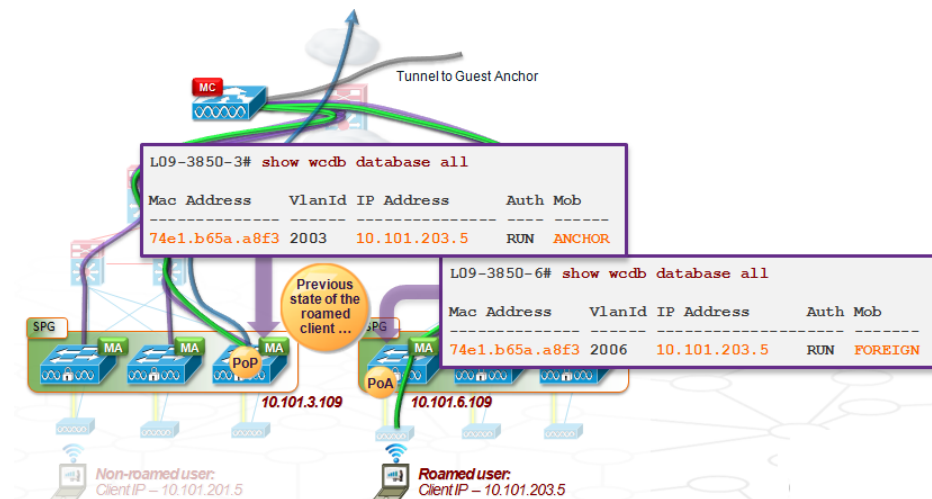
The local non-guest user (after reauthentication) continues to operate as normal, as shown in the command output in the following figure. This is a major enhancement provided by the Converged Access deployment model – as, in this case, an important network element (the MC) has failed, yet traffic for non-roamed users (which are likely to be the bulk of users in many deployments) continues to operate normally following such a failure. This is a significant HA enhancement offered by Converged Access in comparison to alternative deployment modes. However, with the MC down as in this example, any centralized services that depended on the MC, such as guest access, RRM, fast roaming, and key distribution are not available. Inter-SPG roaming involving the MC is also affected.

**Figure 35. MC Failure – Effect on Local (Non-Roamed) Users**



Now consider a previously-roamed client who roamed from the 10.101.3.109 to the 10.101.6.109 switch. The user is still anchored back on the origin switch and foreign on the roamed-to switch.

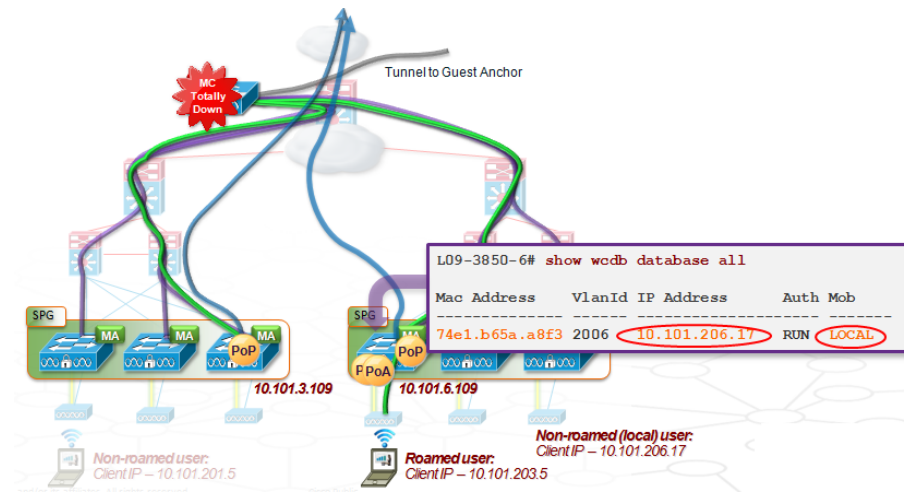
**Figure 36. MC Failure – Effect on Previously Roamed Clients**



If the MC goes down, all of the purple connections (CAPAP tunnels between the MC and the MAs it controls) in the following figure go down. The roamed client authentication information in this case is lost, and such users must reauthenticate for continued network access, obtain a new DHCP IP address from the IP address pool appropriate to them at the switch to which they are now associated, and become local to that new switch. The PoP and PoA for this previously-roamed user are both re-homed on the roamed-to switch. After re-homing and a change of IP address at the client, traffic for this client now flows as for any other local user on that switch.

Prior to the roam, the client in this example had IP address 10.101.203.5 and started on switch 10.101.3.109. After the MC-down event and the client's reauthentication at the roamed-to switch, this previously-roaming client's IP address in this example as shown is now 10.101.206.17 (an IP address obtained from the IP address pool as appropriate at this new access switch for the VLAN involved).

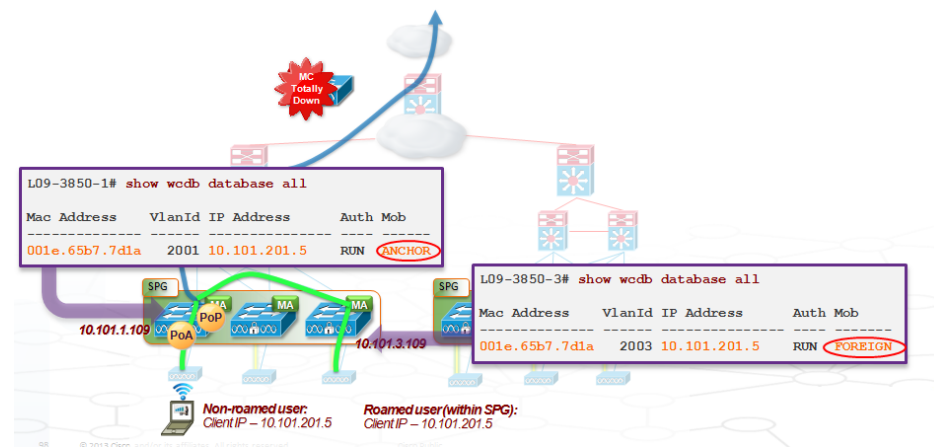
**Figure 37. MC Failure – Effect on Previously Roamed Clients**



Now assume that the MC is still down, and that the previously non-roamed (and non-guest) user tries to roam within the SPG. The roaming event still happens normally. The user’s pairwise master key (crypto key) is already distributed within the SPG, all of the CAPWAP tunnels are still operational, and the user is able to retain the IP address on this roam, even though the MC for this SPG continues to be offline at the time this roam occurs. This is possible since the MC’s involvement in this type of intra-SPG roam is primarily informational only, as all of the switches within the common SPG have a full mesh of CAPWAP tunnels to each other, and can handle this type of intra-SPG roam for a previously-associated client even if the MC for that SPG is offline.

For a properly designed SPG, the majority of the roams should be within the SPG (a building or floor). Although MC is down, which is already a fairly rare event, the user in this case is still able to roam normally. Again, this is a major HA benefit of a Converged Access deployment, since the SPG involved continues to operate normally for previously-associated clients for intra-SPG roams, even though an important network element (the MC for that SPG) is offline.

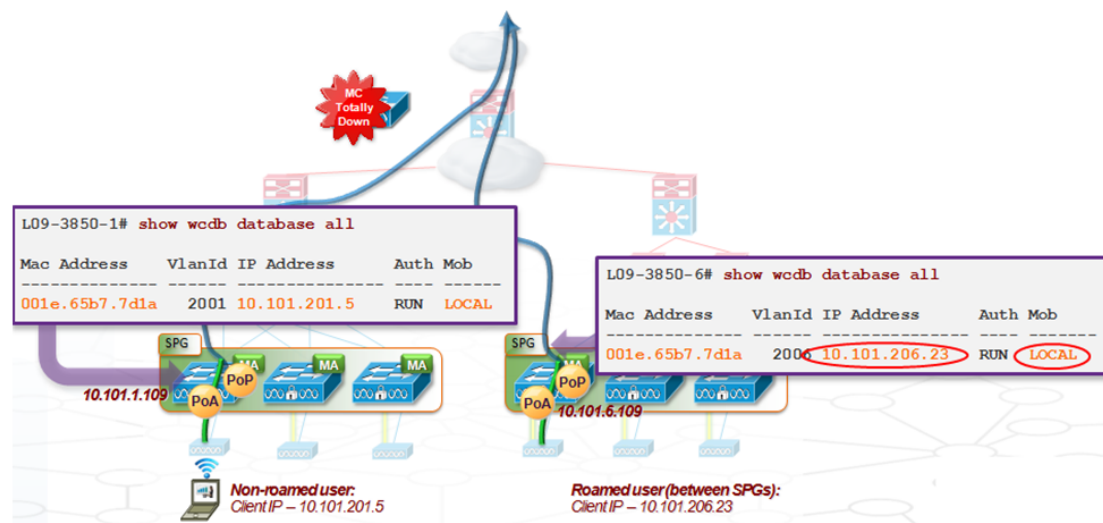
**Figure 38. MC Failure – Effect on Intra-MC Roams Following MC Down**



Now assume that a user roams between SPGs while the MC is down. Because the MC is unavailable, the client authentication information is not able to be passed from the roamed-from MA in the origin SPG, to the roamed-to MA in the target SPG, since the MC is the only connecting point between these two SPGs – and in this example the MC controlling these two SPGs is down. The roaming user between SPGs must reauthenticate, obtain a new DHCP IP address from the IP address pool as appropriate at the new roamed-to switch, and become local to that new switch. The PoP and PoA are both re-homed on this new roamed-to switch in this case.

In the following figure, the user started with IP address 10.101.201.5 prior to this inter-SPG-while-MC-is-down roam. Following the roam, the user has IP address 10.101.206.23, and effectively appears as a new client into the network. Because the only connection between SPGs is the MC that services them both, the user must be treated as a new user in the network following this roam.

**Figure 39. MC Failure – Effect on Inter-MC Roams Following MC Down**



It should be noted that guest traffic is always disrupted in the event that the MC (Catalyst 3850-based or discrete-controller-based) is down. This is because, for scalability reasons, the guest anchor tunnel is originated only from the MC. If an MC is down, guest traffic for users employing that MC is also down (guest users on other MCs are unaffected). This fact should be taken into account when planning a Converged Access deployment.

Because the MC performs an important role within a wireless deployment, it is helpful to plan for redundancy. AireOS-based MCs now support an HA option. The following figure shows the HA mode available in AireOS 7.3 on the WLC 5508 and WiSM-2 controllers. Using this 1:1 HA option, an AireOS controller operates as an MC on a WLC 5508 or a WiSM-2, and is paired up with an identical backup controller in a 1:1 redundancy configuration.

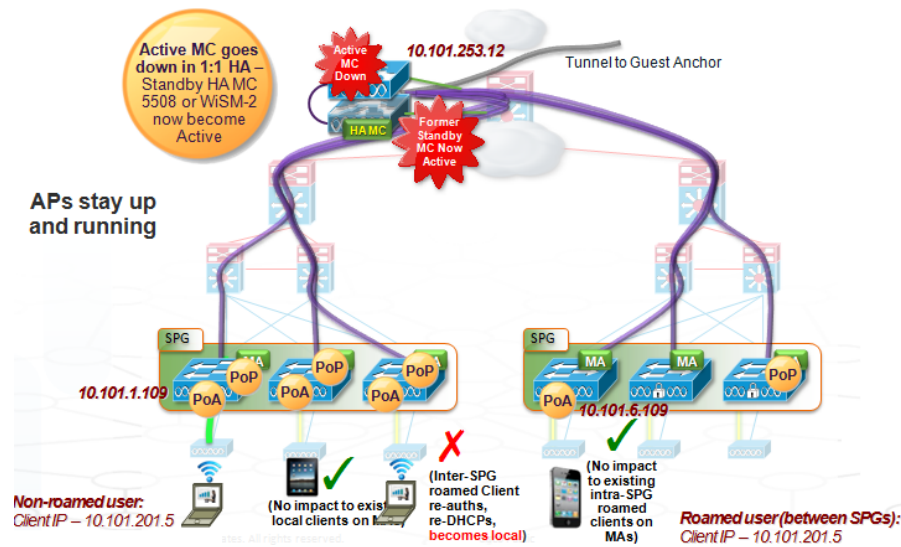
In the event of an MC failure in this type of deployment, local non-guest users on their MAs experience little to no impact following a HA MC failover event, nor do intra-SPG roamed users. Any new intra-SPG or inter-SPG roaming happening after MC HA failover from local MA clients is handled normally. However as before, previously-roamed inter-SPG users do need to reauthenticate, obtain a new DHCP IP address, and become local to their new switch, as their client information at the MC is lost on this type of MC failover.

All APs all stay up and running, and the MC recovers itself rapidly due to the AireOS 7.3 HA functionality. Effectively, the MC redundancy capability provided by AireOS 7.3 on a discrete controller such as a WLC 5508 or WiSM-2 serves to reduce the Mean Time to Repair (MTTR) for the MC component in the network significantly –



and thus serves as an important enhancement for the HA of the resulting solution overall, by restoring the MC function for any associated MAs, and thus restoring normal network operation for those MAs, as rapidly as possible.

**Figure 40. MC Redundancy**



## MC Redundancy Information

In a one-to-one MC redundancy configuration with AireOS 7.3 HA on the WLC 5500 and WiSM-2, one WLC is in the active state and the other WLC is in a hot standby state. The standby system monitors the health of the active WLC. The configuration on the active system is synchronized to the standby WLC through a redundant port. Both of the WLCs share the same set of configurations, including the IP address of the management interface.

The CAPWAP state of the APs that are in the RUN state is also synchronized (the state is retained for any APs that are hosted directly by the WLC 5508 or WiSM-2, so that any such APs do not need to go into the discovery state when the active WLC fails). APs that are hosted by downstream Catalyst 3850s have their AP state retained on the Catalyst 3850 acting as an MA, not on the upstream discrete controller acting as an MC.

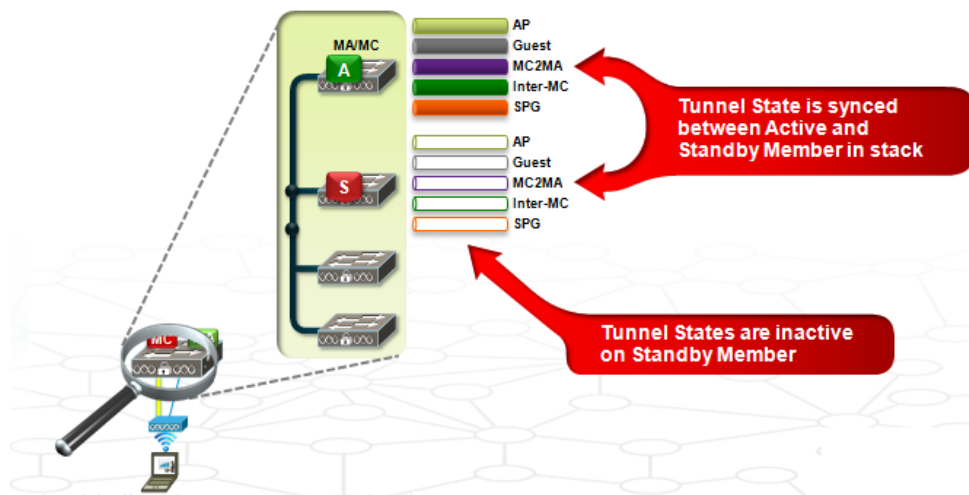
In this configuration, downtime after failover is typically less than one second in the case of controller failover, and up to approximately three seconds in case of network issues.

This configuration can be deployed on a single chassis between two WiSM-2s, or it can be deployed over multiple chassis using VSS, which effectively combines the two into a single virtual switch.

Now consider what happens if an MC fails in an environment where the MC functionality is running on a Catalyst 3850 stack. There is an active switch and a standby switch in the stack. The active switch runs the stack, while the standby is synchronized off the active and is available to take over in the event that the active fails.

The following figure shows how HA works with the Catalyst 3850 stack. A function called tunnel Stateful Switchover (tunnel SSO) runs within the stack. The active switch maintains all the tunnels and their state—AP tunnels, guest tunnels, SPG tunnels, inter-MC tunnels—while the standby switch also has the same set of tunnels that are available. The only difference is that the tunnels on the standby switch are not in the up state, and remain inactive unless and until an intra-stack failure occurs with the active switch within the stack.

**Figure 41. HA with Catalyst 3850 Based MCs**



If the active switch goes down, all of the tunnels become active on the standby switch as the intra-stack HA switchover occurs. As in the earlier examples, there is no impact for existing clients on MAs that are not part of the affected switch stack, assuming these are non-roamed clients, and assuming that the stack involved is not acting as an MC for its SPG. However, clients on the switch stack that experienced the failure need to reauthenticate, and a roamed client on that stack will need to not only reauthenticate, but also obtain a new IP address, and become local to the new switch stack after the failover (the PoA and PoP for such a roaming client both move in this case).

If the entire switch stack operating as an MC for a local SPG is lost due to a reboot or power issue, guest access is lost within the affected SPG, and any CAPWAP mobility tunnels over to the other MCs to which the failed stack may have been associated also go down. Client information for roamed clients is lost in this case, and any previously-roamed clients within the SPG involved need to reauthenticate and become local (change IP addresses) in this event.

Roaming within an SPG is still works while this switch stack operating as an MC is down, since all of the Catalyst 3850 stacks within the common SPG share a full mesh of CAPWAP tunnels to each other. However, as with the failure of a discrete controller operating as an MC, roaming between SPGs while the switch stack operating as the MC is down requires client reauthentication, and obtaining a new IP address for wireless clients following such an inter-SPG roam. Also, similar to the failure of a discrete controller operating as an MC, it is not possible to distribute pairwise master keys while the MC is down. Accordingly, there is no fast roaming available for new wireless clients until the MC is restored. Moreover, while the MC is down, RRM, guest access, and other functions that depend on the MC do not operate within the affected SPG.

Overall, although there is some impact from MC outages or failovers, Converged Access offers a significantly more robust and fault-tolerant deployment option than in many other wireless deployment options. The fact that a Converged Access deployment can largely continue to operate for non-guest traffic even while the MC for a given SPG is entirely down, is a testament to the resilience of the distributed nature of the Converged Access deployment option, and provides a significant HA benefit for Converged Access deployments.



# Multicast, Deployment Options, and Migration

Multicast is important for Converged Access, as it is used by critical applications such as video streaming and music on hold for voice.

The following sections describe multicast in the Converged Access context, including how it works and the theory of operation when associating to a controller and during roaming. The examples describe multicast mainly on the WLC 5760. Multicast is also supported on the Catalyst 3850 platform, and where appropriate, examples are also provided for multicast use on the Catalyst 3850.

## Multicast for Converged Access

By default, multicast forwarding is disabled in the Converged Access environment. Multicast frames received at the controller level are not forwarded to or from APs. Traffic received from the network or from a wireless client, such as IGMP packets, is dropped at the controller. For multicast to work, both multicast forwarding and IGMP snooping are required.

To enable multicast in a Converged Access network, first use the **wireless multicast** global configuration command enable multicast forwarding. Then enable IGMP snooping for IPv4 using the **ip igmp snooping** command and for IPv6 using the **ipv6 mld snooping** command.

As in CUWN, there are two types of delivery mechanisms for multicast traffic from the controller to the APs. With *multicast-multicast* delivery, the multicast traffic is encapsulated in a multicast CAPWAP packet, and all the APs join an internal CAPWAP multicast group that is configured on the controller. With *multicast-unicast* delivery, the multicast traffic is transported inside a unicast CAPWAP header to the AP.

Multicast-unicast is the default delivery method used in Converged Access, as in CUWN. The **ap capwap multicast** command is used to configure multicast-multicast, with the desired multicast group address (usually an internal, private address that is used only between the controller and the AP, not by clients).

The campus LAN and WAN infrastructure must be configured to support multicast traffic prior to enabling multicast forwarding capabilities in existing centralized CUWN and Converged Access deployments. Otherwise, the multicast-multicast operational mode does not function properly throughout the entire network.

The next section provides details on how the multicast process works.

## Multicast Process Details

When multicast forwarding and IGMP snooping are enabled, the WLC intercepts IGMP reports from IPv4 mcast clients and creates a multicast group ID (MGID) based on a (Source, Group, VLAN) tuple. The WLC uses IGMP snooping to determine if an IGMP report should be forwarded to the source router. An MGID to WLAN mapping is sent to the AP, which keeps track of the MGID, WLAN, and client association.

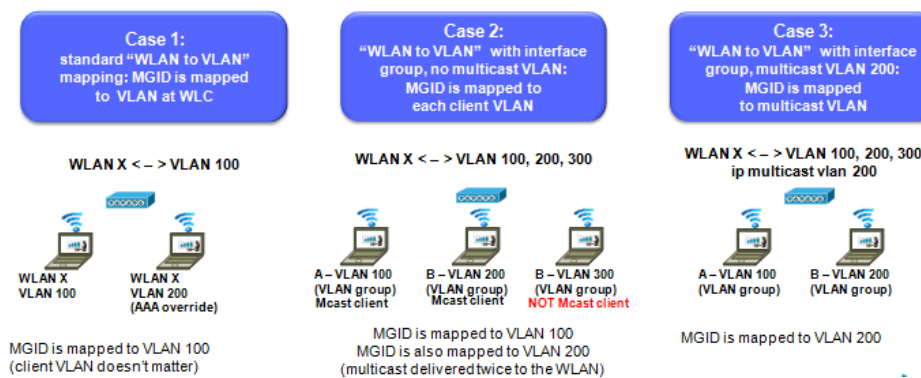
The following figure shows three examples to consider.

Case 1 is the standard case in which the WLAN is defined and mapped to a VLAN, and the MGID is mapped to the same VLAN. In the figure, WLAN X is mapped to VLAN 100. One client is connected to WLAN X and VLAN 100, while the other has had AAA overridden by a policy to be assigned to VLAN 200. From a multicast forwarding point of view, what matters is that the default VLAN 100 is defined on the controller and any multicast that is requested from this client is forwarded to that VLAN.

Case 2 is more complicated, and not optimized. In this case, there is a VLAN group associated with one WLAN, one SSID, and multiple VLANs on the wired side. WLAN X is associated to VLAN 100, 200, and 300. Because there are multicast clients on VLAN 100 and 200, there are two MGIDs, one created for the VLAN 100 and the other for VLAN 200. The challenge is that if clients in these VLANs request the same multicast stream, twice as many MGID messages are forwarded to the router for the two VLANs, and the stream traffic is doubled. Even the AP receives two different streams.

Case 3 shows a better way of handling this situation, using the multicast VLAN feature. In this case, the `ip multicast vlan` command on the controller is used to define a specific VLAN for multicast (VLAN 200 in the figure). The WLAN is mapped only to that VLAN in the interface group. Regardless of how many VLANs are present, all of the IGMP and multicast traffic will be handled by that VLAN.

**Figure 42. Multicast Process Examples**




## Multicast Process Flows

The figures in this section show the process flows for specific multicast operations.

The following figure shows happens when the controller receives the multicast traffic. First it determines if an MGID is associated with the multicast group for the VLAN from which it is receiving traffic. If not, it drops the traffic. However, if an MGID has been created, it forwards the traffic to all the APs who have clients attached that have requested the multicast traffic. If the mode is multicast-multicast mode, the forwarding is performed using multicast. If the mode is multicast-unicast replication, the traffic is replicated and sent by unicast to all the APs with clients attached that requested the multicast traffic. This is a major improvement over the previous process, in which replication was done to all the APs associated to the controller irrespective of whether the AP had clients that had requested the traffic.

**Figure 43. Multicast Forwarding**

### Multicast on Converged Access WLC – Forwarding




When Multicast flows from source to client:

1. WLC checks if a MGID exists for that flow
2. If MGID exists, WLC forwards to all APs, using Multicast
  - No QoS tag by default, no DTLS encryption (if multicast)
  - If Multicast to Unicast replication is used, then traffic is sent only to those Aps that have multicast clients associated
3. APs forward to each WLAN having clients for this MGID
  - Sent at highest mandatory rate, BE queue by default

The following figure shows the process when the multicast client leaves (for example, when stopping a video). The controller uses membership queries to check at defined intervals whether the client is still interested in the traffic. When an IGMP leave is received, the controller checks to see whether any other clients are interested in that traffic. If not, the WLC removes the MGID and does not send a leave to the router. It is a silent drop-off. The controller simply stops replying to the membership query from the router, and the multicast flow stops. The controller also informs the APs that that multicast group ID has been deleted.

**Figure 44. IGMP Leave**

## Multicast on Converged Access WLC – IGMP Leave



- WLC uses membership queries (IGMPv1 general queries or IGMPv2 group-specific queries, depending on config)
- When Leave message is received from last client in a MGID:
  1. WLC removes client from MGID list, removes MGID
  2. WLC does not forward leave to wired infrastructure (timeout is used)
  3. WLC informs the APs about MGID deletion

The following figure shows the process flow for broadcast forwarding. Broadcast is similar to but handled separately from multicast. One or both can be enabled for IPv4 and IPv6. Once activated, the mechanism is the same. When broadcast traffic is received on the VLAN, an MGID is created, and the controller keeps track of the client on the WLAN associated to the VLAN. The MGID to WLAN mapping is sent to all the APs, and the AP keeps track of the association between the multicast groups, the WLAN, and the clients. The traffic is forwarded to the APs using multicast-multicast or multicast-unicast, and no QoS is applied by default. The main difference between multicast forwarding and broadcast is that broadcast traffic is best effort and is sent out over RF at each AP at the lowest mandatory rate configured on that AP.

**Figure 45. Broadcast Forwarding**

## Multicast on Converged Access WLC – Broadcast Forwarding



When Broadcast forwarding is enabled:


1. Broadcast forwarding and multicast forwarding are handled separately (you can activate one, the other or both, for IPV4, and/or for IPV6).
2. For broadcast in new VLAN, MGID is created for IPv4 broadcast (range 0-4095), MGID same as VLAN number
3. WLC keeps track of client WLAN to VLAN mapping (including with AAA override)
4. WLC sends to AP MGID to WLAN mapping (same map for all APs, except with AP groups), AP keeps track of MGID, WLAN and client association
5. Traffic is forwarded to all APs (using Unicast or Multicast, no QoS tag by default, no DTLS encryption (if multicast))
6. APs forward to each WLAN having clients for this VLAN (sent at lowest mandatory rate, Best Effort [AC\_BE] queue by default)

The following figure shows multicast process flow with video streams. The process is the same as with standard multicast forwarding, but with additional steps. A key Cisco differentiator for wireless is VideoStream, which is a video multicast optimization that overcomes one of the limitations of sending standard multicast over the air. VideoStream includes a reservation mechanism and admission control, so a video stream is admitted only if there is the bandwidth on the AP to accept it.

In terms of multicast forwarding, the controller intercepts the IGMP reports from clients, and sends a radio resource control request to the AP, because it needs to check if the AP has available bandwidth. Upon a successful response from the AP, video stream previously defined on the controller is admitted. The WLC creates an MGID, forwards the join report to the router, and traffic starts flowing. The AP receives the MGID and WLAN mapping and keeps the association between this mapping and the client.

When the AP receives the multicast traffic from the controller, it tracks the clients that have been requesting the video, and in hardware converts the multicast traffic to unicast, sending it to each individual client based on the client's data rate. The video is transmitted optimally with Wireless Multimedia Extensions (WMM) video marking and uses the video queue on the radio, resulting in improved video quality when compared with a best effort process.

**Figure 46. Multicast with Video Streams**



Same behavior as standard Multicast forwarding, with additional steps:

1. WLC intercepts IGMP reports from IPv4 mcast clients
2. WLC sends a Radio Resource Control request to the AP
3. APs returns a RRC response for the client and the radio
4. WLC determines if the VideoStream can be admitted
5. WLC creates a new MGID, and forwards join report upstream
6. WLC informs the AP that the stream is admitted and that bandwidth must be reserved
7. WLC also sends to the AP MGID and WLAN mapping
8. AP keeps track of MGID, WLAN and client association

## Multicast with IPv6

IPv6 uses multicast extensively for router and neighbor advertisements and solicitations. Converged Access controllers manage and optimize IPv6 multicast through filtering and throttling.

For router advertisement, the WLC intercepts the router advertisement from the router. It then checks the association between the multicast group and the client VLAN and adds the client to the MGID. As with all other mappings, the information is also sent to the AP using unicast or multicast forwarding. Then the AP forwards the packet to the IPv6 client using unicast or multicast.

If an IPv6 client is on VLAN 100, the client receives router advertisements and knows its IP address and subnet. When the client roams, the foreign controller receives the router advertisement from the home anchor controller, and must forward the message to the roamed client. The foreign controller forwards only the messages from the original client VLAN, in this case VLAN 100, and does not need to send any other router advertisements to the client, for

example, from another VLAN. This process filters and optimizes the delivery of the router advertisements so that the client is not overloaded.

Now consider several roaming scenarios with multicast traffic in use.

For intra-controller roaming, the controller is notified by IOS that the client has moved and updates its table, mapping the client's traffic to the new AP. It checks the configuration on the physical port where the AP is connected to see if multicast is denied. Aside from this, the only other change is the MGID to client mapping.

The next case is that of inter-controller Layer 2 roaming between two switches, for example, in an SPG, where both switches have the same VLAN on the wired side associated to the same SSID. If Layer 2 roaming is non-sticky (Layer 2 roaming is enabled, which is not the default behavior), the process involved is exactly the same as roaming in the existing CUWN network. When a user roams from switch 1 to switch 2, switch 1 moves all the information about this roaming user (unicast and multicast client information) to switch 2. The MGID information on switch 1 is deleted, and a new MGID is created on switch 2. Following the roaming event, unicast and multicast traffic both flow from switch 2.

Recall, however, that the default Layer 2 roaming behavior is sticky, so the unicast traffic is anchored at the home initial switch. When the client roams in this case, switch 1 moves the multicast information to switch 2, but not the unicast information. The unicast information is kept in switch 1 (the user's PoP). In terms of multicast, the behavior is the same as in the non-sticky case – all information is moved to switch 2. By contrast, unicast traffic continues to flow through the initial (anchor) switch – which is where the user's PoP continues to be located after the roam. This process optimizes multicast delivery on such a roam, while still ensuring that the user's unicast traffic is handled as appropriate. Because it is assumed that there is a multicast-enabled network on the roamed-to side, it is preferable to use the switch that the client is actually connected to for the roaming user's multicast traffic.

To summarize, in inter-switch Layer 3 roaming, the associated subnet is different on switch 1 and switch 2 for the same SSID. When a user roams, switch 1 (the roamed-from switch for this user) transfers all multicast group information for this user to switch 2 (the roamed-to-switch for this user), but copies the unicast information. Behavior in this case is thus asymmetric – the unicast traffic for this user is anchored at switch 1, and the multicast traffic for this user flows directly through switch 2. If there are any security policies for multicast traffic flows on switch 1, the same policies must also be present on switch 2. If not, users might gain access to multicast resources that they should be excluded from following a roaming event.

As a last note, if the Converged Access configuration uses an anchor controller for guest traffic, behavior is the same as in the existing CUWN environment. If multicast traffic is received on a guest controller, it is dropped because it is not permissible to forward multicast traffic inside the CAPWAP tunnel from the guest anchor controller to the foreign controller ( multicast forwarding for guest users is not supported).

## Design Options

The following sections evaluate the characteristics and key elements to consider when choosing the best design for a given branch or campus deployment with Converged Access.

### Small Branch

A small branch is typically in a remote location and is limited in the number of APs, say up to 50, which is the limit for a single Catalyst 3850 switch stack with MC functionality enabled. The key characteristic from a design perspective is that this type of deployment uses the Catalyst 3850s as MC and MAs, with no discrete controllers.

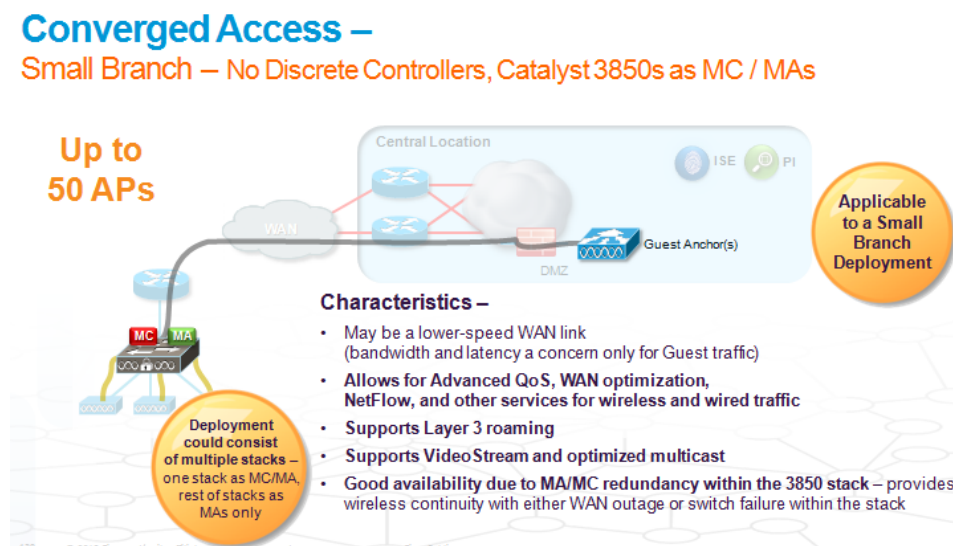
A small branch Converged Access deployment provides features such as Layer 3 roaming at the branch, VideoStream, and optimized multicast – functions that are not supported in a FlexConnect deployment, which would be another alternative deployment option for such small branch use. Connectivity to the branch might be through a relatively low speed WAN link. In such a case, bandwidth and latency are only a concern in this deployment for any traffic transiting the WAN, such as guest access traffic in the deployment shown in the following figure.

This type of deployment allows for advanced QoS, NetFlow, WAN optimization via a WAN router, and other services for wireless and wired traffic, since all traffic (regardless of wired or wireless origin) is treated similarly in this design. MA/MC redundancy is provided within the Catalyst 3850 stack for wireless service continuity should there be a switch failure within the stack or a WAN outage, again providing a significant benefit vs. other possible deployment options. The configuration can be used with a single switch stack, or can include multiple stacks, with one stack operating as MC/MA, and the rest of the stacks configured as MAs only.

The fact that Converged Access is able to optimize wireless and wired traffic flows, provide seamless wired / wireless integration, improve both wired and wireless HA, and provide common wired/wireless traffic visibility and policy control for such a small branch deployment provides significant customer benefits – benefits that may not previously have been available for a small branch such as this.

The following figure shows an example of such a small branch Converged Access deployment.

**Figure 47. Small Branch Deployment**

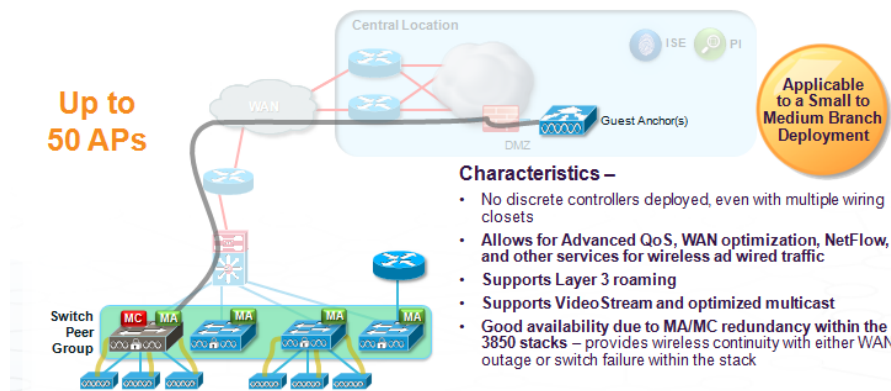


## Small/Medium Branch

The following figure shows a small/medium branch deployment, in which an SPG is created with multiple MAs. One switch or stack can act as an MC, with other stacks acting as MAs within the SPG. Mobility is optimized inside the SPG and roams are kept within the distribution layer. There is no discrete controller to manage. This type of deployment can serve a small or medium branch where there is a larger number of wiring closets. The same functionality and benefits are available as for a small branch deployment. If a Catalyst 3850 is used as an MC for such a deployment, the scalability limit of 50 APs and 2000 clients total applies to the entire SPG formed by that Catalyst 3850 (across all of the MAs it controls), not just to the Catalyst 3850 stack individually.

Figure 48. Small/Medium Deployment

### Converged Access – Small / Medium Branch No Discrete Controllers, Catalyst 3850s as MC / MAs, Single SPG





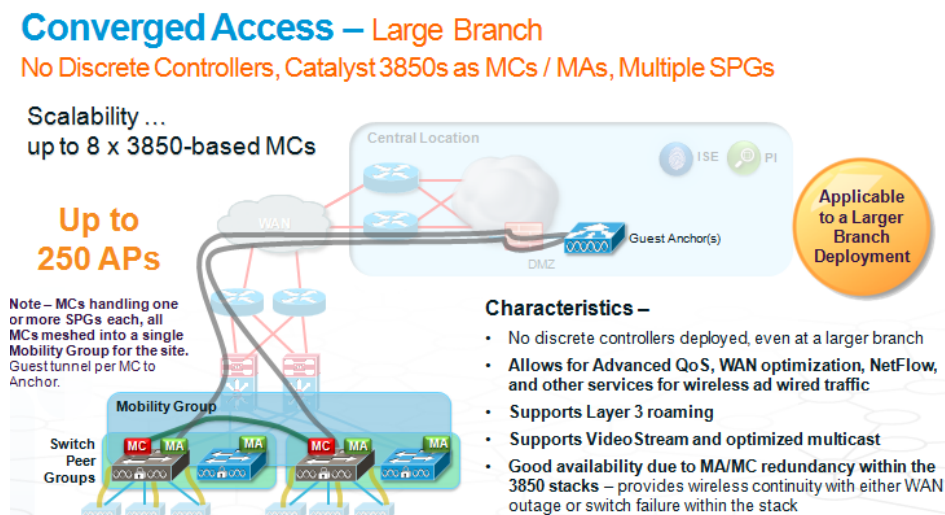
## Large Branch

Moving up to a large branch deployment, the assumption is that more APs are required. This type of deployment can scale up to 250 APs and 16,000 clients, without the requirement to necessarily deploy any discrete controllers. As shown in the following figure, the MC functionality can be distributed across multiple Catalyst 3850 stacks (up to eight in total). These multiple Catalyst 3850-based MCs can then be grouped together into a single MG. As with the earlier branch deployment scenarios, discrete controllers are still not necessary (although they can be used if desired – for example, if even greater scalability is considered to be necessary in future, beyond the 250 AP/16,000 client maximum). The design approach of using Catalyst 3850-based MCs for such a large branch deployment provides seamless and fast roaming between all the APs in the branch.

The deployment could be in a multi-story building where Layer 3 and seamless fast roaming are required throughout the building. The benefits are the same as with the other branch deployments, but with greater scale. When using the Catalyst 3850 as an MC in such a deployment, there is a maximum of eight such MCs possible per MG. The limitation of a maximum of 250 APs (rather than 400) in such a deployment is based on the control-plane scaling capabilities of the MC functionality on the multiple Catalyst 3850s involved. Up to 16,000 clients can be supported in such a deployment (up to 2,000 clients per Catalyst 3850-based SPG, across eight Catalyst 3850 stacks operating as SPGs and grouped together into a common MG).

This deployment option provides for a significant level of scalability, without having to deploy any discrete controllers in the environment. However, to simplify management and allow for greater future scalability, many deployments that begin to approach the 8 x MC limits (250 APs, 16,000 client maximum) opt for discrete controllers to alleviate future scaling concerns. Nevertheless, this use of multiple Catalyst 3850s as MCs remains an important and valuable deployment option for larger branches, and highlights the flexibility and high level of scalability, traffic visibility and control, HA, and wired/ wireless integration.

**Figure 49. Large Branch Deployment**





**Figure 51. Small Campus with Multiple MGs**

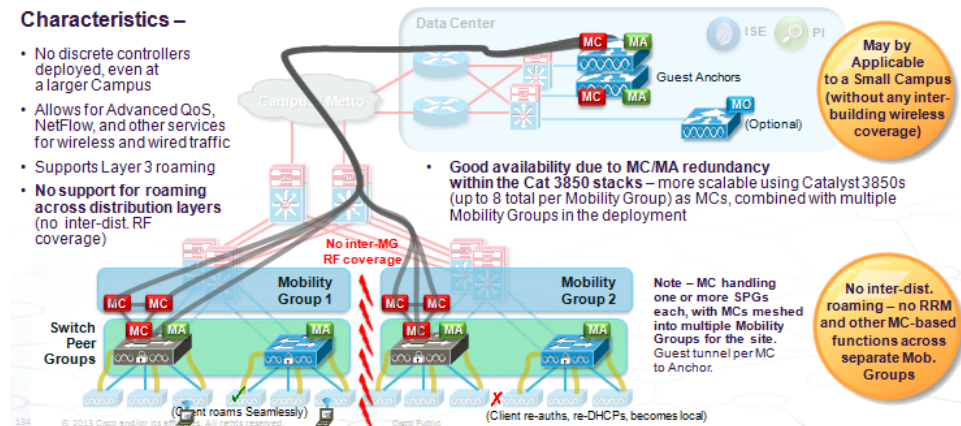
## Converged Access –

**Small Campus – 3850s as MCs / MAs, Multiple Mobility Groups**

Scalability.... > 8 x 3850 MCs, > 250 APs total (w/o inter-dist. roaming)

### Characteristics –

- No discrete controllers deployed, even at a larger Campus
- Allows for Advanced QoS, NetFlow, and other services for wireless and wired traffic
- Supports Layer 3 roaming
- **No support for roaming across distribution layers (no inter-dist. RF coverage)**



## Campus

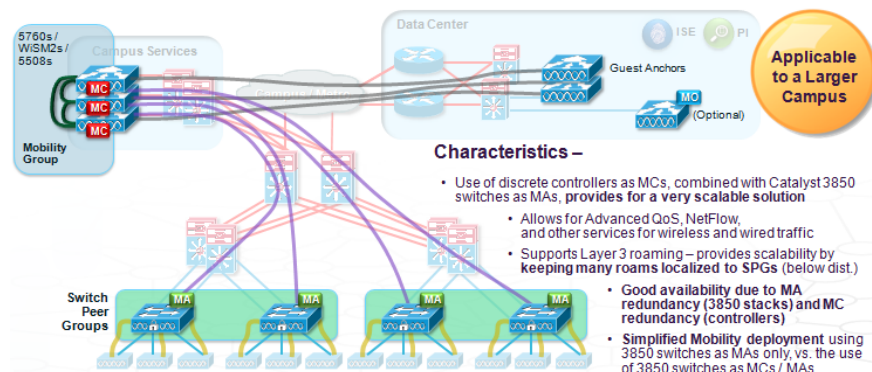
A campus deployment will most likely include a discrete controller or controllers, as shown in the following figure. In this example, these controllers are located in a network block that provides campus services using the MC functionality on the WLC 5760, WiSM-2, or WLC 5508 in these discrete controllers (this functionality could also be distributed in these controllers across the network – which is examined in the alternative campus deployment option provided following this one). Scalability is gained by grouping multiple MCs as needed. By keeping all the discrete-controller-based MCs in the same MG, seamless roaming is provided throughout the entire campus. Functionality is simplified because all of the Catalyst 3850 stacks are operating consistently as MAs only, and HA is supported for the MAs via the inherent HA functionality within the Catalyst 3850 stacks, combined with the HA characteristics of the MG that the MCs are participating in. As with the small campus deployment, QoS, Netflow, and other features are supported, and the excellent traffic visibility, control, and scalability characteristics of the Converged Access solution continue to be provided.

**Figure 52. Campus Deployment**

## Converged Access –

**Campus – Centralized MCs, 3850s as MAs only**

>250 APs



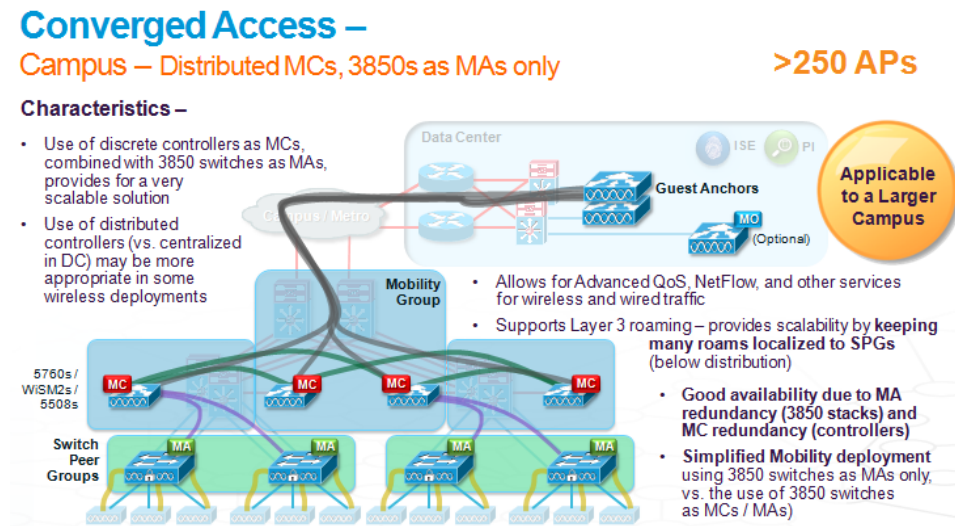
## Large Campus

In some large campus deployments, the network designer might choose to distribute the discrete controllers across the network rather than centralizing them. This could be done to optimize some types of traffic flows within the campus backbone. In such a case, discrete controllers acting as MCs could be attached to the distribution layer switches within the network, and grouped together into a common MGs. In this case, many types of roaming are optimized and kept under the distribution layer within or between SPGs. All the traffic for the guest anchor is sourced from each MC.

Again, functionality is simplified because all of the Catalyst 3850 stacks are operating consistently as MAs only, and HA is supported for the MAs via the inherent HA functionality within the Catalyst 3850 stacks, combined with the HA characteristics of the MG that the MCs are participating in. As with the small campus deployment, QoS, Netflow, and other features are supported, and the excellent traffic visibility, control, and scalability characteristics of the Converged Access solution continue to be provided.

Both of these large campus deployment options (centralized and distributed) offer the greatest mixture of scalability, simplicity, and functionality, and thus would be the most common choices for a Converged Access deployment within a campus environment.

Figure 53. Large Campus Deployment



## IP Addressing

This section describes options for assigning user subnets in Converged Access. The options cover a range of cases and highlight the pros and cons of different design choices that involve dealing with same or different IP address pools for wireless and wired traffic, differentiated policy assignment, and ease of implementation.

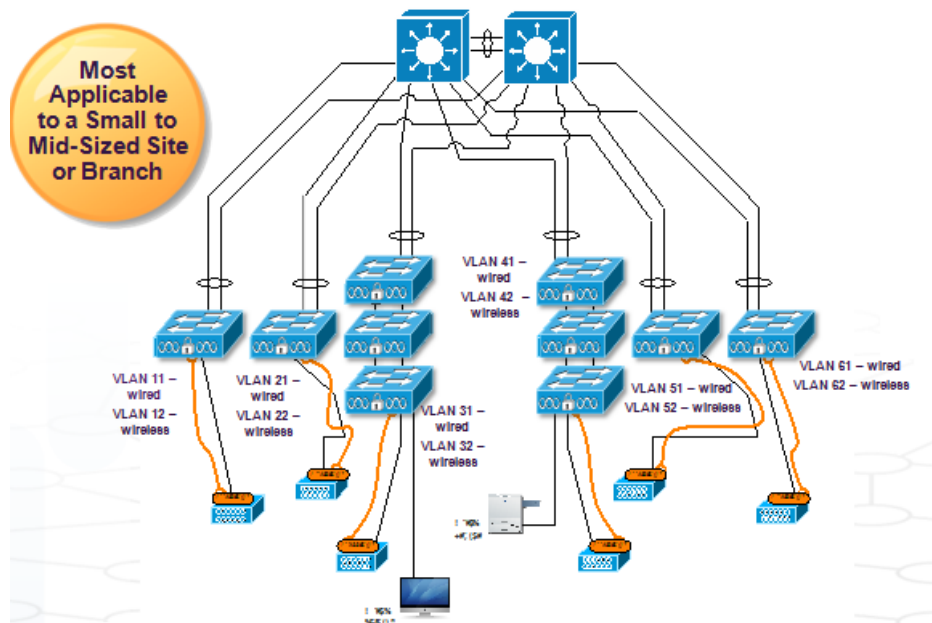
### Option 1. Separate Wired and Wireless VLANs Per Wiring Closet

This option separates wired and wireless VLANs per wiring closet, as shown in the following figure. In this example, there is a pair of VLANs in each closet. This is a simple design that allows the application of separate policies per VLAN to wireless and wired users and eliminates any contention for DHCP between wired and wireless.

However, because wireless clients are moving, it is important to consider how large the subnet must be for that wiring closet to accommodate these non-static clients. For wired connectivity, it is necessary only to count the number of available ports. Wireless usage is much more dynamic, so it is harder to determine the size of the DHCP scope that is required, and thus some of the IP address space as allocated might be wasted simply to accommodate for the maximum possible number of wireless clients that could potentially appear on the network simultaneously.

This approach for IP addressing is applicable mainly to a small or medium sized site or branch, where predicting the maximum size of the wireless subnets needed is easier, based on user and device populations at the small to medium branch involved.

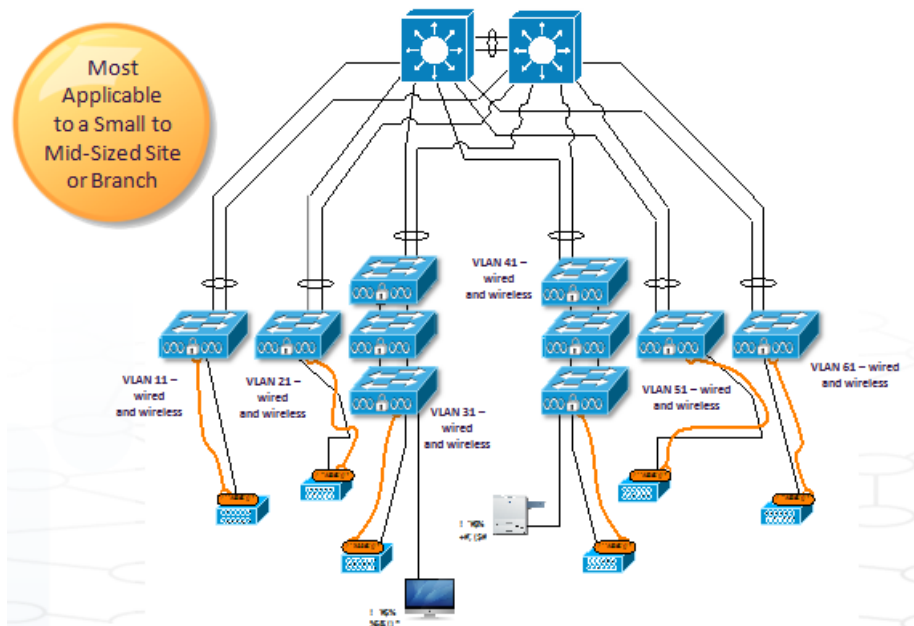
**Figure 54. IP Addressing – Option 1**



**Option 2. Merged Wired and Wireless VLANs Per Wiring Closet**

In this option, the VLANs are merged and the same subnet is used for wired and wireless for each wiring closet, but separated for different wiring closets. For example, VLAN 11 is used for wired and wireless on wiring closet one, VLAN 21 for the second and so on. The main advantage of this option is in saving IP subnets, and thus conserving the associated IP address space to the greatest extent possible. There is still the challenge of sizing subnets, and as well there is the possibility in this deployment option of IP address space contention between wired and wireless clients, since wired and wireless users are mapped into common subnets in this deployment option. Wireless clients could consume all of the IP addresses within a given subnet, resulting in insufficient addresses for wired clients (or vice versa). Moreover, it is not possible to apply separate wired and wireless policies using VLAN based policies alone in this deployment option.

Figure 55. IP Addressing – Option 2



### Option 3. Merged Wired and Wireless VLANs Per Wiring Closet with Spanned Wireless VLAN

This option is a hybrid with separate wired subnets and one wireless subnet spread across multiple wiring closets below a common distribution layer. This deployment option retains the advantage of a separate per-VLAN policy for both wired and wireless users, and avoids IP address space contention between these user communities, as wired and wireless clients are still mapped into separate VLANs. Fewer IP subnets are needed because wireless clients are grouped into a single VLAN (per SSID) below the distribution layer. This deployment option typically requires a VSS deployment at the distribution layer or a single distribution switch with multiple supervisors, to avoid Layer 2 loops and any associated spanning tree blocking / forwarding issues.





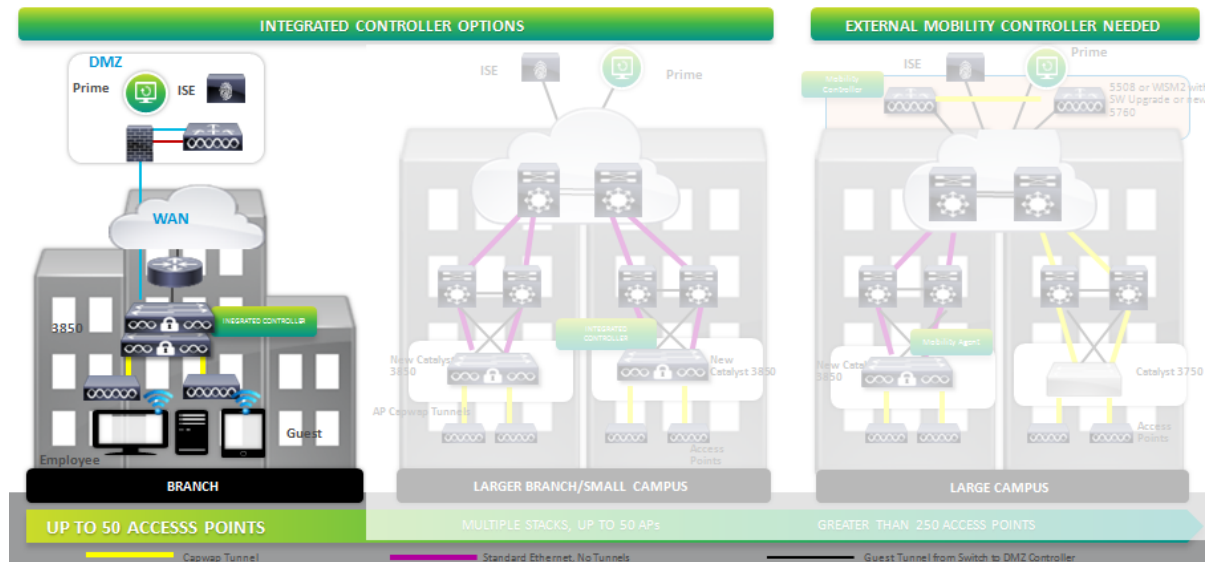
management interface enabled), and would find the centralized controller through DHCP option 43 or other methods.

Notes and important output are highlighted in blue underlined font in the following command examples.

### Basic Deployment Example

The following figure shows a basic branch deployment with fewer than 50 APs. A Catalyst 3850 or a stack acts as MA and MC. This section shows how to configure this switch for a client to join a wireless LAN (WLAN).

**Figure 57. Branch Deployment Example**



The switch acts as Layer 3 switch, so the example defines a management interface, VLAN 31 and client SVIs. The IP address on this VLAN is 192.168.31.42. There are two client VLANs, 32 and 33.

Management VLAN configuration:

```
interface Vlan31
  description MANAGEMENT VLAN
  ip address 192.168.31.42 255.255.255.0
```

SVIs for client VLANs defined locally on the switch:

```
interface Vlan32
  description Client VLAN32
  ip address 192.168.32.2 255.255.255.0
interface Vlan33
  description Client VLAN33
  ip address 192.168.33.2 255.255.255.0
```

Wireless management interface configuration

```
3850(config)#wireless management interface VLAN31 ← This activates the MA functionality
```

```
3850#show wireless Interface summary
```

```
Wireless Interface Summary
  AP Manager on management Interface: Enabled
Interface Name Interface Type VLAN ID IP Address      IP Netmask      MAC Address
-----
Vlan31          Management    31      192.168.31.42  255.255.255.0  2037.06ce.0a55
```



Because this is a branch location, the MC functionality is also enabled so that all functionality is in one system.

The system prompts to save the reconfiguration and to reboot. After reboot, the **show wireless mobility summary** command shows the mobility role of this switch and other important information. In this case, it is the only switch in the MD.

```
3850(config)#wireless mobility controller <= This activates the MC functionality
Mobility role changed to Mobility Controller
Please save config and reboot the whole stack
```

```
3850#sh wireless mobility summary <= After reboot
Mobility Controller Summary:
Mobility Role           : Mobility Controller
Mobility Protocol Port   : 16666
Mobility Group Name      : default
Mobility Oracle IP Address : 0.0.0.0
DTLS Mode                : Enabled
Mobility Domain ID for 802.11 : 0xac34
Mobility Keepalive Interval : 10
Mobility Keepalive Count  : 3
Mobility Control Message DSCP Value : 0
Mobility Domain Member Count : 1
Link Status is Control Path Status : Data Path Status
Controllers configured in the Mobility Domain:
IP          Public IP      Group Name      Multicast IP      Link Status
-----
192.168.31.42  -                default         0.0.0.0           UP : UP
```

The next step is to add an AP. Recall that VLAN 31 is the management VLAN in this example. At the interface level, the **switchport mode access** command is configured on the appropriate switchports that have APs locally attached to the Catalyst 3850. After that is done, the AP is connected and registers to its upstream directly-connected Catalyst 3850 switch, which can be confirmed by the **show ap summary** command.

```
interface GigabitEthernet1/0/15
  description - Access port for Access points
  switchport access vlan 31
  switchport mode access

3850#show ap summary
Number of APs: 1
Global AP User Name: Not configured
Global AP Dot1x User Name: Not configured
AP Name          AP Model  Ethernet MAC      Radio MAC          State
-----
AP3502I          3502I    c47d.4f3a.ed80    04fe.7f49.58c0    Registered
```

The last step is to configure a WLAN and SSID. The example shows a WPA-PSK WLAN in which the client is assigned to VLAN 32.

```
3850(config)#wlan WPA-PSK 4 wpa-psk
3850(config-wlan)#client vlan 32
3850(config-wlan)#no security wpa akm dot1x
3850(config-wlan)#security wpa akm psk set-key ascii 0 Cisco1234
3850(config-wlan)#no shut
```

Finally, the **wireless client summary** command verifies that the client is connected and displays the WLAN and wireless band. The **show wcdb database** command displays important information regarding mobility. In this case, the client is registered locally to that controller, so the mobility is shown as local.

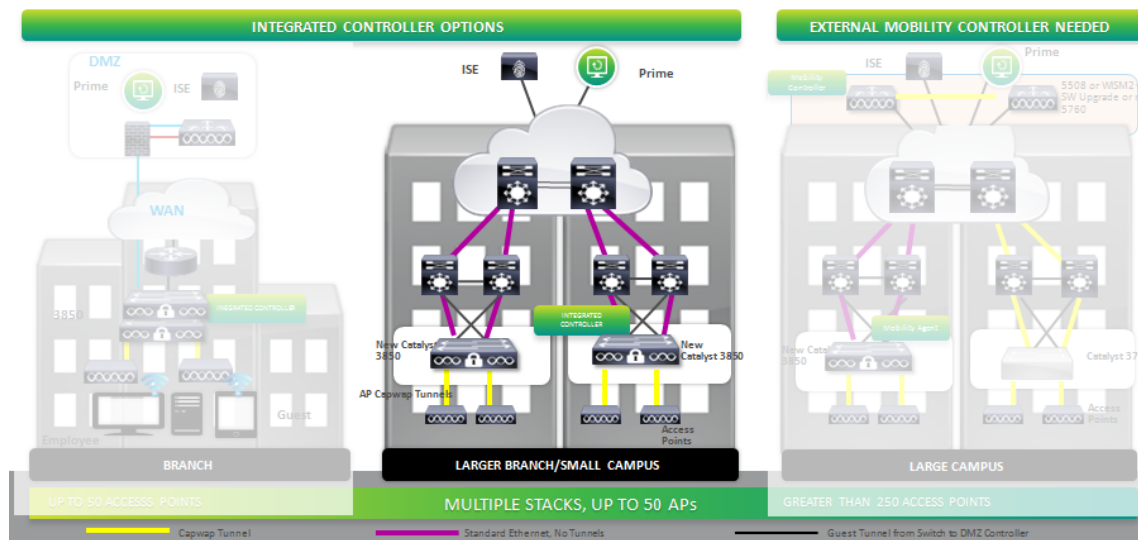
```
3850r#sh wireless client summary
Number of Local Clients : 1
MAC Address      AP Name
-----
f81e.dfe2.e80e  AP3502I
WLAN State      Protocol
-----
4              UP              11n(5)

3850#sh wcdb database all
Total Number of Wireless Clients = 1
Clients Waiting to Join = 0
Local Clients = 1
Anchor Clients = 0
Foreign Clients = 0
MTE Clients = 0
Mac Address      VlanId IP Address      Src If      Auth      Mob
-----
f81e.dfe2.e80e   32 192.168.32.57  0x00FF5BC00000011  RUN      LOCAL
```

### Small Campus Deployment Example

The following figure shows a large branch or small campus deployment. The number of ports and switches is increased, but from a wireless perspective there are still only 50 APs and one controller.

Figure 58. Small Campus Deployment Example



Setting up this example requires connecting and configuring another switch in the same SPG. The **wireless mobility controller peer-group** command assigns a name to the SPG. An MA is then assigned to the SPG with the IP address of the new switch. VLAN 21 is configured, and the controller is identified to the MA. The show wireless mobility summary shows the result with the data plane and control plane both up.

```
3850-MC(config)#wireless mobility controller peer-group GroupABC
3850-MC(config)#wireless mobility controller peer-group GroupABC member ip 192.168.31.44 public-
ip 192.168.31.44
```

On the MA side, all that is needed is to turn on MA and point to the MC.

```
interface Vlan21
  description MANAGEMENT VLAN
  ip address 192.168.31.44 255.255.255.0
```

```
3850-MA(config)#wireless management interface VLAN 21 ← This activates the MA functionality
3850-MA(config)#wireless mobility controller ip 192.168.31.42 ← This points MA to the MC
```

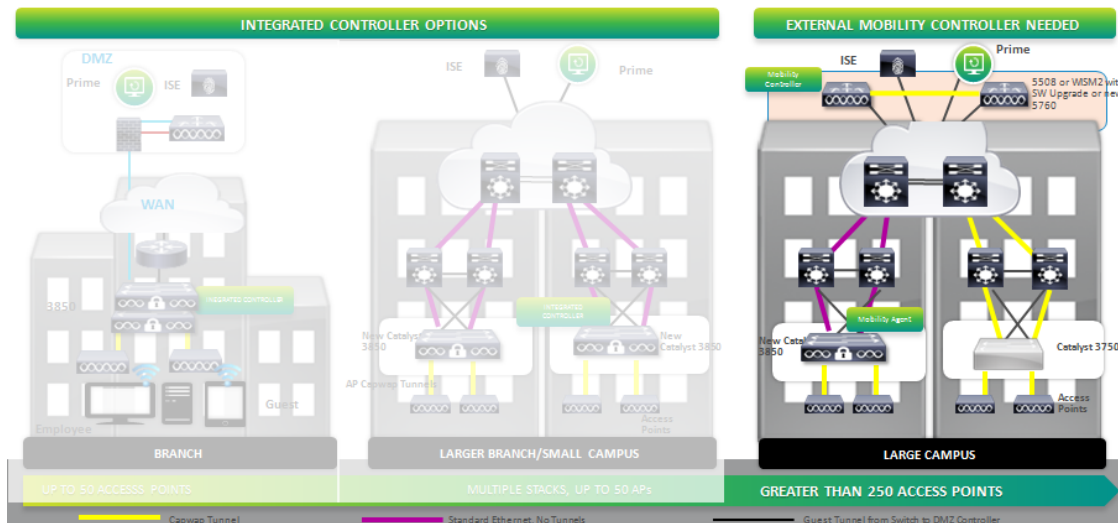
The configuration can be verified on the MC:

```
3850-MC#sh wireless mobility summary
Mobility Controller Summary:
Mobility Role           : Mobility Controller
Mobility Protocol Port  : 16666
Mobility Group Name     : default
Mobility Oracle IP Address : 0.0.0.0
DTLS Mode               : Enabled
Mobility Domain ID for 802.11 : 0xac34
Mobility Keepalive Interval : 10
Mobility Keepalive Count  : 3
Mobility Control Message DSCP Value : 0
Mobility Domain Member Count : 1
Link Status is Control Path Status : Data Path Status
Controllers configured in the Mobility Domain:
IP                Public IP          Group Name          Multicast IP      Link Status
-----
192.168.31.42    -                    default            0.0.0.0          UP : UP
Switch Peer Group Name : GroupABC
Switch Peer Group Member Count : 1
Bridge Domain ID      : 0
Multicast IP Address  : 0.0.0.0
IP                Public IP          Link Status
-----
192.168.21.44    192.168.31.44    UP: UP
```

## Large Campus Deployment Example

The following figure shows a large campus deployment with more than 250 APs. As recommended, it includes a discrete controller, such as WLC 5508 or WiSM-2 or WLC 5760.

Figure 59. Large Campus Deployment Example



The following commands configure the WLC 5760 as an MC and member of an SPG.

```
interface Vlan21
  description MANAGEMENT VLAN
  ip address 192.168.21.42 255.255.255.0

5760(config)#wireless management interface VLAN21
5760(config)#wireless mobility controller peer-group GroupABC
5760(config)#wireless mobility controller peer-group GroupABC member ip 192.168.21.44 public-ip
192.168.21.44
```

The following commands configure the Catalyst 3850 as an MA.

```
interface Vlan21
  description MANAGEMENT VLAN
  ip address 192.168.21.44 255.255.255.0
3850(config)#wireless management interface VLAN21
3850(config)#wireless mobility controller ip 192.168.21.42
```

The next command examples show the addition of another MG, for example, if a WLC 5508 is acting as a centralized controller for CUWN employment, and it is required to merge the two networks. The configuration defines the wireless MG name and member. If required, NAT is supported between controllers acting as MCs, as well between MC and MA. NAT is not supported between MAs.

After the configuration is done, the **show wireless summary** command displays the changes: one switch has the MC role and there are two members of the MG.

MG configuration:

```
5760(config)#wireless mobility group name sevt-lab
5760(config)#wireless mobility group member ip 10.1.1.5 public-ip 10.1.1.5
```

Verify the configuration:

```
5760-simo# sh wireless mobility summary
```

```
Mobility Controller Summary:
Mobility Role                : Mobility Controller
Mobility Protocol Port       : 16666
Mobility Group Name          : sevt-lab
Mobility Oracle               : Disabled
Mobility Oracle Ip Address   : 0.0.0.0
DTLS Mode                    : Enabled
Mobility Domain ID for 802.11 : 0x2fee
Mobility Keepalive Interval  : 10
Mobility Keepalive Count     : 3
Mobility Control Message DSCP Value : 0
Mobility Group Members Configured : 2
```

Controllers configured in the Mobility Domain:

IP Address	Public IP Address	Group Name	Multicast IP	Status
192.168.21.42	-	sevt-lab	0.0.0.0	UP
10.1.1.5	10.1.1.5	sevt-lab	0.0.0.0	UP

Switches configured in Group20 switch Peer Group: 1

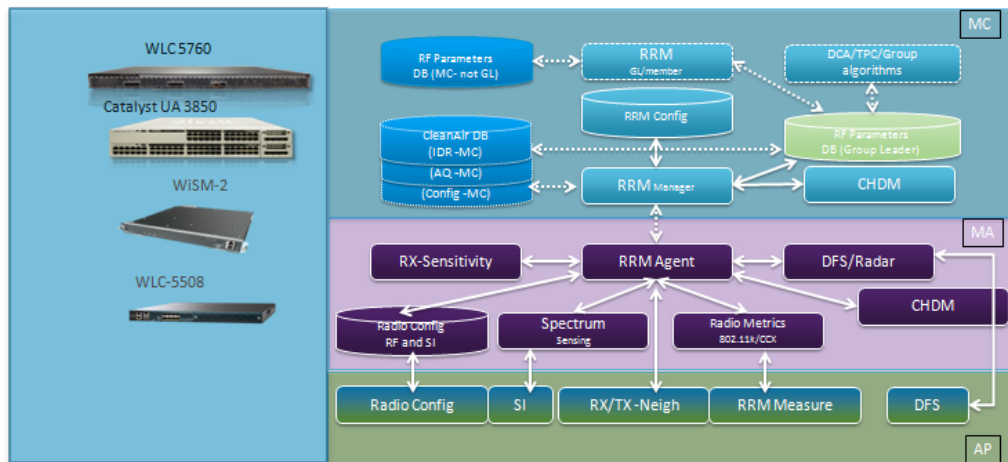
IP Address	Public IP Address	Status
192.168.21.44	192.168.21.44	UP

## RRM and CleanAir

This section describes the RRM and CleanAir implementations for Converged Access. The following figure shows how RRM and CleanAir functionality are split between the three different entities: MC, MA, and AP.

All RRM, Dynamic Frequency Selection (DFS), and CleanAir configuration is done at the MC level. Whereas the MA has an RRM agent that talks to the RRM manager on the MC, no RRM configuration is needed on the MA. A spectrum analysis process operating on the AP provides needed instrumentation for CleanAir. All measurements are collected and sent to the MC, where the RRM process coordinates channel selection and power levels across all APs.

Figure 60. Overview of RRM and CleanAir



## RRM

In RRM as in CUWN, an MC is the group leader for the RF group. If the CUWN AireOS controller is running 7.3.112.0 or 7.5, it can participate in the RF group with a Converged Access controller.

For RRM interoperability between controllers, the RRM version is important for compatibility. If a controller is not running a compatible release, an error message such as the following is generated. This is important because if the two controllers cannot form a group, they see each other's APs as rogues.

```
(Cisco Controller) >*rrm Socket Task: Sep 12 16:59:30.520: Airewave Director:  
Received incompatible RRM protocol or header version (30,2) from 192.168.10.101, our version  
(30,1)
```

When implementing RRM and CleanAir, one may want to check configuration parameters that would not normally need to be checked following an upgrade. For example, 1 Mbps is the mandatory minimum data rate by default. Default channels should be fine for most cases, but it is helpful to check and verify.

All pertinent RRM information is listed in the output of the **show tech-support wireless** command. Both static and dynamic RF groupings are supported. The following commands statically define the group leader MAC address and IP address.

```
(config)#ap dot11 24 rrm group-mode leader  
(config)#ap dot11 24 rrm group-member <mac_addr> <IP addr>
```

ClientLink is disabled by default and can be enabled using the following command.

```
(config)#ap dot11 24/5ghz beamforming
```

Use the following command examples to obtain information about the RRM configuration. The following command shows AP-specific RRM metrics at the MA (local switch).

```
3850#sh ap dot11 2 channel
Automatic Channel Assignment
Channel Assignment Mode           : AUTO
Channel Update Interval          : 600 seconds
Anchor time (Hour of the day)    : 0
Channel Update Contribution      : SN..
Channel Assignment Leader        : 5760 (192.168.10.101)

DCA Sensitivity Level            : MEDIUM (10 dB)
Channel Energy Levels
  Minimum                        : -82
  Average                        : -82
  Maximum                        : -82
Channel Dwell Times
  Minimum                        : 4 hours 0 minutes 13 seconds
  Average                        : 4 hours 0 minutes 13 seconds
  Maximum                        : 4 hours 0 minutes 13 seconds
802.11b Auto-RF Channel List
802.11b Auto-RF Allowed Channel List : 1,6,11
Auto-RF Unused Channel List      : 2,3,4,5,7,8,9,10
```

Use the **show ap dot11** commands to display information on the RRM RF group functions such as grouping, channel, and TX power on the MA or MC. Only devices with local APs show statistics.

```
5760#sh ap dot11 2 channel
Automatic Channel Assignment
Channel Assignment Mode           : AUTO
Channel Update Interval          : 600 seconds
Anchor time (Hour of the day)    : 0
Channel Update Contribution      : SN..
Channel Assignment Leader        : 5760 (192.168.10.101)
Last Run                         : 21 seconds ago

DCA Sensitivity Level            : MEDIUM (10 dB)
Channel Energy Levels
  Minimum                        : unknown
  Average                        : unknown
  Maximum                        : unknown
Channel Dwell Times
  Minimum                        : unknown
  Average                        : unknown
  Maximum                        : unknown
802.11b Auto-RF Channel List
802.11b Auto-RF Allowed Channel List : 1,6,11
Auto-RF Unused Channel List      : 2,3,4,5,7,8,9,10
```

```
3850#sh ap dot11 2 group
Radio RF Grouping
802.11b Group Mode               : AUTO
802.11b Group Update Interval    : 600 seconds
802.11b Group Leader            : 5760 (192.168.10.101)
802.11b Group Member            : Cisco_69:9a:64(192.168.10.8)
                               : 5760(192.168.10.101)
```

```
5760#sh ap dot11 2 group
Radio RF Grouping
802.11b Group Mode               : STATIC
802.11b Group Update Interval    : 600 seconds
802.11b Group Leader            : 5760 (192.168.10.101)
802.11b Group Member            : 5760(192.168.10.101)
                               : Cisco_69:9a64(192.168.10.8)
                               : 3850 (192.168.10.100) (*Not a Manager)
802.11b Last Run                 : 506 seconds ago
Mobility Agents RF membership information
```

```

-----
No of 802.11b MA RF-members : 1
MA Member name                IP address
-----
3850                          192.168.10.100

```

## CleanAir

CleanAir works the same as it does in CUWN, although some defaults need to be modified. PI 2.0 is required for any upper level display features. Information is available for all functions from the command line. For Spectrum Expert Connect (SE Connect) mode (connecting directly to the AP), CleanAir information can be displayed using the SE Connect application. All of the information and configuration is done at the MC, but the AP interface and radio for CleanAir can be enabled or disabled at the MA level.

Use the following commands to enable or disable ClearAir on a device (the example is for a 2.4GHz radio, it is similar for 5GHz).

```

(config)# ap dot11 24Ghz cleanair <= Enable
(config)# no ap dot11 24Ghz cleanair <= Disable

```

The following is an example of an AP radio level command at the MA.

```

3850#sh ap sum (use sh ap dot11 24Ghz/5GHz to see interfaces)
[snip]
AP Name                AP Model  Ethernet MAC  Radio MAC  Port
-----
AP0022.bd18.87c0      3502E    0022.bd18.87c0 0022.bdcc.d570 Gi1/0/1
(config)#ap name AP0022.bd18.87c0 dot11 24ghz/5ghz cleanair (cr) - enables cleanair on radio
2.4/5 GHz
(config)#ap name AP0022.bd18.87c0 no dot11 5ghz cleanair (cr) - to disable it

```

All CleanAir commands are processed on MC and passed to the MA.

```

5760#sh ap dot11 2 cleanair config
Clean Air Solution..... : Disabled
Air Quality Settings:
  Air Quality Reporting..... : Disabled
  Air Quality Reporting Period (min)..... : 15
  Air Quality Alarms..... : Enabled
  Air Quality Alarm Threshold..... : 35

```

Use the following command to display information from the MC.

```

5760# sh ap dot11 24Ghz cleanair ?
  air-quality  no description
  config      Displays CleanAir Configuration for 2.4GHz band
  device      no description
5760# h ap dot11 24Ghz cleanair device type ?
  all         Displays all CleanAir Interferers for 2.4GHz band
  bt-discovery Displays CleanAir Interferers of type BT Discovery for 2.4GHz band
  bt-link     Displays CleanAir Interferers of type BT Link for 2.4GHz band
<snip>
5760# sh ap dot11 24Ghz cleanair air-quality summary/worst
AQ = Air Quality
DFS = Dynamic Frequency Selection
AP Name                Channel  Avg AQ  Min AQ  Interferers  DFS
-----
AP0022.bd18.87c0      11      99     99     0           No

```

The CleanAir configuration is generally under the CleanAir tag. The only exceptions are Event Driven RRM and Persistent Device Avoidance, which are dynamic channel assignment (DCA) functions and are found under the RRM channel configuration.

```
5760(config)#ap dot11 2 rrm channel ?
cleanair-event  Configure cleanair event-driven RRM parameters
dca             Config 802.11b dynamic channel assignment algorithm parameters
device         Configure persistent non-WiFi device avoidance in the 802.11b channel
assignment
foreign        Configure foreign AP 802.11b interference avoidance in the channel assignment
load           Configure Cisco AP 802.11b load avoidance in the channel assignment
noise         Configure 802.11b noise avoidance in the channel assignment
```

## Migration

This section considers migration issues when moving from a WLC 5508 in a centralized deployment mode to a WLC 5760 with the same functionality as a centralized controller for remote connected APs.

The following table compares a centralized WLC 5508 and WLC 5760. Scale grows significantly with the WLC 5760. Deployments using the WLC 5760 are restricted to local mode only (no support for FlexConnect or other deployment options are offered with the WLC 5760).

**Table 2. Comparison of WLC 5508 and WLC 5760**

Feature	WLC 5508	WLC 5760
Throughput	8 Gbps	60 Gbps Line-rate
Scale	500 APs, 7000 Clients	1000 APs, 12000 Clients
Data forwarding Modes	Local, Flex, Bridge, OEAP	Local Mode
Resiliency	AP SSO, N+1, HA SKU	N+1, Multiple LAG, HA SKU
QoS	Alloy (precious metal: Platinum, Gold, Silver, Bronze) QoS	Granular QoS (MQC)
Security	Dynamic ACLs, SGA (SXP), 802.11w	Downloadable ACLs, Dynamic ACLs
Roaming	Layer 2 and Layer 3 CCKM Fast Secure Roaming, Neighbor List, 802.11r, 802.11k	Layer 2 and Layer 3 CCKM Fast Secure Roaming (FSR)
Services	Bonjour, AVC, Static Netflow	Flexible Netflow
IPv6	IPv6 Client Mobility, First Hop Security	IPv6 Client Mobility, First Hop Security
BYOD	ISE 1.2, Single SSID, Device Sensor	ISE 1.2, Single SSID
CLI	Available	IOS CLI, Secure Shell, EEM/TCL/TK
Licensing	License PAK based on serial number	Right to use (RTU)

The deployment design is similar for both platforms. There is one wireless management interface per box, with up to six if the wireless interfaces are managed with the six physical ports. The WLC 5760 does have a routable service port (which is also available on the WLC 5508) and supports multiple LAGs.

IPv6 can be supported natively on the client VLAN interfaces and on the management interface; however, management is not supported for the wireless management interface. Some IPv6 features are supported, such as Layer 3 client mobility and DHCP. IPv6 addresses are currently not supported on APs.



It is possible to filter and adjust some multicast IPv6 messages such as Router Advertisements (RAs) so as not to overload the wireless clients. By default, multicast messages are dropped, so the router advertisement is blocked. The following example demonstrates how RA messages can be throttled to a maximum of five in every throttle period.

```
5760(config)#ipv6 nd ra-throttler policy Mythrottle
5760(config-nd-ra-throttle)#throttle-period 20
5760(config-nd-ra-throttle)#max-through 5
5760(config-nd-ra-throttle)#allow at-least 3 at-most 5
```

The interface must be defined towards the network where the RA guard is received as a trusted interface and it is necessary to specify its connection to a router.

```
5760(config)#ipv6 nd rguard policy Mypolicy
5760(config-nd-rguard)#trusted-port
5760(config-nd-rguard)#device-role router
```

When policy is applied to the interface towards the network, RAs start to flow.

```
5760(config)#interface tenGigabitEthernet 1/0/1
5760(config-if)#ipv6 nd rguard attach-policy Mypolicy
5760(config-if)#ipv6 nd ra-throttler attach-policy Mythrottle
```

It is also possible to configure neighbor discovery suppression. For example, suppression can be configured at the client VLAN level on the MC and enforced at the AP level.

```
5760(config)#vlan configuration 19-21,23
5760(config-vlan-config)#ipv6 nd suppress
```

Hybrid deployments involve a mix of platforms. The new Converged Access mobility features are supported on AireOS 7.3 MR1 on the WLC 5508 and WiSM-2 platforms. MC (and optionally, MO capabilities) are supported on the WLC 5508 and WiSM-2 in this release.

The following conditions apply:

- Converged Access mobility is supported on AireOS 7.3.112.0 for the WLC 5508 and WiSM-2.
- MC and MO functions are supported on the WLC 5760. MA-only functionality for Converged Access APs is supported only on the Catalyst 3850.
- Seamless and CCKM fast roaming are supported between Converged Access and CUWN. All controllers need to be in the same MG. Roaming is always treated by default as a Layer 3 roam, and traffic is anchored at the home switch/controller.
- The WLC 5760 can terminate CAPWAP tunnels from APs connected to non-MA switches.
- Catalyst 3850 (acting as an MA) allows only APs to terminate CAPWAP locally. An AP cannot be connected to a Catalyst 3850 (acting as an MA) and be registered to a CUWN controller, although a Catalyst 3850 can have its MA functionality locally disabled, if desired, to allow for CAPWAP passthrough from an attached AP to an upstream discrete controller located elsewhere in the network.
- An AP failing over between a WLC 5760 and a WLC 5508/WISM-2 always reloads before becoming active.
- Seamless roaming is supported between a WLC 5508/WiSM-2 running 7.3.112.0 and a WLC 5760 running in Centralized mode. Seamless roaming is also supported between a WLC 5760 running in centralized mode and WLC 5760 running as an MC with a Catalyst 3850 running as an MA.

The following figure summarizes features and compatibility for Inter-Release Controller Mobility (IRCM) in a hybrid environment.

**Table 3. Hybrid Deployment Compatibility– IRCM**

CUWN Service	Y = Compatibility in Classic Flat Mobility							O = Compatibility in Hierarchal Mobility	
	4.2.x.x	5.0.x.x	5.1.x.x	6.0.x.x	7.0.x.x	7.2.x.x	7.3.101.0	7.3.112.0 Note: 1	IOS-XE 3.2.1 SE
Layer 2 and Layer 3 Roaming	Y	–	–	Y	Y	Y	Y	0	0
Wireless Guest Anchor/Termination	Y	Y	Y	Y	Y	Y	Y	0	0 <sup>2</sup>
wIPS & AwIPS Rogue Detection	Y	–	–	Y	Y	Y	Y	0	0 <sup>3</sup>
Fast Roaming (CCKM) in a mobility group	Y	–	–	Y	Y	Y	Y	0	0
Location Services	Y	–	–	Y	Y	Y	Y	0	0
Radio Resource Management (RRM)	Y	–	–	Y	Y	Y <sup>4</sup>	Y <sup>4</sup>	0 <sup>5</sup>	0 <sup>5</sup>
Management Frame Protection (MFP)	Y	–	–	Y	Y	Y	Y	0	0
AP Failover	Y	–	–	Y	Y	Y	Y	0 <sup>6</sup>	0 <sup>6</sup>

**NOTES:**

1. New Mobility is only supported on AireOS CT5508 & WISM-2 platforms but **does not** form any IRCM or GA with CT2500/CT7500/CT8500/v-WLC
2. Guest Anchor Termination is only supported on CT5760/CT5508/WISM-2. CT5760/CT5508/WISM-2/Cat3850 all supported as a Foreign
3. Rogue Detector Mode not supported
4. In Release 7.2 RF Profiles and groups was introduced. RRM for release 7.2 and later is not backwardly compatible with previous releases.
5. RRM Converged Access is compatible with CUWN release 7.3.112.0 but **does not** support RF Profiles and Groups introduced in 7.2
6. No AP SSO in IOS for CT5760. AP Intra-OS Platform Fast Failover Supported. AP Inter-OS Platform Image Download & Reboot performed.

The following table shows resiliency considerations.

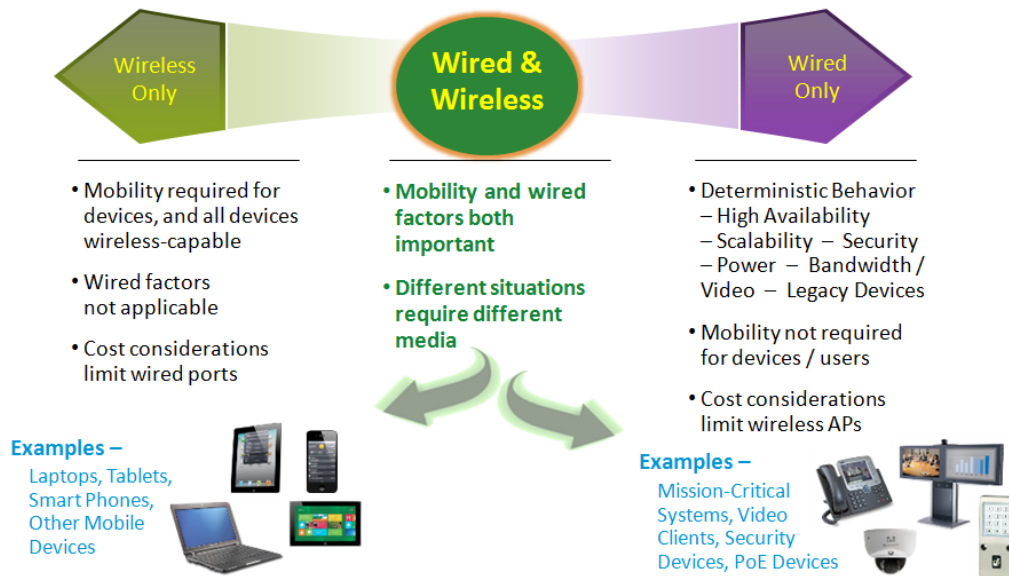
**Table 4. Resiliency Considerations**

Item	AireOS Before 7.3	AireOS 7.3	WLC 5760 in Centralized	WLC 5760 as MC	Catalyst 3850 as MA	Catalyst 3850 as MC
N:1, N:N:1 redundancy (AP related)	Yes	Yes	Yes	na	na	na
AP SSO	No	Yes	No (future)	na	na	na
AP SSO within the stack	na	na	na	na	na	Yes
Multiple LAGs	No	No	Yes	Yes	Yes	Yes

# Quality of Service

The following sections describe the capabilities and benefits of Quality of Service (QoS) in the Converged Access environment. The goal of QoS in Converged Access is to apply differentiated traffic handling as appropriate, regardless of the medium (wired or wireless). The following figure provides an overview of the characteristics on the wired and wireless sides and the need for coordination to implement consistent and appropriate policies.

Figure 61. Wired and Wireless Considerations for QoS



## QoS Elements

QoS is made up of the following elements:

- **Classification.** Classification is identification of traffic based on a number of possible factors, including ACL, Differentiated Services Code Point (DSCP), and NBAR. For example, the Catalyst 3850 classifies traffic based on ACL and DSCP.
- **Marking or mutation.** When traffic is marked, QoS operations on that traffic can be applied. This can be accomplished directly with the `set` command or through a table map, which takes input values and translates them directly to values on output.
- **Shaping and policing.** When a burst of traffic goes above a specific rate, shaping takes the traffic and plays it out over time instead of dropping it, as is done with policing.
- **Queuing.** This is the ability to take traffic and fill up a queue for servicing and scheduling based on bandwidth allocation. The traffic is then scheduled or sent out through the port.
- **Bandwidth allocation.** This determines the available capacity for traffic that is subject to QoS policies.
- **Trust.** In some cases, a policy requires that certain traffic is allowed to pass through, regardless of its QoS value.

## Existing QoS deployments and Challenges

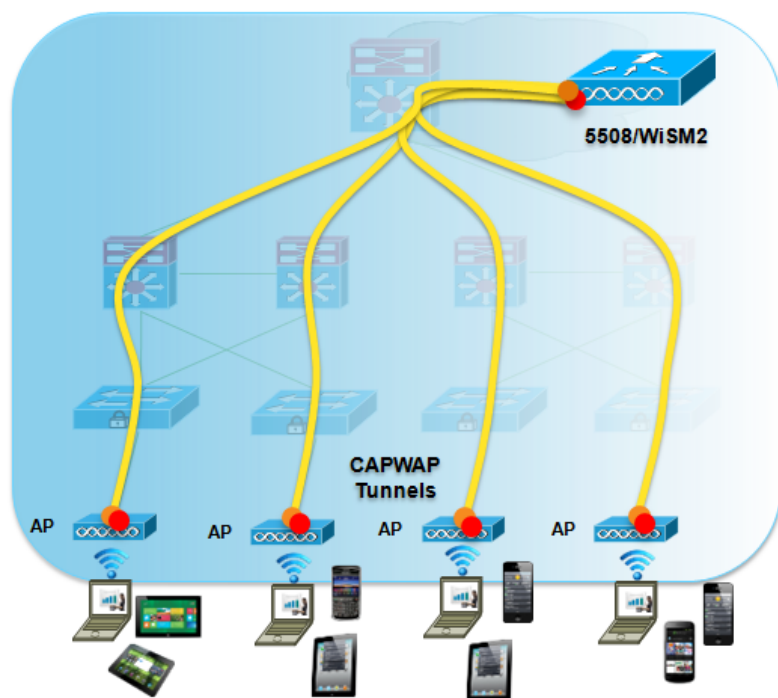
The advent of new devices and BYOD brings challenges for QoS. It is necessary to provide more applications at the access edge, especially in the wireless domain. In addition to voice and video, this can include data-based applications like Citrix and SAP.

The existing CUWN mobility architecture uses a CAPWAP tunnel between the APs, for example, directly tied to the WLC 5508. There are challenges with this overlay approach, because it does not provide full visibility into the applications that should be segmented and separated for proper QoS treatment, as the traffic flows through the network infrastructure.

The existing approach is to look at the WMM value (user priority, or UP) coming in from the wireless client, translate this to a DSCP value, and pass it into the network. However, some devices such as the iPhone 5 do not mark their traffic in a way that is useful from a business perspective. For example, FaceTime might be given a user priority value, whereas a Jabber client might be given a value based on the endpoint. It is preferable to examine the traffic coming from the end device, and mark it appropriately based on the business QoS policy.

Limited visibility in the existing environment leads to a lack of granular classification. Moreover, the WLC 5508 and WiSM-2 perform QoS in software. This is different from the wired world, which typically features hardware-based QoS.

**Figure 62. Existing Wired QoS Environment**



In the existing environment from the classification/marketing perspective, the WLC chooses a profile (platinum, gold, silver, or bronze). The profile is attached to a WLAN or an SSID and pushes the policy down to the AP, which enforces the classification and does marking based on the UP value that was set in the profile. WMM (user priority) client marking is allowed up to the profile value and non-WMM client traffic is marked to the profile value. The DSCP value is set in the CAPWAP header corresponding to the marking.

For policing, per-user bandwidth contracts are applied downstream at WLC and upstream at the AP, and per SSID (per AP and per radio) bandwidth contracts are applied upstream and downstream at the AP.

### Existing Campus QoS Architecture

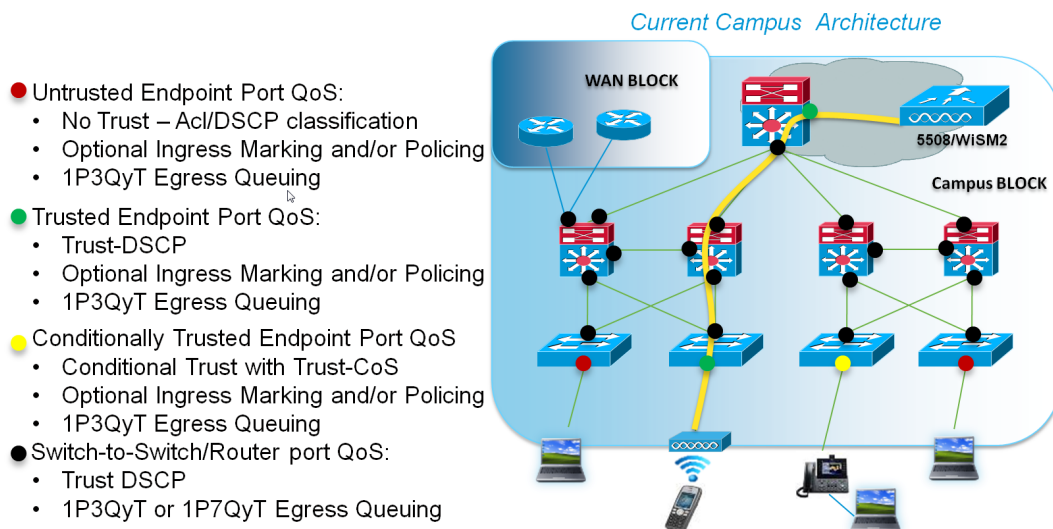
The wireless side is treated differently from the wired side in the CUWN architecture. On the wireless side, the endpoint is first examined and classified based on the endpoint.

In the following figure, the endpoints on the far right and left are untrusted (red). A QoS policy might remark this Citrix, SAP, voice, or video based traffic to a certain DSCP value, but the original values should not be simply be trusted and passed into the network.

The trusted endpoints (green) are set up for wireless APs and WLCs today. There is no visibility to differentiate different traffic types, so it is necessary to trust the AP and the WLC to mark the traffic appropriately. When that is done, the outer DSCP value in the CAPWAP-encapsulated packet is passed further into the network, unchanged.

The conditionally trusted endpoints (yellow) such as IP phones are validated with CDP and trust is extended, allowing them to send the DSCP values, the CS3 and EF in for their voice bearer traffic to move through the network. Everywhere else space is queued based on the DSCP values. A QoS configuration must be created at each point in the figure.

**Figure 63. Wired QoS Environment**



The following commands show how QoS is enabled in this wired environment. There are numerous detailed commands to set thresholds and other parameters. The commands are too granular for many users, an issue that is resolved with the Catalyst 3850. (This platform uses MQC, unlike the preceding Catalyst 3750 switch platforms, which used the older MLS QoS configuration constructs.)

```
C3750-X(config)#mls qos
C3750-X(config)#interface GigabitEthernet 1/0/1
C3750-X(config-if)#mls qos trust dscp

C3750-X(config)#mls qos queue-set output 1 buffers 15 30 35 20
C3750-X(config)#mls qos queue-set output 1 threshold 1 100 100 100 100
C3750-X(config)#mls qos queue-set output 1 threshold 2 80 90 100 400
C3750-X(config)#mls qos queue-set output 1 threshold 3 100 100 100 400
```

```

C3750-X(config)#mls qos queue-set output 1 threshold 4 60 100 100 400
C3750-X(config)#mls qos srr-queue output dscp-map queue 1 threshold 3 32 40 46

C3750-X(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 16 18 20 22
C3750-X(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 26 28 30 34 36 38
C3750-X(config)#mls qos srr-queue output dscp-map queue 2 threshold 2 24
C3750-X(config)#mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
C3750-X(config)#mls qos srr-queue output dscp-map queue 3 threshold 3 0
C3750-X(config)#mls qos srr-queue output dscp-map queue 4 threshold 1 8
C3750-X(config)# mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14

C3750-X(config)#interface range GigabitEthernet1/0/1-48
C3750-X(config-if-range)# queue-set 1
C3750-X(config-if-range)# srr-queue bandwidth share 1 30 35 5
C3750-X(config-if-range)# priority-queue out

```

## QoS Policies

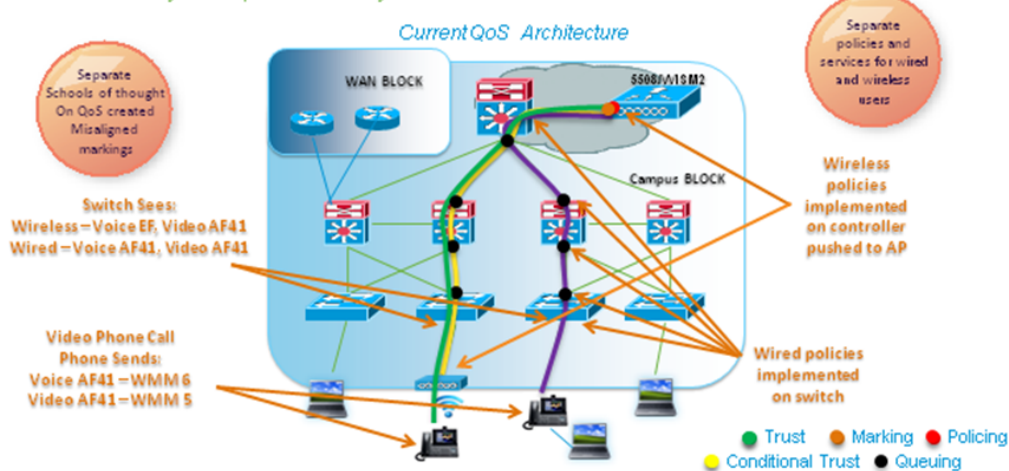
The following figure shows how policies are overlaid in the CUWN environment. Consider a wireless call made from the phone at the bottom left to the wired phone on the right. In the example, traffic coming in from the wireless phone is trusted, and the wired IP phone is trusted conditionally. So, different policies must be configured in the switch network even though they are protecting the same traffic over the same path and passing through essentially the same devices. Management, configuration, and deployment are all distributed.

In addition, if both the AP and the wired phone are attached to the same switch, an example can be seen of the major difference between wired and wireless QoS. Traffic from the device (9971 phone) is marked with WMM 5,6 for video and voice streams and DSCP AF41 for voice and video streams if video is configured in the Cisco Unified Communications Manager (CUCM) for these endpoints. Because the wireless side uses WMM to mark the outer DSCP value and wired uses DSCP, the voice and video streams from the same device are marked differently depending on how the device is attached. Wired and wireless QoS in this example thus have divergent treatment.

**Figure 64. Policies Overlay in Existing QoS Environment**

## Existing QoS deployments

How we overlay QoS policies today



## QoS with Converged Access

Converged Access provides enhanced QoS on the wired and wireless sides. On the wired side, the modular QoS based CLI (MQC) supports alignment with 4500E series (Sup6, Sup7) and provides for class-based queuing, policing, shaping, and marking. The number of queues and queue sets is increased and flexible MQC provisioning abstracts the queuing hardware.

On the wireless side, Converged Access supports granular QoS control at the wireless edge. Tunnel termination allows QoS treatment per SSIDs and per client, and common treatment of wired and wireless traffic throughout the network. Enhanced bandwidth management with Approximate Fair Drop (AFD) is an innovative and powerful new QoS enforcement technique pioneered by the Catalyst 3850 platform. AFD ensures fairness at the client, SSID, and radio levels for Non-Real-Time (NRT) (non-strict-priority) traffic, and wireless specific interface control allows policing capabilities per-SSID and for upstream and downstream clients. AAA support is provided for dynamic client-based QoS and security policies for per SSID bandwidth management.

Many of the QoS features are supported because Converged Access provides local termination of the CAPWAP tunnel from the AP at the ingress Catalyst 3850 switch port to which that AP is attached, thereby providing the needed visibility to do granular ACL-based DSCP classification. Tunnel termination allows customers to provide QoS treatment per SSID and per-client, and common treatment of wired and wireless traffic throughout the network. SSID-based QoS policies are actually done at the BSSID level, not the SSID level, meaning that these policies can be applied per-AP and per-radio.

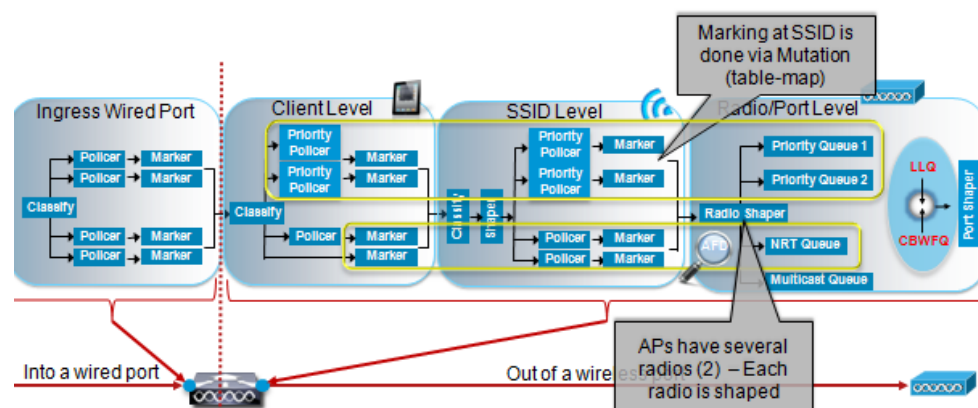
With the existing arrangement, a client in one SSID who decides to access bandwidth-intensive video streaming can overload the radio and cause drops for others who are doing business based applications. There is no fairness among those clients. With AFD, each SSID provides client-based fairness by default. The client who is demanding excessive bandwidth has traffic dropped at a higher proportional rate.

Wireless-specific interface control means that policing policies can be done per-SSID and per-client, upstream and downstream. AAA support is also available for dynamic client based QoS and security policies.

Bandwidth management can be done on a per-SSID basis, if needed (for instance) to manage guest traffic relative to other business traffic. As an example, the guest SSID could be allocated 10 percent of the bandwidth, while the enterprise clients are allocated 90 percent of the bandwidth. Client fairness is applied within the SSID, and bandwidth management is also provided between SSIDs on the same radio.

The following figure is a conceptual view of how wired to wireless policies can be applied with Converged Access.

**Figure 65. QoS Policies – Wired to Wireless**



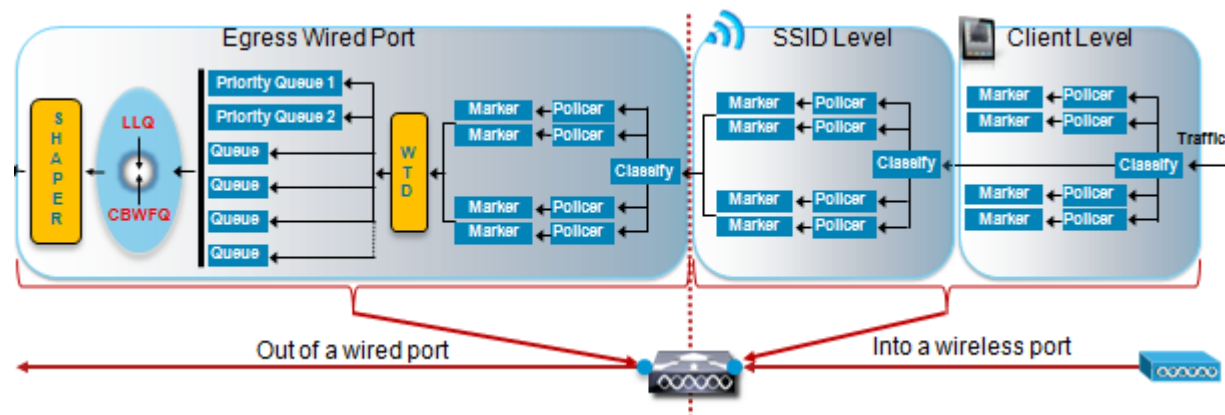


On the wired side, MQC provides a unified provisioning language on the access edge with class maps, policy maps, set commands, and policers, synonymous across platforms. The following figure shows the levels at which wireless to wired policies can be applied.

Policies can be applied at the ingress wireless and egress wired levels to classify, police, and mark traffic from the client. For granular policies on ingress, the client level can be used without a policy at the SSID level. The SSID level can be used to aggregate and police a group of users, for instance, for priority traffic or non-real time traffic. On egress for wired ports, traffic can be classified and marked. Weighted tail drop is used as the congestion control mechanism. New to the access layer switch portfolio is the ability to configure a hierarchical shaper. This provides the ability to limit upstream traffic for sub-rate circuits or to simply limit bandwidth while allowing for granular queuing upstream.

The following figure shows how wireless to wired policies can be applied.

**Figure 66. QoS Policies – Wireless to Wired**



## QoS Behavior

It is important to note the default behavior of any platform and consider the wired-wireless integration of the Catalyst 3850. The default QoS behavior for traffic traversing from wired-to-wireless or vice versa is untrusted. This means that traffic QoS markings are all reset to zero as they pass this boundary. Trust is provided only on traffic sent from wired-to-wired interfaces.

Each radio has a default shape rate based on the maximum rate of the specific radio, and the aggregation of the radio rates per AP is provided by an additional shaper set on that wireless port. As an example, if two radios are present within an AP, the aggregate bandwidths of the two radios are added together and that shape rate is placed on the physical port. The service policy, **port-child-policy**, is a statically defined default policy on the platform that is used to configure egress queuing for wireless ports. The **port-child-policy** has by default two configured classes or queues. One of them is the **class-default**, the other is the **non-client-nrt** queue. Class-default is not seen in the CLI, as it is a catch-all class, but the **non-client-nrt** class is shown in **port-child-policy**. Priority traffic is handled by internally using the first priority queue. The priority queues must be enabled on the physical port if they are to be used, similar to the implementation on the Catalyst 3750.

The Cisco Wireless Control Module (WCM) is the software functionality within the Catalyst 3850 or WLC 5760 platform that provides the bulk of the wireless functionality, operating under IOS-XE. WCM does all the required work and installs these policies as needed.

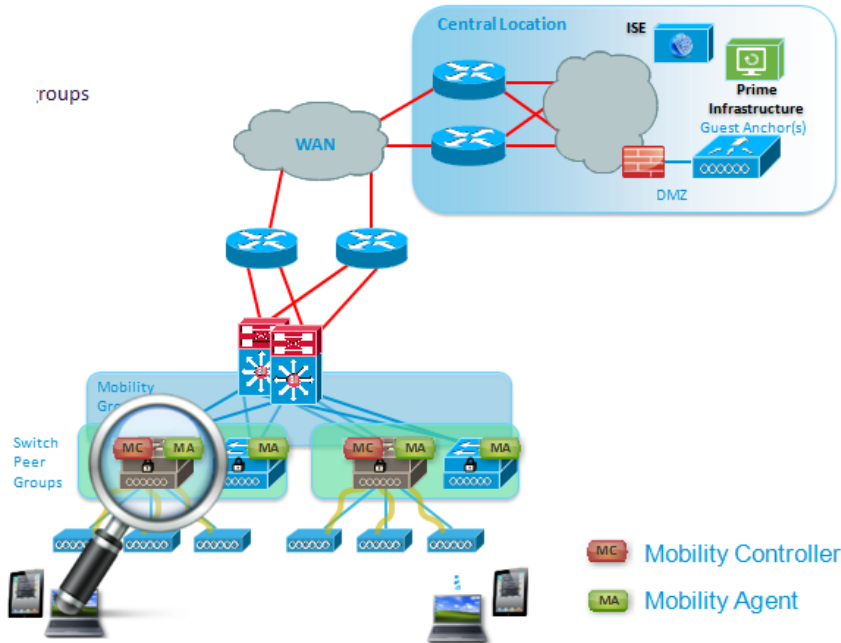


## QoS Design Options and Policies

From a QoS standpoint, the goal is to simplify the transition to MQC by using ISE, for instance, to deploy QoS policies. The QoS policies themselves are distributed just as they are today at the access edge on the Catalyst 3850. All of the client-based policies are in the global CLI, and ISE pushes the name of the policy down to the Catalyst 3850. It is a best practice to keep policies simple and reuse policies as often as possible.

The following figure shows an example deployment that illustrates some of the design options.

**Figure 67. Example of Design Options to Support QoS**



This campus example includes two SSIDs: faculty SSID and student SSID. For users, the QoS policy name is pushed from the ISE down to the Catalyst 3850 when a user authenticates and is authorized on that SSID. User groups (faculty/student) are separated by allocating bandwidth to each SSID. Users are provided fairness within the SSID courtesy of AFD, and specific per-user parameters are allocated to protect voice, video, and data applications for both faculty and students.

For a trust boundary, the following example illustrates how a policy is applied at the WLAN with a DSCP value so that traffic with a DSCP value that comes in is allowed to go out at the same DSCP value. Traffic is sent upstream without having to remark the DSCP value to zero.

The same policy is applied downstream with the addition of user-priority, which is necessary for wireless queuing for faculty and students, allowing traffic to be passed within the WLAN from wired to wireless and wireless to wired. A bandwidth allocation ratio of 90 to 10 percent allows limits students bandwidth for the SSID to 10 percent of the radio. The ratio is 90 percent for the faculty, to allocate a larger amount of bandwidth for the faculty than for the students. The policy can be specified for the WLAN for each of the VLANs.

```
table-map dscp2dscp  
default copy
```

```
Policy-map TRUST-BW-FACULTY
```

```

Class class-default
  set dscp dscp table dscp2dscp
  set wlan user-priority dscp table dscp2up
  bandwidth remaining ratio 90

table-map dscp2dscp
default copy

Policy-map TRUST-BW-STUDENTS
Class class-default
  set dscp dscp table dscp2dscp
  set wlan user-priority dscp table dscp2up
  bandwidth remaining ratio 10

```

When it comes to user-based policies, the user must first be authorized to associate to the SSID in question. This can be accomplished through ISE, but the credentials do not have to be defined by ISE (they could, for example, originate in Active Directory or another credential store to which ISE has access).

After the user is authenticated, the QoS policy name is pushed down to the Catalyst 3850 and tied to the client-specific MAC address, so that the client receives the treatment that is specified in the policy. This is a wireless policy based on MQC, which involves policing and marking. There are multiple classes, including voice, video, signaling, and transactional, just as in the wired domain today.

```

policy-map FACULTY
  class VOIP
    set dscp ef
    police 128000 conf transmit exceed drop
  class VIDEO
    set dscp AF41
    police 384000 conf transmit exceed drop
  class SIGNALING
    set dscp cs3
    police 32000 conf transmit exceed drop
  class TRANSACTIONAL-DATA
    set dscp af21
  class class-default
    set dscp default

```

For wired port uplinks, just as in any other campus based platform, queuing must be configured on a per-port basis. Priority levels are policed to a specific rate. There is a control and management queue with bandwidth allocation for the transactional, scavenger class, and class default traffic. All of this is marked at the client level. The traffic that comes in from the client level has markings treated and queued out to the uplink. Queuing also applied to wireless ports. There are four queues on the wireless side, all of which are subject to bandwidth allocation.

For mobility, the QoS policy is retained by, and enforced at, the PoA rather than the PoP. The client who associates gets a policy, which is already configured at the PoP. When the client roams, since the new policy is already on the other switch, a fast roam can occur between the Catalyst 3850s. The MC sends over the policy name associated with the MAC address of the client, and it applies the new QoS policy to the client on the new switch that it is attached to.

In summary, the QoS with Converged Access provides the following important features and benefits.

- Differentiated policy – per user
- Converged user policies – wired/wireless
- Simplified, more flexible CLI
- CLI alignment with other MQC-based switches

- Improved resource allocation
- QoS-based control at the access edge
- No change to wireless side

## Management

The following sections describe changes in PI to make Converged Access deployments easier.

### Cisco Prime Infrastructure

#### Features

The goal of Cisco Prime Infrastructure (PI) is to support the promise of one management, including lifecycle, assurance, and best practices and compliance. Lifecycle means that every device from Cisco can be managed, including wired to wireless, from campus to branch. Assurance encompasses how the user experience is driven, and how to keep track of what applications and services are running on the network.

Best practices and compliance include regulatory checking and best practice checking with Cisco Validated Designs for better analysis of what has been deployed in the network infrastructure. Auditing and reporting help administrators understand what has been done, how that relates to Cisco best practices, and what remediation should be done when issues arise. The result is better operational productivity services, improved user productivity, with effective best practices.

Figure 68. Cisco PI



#### Lifecycle Management

Cisco PI now includes design, deploy, operate, and administer services to support lifecycle management. The design process includes assessment and planning for new services, policies, and metrics. Deploy including scheduling the rollout of network changes, bringing up new sites, and defining audit baselines. Operation is day-to-day management with the aid of centralized health and alert monitoring and one-click workflows. Administration includes maintenance and update management services, HA, and integration.

## Platform and Technology Support

These capabilities are supported in the hardware platforms and are being added to PI 2.0 to support the new Catalyst 3850 and the new WLC 5760 controller. Support includes simplification of configuration and management and out-of-box templates and best practices for quick error-free deployment. There are new workflows for mobility architecture support and reports to improve lifecycle and network visibility on TrustSec 802.1x deployment, AVC, and IPv6 readiness.

## Simplifying Deployment

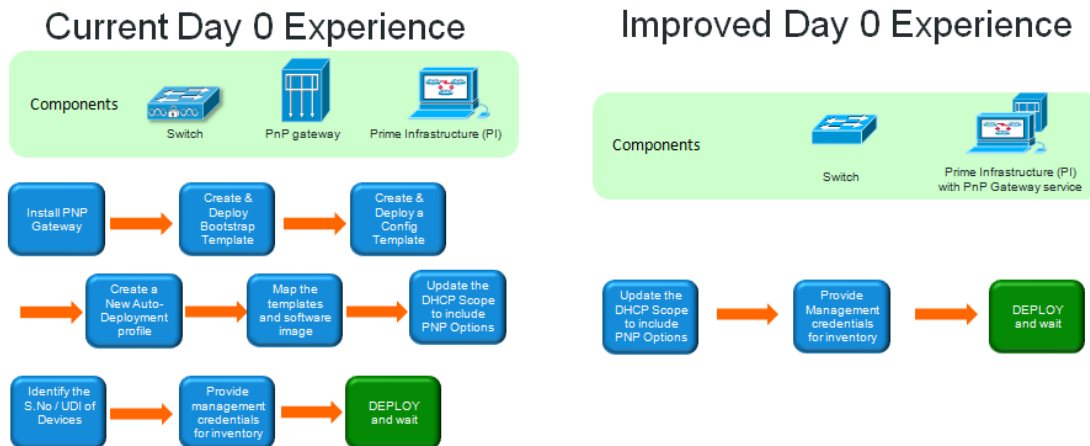
The 360 views in PI pull different information to show CPU and memory utilization in one window. The user 360 view helps to quickly isolate and fix end-user or end-point issues (such as response time, network access, and configuration). The device 360 view helps identify and fix device related problems (performance, faults, interface, and modules). And the application 360 views helps identify and fix network issues related to app delivery (app discovery, utilization, user/device/site, and association).

With these views tied together in one system that handles both wired and wireless services, Converged Access deployments become much easier. The single entry point allows administrators to solve access and experience problems, while embedded best practices simplify and improve management and troubleshooting for increased deployment and operational simplicity.

## Improved Day Zero Experience – Plug and Play

The following figure shows the existing day zero experience (left) and the much improved day 0 experience on the right. With PI 2.0, the plug and play (PNP) process has been improved to be simpler and easier to deploy. The PNP service is now built in to PI, where as in the past it was a separate device. The formerly nine step process is now a three-click process. There are still three ways to deploy a bootstrap configuration to a device on the network: through TFTP using a DHCP process, having PI email the bootstrap to the user, or saving the file on the local machine the administrator is signed in from.

Figure 69. Existing Day 0 Experience Versus Converged Access Day 0 Experience



When PI is up and running, the PNP setup can be initiated. The user selects **Workflows > Plug and Play Setup**. Completing this process will create a bootstrap configuration file.

The following example shows how to deploy the file within the TFTP service on PI.

The screenshot displays the Cisco Prime Infrastructure interface. At the top, the navigation bar includes 'Home', 'Design', 'Deploy', 'Operate', 'Report', 'Administration', and 'Workflows'. The 'Workflows' menu is expanded, showing 'Plug and Play Setup' and 'Initial Device Setup'. The main content area is titled 'Workflows: Plug and Play Setup' and features a progress bar with three steps: 'Before you Begin', 'Create Profile', and 'Save Profile'. Below the progress bar, the 'Before You Begin' section contains a detailed description of the wizard and a list of three methods for profile delivery: 1. DHCP-based Autoinstall, 2. Prime Utilities, and 3. File Transfer. A final note states that after devices 'call home', they can be found in the Notifications drawer.

In the next step, the administrator defines credentials and the location of the PNP gateway service. With the new changes to PI 2.0, this is actually the IP address of the PI server. Preliminary SNMP information is provided along with Telnet credentials to be set on the new network devices, to enable communication between the devices and PI. When the new network device has found PI, the initial device setup guided workflow can carry the administrator through a process to change these as needed.

The screenshot shows the 'Create Profile' step of the 'Workflows: Plug and Play Setup' in Cisco Prime Infrastructure. The interface includes a progress bar with three steps: 'Before you begin', 'Create Profile', and 'Save Profile'. Below the progress bar, the 'Create Profile' section contains the following fields and options:

- Pre-deploying on:** Campus Devices (selected), Configuration Details, Factory Defaults.
- Properties:**
  - Credentials:**  Show Clear Text
  - SNMP Fields:**
    - Read-Only Community String: [\*\*\*\*\*] Confirm: [\*\*\*\*\*]
    - Read-Write Community String: [\*\*\*\*\*] Confirm: [\*\*\*\*\*]
  - SSH/Telnet Credentials:**
    - Enable SSH  Enable Telnet
    - User Name: [joewill] Confirm: [\*\*\*\*\*]
    - Password: [\*\*\*\*\*] Confirm: [\*\*\*\*\*]
    - Enable Password: [\*\*\*\*\*] Confirm: [\*\*\*\*\*]
  - Plug and Play Gateway Location:**
    - PnP Gateway Host Name: [PI20]
    - PnP Gateway IP Address: [10.1.1.16]

At the bottom of the form, there are 'Cancel', 'Previous', and 'Next' buttons.

The next step is to save and deploy the profile.

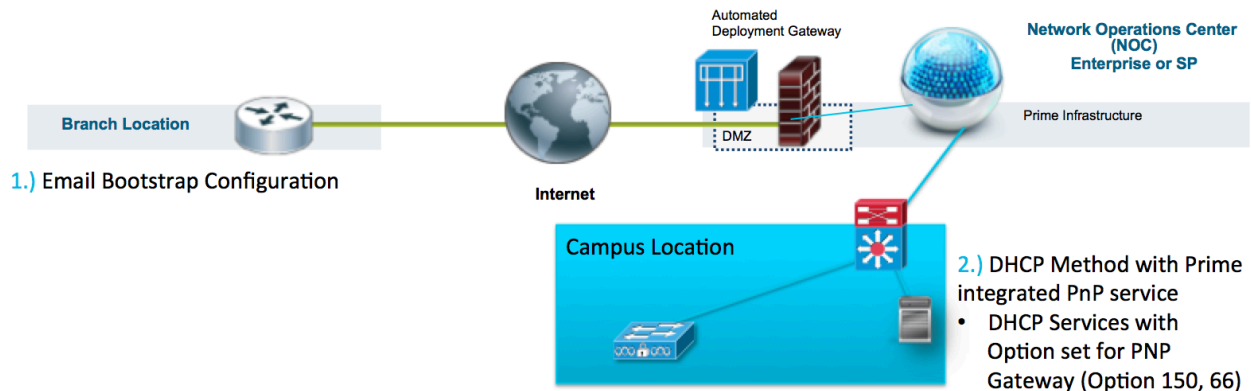
The screenshot shows the 'Save Profile' step of the 'Workflows: Plug and Play Setup' in Cisco Prime Infrastructure. The progress bar now shows 'Before you begin' and 'Create Profile' as completed steps with green checkmarks, and 'Save Profile' as the current step. The 'Save Profile' section contains the following options and fields:

- Deploy via:**  TFTP
- Email to other operators
- Export the file

Additional information and fields include:

- Deploy via TFTP:** Deploy via TFTP - means that the new IOS device can netboot and quickly call home to Prime Infrastructure for further configuration. Cisco recommends using this method when possible.
- Email to other operators:** Email to other operators - allows a bootstrap profile to be emailed to other operators who may then use with the iPhone or laptop utilities to configure new devices. This is useful when TFTP isn't available, for example at a branch site. Alternatively a PIN may be emailed. This PIN is used to download the bootstrap configuration from the Prime Infrastructure PnP gateway and then apply the associated bootstrap profile.
- What to email:**  Bootstrap  PIN
- To:** [example1@cisco.com, example2@cis]
- Export File:** Export File - allows the user to export a bootstrap profile as a file to the local system.

If a device has a defined template, that PNP option can be used.



One scenario is if the device is outside the DMZ. For this case, PNP or automated gateway service can be deployed in the DMZ. The bootstrap information can be defined on the device and the device can come up, find the deployment gateway outside the DMZ, and pull its software and configuration down. Internally, DHCP can be set to point to the TFTP server. When a device is plugged in to the network, it obtains an IP address through DHCP, and from option 150/66 the device will learn the IP address of the TFTP service. The device will request the bootstrap configuration using that IP address and the bootstrap configuration will direct the device to the PI server. If there is a configuration predefined for that devices based on the device ID, it will be downloaded. If not, the device will add itself to PI and the administrator will see on the bottom bar that a new network device has been added to PI.

Sample bootstrap configuration file:

```
ip host PI20 10.1.1.16
cns trusted-server all-agents PI20
cns id hardware-serial
cns id hardware-serial event
cns id hardware-serial image
cns event PI20 11013 keepalive 120 2 reconnect-time 60
cns exec 80
cns image server http://PI20/cns/HttpMsgDispatcher status http://PI20/cns/HttpMsgDispatcher
cns config partial PI20 80
cns config initial PI20 80
```

## Design and Deploy

Design templates incorporate configuration information which is then pushed to devices. Model-based templates are predefined within PI and can be focused on security, routing, or wireless controller services with parameters specified based on the device IDs. A second option is to use CLI-based templates that allow custom information, special banners on devices, and custom routing configuration. Special ACLs can be defined to be consistent with other devices within the network. The CLI template can be pushed to the network with device-specific information or generic information for the full network. Finally, a composite template can be created as a template of templates. If several different templates are created, for security, wireless, wired, routing, Layer 2, and Layer 3 services in your network, they can be combined into a composite based template which is then pushed to the network.

The templates can be pushed through the deploy process to specific devices, or PNP can be used. With PNP, the templates are provided to the plug and play gateway service. The gateway service checks for the target device IDs and deploys the templates on those devices.

PI 2.0 supports an alternative mechanism that does not use templates, but instead allows creation of configurations for the devices that are pulled into the network infrastructure.

The following figures show how initial device configuration works through the wizards. The figures assume that there is a new device on the network that has been brought into PI.

Initial device setup starts by choosing **Workflows > Initial Device Setup**.

**Workflows: Initial Device Setup**

Before You Begin → Assign to Site → Choose Other Devices → Configuration → Confirmation

**Before You Begin**

This wizard will help configure devices discovered via Plug and Play Setup or other Prime discovery mechanisms. The wizard currently supports the Catalyst 2xxx, 3xxx, and 4xxx switches, and 5760 controller. It does not include routers. Additional device support will be added in future releases.

The following features can be configured

Wired	Wireless (Currently only 3850 and 5760)
Management Details	Mobility Domain, Mobility Group and Swi...
Authentication Settings	Mobility Controllers and Mobility Agents
VLAN and Switching	Provision Enterprise Access
Uplink	Provision Guest Access
Enable Basic Services	Security and QoS policies

**More Details**

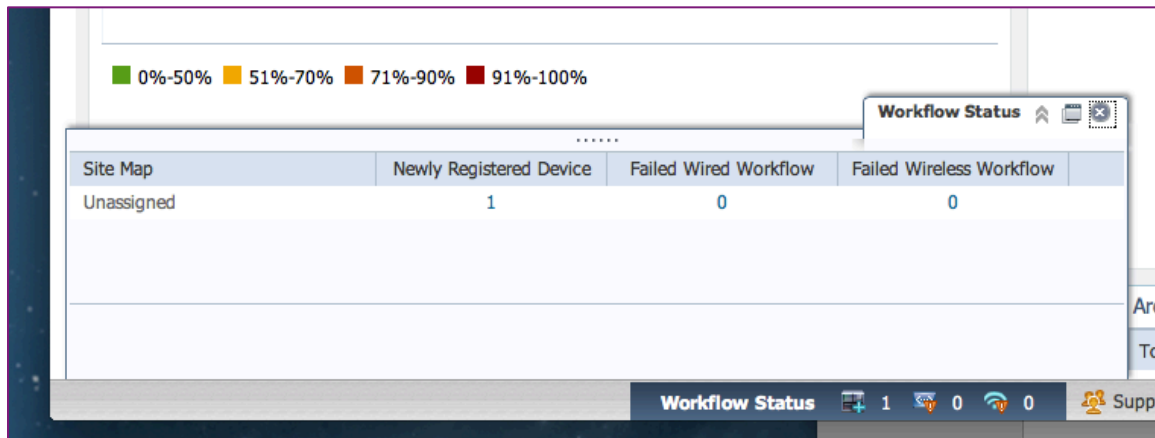
Downlink port configuration is done using Cisco Auto Smartports. All downlink ports on the target device are assumed to be Data Ports by default. If a Cisco IP Phone or access point is discovered then appropriate port configuration is automatically deployed.

**Before You Continue**

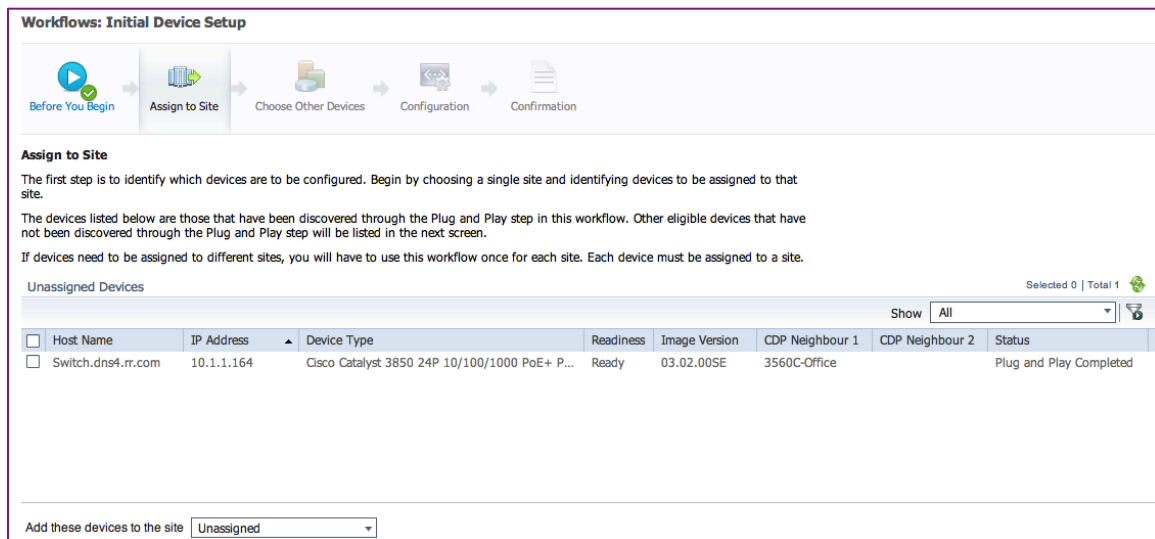
You will be assigning devices to a site, but sites must be defined before using this wizard. You may create new sites in Device Work Center by selecting 'Create Group' in the Device Group panel. You can also go directly to Site Map Design by using this [Link](#). A device can only be associated with a single site. This wizard configures one site at a time. So to assign devices to multiple sites you will need to use this wizard once per site. This wizard requires that Cisco Discovery Protocol (CDP) is enabled on the uplink switch or router that will be connecting to the new devices. To quickly configure using recommended values, use the default Guided mode. For additional configuration options, use Advanced mode.



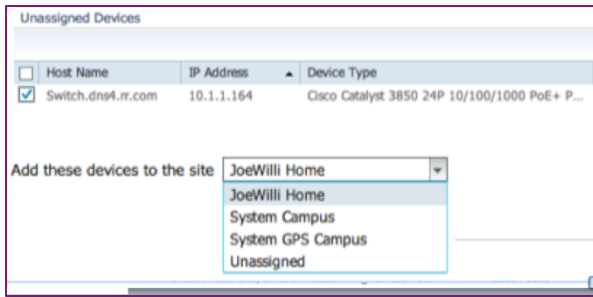
Sites are used in this workflow. The administrator should have the sites defined before continuing with the process. The first step is to select devices and place them in a site. When a new device is added to the network, the workflow for initial device setup can begin. The following example shows that a new device has been found by PI. In the bottom alert bar within PI, the workflow status has automatically updated and a new device is ready to be added to the system.



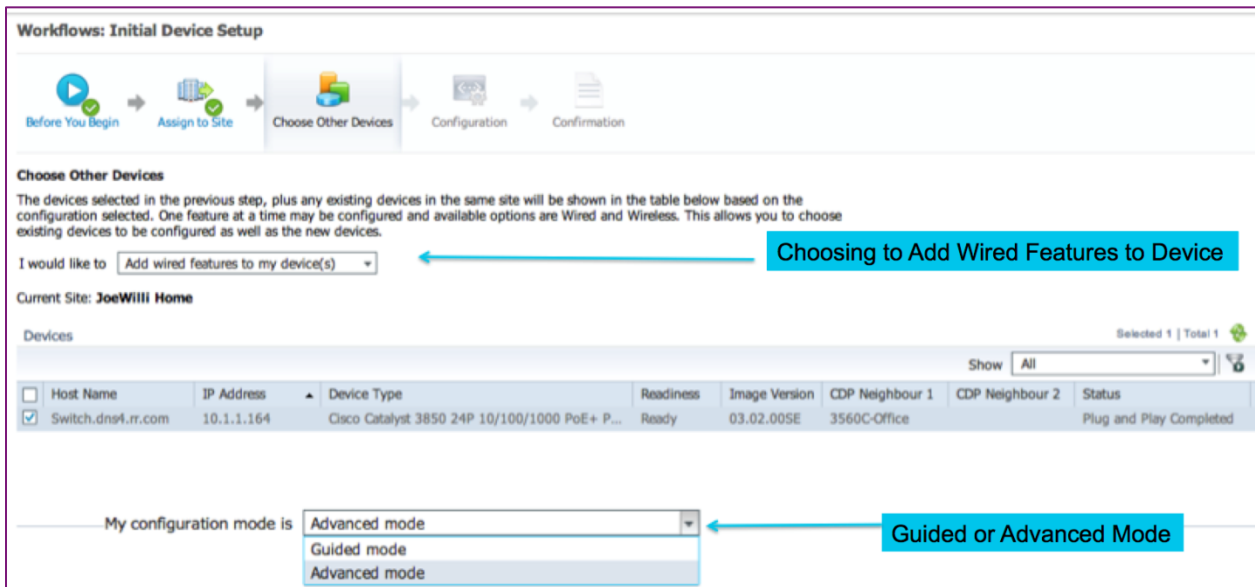
By moving the cursor over the workflow status, the menu comes up and shows one newly register device that is not assigned to a site map. Clicking this brings up the device setup workflow.



Now one or more devices can be selected and added to a site within PI (the site must be already set up).



Now that the device has been successfully added to a site within PI, the administrator is asked which workflow to perform. The **Add wired features to my device** option must be run before adding wireless features. If this has not been done, it is not possible to select the new device after selecting **Add wireless Features to my device**. When the **Add wired features** option has been selected and the device or devices has been chosen, either of two configuration modes that can be selected: Guided Mode, or Advanced Mode.



In the IP Address setup area, the administrator can now change the address that was handed out by DHCP along with the default hostname.

Changes are not saved until the **Save** icon is clicked directly under that line.

**Workflows: Initial Device Setup**

Before You Begin → Assign to Site → Choose Other Devices → **IP Address setup** → Credentials → VLAN and Switching → Auto Smartports, Uplin... → Confirmation

**IP Address Setup**

This step allows you the option to manually assign IP Addresses for previously selected devices. You may also use previously selected options - and skip this step. If you have lots of devices you may find it easier to edit a spreadsheet instead of this table. To do this, export the list of devices as a CSV file, edit that file and then import the file to overwrite this table.

You are editing the following device(s) *Switch.dns4.rr.com...*

I would like to

**Device Management Options**

Devices Total 1

Show

Serial Number	Device Type	Host Name	IP Address	Subnet Mask	Gateway
FOC1634V0MT	Cisco Catalyst 3850 24P 10/100/1000 PoE+ Ports Lay	3850poedns4.rr.com	10.1.1.221	255.255.255.0	10.1.1.2

Now that the IP address and Hostnames have been configured, the administrator can set the credentials on the new system. In this window the user can accept the given values from the original bootstrap config or new ones can be defined.

**Workflows: Initial Device Setup**

Before You Begin → Assign to Site → Choose Other Devices → IP Address setup → **Credentials** → VLAN and Switching → Auto Smartports, Uplin... → Confirmation

**Credentials**

Devices that have been "discovered" are using default credentials assigned via the Plug and Play Setup Wizard. You may define new credentials that will be used to access SNMP, Telnet and SSH. These new credentials will be pushed to the selected devices.

You are editing the following device(s) *Switch.dns4.rr.com...*

I want to

**Settings**

Show Clear Text

Use Credentials

Use RADIUS  Use TACACS  None

**SNMP Credentials**

\* Read Only Community String  \* Confirm

\* Read-Write Community String  \* Confirm

All fields are required. These SNMP v2 community settings will be configured on the device and used by Prime Infrastructure for discovery purposes.

**SSH/Telnet Credentials**

Enable SSH  Enable Telnet

\* User Name  \* Confirm

\* Password  \* Confirm

\* Enable Password  \* Confirm

All fields are required. Telnet is enabled by default. If you prefer SSH ensure that you have the K9 image. The same credentials will be used for both Telnet and SSH.

The administrator can now specify VLAN and switch settings. At least one VLAN must be defined on the system for the device to continue being monitored. If a wired management VLAN and a data VLAN are created, the IP address defined on last page will be placed on the wired management port.

**Workflows: Initial Device Setup**

**VLAN and Switching**

This allows you to configure a VLAN - useful when your network becomes large. It also allows you to edit specific switch settings.

You are editing the following device(s) [Switch.dns4.rr.com...](#)

**VLAN**

VLANs segment network traffic. Enter Data, Voice, Wired Management, and Wireless Management VLANs. Wireless access points will use the Wireless Management VLAN.

\* Data  \* Voice

\* Wired Management  \* Wireless Management

**Switching**

By default the following switching features are configured to Cisco Recommended settings. You may review and change if needed.

**Settings**

- Basic Services**  
This option enables timestamp for logging, debug, password encryption, select debug commands and compresses the config file in NVRAM.  
 Enable switch basic services
- Neighbor Discovery**  
Cisco recommends enabling Neighbor Discovery.  
 Enable CDP  
 Enable LLDP
- Spanning Tree**  
This enables the spanning tree algorithm that the device should use for spanning tree convergence.  
 Rapid PVST
- Power Sharing**  
When applicable, the power supplies will act in a power shared mode.  
 Enable power sharing
- System Redundancy**  
When applicable, this enables System Redundancy using stateful switch over.  
 Enable System redundancy

In addition to setting the VLANs from here, several recommended configuration options are enabled here to comply with Cisco best practices. The administrator can change these by removing the checkmark next to any of the items listed in the settings area.

In the guided workflow, Smartports are applied to guarantee that APs are defined in the appropriate management VLANs. The administrator can also choose how the uplinks for this device are going to be configured. The options include enabling Etherchannel, enabling layer 2 trunking, or both.

**Workflows: Initial Device Setup**

**Auto Smartports and Uplink**

This workflow enables Cisco Auto Smartports and Cisco Auto Quality of Service (Auto QoS) on downlinks by default. This ensures that devices attached to ports automatically receive the correct configuration and quality of service parameters. For example, Access points will be mapped to the correct VLAN and receive the correct quality of service. This is a requirement for this workflow.

You are editing the following device(s) [Switch.dns4.r.com...](#)

**Cisco Auto Smartports**

Cisco Auto Smartports and Cisco Auto Quality of Service (Auto QoS) are automatically enabled by default on downlinks. This ensures that devices attached to ports will automatically receive the correct configuration and quality of service parameters. For example, Access Points will be mapped to the correct VLAN and receive the correct quality of service, Phones will do the same, and so on. VLANs are inherited from earlier in this workflow.

**Enabling Uplink Features**

It is recommended that once your devices are deployed that you consider enabling uplink-specific features like EtherChannels and Trunking. This requires knowledge of your intended topology and cannot be determined by this wizard. These settings can be configured in either the Advanced Mode of this wizard or later using the Device Work Center. Cisco recommends enabling EtherChannels and Trunking on uplinks.

I would like to

The final step is to confirm the settings that will be deployed to the new switch or switches that were chosen during the original deployment screen. The administrator should verify that the correct IP address is being assigned along with subnet mask and gateway. If they is not correct, it could be that the **Save** icon was not selected after entering the values.

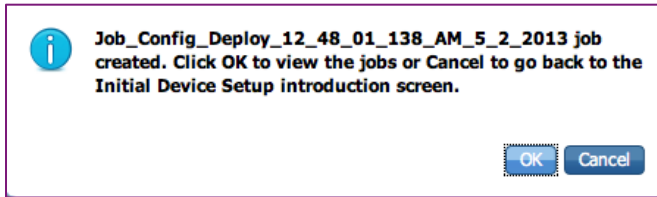
**Workflows: Initial Device Setup**

**Confirmation**

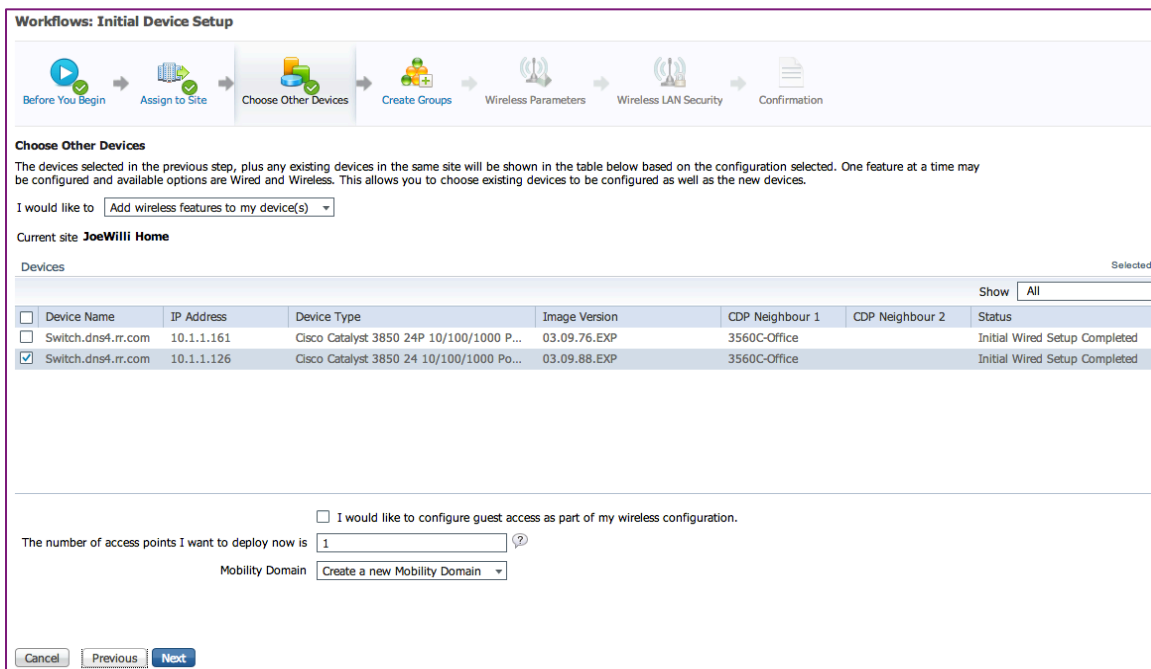
The following devices will be configured with the information shown below.

Serial Number	IP Address	Device Type	Readiness	Image Version	CDP Neighbour 1	CDP Neighbour 2	Status
▼ FOC1634V0...	10.1.1.164	Cisco Catalyst 3850 24P 10/100/1000 PoE+ Ports ...	Ready	03.02.005E	3560C-Office		Plug and Play Comp...
Management IP Address 10.1.1.221 Subnet Mask 255.255.255.0 Gateway 10.1.1.2 Credentials Use existing credentials Authentication Method None VLANs Data VLAN:1, Voice VLAN: 20, Wired Management VLAN: 30, Wireless Management VLAN: 40 Uplink Enabled EtherChannel and Layer 2 Trunking on all of the following uplink ports. <a href="#">Gi1/1/1, Gi1/1/2....</a> Downlink Enabled Auto SmartPorts on all of the following downlink ports. <a href="#">Gi1/0/8, Gi1/0/9....</a>							

When settings are confirmed a job notification pops up for the administrator with the job number assigned so the administrator can check the status. When the job has completed successfully the administrator can deploy wireless service to the switch.



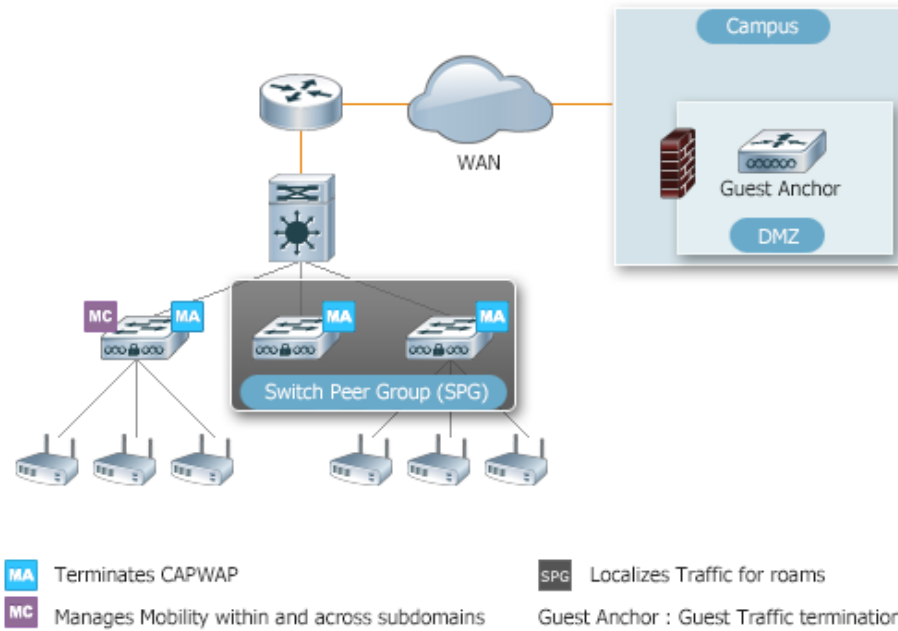
So now that the device has a wired configuration and is in PI, the **Add wireless features to my device** workflow can be deployed. Notice Initial Wired Setup Completed status in the following screenshot. The status must be complete to be able to continue. PI requires an initial wired configuration to be deployed on this device.



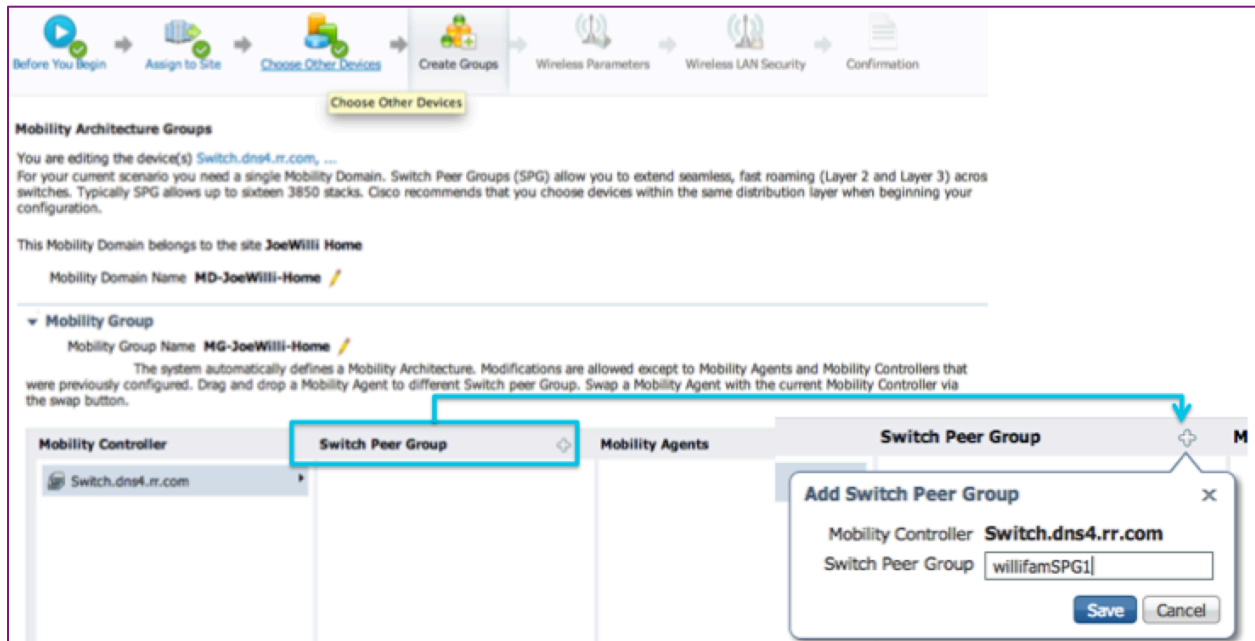
The Catalyst 3850 that will run as either an MA or MC must have a wireless management VLAN defined on it. If planned properly, this was done already using the wired guided workflow.

Having completed that step the guided workflow now determines what devices are MAs and MCs and how they are connected through SPGs.

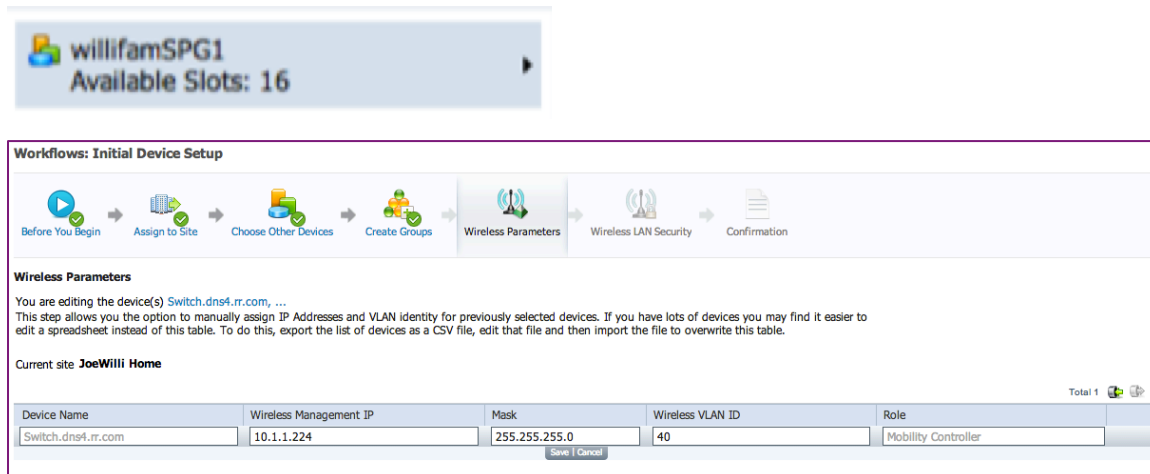
**Figure 70. Design and Deploy Workflow**



In this example, the administrator has defined a mobility domain named MD-Joewilli-Home. A new Catalyst 3850 is being defined as an MC. This device is also being placed in an SPG named willifamSPG1. The SPG does not exist yet, so the administrator can click + to create a new one.



PI also keeps track of how many MAs can exist in an SPG based on the capabilities of the controller defined within that SPG. Here the MA limit is 16 based on the capabilities of the Catalyst 3850 acting as the MC for that SPG. Now once the SPG has been defined, additional MA devices can be added into the willifamSPG1 peer group. When the domain, controller, agents and peer groups have been defined, the administrator can continue to complete the wireless portion of the network, including specifying the wireless management VLAN across the devices.





After this has been defined, secure WLANs can be defined for rollout across the network.

**Workflows: Initial Device Setup**

**Wireless LAN Security**

You are editing the device(s) [Switch.dns4.rr.com, ...](#)  
 This step allows you the option to add secure wireless for LAN connectivity.  
 Current site **JoeWillii Home**

▼ **Secure wireless LAN Properties** \* Indicates required fields

- \* Enterprise WLAN Name:
- \* Enterprise WLAN ID:
- \* VLAN for client connectivity:  ?
- \* Security Profile:  ?
  - Open WLAN
  - DOT1X
  - WPA12+DOT1X+CCKM
  - WPA12+PSK
  - WPA12+DOT1X

Similar to the wired workflow, the wireless workflow completes with a summarization screen before the changes are committed and sent out to the network devices.

PI still also supports adding devices manually from the Operations menu, but for Converged Access devices, they must be added through the lifecycle view, not the classical view.

## Single Pane Glass Visibility

The Mobility Work Center in PI has now have visibility into the roles and responsibilities of Converged Access devices in the network. An example is shown in the following figure. The left side shows the build-out hierarchy, including the MD, SPGs, and device relationships. The interface lists the devices, their MG, what is the role of the device in the MG. Is it an MA, MC, or does it provide multiple functions? The device role can also be changed on this page. It is not necessary to use the wizard or redeploy a template. Click the device and indicate the desired change.

**Device role can be changed to MO/MC/MA**

**Lists devices and their role**

**Tree hierarchy displays mobility domain, SPG and device-level relationship**

Device Name	Management IP	Wireless Interface IP	Mobility Group	Mobility Role
<input type="checkbox"/> Edison106	172.19.28.106	10.1.1.3	default	Admin - MC, Operational - MA
<input type="checkbox"/> MC	172.18.136.161	35.1.1.6	default	Admin - MC, Operational - MC
<input type="checkbox"/> katana115	172.19.28.115	10.1.1.4	g-115c	Admin - MC, Operational - MC, MO

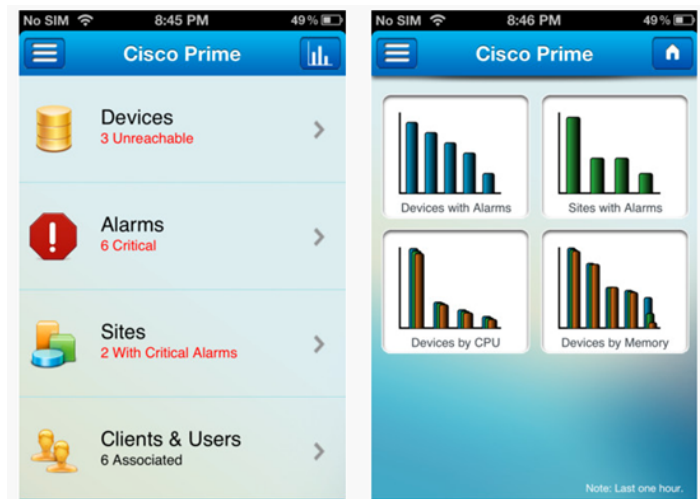
The Client List page shows the devices each client is associated to. Recall that a device can be anchored to a particular switch. It can be associated with an MC, and might have an MA or MC that it is communicating with. The client listing shows the anchor controller, what the device connects to, whether there is an MO in the path, the MG, and the MA.

**Figure 71. Client Listing Interface**

The screenshot shows the 'Clients and Users' interface. A table lists client information with columns for MAC Address, IP Address, IP Type, User Name, Type, Anchor Controller, Mobility Oracle, Mobility Group, Mobility Controller, Anchor Mobility, Switch Peer Group, and Anchor Switch. A blue callout box points to the last four columns, labeled 'Additional columns for Converged Access-specific mobility information'.

MAC Address	IP Address	IP Type	User Name	Type	Anchor Controller	Mobility Oracle	Mobility Group	Mobility Controller	Anchor Mobility ...	Switch Peer G...	Anchor Switch P...
c4:71:fe:d7:1e:19	10.32.35.245	IPv4	ajtdmw		N/A	10.34.150.69	default	10.34.150.69		wmbu-alpha-ng...	
00:22:90:fd:d9:10	10.32.35.173	IPv4	ajtdmw		N/A	10.34.150.69	default	10.34.150.69		wmbu-alpha-ng...	
00:1c:58:cd:4d:46	10.32.35.146	IPv4	ndoshi		N/A	10.34.150.69	default	10.34.150.69		wmbu-alpha-ng...	
c4:71:fe:d7:1f:6c	10.32.35.172	IPv4	ajtdmw		N/A	10.34.150.69	default	10.34.150.69		wmbu-alpha-ng...	
00:1e:7a:ba:d7:ac	10.32.35.151	IPv4	bkudipud		N/A	10.34.150.69	default	10.34.150.69		wmbu-alpha-ng...	
00:22:90:fd:c9:0b	10.32.35.242	IPv4	ajtdmw		N/A	10.34.150.69	default	10.34.150.69		wmbu-alpha-ng...	wmbu-alpha-ng...
c4:71:fe:d7:2b:0f	10.32.35.195	IPv4	ajtdmw		N/A	10.34.150.69	default	10.34.150.69		wmbu-alpha-ng...	wmbu-alpha-ng...

The PI iPhone app provides the ability to look at the devices on the network directly from an iPhone or an iPad. The interface shows alarms other information from a client user perspective. So it is not necessary to be physically at a desktop to get the information. This is yet another way that Cisco makes it easy to deploy Converged Access services.



# Security and Guest Access

The next sections describe the solutions for security and guest access for Converged Access. As millions of devices connect to networks today, the challenge is to provide an enhanced level of security that is consistent for wired, wireless, and VPN while also ensuring end-to-end scalability

Converged Access meets this challenge with a common platform, contractor access, employee access, wired, wireless, and guest access, including BYOD. The solution scales into very large networks with the Catalyst 3850 and leverages the WLC 5760 to provide secure access for all different classifications of users and their devices.

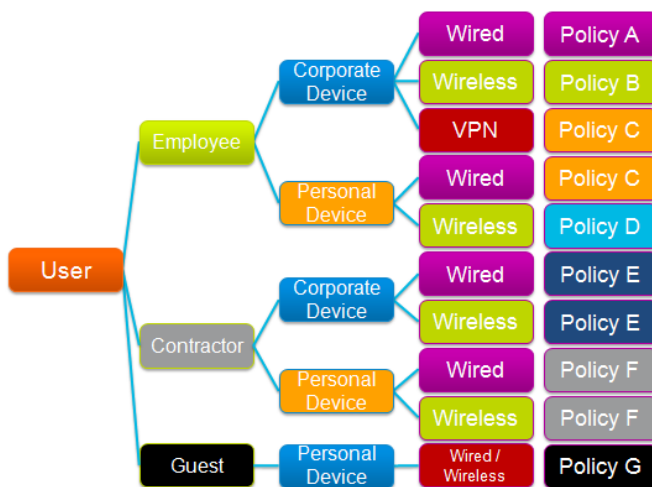
## Integrated Policies

Policy integration first requires a clear definition of the policies to implement for employees, contractors, and guests, along with the devices they may be using. Each type of user might have a corporate device or personal devices, or both. Historically the employer has provided the devices, or had control over them in the corporate network. But now with the BYOD explosion, personal devices are now becoming more of a norm and requirement for network connectivity. Moreover, BYOD device can be wired or wireless and might or might not involve VPN.

The example in the following figure shows how the policies might be structured. The figure shows policies A, B, C, and D that apply to employees across network connectivity and device types. A common policy C covers VPN connections for an employee who uses a corporate device or a personal device over a wired connection. If an employee comes in with a personal wireless device, a different policy might be required (D in the figure) due to regulatory HR compliance issues, or other restrictions needed by the organization. For contractors, the E and F policies cover corporate devices issued to contractors, while a different policy (F in the figure) covers the personal devices of contractors. Finally, guests are subject to policy G, whether they come in wired or wireless.

This is just one example of how to structure and go about the decision making process for defining policies. Actual policies can be simpler or more elaborate, with reuse of some common policies and additional policies where needed.

**Figure 72. Users, Devices, and Policies**



## Policy Definition and Application

Policies need to be defined, reside, and be enforced on specific devices in the network. On-device policies, whether on a WLC or a wired switch, or on a Catalyst 3850 platform, can provide both wired and wireless device termination. AAA services, local policy objects, local policy rules, and users can all be defined on the local switch, but while acceptable in small networks, that solution is not scalable.

Scaling networks up to millions of devices and users requires centralized policies with centralized user databases such as Lightweight Directory Access Protocol (LDAP) servers, Remote Authentication Dial In User Service (RADIUS) servers, and Active Directory servers. With these centralized servers, centralized policy objects and central policy control can be stored and enforced, and inspected profiling can be done from the central location leveraging ISE.

In some cases, a combination of local and centralized policies is optimal. An emergency or fallback policy on the device, some localized QoS policies, and native ACLs can work in conjunction with each other to scale the network and provide common policy definitions throughout the network, while retaining the ability to enforce some policies if the link to the centralized policy engine fails.

Prior to Converged Access, policy application was applied at different places for wired, wireless and guests. With Converged Access, policy application is distributed, allowing for better scalability

Inconsistent policy definition has been a challenge across different platform types, whether a firewall providing VPN access, a VPN concentrator, a router, or a switch. Inconsistencies can involve downloadable ACLs, different filter IDs, different AV-pairs, WLC providing Airespace ACL names, different methodologies of assigning VLAN assignments (802.1, VLAN name, ID number, or Airespace VLAN names), and an inability to provide common QoS templates or assignments across those platforms. With Converged Access, this becomes much simpler – all WLC attributes and configurations can now be applied at a converged point in the network.

Consider how a single centralized policy affects flow through the network. Employees using the same SSID can be associated to different VLAN interfaces and policy after an 801.1x Extensible Authentication Protocol (EAP) authentication. Employees using corporate wired and wireless devices with their ISE and Active Directory User ID can be assigned to same VLAN to have full access to the network, while employees using personal devices with the same Active Directory user ID can be assigned to another VLAN for Internet access only.

## Session-Aware Networking

Session aware networking (SaNet) is a new policy engine for TrustSec. It is a simplified but more powerful implementation for policy enforcement and configuration through Cisco's common classification policy language (C3PL). This structure allows definition of events and different classes, and how events are merged into classes and trigger different actions. Sessions and interfaces can be flexibly defined. Functions such as AAA, WebAuth MAB, and MAC Authentication Bypass (MAB) are pulled together in IOS in the Converged Access deployment model.

## Policy Enforcement

In Converged Access, policy definition is done in IOS, and IOS inspects and enforces the policy. For wireless clients inside the Catalyst 3850, a wireless process called WCM runs and interacts with IOS to decide which policy will be applied after the decision has been made in IOS. For Layer 3 roaming, ACL policies are applied at the anchor switch or the PoP in the network where that client's traffic is associated to in the network. As mentioned previously, this is the default behavior of a Converged Access network when the system is configured and deployed. However, also as mentioned previously, this type of sticky user roaming can be disabled to permit Layer 2 roaming (assuming the associated wired network deployment supports this, as outlined previously in this document). In

Layer 2 roaming, any ACL assignments for users are handed off between the involved switches, and the Layer 2 ACL roamed policy (after being retrieved from the AAA server by the roamed-to switch, and reapplied there in hardware) is enforced at the now-moved PoP on that roamed-to switch, where the roamed client and their newly-associated AP directly connect and attach to the network.

### **Per-Session VLAN Assignment**

Prior to Converged Access with the Catalyst 3850, MAC address bypass VLAN assignment was done on a per user and per device basis. With Converged Access, VLAN assignment can be done on a per session basis. For example, consider a machine that boots up, authenticates, and authorizes to the network, and is assigned to a VLAN. If a virtual machine is started on the machine and is bridged through the network, there is a second MAC address. Based on the second MAC address, the virtual machine session can be assigned to a different VLAN.

### **Service Templates**

Service templates are typically defined on the switch. They can also be defined on the RADIUS server and downloaded dynamically as needed per authorization or during CoA (ISE 1.2 feature). The service template can be used as for actions per control policy or as part of the RADIUS authorization (AV-pair). It is similar to applying a port ACL using filter-id.

### **ACLs**

ACLs can be downloaded from an AAA server. For example, assume a wireless client connects to the network. Authentication is processed in the access device, in this case on the Catalyst 3850 acting as an access switch. The authentication request goes to the RADIUS server, which passes the ACL information for this user (if any) back down to the access switch for enforcement. The Catalyst 3850 applies the downloadable ACL for this user. The traffic flow, and associated security enforcement, is the same for both wireless and wired users.

The process works as follows:

1. A wireless client requests association.
2. The MA responds back with association.
3. The WCM triggers the IOS module to do authentication.
4. IOS starts the authentication process for the client with the AAA server.
5. The AAA server responds with 'access accept,' including discretionary ACL (dACL) name and version number in policy attributes.
6. If the switch has previously downloaded the dACL and has the current version, it uses the cached version.
7. If the switch does not have current version, it queries the AAA server for the latest dACL version.

### **Device Enrollment and Provisioning**

A key characteristic of BYOD is the ability to do device enrollment and provisioning in the network. The process works as follows:

1. Employee associates to BYOD-Secure SSID
2. Employee enters username and password
3. MA does PEAP authentication

4. Server authenticates
5. MA does client URL redirection
6. Device registration page load & MAC get prepopulated
7. Employee registers device
8. Supplicant Provisioned and certificate installed
9. CoA occurs and supplicant authenticates using EAP-TLS
10. dVLAN, dACL, QoS policy for Employee pushed to MA

### Catalyst Integrated Security Features

Cisco Integrated Security Features (CISF) are provided for IPv4 in the common platform for both wired and wireless. IP device tracking and DHCP snooping are fundamental for identity. IP Source Guard (IPSG) is enabled on a per-WLAN basis while IP theft is enabled globally.

For IPv6 first-hop security includes IPv6 RA guard, IPv6 DHCP guard, RA throttle, and multicast suppression. The current IPv6 capabilities provide a foundation that will continue to be expanded in future releases.

Dynamic ARP inspection provides man-in-the-middle attack prevention and enforcement. It is an IOS feature that validates ARPs and then forwards them once they are validated from the port or the AP, depending on the clients connected. It is supported for both wired and wireless clients. Traditionally, this has been done just on the edge, or separately in the WLC. With Converged Access there is a common enforcement point at the edge of the network that can provide rate limiting and error disabling.

IP Source Guard is an IOS security feature that prevents IP spoofing attacks. It can be enabled on the wireless side on a per WLAN basis per SSID, or it can be enabled globally on the switch for wired and wireless. The spoofed traffic is dropped in hardware where it is implemented and enforced.

### Security Features Summary

The following table summarizes security features on the Catalyst 3850, WLC 5760, and WLC 5508.

**Table 5. Summary of Security Features**

Feature	Catalyst 3850	WLC 5760	WLC 5508
BYOD Functionality	YES	YES	YES
Rogue detect / classify / contain, RDLP	YES	YES	YES
Port Security	YES	YES	NO
IP Source Guard	YES	YES	NO
Dynamic ARP Inspection	YES	YES	NO
LDAP, TACACS+, RADIUS	YES	YES	YES
LSC and MIC	YES	YES	YES
AP dot1x EAP-FAST	YES	YES	YES
Secure Fast Roaming	YES	YES	YES
802.1X-rev-2010 (MACsec / MKA)	H/W Ready	H/W Ready	NO

Feature	Catalyst 3850	WLC 5760	WLC 5508
IP Theft, DHCP Snooping, Data Gleaning	YES	YES	YES
IOS ACL	YES	YES	YES
Adaptive wIPS, WPS	YES	YES	YES
CIDS	YES	YES	YES
TrustSec SGT / SGACL	H/W Ready	H/W Ready	SXP
Guest Access	YES	YES	YES
IPv6 RA Guard	YES	YES	NO
MFP	YES	YES	YES
IP Device Tracking	YES	YES	NO
CoPP	Static	Static	NO

## Guest Access

This section presents examples of how guest access works with Converged Access. The Converged Access deployment option supports flexible security features for wired and wireless with integrated policies and increased scalability through partially centralized and partially distributed configuration. Configuration is simplified, and the guest anchor component described in the following sections is consistent with existing CUWN networks.

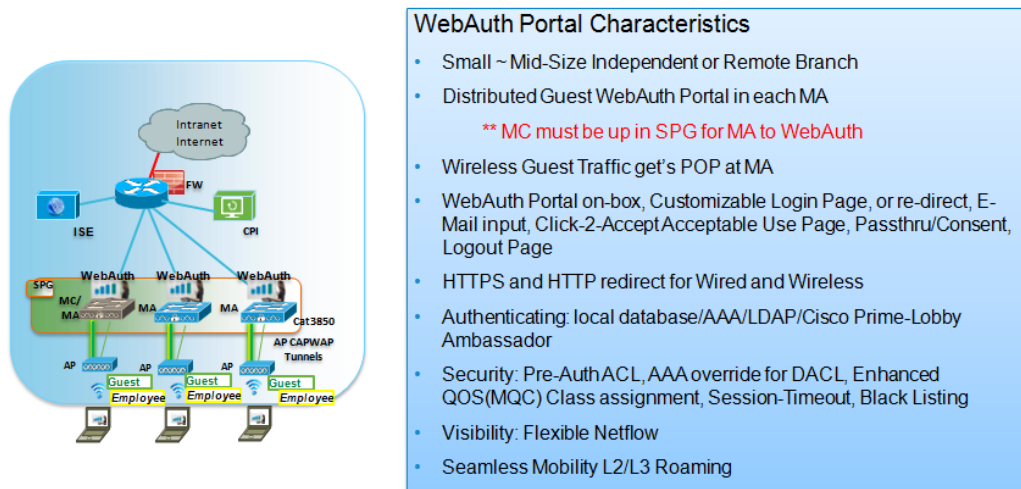
### Deployment Examples

The following figure shows a small or mid-sized branch with a Catalyst 3850 supporting fewer than 250 APs with no guest anchor. This solution provides wired and wireless authentication using distributed WebAuth. The WebAuth portal supports functions such as customization, customization of the login page, a redirect page, and email input.

Redirect for HTTPS and HTTP is supported for wired and wireless users. Clients can be authenticated against a local database, AAA, LDAP server, or lobby ambassador using PI. Available security features include ACLs, AAA override, downloadable ACLs, QoS with MQC, session time-outs, and client blacklisting. Edge of the network visibility is available with Flexible Netflow, and Layer 2 and Layer 3 roaming can occur after clients are authenticated to the network.

For this functionality to work for wireless access, the MC function for the SPG must be up and active so that the traffic is terminated at the PoP in the network.

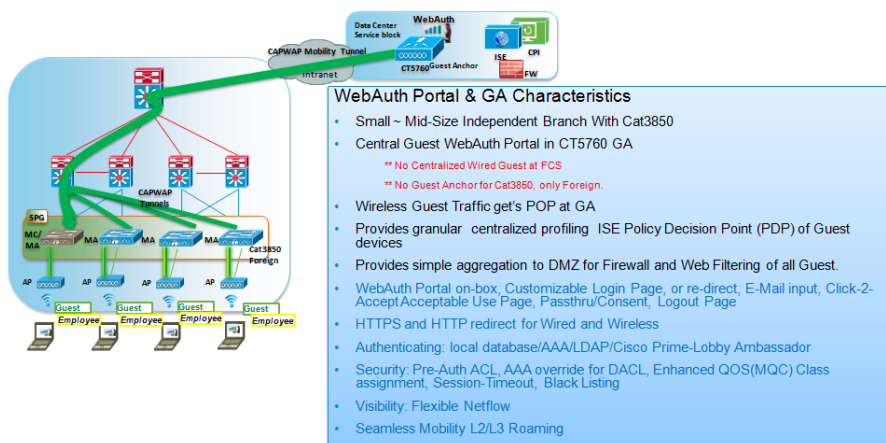
**Figure 73. Guest Access – Small to Mid-Size Branch**



In the following figure, the site still has fewer than 250 APs, but a guest anchor is included. Connectivity is provided up to a guest SSID, which redirects the guest traffic up to a mobility anchor (in this case, a WLC 5760 in the datacenter or DMZ). The guest anchor controller performs the WebAuth and provides Layer 3 credential pages to the clients, passing through the MC and out to the MAs.

This configuration provides centralization, in this case for a WebAuth portal, but it could also provide centralization for other types of client, guest, or contractor access, and even be used for employee traffic centralization. All the traffic is terminated at the PoP at the guest anchor. Granular centralized profiling inspection and inline posture assessment is supported, along with simple aggregation to the DMZ and web filtering through the firewall for all guest users.

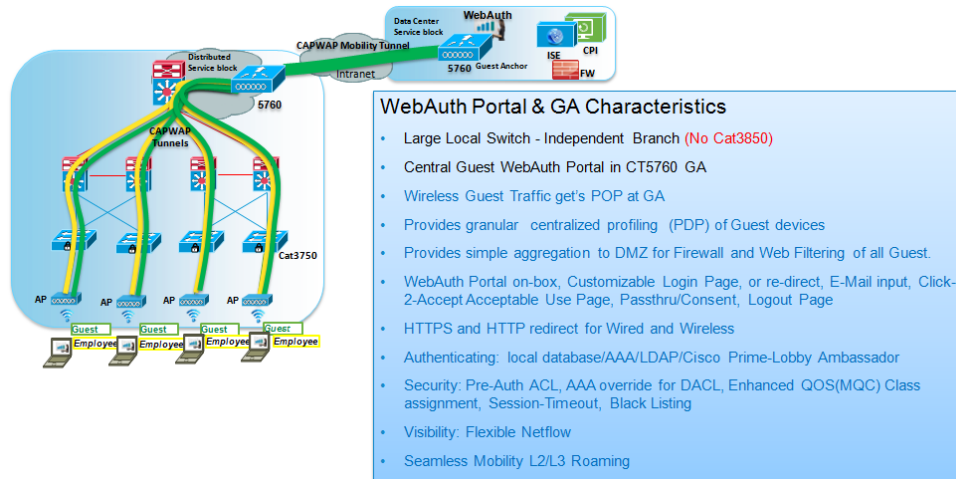
**Figure 74. Guest Access – Small to Mid-Size Branch with Guest Anchor**





The following figure adds centralization to the guest access solution using a WLC 5760. The APs are registered through the yellow tunnels up to a controller in the large campus, which provides the MC and MA functionality. The MC redirects the guest anchor tunnel into the DMZ and to another WLC 5760 that serves as the guest anchor controller. (The guest anchor controller cannot be a Catalyst 3850.) This approach provides all of the capabilities of the other guest access solutions, but supports additional scaling with up to 71 foreign anchor controllers.

**Figure 75. Guest Access – Large Campus**



## Guest Access Configuration

This section provides configuration examples for guest access. The examples are based on the 7.0 release train from AireOS, so the feature and functionality level interoperability is with the code levels of the 7.0 release.

The following sample CLI configuration defines a virtual IPv4 address (192.0.2.1) and two types of web authentication. WEBAUTH1 performs WebAuth against a central internal database, while WEBCONSENT provides a click-to-accept WebAuth page with no user authentication. As a best practice, in addition to enabling HTTP and HTTPS, the WEBCONSENT configuration enables the HTTP secure server (ip http secure-server) and disables HTTP port 80 connections into the controller (ip http active-session-modules none). Although this disables port 80, the controller can still capture the client's port 80 webpage request and redirect it to the internal WebAuth page (an HTTPS secure page). When the user's credentials are established, policy is pushed down, and the client can pass traffic through the secure or unsecure port 80; however, the user is still not able to directly access the Catalyst 3850 or the WLC 5760 local WebAuth agent.

```
! First section is to define our global values and the internal Virtual Address.
! This should be common across all WCM nodes.
parameter-map type webauth global
  virtual-ip ipv4 192.0.2.1
! This is for generic WebAuth and will authenticate against internal user database
parameter-map type webauth webauth1
  type webauth
  banner text ^C WEBAUTH1^C
! This is for generic WebAuth with Consent form Click-2-Accept, no Authentication
parameter-map type webauth webconsent
  type webconsent
banner text ^C WEBCONSENT^C
! Configure http server in global config. These are needed to enable Web Services in IOS
ip http server
ip http secure-server
ip http active-session-modules none
```

!

The following sample configuration shows wireless WebAuth. It specifies a client VLAN and local authentication using the WebAuth1 profile. Other authentication options, including RADIUS, could be specified instead.

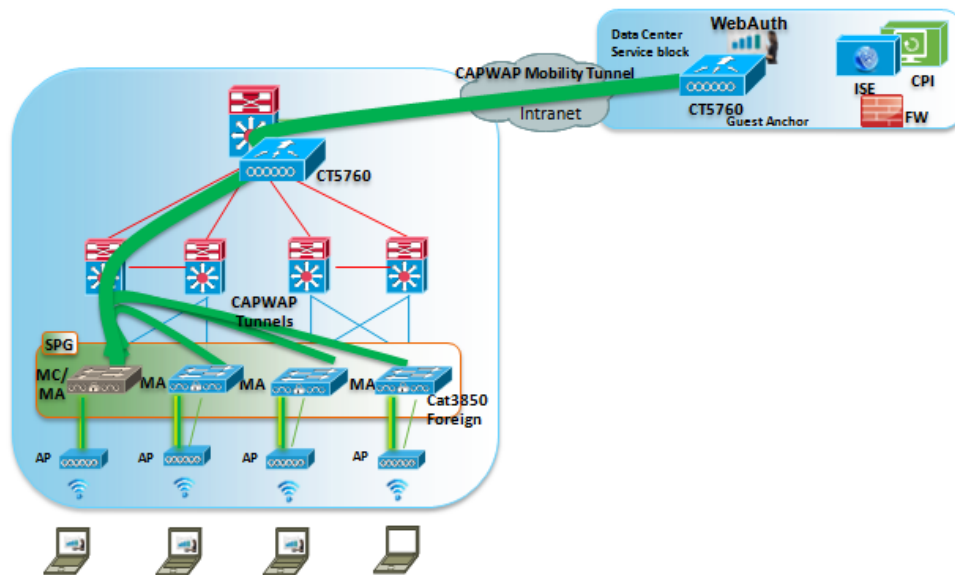
```
! This WLAN "web1" will advertise an SSID called "web1",
! Place the user in VLAN 21,
! Disable default WPA authentication and Enable web-auth security
! Use wcm_local authentication for this security from global AAA Setup
! Associate earlier defined parameter-map "webauth1"
!
Wlan web1 11 web1
  client vlan 21
  no security wpa
  security web-auth
  security web-auth authentication-list wcm_local
  security web-auth parameter-map webauth1
  no shutdown
!

! Sample AAA Global setup for wcm_local
!
username abc password 7 08204E4D
aaa new-model
aaa local authentication wcm_local authorization wcm_author
!
aaa user profile local
!
aaa authentication login wcm_local local
aaa authentication dot1x wcm_local local
aaa authorization network wcm_local local
!
```

The next example adds centralization for the campus deployment. Guest anchoring is supported in the DMZ by a WLC 5760 (could also be a WLC 5508 or WiSM-2 running in the new mobility mode for hierarchical mobility).

The foreign role (bottom right in the figure) is on a Catalyst 3850, but could also be on a WLC 5760, WLC 5508, or WiSM-2. When a client authenticates to the network, the Layer 3 WebAuth happens from the PoP in the network. The WLC 5760 in the DMZ does the transactions to a RADIUS server on the same box for authentication. Layer 2 functions such as MAC address bypass, MAC filtering, and Dot1x, happen at the edge of the network at the MAs.

**Figure 76. Guest Access Example for Large Campus**



The following sequence shows the anchoring configuration. The mobility anchor points back to the anchor controller in the DMZ. The foreign MA has IP address 192.168.21.44, and the guest anchor controller has IP address 192.168.21.43.

```
! Config on Foreign MC/MA (192.168.21.44)
! All Mobility Group Configuration must be completed prior to these steps
! Place the user in dummy VLAN 1 and establish (GA) Tunnel
! to Anchor (GA) controller (192.168.21.43) , Disable Snooping on foreign VLAN
! Disable default WPA authentication and Enable web-auth security
! Use wcm_local authentication for this security from global AAA Setup.
! Associate earlier defined parameter-map "webauth1"
!
no ip dhcp snooping vlan 1
wlan web1 11 web1
  client vlan 1
  mobility anchor 192.168.21.43
  no security wpa
  security web-auth
  security web-auth authentication-list wcm_local
  security web-auth parameter-map webauth1
  no shutdown
!
```

On the anchor side, the configuration specifies that the anchor controller will accept mobility anchor configurations from other foreign places in the network on this WLAN and this SSID.

That is all that is required for the configuration. The mobility anchor commands, one on the foreign MC/MA and one on the guest anchor controller, implement the guest anchoring virtualization in the network. There is no need for multiple policy-based routing, Multiprotocol Label Switching (MPLS), or Virtual Private LAN Services (VPLS) configuration or use.

```
! Config on Anchor GA (192.168.21.43)
! All Mobility Group Configuration must be completed prior to these steps
!
! Place the user in VLAN 24 and establish (GA) Tunnel to a local GA controller (192.168.21.43)
! Disable default WPA authentication and Enable web-auth security
! Use wcm_local authentication for this security from global AAA Setup.
```

```

! Associate earlier defined parameter-map "webauth1"
!
wlan web1 11 web1
  client vlan 24
  mobility anchor 192.168.21.43
  no security wpa
  security web-auth
  security web-auth authentication-list wcm_local
  security web-auth parameter-map webauth1
  no shutdown
!

```

## Release Compatibility for Guest Access

The following table shows the Inter Release Controller Mobility Compatibility (IRCM) matrix.

**Table 6. Release Compatibility for Guest Access**

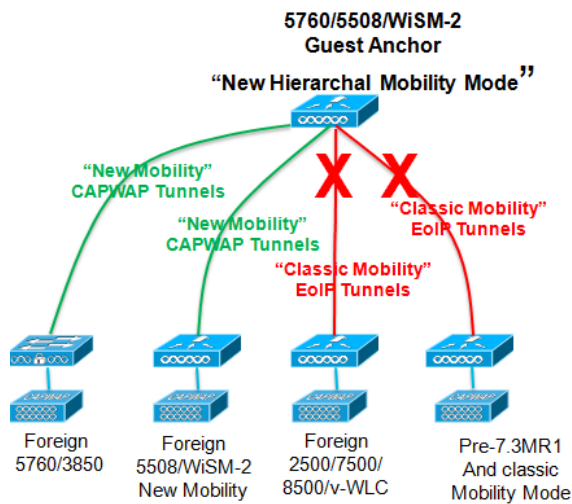
CUWN Service	Y = Compatibility in Classic Flat Mobility							O = Compatibility in Hierarchal Mobility	
	4.2.x.x	5.0.x.x	5.1.x.x	6.0.x.x	7.0.x.x	7.2.x.x	7.3.101.0	7.3.112.0 Note: 1	IOS-XE 3.2.1 SE
Layer 2 and Layer 3 Roaming	Y	-	-	Y	Y	Y	Y	0	0
Wireless Guest Anchor/Termination	Y	Y	Y	Y	Y	Y	Y	0	0 <sup>2</sup>
wIPS & AwIPS Rogue Detection	Y	-	-	Y	Y	Y	Y	0	0 <sup>3</sup>
Fast Roaming (CCKM) in a mobility group	Y	-	-	Y	Y	Y	Y	0	0
Location Services	Y	-	-	Y	Y	Y	Y	0	0
Radio Resource Management (RRM)	Y	-	-	Y	Y	Y <sup>4</sup>	Y <sup>4</sup>	0 <sup>5</sup>	0 <sup>5</sup>
Management Frame Protection (MFP)	Y	-	-	Y	Y	Y	Y	0	0
AP Failover	Y	-	-	Y	Y	Y	Y	0 <sup>6</sup>	0 <sup>6</sup>

**NOTES:**

1. New Mobility is only supported on AireOS CT5508 & WISM-2 platforms but **does not** form any IRCM or GA with CT2500/CT7500/CT8500/v-WLC
2. Guest Anchor Termination is only supported on CT5760/CT5508/WISM-2. CT5760/CT5508/WISM-2/Cat3850 all supported as a Foreign
3. Rogue Detector Mode not supported
4. In Release 7.2 RF Profiles and groups was introduced. RRM for release 7.2 and later is not backwardly compatible with previous releases.
5. RRM Converged Access is compatible with CUWN release 7.3.112.0 but **does not** support RF Profiles and Groups introduced in 7.2
6. No AP SSO in IOS for CT5760. AP Intra-OS Platform Fast Failover Supported. AP Inter-OS Platform Image Download & Reboot performed.

There are some constraints, especially when migrating from centralized access into the converged access model. The green lines in the following figure connect back to the three supported platforms running in the hierarchical mobility mode, while the red lines indicate that classic mobility is not supported in the Converged Access hierarchical model.

**Figure 77. Constraints**



## Troubleshooting

This section shows some useful commands for troubleshooting issues with guest access.

```
! Client Specific
debug client mac-address <mac>
! Mobility / Guest Anchoring
debug mobility keep-alive
debug mobility handoff
! DHCP
debug ip dhcp snooping packet
debug ip dhcp snooping event
! Client Sessions / States
show access-session
show wireless client summary
show wireless client mac-address <mac> detail
show ip dhcp binding
show ip arp <mac>
! WebAuth Session
debug ip admission page
```

# Catalyst 3850 Hardware Architecture

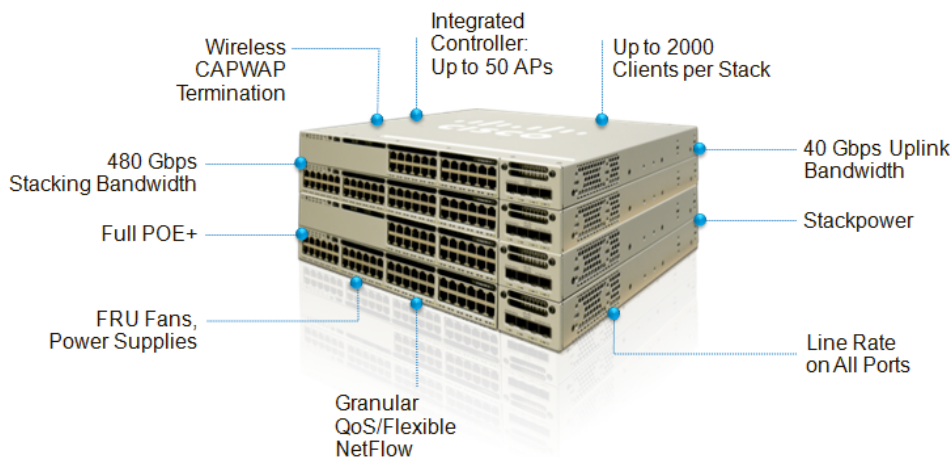
The following sections provide an overview of the Catalyst 3850 hardware architecture.

## Catalyst 3850 Platform Overview

The following figure shows many of the key Catalyst 3850 hardware features. The top portions relate to wireless capabilities and scale, including local CAPWAP termination and support of up to 50 APs and 2000 clients. The bottom portion relates to wired capabilities of the switch. The Catalyst 3850 has 40G of uplink bandwidth and 480G of stack bandwidth. The platform has full support for PoE+ on all 48 ports, when appropriate power supplies are deployed.

Flexible NetFlow is available on, and native to, all the ports on the Catalyst 3850 switch (a separate service module for NetFlow support, required on previous Catalyst 3750 platforms, is not needed). All of these features are built into Cisco's innovative and powerful new UADP ASIC, which forms the heart of the both the Catalyst 3850 and WLC 5760 platforms.

**Figure 78. Catalyst 3850 Platform**



## Network Modules

The following figure shows the available network modules for the Catalyst 3850 platform. The 4x1G SFP module is shown on the left, and is supported on both the 24-port and the 48-port versions of the Catalyst 3850. The middle image is the 2 x 10G module, which supports multiple modes: 2 x 10G ports, 1 x 10G and 2 x 1G ports, or 4 x 1G ports, with support for both SFP and SFP+ optics. The right-most network module shown is the 4x10G module. On this module, all four ports can operate in 1G mode or 10G mode independently, with auto-sensing for SFP and SFP+ optics supported. It is important to note that the 4-port 10G module is supported only on the 48-port version of the Catalyst 3850 switch.

**Figure 79. Catalyst 3850 Network Modules**



### Power Modules

The Catalyst 3850 provides two power supply bays for modular, field-replaceable power supplies, thus providing the option for power redundancy and scalability within the Catalyst 3850 platform. The following figure shows options available for the Catalyst 3850 power supplies, which follow the same design as in the Catalyst 3750-X platform.

Power supply modules from the Catalyst 3750-X can be used in the Catalyst 3850. Cisco’s innovative and powerful StackPower capability, which was introduced with the Catalyst 3750-X platform, remains the same, and the StackPower cables from existing 3750-X switches can be used with the Catalyst 3850 (StackPower pools consisting of a mixture of Catalyst 3750-X, Catalyst 3750, and Catalyst 3850 switches are not supported). Versions of the power supplies available for the Catalyst 3850 include 350-watt AC, 715-watt AC, 1100-watt AC, and 440-watt DC. The use of two of the 1100-watt AC power supplies with a single Catalyst switch provides for full PoE+ capability (30 watts per port) on all 48 ports, and StackPower can be used to pool power across multiple Catalyst 3850 switches.

**Figure 80. Catalyst 3850 Power Modules**

PWR Modules
PWR-C1-350WAC
PWR-C1-715WAC
PWR-C1-1100WAC
PWR-C1-440WDC

Same as 3K-X Series




- Power Modules is same as 3K but with a new PID
- Classic 3K Power Module can work on Catalyst 3850s
- No Interworking with classic 3Ks for StackPower

## Stacking and Stacking Cables on the Catalyst 3850

The Catalyst 3850 platform employs a new, higher-performance stacking architecture known as StackWise-480. This stacking architecture operates at up to 480Gbps, providing major improvements in stacking performance and capability as compared with the StackWise/ StackWise Plus architecture used with the Catalyst 3750 and Catalyst 3750-X platforms.

The following figure shows the stacking cable used with the StackWise-48 system, which is different from the StackWise cable used with the Catalyst 3750 / 3750-X platforms. Due to their different stacking architectures and cabling, the Catalyst 3750 / 3750-X and Catalyst 3850 platforms cannot be stacked together. The cable lengths available for StackWise-480 are 0.5 meter, 1 meter, and 3 meters. The advanced StackWise-480 system employs three high-performance bidirectional stacking rings, and is described in greater detail in the following text.

**Figure 81. Catalyst 3850 Stacking Cables**



## IOS-XE on the Catalyst 3850

The Catalyst 3850 runs IOS-XE, which enables multi-core CPU support. IOS itself runs as a daemon over the underlying IOS-XE Linux-based kernel, with some supporting processes running outside the IOS context. IOS-XE also provides support for other selected hosted applications. For example, the WCM runs as a hosted application on the Catalyst 3850 platform with IOS-XE, in a different address space than IOSd within the Linux-based platform infrastructure.



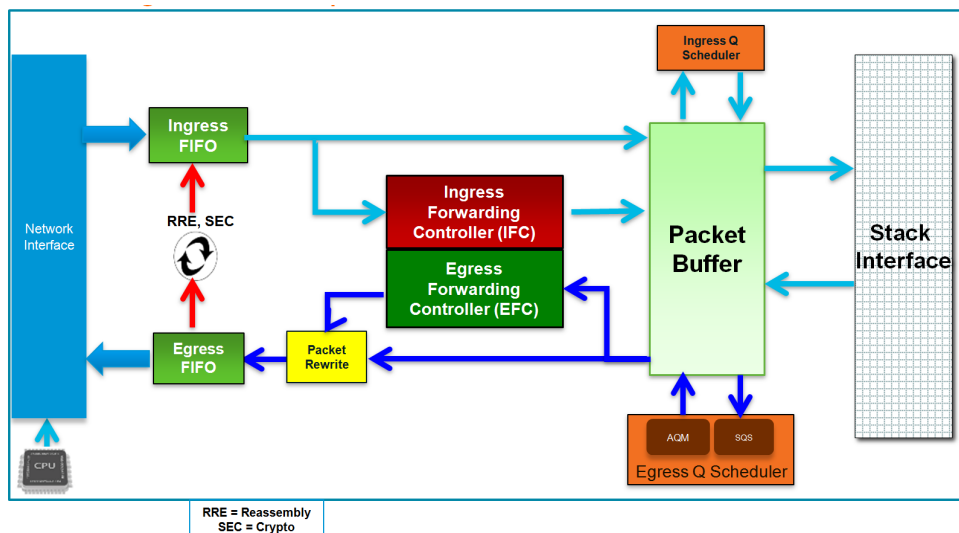


## ASIC Details

The following figure is a block diagram of the UADP ASIC. The Network Interface is the traffic multiplexer for all the traffic ingressing front panel ports and the CPU. The ingress FIFO provides temporary holding for packets as they are played toward the Ingress Forwarding Controller (IFC) and also to the packet buffer.

The forwarding controller is the UADP ASIC, and is the heart of the Catalyst 3850. In this logical representation ingress and egress are separated, but physically they are in the same ASIC. The 6MB packet buffer is used for all the tuning and buffering inside the system. The ingress queue scheduler schedules the packets going toward the 480G stack. The 480Gbps stack bandwidth is achieved by spatial reuse. The Stack Queue Scheduler (SQS) queues packets that are locally switched within the ASIC as well as packets ingressing over the stack interface. The Egress Queue Scheduler (EQS) is used for queuing packets to local front panel ports and the packet recirculation block.

The Active Queue Management (AQM) block handles queuing and replication of multicast packets or egress SPAN packets. The logical egress forwarding controller looks up all the egress features (such as ACL, QoS, and NetFlow lookups) on the egress interface. When packet processing is complete and before the packet is switched out, any rewrites that need to be performed on the packet (Layer 2 rewrites, Layer 3 rewrites, encapsulation, or fragmentation) are formed as part of an instruction set. These instructions are passed to the packet rewrite block where the actual rewrites are completed. Before the packet egresses, the egress FIFO acts as a temporary staging area. Any fragment reassembly is done in the recirculation block, as well as any encryption or decryption of Datagram Transport Layer Security (DTLS) packets. Packets that require recirculation are fed into the ingress FIFO again.



## Packet Walks

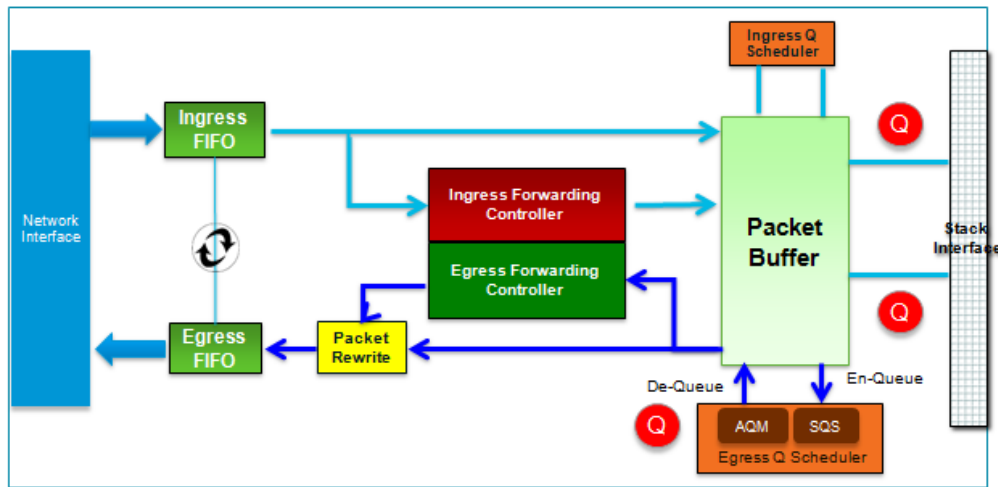
This section describes how packets traverse the system for local switching, remote switching, multicast local replication, multicast remote replication, and recirculation.

Different paths are available through the switch. Local switching occurs inside a single ASIC, while remote switching is for packets where the ingress is one member and the egress is another member in the stack or even if the destination port lies within another ASIC in the same system. As with local switching, local replication for multicast is within the ASIC, and remote replication is across ASICs. The recirculation path is explained later.

The following figure shows the local switching flow. Consider the situation in which a packet ingresses from the first uplink port destined to one of the downlink ports in the range 1-24. Packets enter through the network interface from the front panel port, and some fields of the packet are snooped into the ingress forwarding controller, with the rest of the packet held in the packet buffer. The ingress forwarding controller does the necessary ingress lookups: Layer 2 source, Layer 3 source, Layer 2 destination, Layer 3 destination, multicast, ACL, QoS, NetFlow, and so on. It then creates a control header, or descriptor, which is passed to the packet buffer. The packet buffer prepends the descriptor to the packet. Since this is a locally switched packet, it sends an enqueue signal to the Stack Queue Scheduler (SQS) that passes the queuing to the Egress Queue Scheduler (EQS). When the packet is ready to be transmitted, the EQS sends a dequeue signal to the packet buffer, which relays the packet to the egress forwarding controller (EFC), and also to the rewrite block. The egress forwarding controller does all of the feature processing (ACL, QoS, NetFlow) for the egress interface, and the final rewrite instruction is provided to the packet rewrite block. The packet is rewritten accordingly and switched to the egress FIFO towards the network interface, delivering to the front panel port.

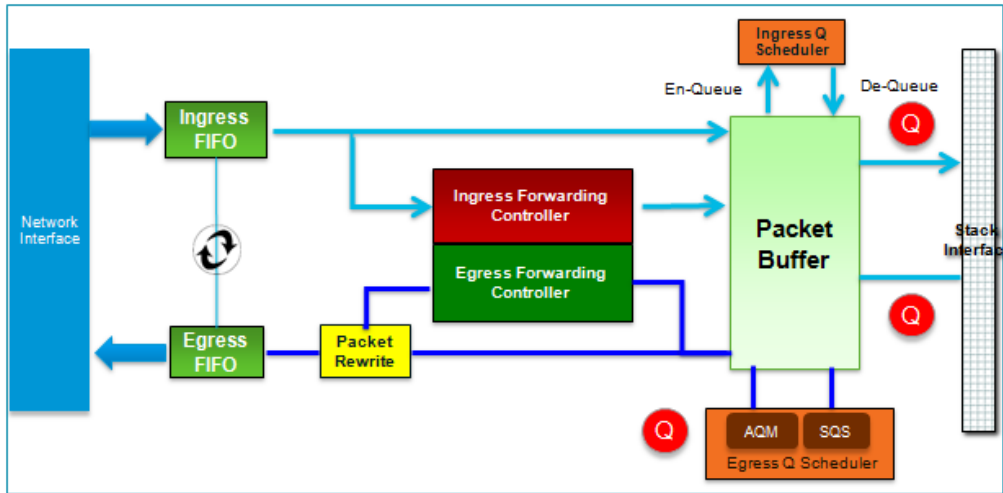
The egress forwarding controller does all of the feature processing (ACL, QoS, NetFlow) for the egress interface, and the final rewrite instruction is provided to the packet rewrite block. The packet is rewritten accordingly and switched to the egress FIFO towards the network interface, delivering to the front panel port.

**Figure 84. Packet Walk for Local Switching**



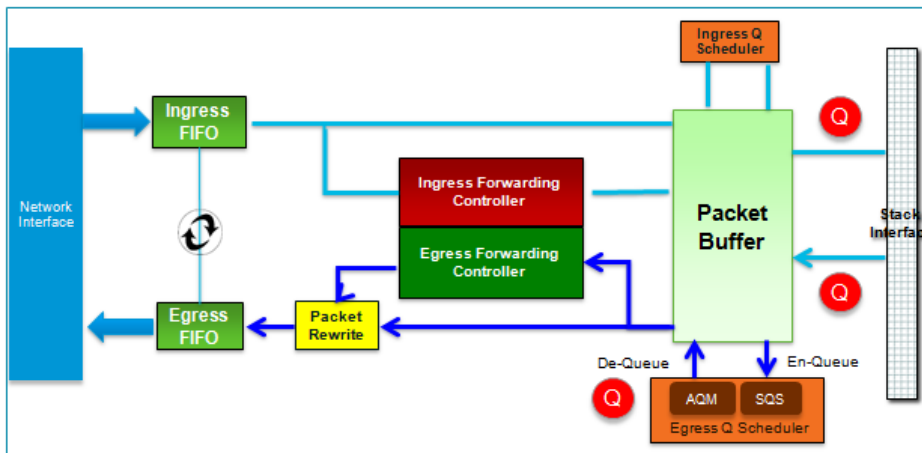
The following figure shows the remote switching flow for ingress. In this case, the ingress is on one switch and the egress could be on another switch or across a different ASIC. As with local switching, the packet enters from the front panel ports into the network interface, through the ingress FIFO, and is forwarded out to both the packet buffer and the IFC. In this case, after the descriptor is prepended, the ingress queue scheduler (IQS) comes into play. The packet buffer sends an en-queue signal to the IQS, and when the packet is ready for transmission over the stack, the IQS sends a de-queue signal to the packet buffer. The packet buffer then switches the packet out to the stack interface.

**Figure 85. Packet Walk for Remote Switching – Ingress**



The following figure shows the remote switching flow for egress. In this case, when the packet arrives from the stack interface, the packet buffer accepts the packet, and it follows the the same path as before. The queuing occurs in the SQS, and then EQS. An en-queue signal is sent to the EQS, which sends a de-queue signal back when the packet is ready to be forwarded to the EFC. The egress lookups are done along with final instruction to the packet rewrite block. The packet is rewritten and switched to the appropriate front panel port.

**Figure 86. Packet Walk for Remote Switching – Egress**

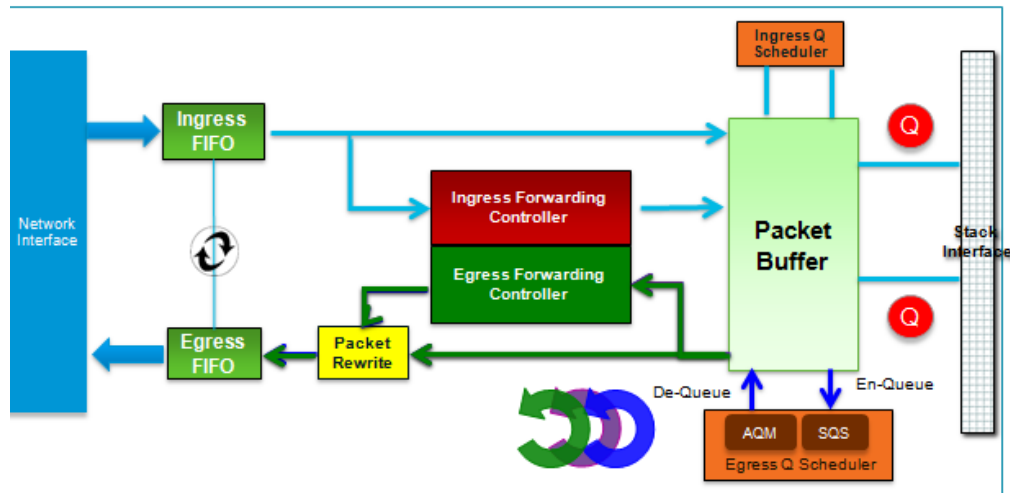


The following figure shows multicast flow with local replication, which means that the source and receivers are in the same ASIC. For example, the source could come in uplink port 1, with receivers on ports 9, 13, and 19 – ports that are in the same ASIC. Flow is similar to the previous examples. The packet comes in and ingress lookups are done in the IFC. The difference is that because it is a multicast packet, active queue management (AQM) comes into play to provide the replication. The AQM checks the list of outgoing interfaces for which packets need to be replicated.

When the en-queue signal comes, the AQM indicates that it is a multicast packet and de-queues the first copy. It sends the signal to the packet buffer to de-queue the packet to, say, port 9. The first copy goes through the EFC, where it checks the configuration on the egress interface (port 9, for example), does the necessary processing, and

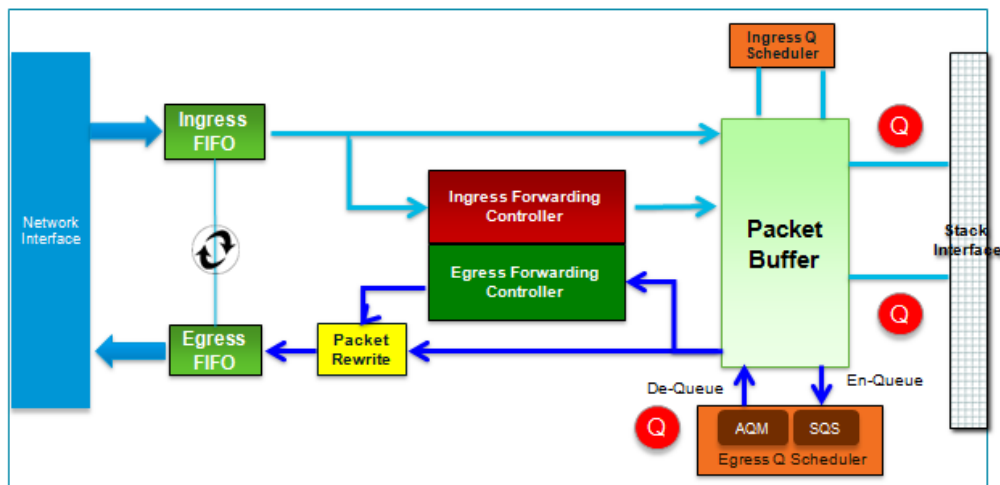
delivers the rewrite instructions to the packet rewrite block. The packet is switched towards the egress FIFO, and subsequently toward port 9. The AQM now indicates that the packet is gone and schedules the next copy.

**Figure 87. Packet Walk for Multicast Local Replication**



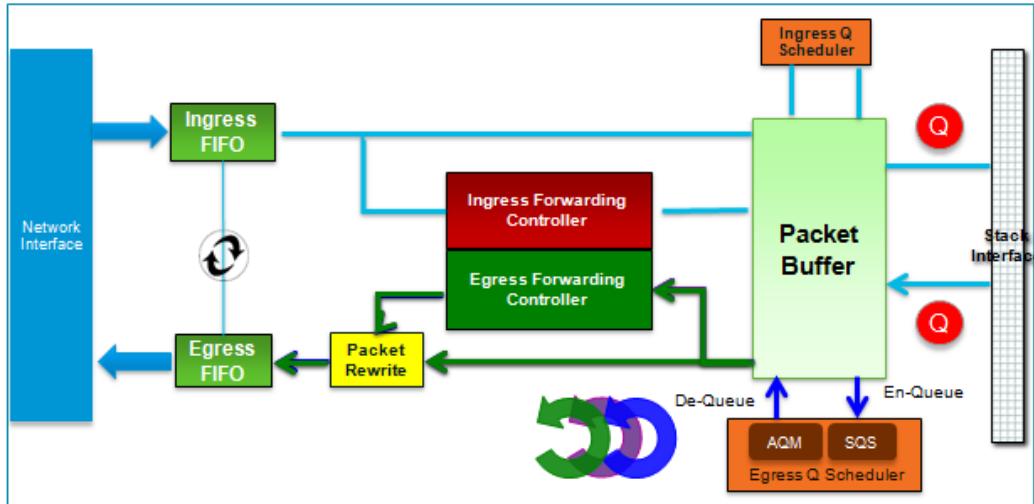
The following figure shows the flow for multicast across stacks on the ingress ASIC. In this case, the source is on one member, while the receivers can be on another member or even an adjacent ASIC on a 48 port system. The source packet enters, goes through the IFC, which does the processing and delivers a descriptor to the packet buffer. Because it is necessary to go through the stack interface, the IQS handles the replication for this source packet. Only one copy is sent over the stack interface, even though there might be many receivers on the stack ring.

**Figure 88. Packet Walk for Multicast Across Stacks – Ingress**



The following figure shows the multicast across stacks for egress. On the egress ASIC/switch, the packet enters from the stack interface. It is the responsibility of the egress ASIC/switch to replicate it to all its local ports. It is an efficient mode of switching multicast packets, where only one copy is sent over the stack ring, while every egress ASIC/switch performs replication for its own set of local ports.

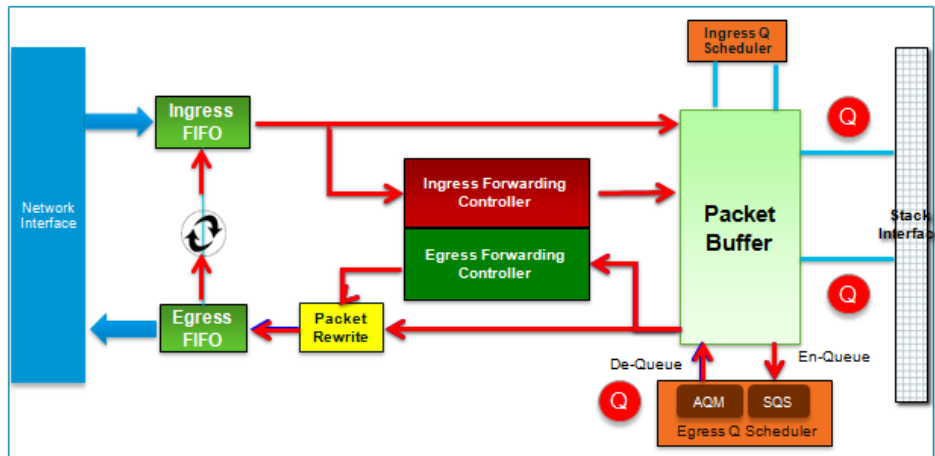
**Figure 89. Packet Walk for Multicast Across Stacks - Egress**



The following figure shows the flow for recirculation. Assume that a packet comes in on a wired port, destined to a wireless client which is associated to an AP on another port. The packet goes through the ingress forwarding controller. Processing is similar to the previously examples. The packet goes to the packet buffer and then the ingress forwarding controller determines that it is destined for a wireless client and needs to be CAPWAP-encapsulated. That information is passed in the descriptor, and the packet is queued as normal, and sent to the EFC.

The EFC applies the egress VLAN or the client policies and gives a final rewrite instruction to the packet rewrite block to encapsulate the packet in CAPWAP. The packet is encapsulated in CAPWAP, delivered to the egress FIFO, and redirected to the recirculation block. This recirculated packet passes through the IFC, where a destination lookup is performed and passed to the EFC for final rewrites and subsequently switched to the correct transmit front panel port.

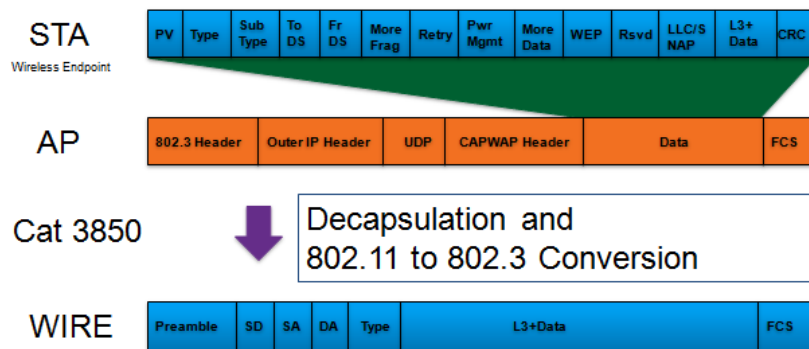
**Figure 90. Packet Walk for Recirculation**



**CAPWAP Frames**

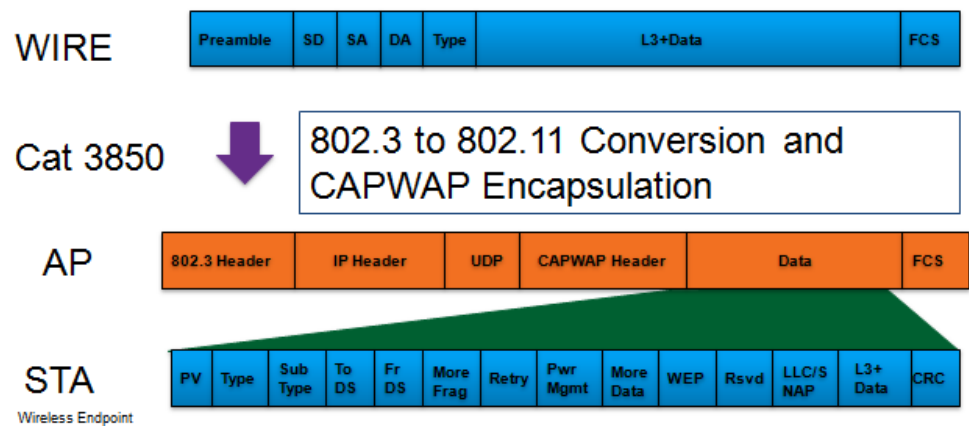
The following figure shows the structure of a CAPWAP frame from a wireless LAN to a wired post. The packet from the end station is transmitted to the AP in 802.11 format, which encapsulates it in CAPWAP, puts it into the data payload area of the CAPWAP header and sends it to the Catalyst 3850, where the decapsulation and the conversion happens inside the switch. In this way, the 802.11 frame is converted to a normal looking Ethernet frame and switched off into the wired world.

**Figure 91. CAPWAP Frames – Wireless to Wired**



The following figure shows the CAPWAP frame going from wired to wireless. It is similar to an Ethernet frame. Conversion and CAPWAP encapsulation happens inside the switch. The AP gets the frame, which is transmitted back to the wireless clients.

**Figure 92. CAPWAP Frames – Wired to Wireless**



## Traffic Walks

The examples in this section show how traffic flows through the architecture.

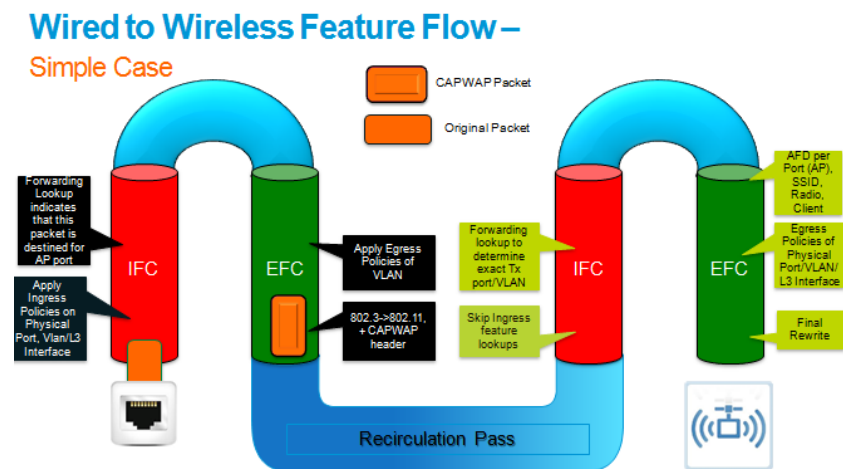
In the simple case shown in the following figure, a packet comes from a wired to a wireless client. There is no fragmentation or encryption of the data. Packets enter the front panel Ethernet port into the IFC where all the ingress lookups are done based on the physical port.

The forwarding controller determines that the packet is destined for an end client off an AP port, and it needs to be CAPWAP-encapsulated. The egress policies of the VLAN and the client are applied and provided to the packet rewrite block, which converts the packet from an Ethernet packet, to a wireless packet and encapsulates the packet in CAPWAP header. It is then shipped to the recirculation pipe, and again fed in through the IFC.

All ingress features are ignored but forwarding lookup is performed on the outer IP header. All the egress policies of the physical interface for port, VLAN, and Layer 3 interface are applied, the final rewrite is done, and the packet is shipped to the AP.



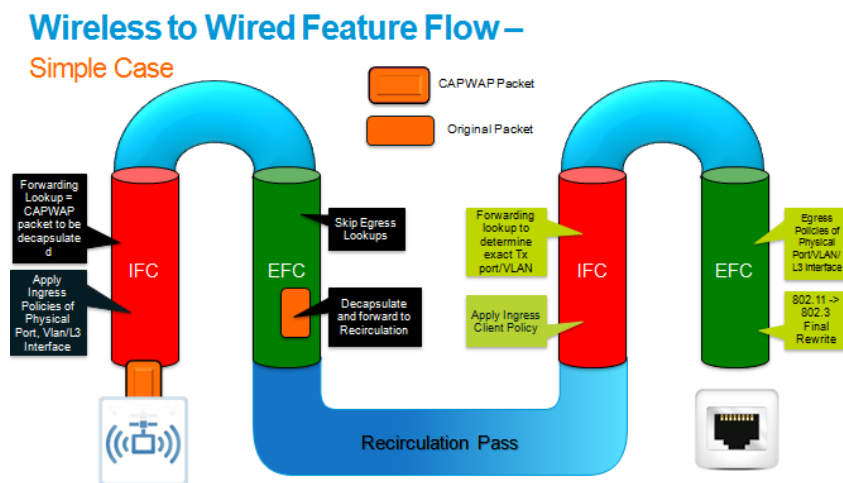
**Figure 93. Wired to Wireless Feature Flow - Simple**



The following figure shows the wireless to wired. An encapsulated packet enters from the wireless LAN, and the policies of the physical port are applied on the outer header. Because the packet is CAPWAP encapsulated, it must be decapsulated for the system to process it further. The packet is decapsulated – at this stage the original wireless packet is exposed to the switch hardware and the rich client features are now applied to this packet after it is forwarded to the IFC via recirculation.

It applies the ingress policy (ACL, QoS, or NetFlow) and does a forwarding/VLAN lookup to determine the egress port. It is sent to the EFC, which applies egress policies and converts the packet from an 802.11 to an Ethernet packet. After a final rewrite, it is shipped to the wired side.

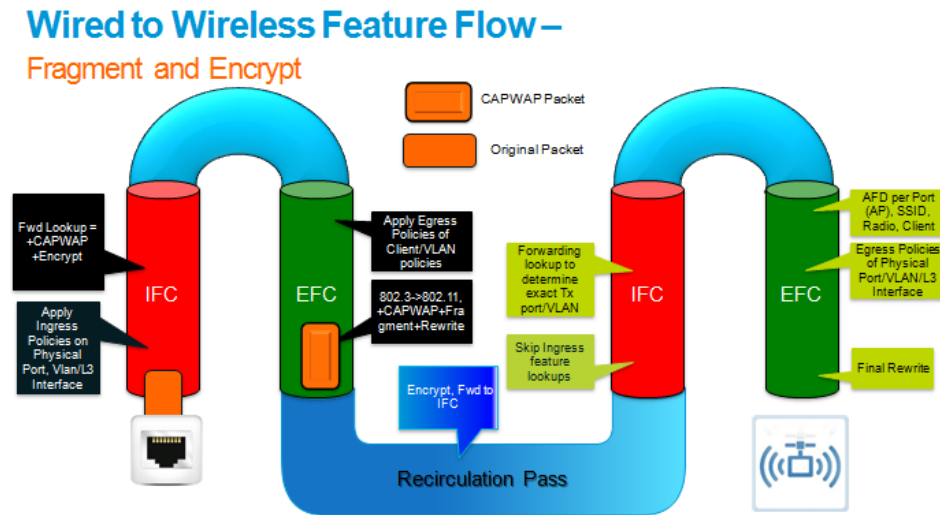
**Figure 94. Wireless to Wired Feature Flow - Simple**



The following figure shows the wireless to wireless flow, which could be headed to another end station associated to the same AP, or it might be destined to a CAPWAP mobility tunnel. So when the packet comes in as CAPWAP encapsulated, there are two passes. First, policies are applied to the outer header. Egress lookups are ignored and the packet is decapsulated and recirculated.

There is now visibility into the original packet, and ingress policies of the client are applied. Forwarding lookup indicates that the packet needs to go out to an AP or a mobility tunnel. The packet is sent to the EFC for CAPWAP encapsulation and egress client policies are applied. The packet is converted from a 802.11 to 802.3 if it is going to head out a CAPWAP mobility tunnel. If it is going to go out back to the AP or the end station, the 802.11 format is retained, the packet is CAPWAP encapsulated, and is fed back into the recirculation. The forwarding lookup determines the destination and then feeds the packets into the EFC pipe where AFD (Approximate Fair Drop) – hierarchical bandwidth management - for wireless is applied. The final rewrite is performed and the packet is shipped back to an AP or to a CAPWAP mobility tunnel.

**Figure 95. Wired to Wireless Feature Flow – Frame and Encrypt**




Consider a situation where the wired to wireless requires DTLS encryption as well as fragmentation. The UADP ASIC is flexible in this case, since it can encrypt and fragment in the same pass and does not require a separate recirculation pass. The original wired packet enters the port. The normal policies on the ingress port, ACL, QoS, are applied on this packet. The forwarding lookup will indicate that this packet needs to be CAPWAP encapsulated, fragmented as well as encrypted. The egress policies of the client or VLAN are applied to this packet, and the packet is converted from 802.3 to 802.11, and the CAPWAP encapsulation is added. The packet is fragmented in the rewrite block as well and delivered to the recirculation block. The recirculation block performs DTLS encryption, and forwards them to IFC. A forwarding lookup based on the outer header is performed on the pass through IFC and sent to EFC to be forwarded to the appropriate port. The feature policies, or queuing policies of the egress interface is applied. A final rewrite is done before transmitting it off to the AP.

## Catalyst 3850 Stacking Architecture

The following figure shows the basic stacking architecture for the Catalyst 3850. There are six rings in total, with three rings going east, and three rings going west. Each ring runs at 40G, so the total stack bandwidth is 240G. With spatial reuse the bandwidth can be doubled to 480G. Before transmission on the stack, packets are segmented into 256 byte segments, which are again reassembled and reordered at the destination switch within the stacking ring.

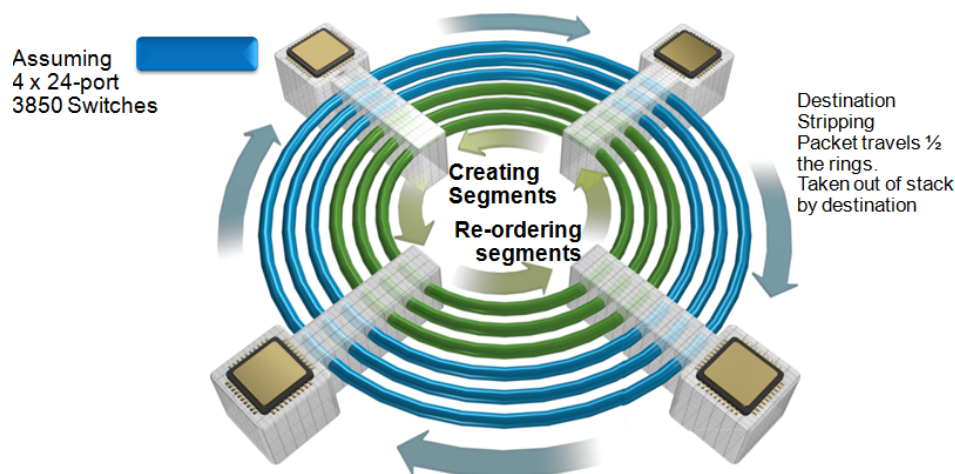
**Figure 96. Stack Ring**



The unicast path is shown in the following figure. It is a token-based access scheme, with a token for each of the six rings. Depending upon the credits which each specific ASIC has, it accesses those rings, it can access some or all of the rings.


In this case the ASIC has credits on four rings so it segments this packet into segments of 256 bytes and accesses the two rings - it sends out the segments, one is a blue ring one is a green ring, which is again east, west. The segments reach the destination, where the packet is reassembled and then sent back into the EFC. The unicast packet travels half the ring, from source to destination, and the destination takes the packet out of the stacking ring – the segments do not come back to the source—the destination takes the packet out of the ring. So a unicast packet travels half the ring where the destination strips off the packet from the ring.

**Figure 97. Unicast Packet Path**



With spatial reuse, multiple communications is possible among different ASICs. In this case, frames enter the two ASICs to the left, each meant for the two ASICs to the right. The transmitting ASIC accesses the rings based on credits it has, segments the frames, and transmits them onto the respective rings. The same occurs on the other transmitting ASIC that accesses different rings than before, segments the frame and send them across. At the respective destinations, the segments are received and taken off the ring. Reordering of these segments occurs, and they are forwarded to the EFC subsequently for further processing.


**Figure 98. Unicast Packet Path – Spatial Reuse**



Multicast behavior on the ring is different. In multicast the segments circle back to the source where the source strips them off the ring – not the destination/s. If a multicast packet is transmitted onto the ring by a source ASIC, it is segmented and sent across the ring for interested receivers to make a copy and send it to their EFCs respectively. The segments are not taken off the ring by these destinations. These segments come back to the source on the ring, and the source detecting that it sent them, takes them off the ring. In this way, efficient multicast replication occurs,


where one copy is sent over the stacking ring, with multiple ASICs grabbing these off the ring and replicating for their local interested ports.

**Figure 99. Multicast Packet Path on the Stack Ring**



Any faults on the stacking ring that occur are detected immediately by hardware. The hardware notifies the software that wraps the rings around the fault immediately. This takes the system stack bandwidth to 240Gbps. When the fault is rectified, the hardware again detects the rectification and signals the software. The software then instead of recovering the ring immediately, validates the stack connections and ensures they are reliable. After the software checks are performed, the stack then returns to its full stack bandwidth to 480Gbps. The point is, the fault detection and isolation (ring wrap), happens instantaneously, compared to the recovery (unwrap) that is graceful and reliable.

**Figure 100. Stacking Ring Healing**



## Catalyst 3850 High Availability

The following section describes the advanced HA capabilities on the Catalyst 3850, provided by the IOS-XE infrastructure, and the stacking infrastructure of the Catalyst 3850.

### HA Redundancy

The HA model on Catalyst 3750-X used hybrid control-plane processing with N:1 stateless control-plane redundancy, distributed Layer2/Layer3 forwarding redundancy, and stateless Layer 3 protocol redundancy.

In contrast, the Catalyst 3850 uses centralized control-plane processing. There is 1:1 stateful redundancy with an active switch and standby switch. Information is synchronized between the two based on SSO capability. Distributed hardware forwarding is used on the Layer 2 and Layer 3 side. IOS HA framework alignment works for the Layer 3 protocols.

### Stacking vs. Catalyst 6500

From a HA perspective, the Catalyst 3850 is similar to the Catalyst 6500. In a stack of four switches, there is an active switch, a standby switch, and other two are member switches. The active and standby switch members run IOSd and WCM, synchronizing information such as the running configuration, startup configuration, licensing, and boot variables. The active switch programs all the hardware resources such as ternary content-addressable memory (TCAM) and tables on the stack system according to the configuration.

The key difference between the Catalyst 6500 and Catalyst 3850 is that the distributed forwarding cards (DFCs) on the Catalyst 6500 are in a chassis. The DFCs cannot be a potential supervisor on the Catalyst 6500 because they lack the supervisor hardware capabilities. This is different from a Catalyst 3850 stack, where each Catalyst 3850 line card switch in the stack is also a potential active or standby switch. Although they are not in a chassis, they are still connected by the high-performance stacking deployment and cabling system provided by the Catalyst 3850's StackWise-480 architecture.

### Catalyst 3850 Software HA Processes

As mentioned above, not all processes run on each switch in the stack. There are some main processes that are supported by other infrastructure and distributed processes. The RP domain is the command center of the stack running software processes, including IOSd and WCM, on the active and standby switches of the stack. The line card domain is a set of software processes, including Forwarding Engine Driver (FED) and Platform Manager that implement the distributed line card portions of the stack control plane. The Infra domain supports the Route Processor and Linecard domains.

The active switch runs the Active RP domain, Linecard domain and Infra domain. The standby switch runs the Standby RP Domain, Linecard domain and Infra domain. The member switches only run Linecard and Infra domains.

### Stack Discovery and Formation

When the stack ports come online, they run a stack discovery process to discover the topology on the stack interfaces. The process is first done by sending broadcast, and then by neighborcast (packets with just one hop), where each switch recognizes its neighbors on the left and right. Then when the topology is complete, the active election starts. The intra and line card domains boot in parallel during the stack topology phase.

In a full ring, discovery exits after all the members are discovered. In a half ring system, the hardware waits for two minutes before exiting, waiting for any switch to initialize. After the discovery exits, the active election begins. Active election is straightforward in the Catalyst 3850. It happens on the highest priority that is configured, or the lowest MAC address.

When one of the switches wins the election, it signals the software module to initialize the active Route Processor domain. The RP domain comes online to run IOSd, and reads through the configuration. The RP domain then programs the hardware on all the member ports on all the member switches. Traffic resumes when programming is complete. The active switch also starts a two minute timer in which it elects its own standby switch. It signals the RP domain on the standby switch to boot up and initialize. Following initialization, there is bulk synchronization between the active and standby RP domains, culminating with the switches established in the active and standby SSO-Hot states.

The following output of the **show switch** command shows the active and standby switch indication in a stack when SSO state is reached.

```
Lightning-II#show switch
Switch/Stack Mac Address : 2037.06cf.0e80

Switch#   Role      Mac Address      Priority  H/W   Current
-----
*1        Active   2037.06cf.0e80   10       PP    Ready
2         Standby  2037.06cf.3380   8        PP    Ready
3         Member   2037.06cf.1400   6        PP    Ready
4         Member   2037.06cf.3000   4        PP    Ready
```

\* Indicates which member is providing the "stack Identity" (aka "stack MAC")

## Stack Addition and Deletion

Every time the stack cable is disconnected or connected, the stack discovery process runs to determine the stack topology. When discovery is complete, the active switch detects the new member and programs its hardware with the relevant configuration. Traffic resumes on the new switch after programing is complete. If the new switch that is added to the already established stack has a higher priority than the current active switch in the stack, the active status is not pre-empted. The current active switch stays active unless the stack reboots or if there is a failover to the standby switch.

Deletion of a switch from the stack is similar to removing a line card from the Catalyst 6500 system. The removal of a switch (intentional or otherwise) triggers a stack topology discovery that signals to the active switch that a neighbor switch was lost. The active switch initiates a cleanup process for that switch member, removing TCAM entries, MAC addresses, CDP tables, shutting down ports, and so on. The result will be half ring, with bandwidth decreasing from 480Gbps to 240Gbps. The configuration of the removed switch is moved to a preprovisioned state. This is done in case the switch returns to its membership following a Stack cable fault; the configuration is still present and applied to the switch.



## Performance and Scalability

The following sections describe the performance and scalability characteristics of the Catalyst 3850 hardware platform.

### UADP Performance

The UADP ASIC makes the Catalyst 3850 a 64-byte line rate switch with 56G of raw capacity on each ASIC for a total of 84 million packets per second switching capacity. Latency through the UADP ASIC is minimal, with recirculation (if needed for a given packet) operating with a recirculation latency of 1-2 microseconds. The UADP ASIC supports a total of 24G downlink ports and 2 x 10G uplink ports, plus 2G of bandwidth available to the CPU and 10G expandable, where 8G is given to recirculation bandwidth. The total is 56G of total bandwidth available per UADP ASIC, which provides significant headroom available for the bandwidth to accommodate next-generation deployments, such as 802.11ac wireless, and beyond.

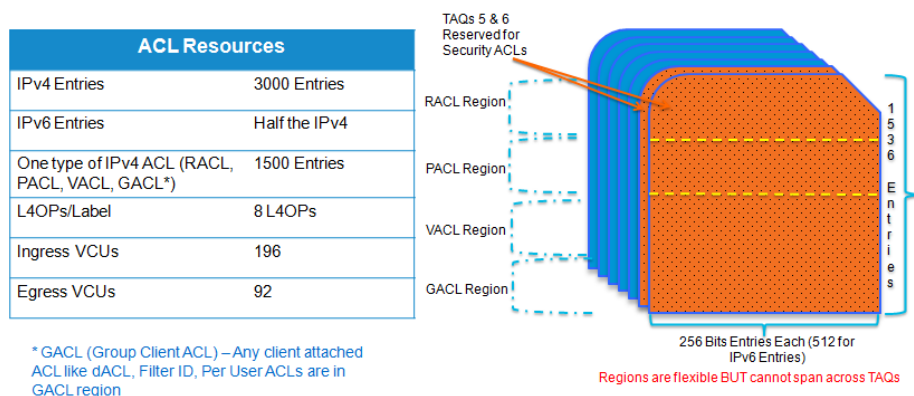
### TCAM and ACL Sizing

The following figure shows data for TCAM and ACL sizing. TCAMs are divided into banks with approximately 256 bit entries for each. Two banks are reserved for ACLs. Quality of Service (QoS) uses another TCAM. A Group ACL (GACL) is used for programming all the policies on the client entities. A Client is defined as a wireless host, or even a wired host passing 802.1x authentication.

The TCAM is divided into regions for Packet ACL (PACL), Router ACL (RACL), VLAN ACL (VACL) and Group ACL (GACL). This arrangement is flexible, so if there are more RACLs, for example, they can take up some of the space in other regions.

However, IPv4 Access-control List Entries (ACEs) in an Access-Control List (ACL) cannot span from one bank to another bank. Hence, an IPv4 ACL can have a maximum of 1500 ACEs. IPv6 takes double the space of IPv4, so the ACEs in an IPv6 ACL are half the size of those for IPv4 entries. There are eight Layer 4 operators that are needed for ACEs containing “less than”, “greater than”, and “range” of Layer-4 (L4) ports. Those parameters can be handled by the Layer 4 operators along with ingress and egress Value Comparison units (VCUs), which help in programming the ACEs containing Layer 4 operations in the ACL.

**Figure 101. TCAM and ACL Sizing**





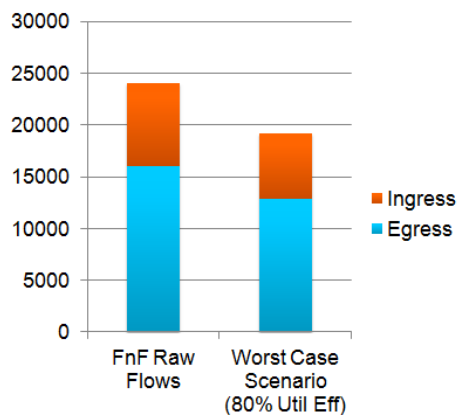
## FNF Scale and Performance

Flexible NetFlow (FNF) scale and performance is fully integrated into the ASIC. There's no performance impact with FNF, as all the calculation and statistics are implemented in hardware. Ingress and egress NetFlow are supported on all ports as is NetFlow sampling. Sample can be one packet out of two, or one packet in 1024 packets. FNF also supports IPv6 and Layer 2.

The UADP ASIC can store up to 24K of flow entries and statistics in hardware. However, due to hash utilization efficiency which is 80%, the actual capacity might be 19K (worst case), although it is possible to get as high as 22K, depending upon how the flows are distributed by the hash. The 24K space is distributed into 8K for ingress NetFlow and 16K for egress NetFlow.

The following figure shows FNF raw flows and the expected worst case flow support.

**Figure 102. TCAM and ACL Sizing**



## QoS Scalability

The following table shows QoS values for scalability.

**Table 7. QoS Scale**

QoS Scale Numbers	
Class-maps (Ingress)	1024
Class-maps (egress)	512
Table-maps (ingress)	16
Table-maps (egress)	16
Aggregate Policers	2000
Microflow Policers (wireless)	24000
Wired Queues/port	8 queues
Wireless Queues/port	4 queues
Buffer/ASIC	6 MB

## Scalability Comparison

The following table compares scalability information for the Catalyst 3850 with the previous generation Catalyst 3750-X. There are performance increases built across the board into the new Catalyst 3850 switch. The bandwidth jumps from 64G to 480G on the stacking side, while the uplinks double from two to four 10G uplinks. Addressing spaces increases significantly, and ACEs are also increased.

**Table 8. Scalability Comparison Between Catalyst 3750X and Catalyst 3850**

Field	Catalyst 3750X	Catalyst 3850
Stacking BW	64G	480G
Uplinks	2x10G	4x10G
MAC addresses	4k	32k
Unicast Routes	6k	24k
IGMP Groups and Multicast routes	1k	4k
Security ACEs	2k	3k
QOS ACEs	0.5k	2.8k
PBR ACEs	0.5k	1.2k

## Migration

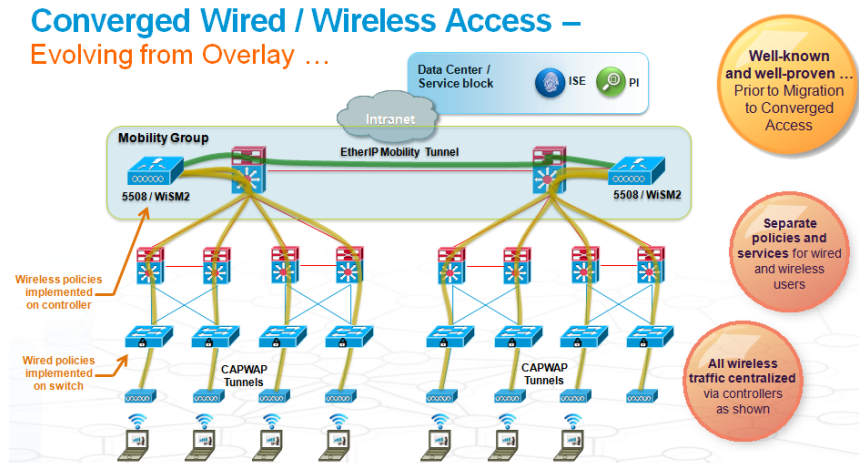
Converged Access splits the Cisco wireless controller architecture into two pieces – a control plane that can run on a WLC 5760, WLC 5508, or WiSM2 controller (or, for smaller deployments, on the Catalyst 3850 switch itself), and a data plane that consists of CAPWAP termination on the Catalyst 3850 switch. This capability evolves the Cisco wired and wireless portfolio, and enables several important advantages, including greater scalability, better traffic visibility, and enhanced traffic control – all achieved through an integrated wired and wireless network deployment.

One area that will need to be undertaken in many deployments is the in-place integration and migration of a wireless network from the CUWN deployment model to a Converged Access deployment model. This might involve a hybrid deployment consisting of both AireOS-based and IOS-based wireless controllers/ switches as an integrated model, or alternatively can involve the wholesale migration towards a Converged Access-only deployment.

For small to mid-sized branch deployments, such integration or migration models are fairly straightforward. For example, for a small branch, all that might be needed is the deployment of the appropriate Catalyst 3850 switches or stacks, and the movement of the associated APs, users, and devices onto the new network at the site. This might be done as a “flash cut” for small deployments, and is fairly straightforward (although a feature comparison should be done for the deployment to compare the capabilities of the existing CUWN deployment with those offered by the Converged Access deployment, as these might be using different software versions and have different features).

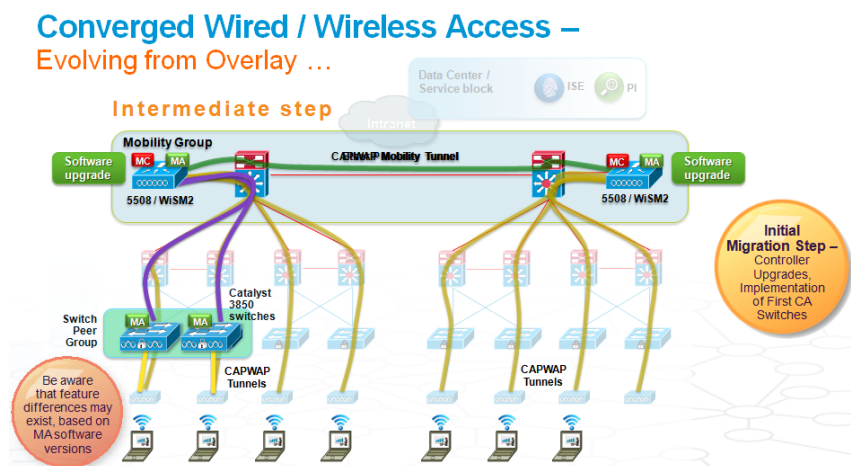
A more complex deployment scenario involves migrating a larger wireless deployment, such as a large campus implementation beginning with a configuration such as that shown in the following figure. This type of migration scenario is explored in more detail in the following text.

**Figure 103. Example Deployment Prior to Migration**



The migration to Converged Access deployments from an existing CUWN deployment for a larger campus site can begin with a software upgrade onto an existing WLC 5508 or WiSM2 controller to prepare it for integration with downstream Catalyst 3850 switches. The minimum AireOS software level to provide the level of integration is AireOS 7.3 MR1. This step of upgrading to this (or a subsequent compatible) AireOS code on any existing discrete wireless controllers provides the ability to enable seamless roaming across the entire wireless domain, with discrete controllers and new Catalyst 3850 switches providing wireless termination capability. At this stage, differences in wireless functionality might exist between the wireless code versions that exist on the MAs (Catalyst 3850 switches) and MCs (WLC 5508s/ WiSM2). A thorough review of the wireless functionality within the wireless network, vs. the functionality provided by the associated Catalyst 3850 switch software at the time of the upgrade, should be undertaken to ensure that all needed wireless functionality is capable of being provided.

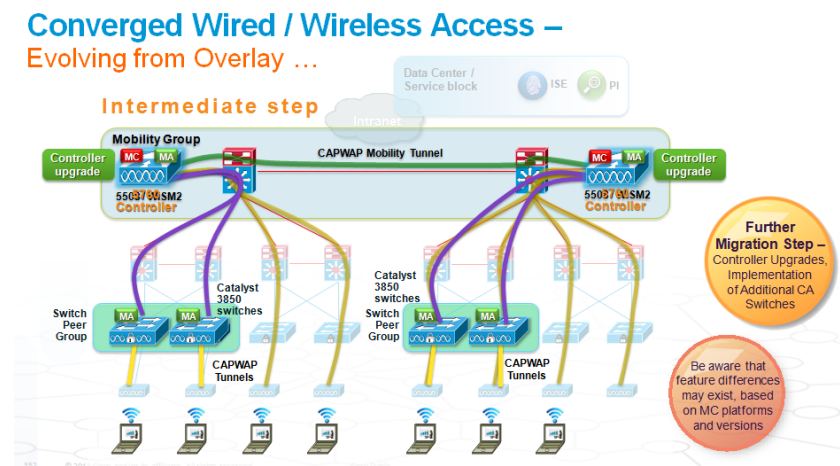
**Figure 104. Beginning Migration to Converged Access**



Integration of an existing AireOS deployment with Catalyst 3850 switches requires the use of the new hierarchical mobility mode throughout the entire MG within the campus, which is one of the reasons that the AireOS 7.3 MR1 code release is required. Any associated release notes for the targeted version of AireOS should be checked to ensure that the hierarchical mobility mode is supported in that release. Also, enabling this mode of MG operation (which uses CAPWAP rather than EoIP tunnels between the discrete controllers within the common MG) requires all controllers within the MG to be restarted (even if they are already running a compatible AireOS code revision).

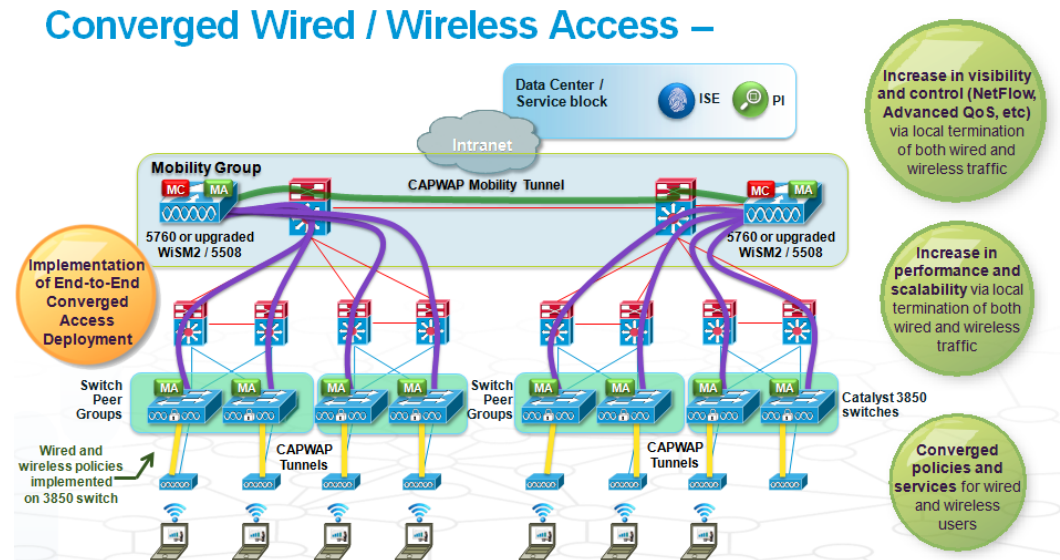
To continue the evolution, additional Catalyst 3850s can be deployed over time. It might be desirable at some point to upgrade or integrate WLC 5760 controllers as MCs to allow for greater system scalability and overall performance. Existing WLC 5508s/ WiSM2s can be deployed in a common MG with WLC 5760s provided that compatible AireOS and IOS-XE releases are employed on the respective platforms. Differences might still remain between the wireless capabilities of the WLC 5760s and the WLC 5508s/ WiSM2s based on software versions and associated functionality. The differences should be thoroughly examined prior to implementing the WLC 5760 into an existing CUWN wireless network deployment.

**Figure 105. Continued Migration to Converged Access**



The final step for such a large campus deployment is the implementation of a complete end-to-end Converged Access deployment across the campus, with all of the associated benefits and capabilities. These include the convergence of wired and wireless QoS and security policy enforcement at the access switch, increase in performance and scalability offered by the distributed nature of the Converged access deployment, and the increase in traffic visibility and control for both wired and wireless users.

**Figure 106. Migration to Converged Access**



Integration of an existing CUWN system with, and migration to, a Converged Access system is fully supported within the Cisco wired/wireless system architecture. The significant compatibility and investment protection offered by Converged Access Group provides a major benefit for Cisco customers as they look to adapt and grow their existing wireless networks.

As with all upgrades, appropriate attention to detail and planning are required. However, the benefits that accrue to a Converged Access deployment within an organization are likely to outweigh the work required for planning and implementation of system integration and migration.

## Conclusion

Converged Access is an exciting and compelling new option for integrated wired and wireless deployments. By addressing the business issues and opportunities provided by true wired/wireless convergence, Converged Access provides many significant benefits.

By terminating wireless traffic at the access edge of the network, Converged Access offers these major advantages:

- Immediate traffic visibility for both wired and wireless traffic flows at the same place in the network and at the same time, using functions such as Flexible NetFlow.
- Traffic control with common policy for QoS and security for both wired and wireless traffic flows, applied at the same place in the network (the access edge).
- Significantly greater scalability across multiple dimensions (including total users, traffic load, flows, and multicast handling) through a distributed wireless deployment model.

By tightly integrating wireless and wired networking at the access layer, Converged Access directly addresses critical business requirements – requirements that are generated by the ever-increasing proliferation of wireless-only devices (smartphones, tablets, and many laptops) on the networks of many organizations. At the same time, Converged Access allows increased bandwidth for these devices (with new wireless technologies such as 802.11ac) to be provided in a scalable fashion.

By using an integrated rather than an overlay approach, Converged Access allows for greater traffic visibility and control. Organizations can harmonize how they handle user traffic at the access edge, regardless of the access method (wired or wireless), with a common set of policies and solutions. With so many new devices attaching to corporate networks and with wireless access speeds increasing with 802.11ac, Converged Access becomes a critical element within an organization's technology portfolio.

With one of the largest and most comprehensive wired and wireless product suites in the networking industry, Cisco is uniquely able to offer this solution. The innovations driven by the Cisco UADP ASIC and the IOS-XE operating system – leveraged as the common foundation of the Catalyst 3850 and WLC 5760 platforms – provide a solid and scalable base upon which Cisco customers can build Converged Access solutions in their own environments. In addition, the advanced, multi-level set of QoS capabilities and integrated security functionality offered by these innovative platforms provide significant benefits from both a business and technical perspective across many different types and sizes of deployments.

By bringing wired and wireless networking together into a true, continuous whole, and by providing seamless interoperability with existing Cisco Unified Wireless deployments, the Converged Access solution offers compelling benefits today, with a clear upgrade path towards the future. Given the ever-growing scope of wireless usage and the ever-increasing capabilities and performance offered as wireless technologies evolve, the Converged Access solution serves as an important adjunct to, and expansion of, Cisco's offerings within the wireless realm.

Converged Access allows organizations to plan and implement their networks today with a keen eye towards the future – a future in which common policies, services, and capabilities are applied in a scalable, secure fashion to all wired and wireless traffic flows.