CHAPTER 2

Cisco Unified Wireless Technology and Architecture

The purpose of this chapter is to discuss the key design and operational considerations in an enterprise Cisco Unified Wireless Deployment.

This chapter examines the following:

- LWAPP
- Roaming
- Broadcast and multicast handling
- Product choices
- Deployment considerations

Much of the material in this chapter is explained in more detail in later chapters of the document. Recommended reading for more detail on the Cisco Unified Wireless Technology is *Deploying Cisco 440X Series Wireless LAN Controllers* at the following URL: http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html.

LWAPP Overview

Lightweight Access Point Protocol (LWAPP) is the core protocol for the centralized WLAN architecture that provides for the management and configuration of the WLAN, as well as the tunneling of the WLAN client traffic to and from a centralized WLAN controller (WLC). Figure 2-1 shows a high level schematic of the basic centralized WLAN architecture, where LWAPP APs connect to a WLC.

Note The term WLC is used as a generic term for all Cisco WLAN Controllers in this document, regardless of whether the WLAN controller is a standalone appliance, an ISR or switch module, or integrated, because the base WLAN features are the same.

Figure 2-1 LWAPP APs Connected to a WLC

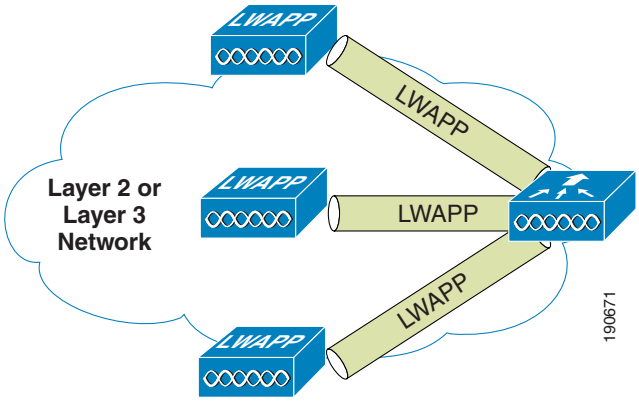

Although the LWAPP protocol has a number of components, only the components of the LWAPP protocol that impact the network design and operation are discussed in this document.

The key features are the LWAPP split MAC tunnel, the various tunnel types, and the WLC discovery process.

Split MAC

One of the key concepts of the LWAPP is concept of split MAC, where part of the 802.11 protocol operation is managed by the LWAPP AP, and other parts of the 802.11 protocol are managed by the WLC.

A schematic of the split MAC concept is shown in Figure 2-2. The 802.11 AP at its simplest level is the 802.11 radio MAC layer providing bridging to a wired network for the WLAN client associated to the AP Basic Service Set Identifier (BSSID). as shown in Figure 2-2a.

The 802.11 standard extends the single AP concept to allow multiple APs to provide an extended service set (ESS), where multiple APs use the same ESS identifier (ESSID; commonly referred to as an SSID) to allow a WLAN client to connect to the same network through different APs.

The LWAPP split MAC concept breaks the APs making up the ESS into two component types: the LWAPP AP, and the WLC. These are linked via the LWAPP protocol across a network to provide the same functionality of radio services, as well as bridging of client traffic in a package that is simpler to deploy and manage than individual APs connected to a common network.

Note Although the split MAC provides a Layer 2 connection between the WLAN clients and the wired interface of the WLC, this does not mean that the LWAPP tunnel passes all traffic; the WLC forwards only IP Ethertype, and its default behavior is not to forward broadcast or multicast traffic. This becomes important when considering multicast and broadcast in the WLAN deployment.

Figure 2-2 Split MAC Concept

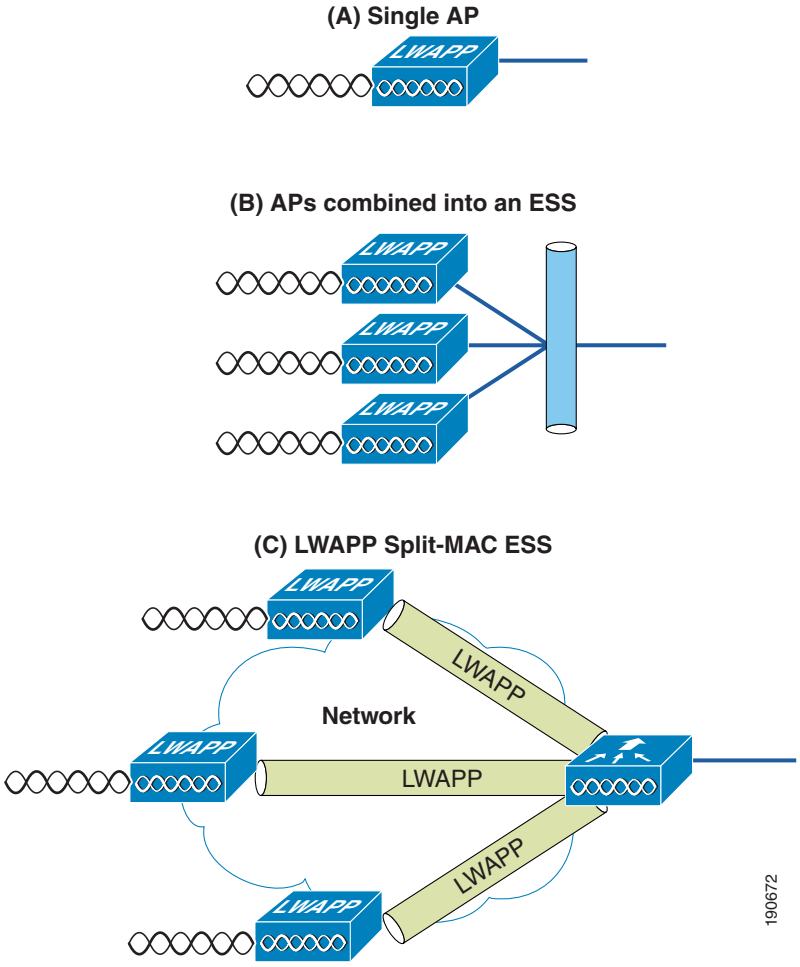

The simple timing-dependent operations are generally managed on the LWAPP AP, and more complex and less time-dependent operations are managed on the WLC.

For example, the LWAPP AP handles the following:

- Frame exchange handshake between a client and AP
- Transmission of beacon frames
- Buffering and transmission of frames for clients in power save mode
- Response to probe request frames from clients; the probe requests are also sent to the WLC for processing
- Forwarding notification of received probe requests to the WLC
- Provision of real-time signal quality information to the switch with every received frame
- Monitoring each of the radio channels for noise, interference, and other WLANs
- Monitoring for the presence of other APs
- Encryption and decryption of 802.11 frames

Other functionality is handled by the WLC. Some of the MAC-layer functions provided by the WLC include the following:

- 802.11 authentication
- 802.11 association and reassociation (mobility)
- 802.11 frame translation and bridging
- 802.1x/EAP/RADIUS processing
- Termination of 802.11 traffic on a wired interface, except for the REAP and H-REAP, which are discussed later in this guide

When the WLAN LWAPP tunnel traffic reaches the WLC, it is mapped to the matching VLAN interface configured on the WLC that defined the SSID, operational state, and WLAN security and quality parameters for that WLAN. WLC WLAN parameters define the wired interface to which the WLC WLAN is mapped. The wired interface on the WLC is typically a VLAN configured on a WLC port, but a WLAN client can be mapped to a specific VLAN interface on the WLC based on parameters sent by the AAA server after successful EAP authentication.

Layer 2 and Layer 3 Tunnels

LWAPP allows tunneling within Ethernet frames (Layer 2) and within UDP packets (Layer 3). This is configurable on the WLC, but not all WLCs support Layer 2 tunneling, and a WLC can support only one tunnel type at a time.

Layer 2 Tunnel

When using Layer 2 LWAPP, the WLC and the LWAPP APs still require IP addresses, but the Layer 2 LWAPP connection uses Ethertype 0xBBBB to encapsulate the LWAPP traffic between the AP and the WLC, and all interaction between the LWAPP AP and the WLC are within the Ethertype 0xBBBB.

Although Layer 2 LWAPP is one of the simplest ways to establish LWAPP connection, and is sometimes the easiest way for the initial configuration of APs or troubleshooting AP WLC connectivity, it is not generally recommended for enterprise deployment, and is not discussed in detail in this document.

The primary reasons for Layer 2 LWAPP not being recommended are the following:

- The need to provide a Layer 2 connection between the LWAPP APs and the WLC limits the location of the APs or WLC, unless Layer 2 connections are extended across the enterprise network, which goes against current networking best practice.
- Layer 2 LWAPP is not supported on all LWAPP AP and WLC platforms.
- Layer 2 LWAPP does not support CoS marking of the Ethertype frames, and therefore is not able to provide end-to-end QoS for tunnelled traffic, although the client traffic DSCP is maintained within the tunnel.

Layer 3 Tunnel

Layer 3 LWAPP tunnels are the recommended LWAPP deployment type, and use IP UDP packets to provide communication between the LWAPP AP, and the WLC. The LWAPP tunnels between the LWAPP APs and the WLC perform fragmentation and reassembly of tunnel packets; allowing the client traffic to use the full 1500 byte MTU and not have to adjust for any tunnel overhead.

Note To optimize fragmentation and reassembly, the number of fragments that the WLC or AP expect to receive is limited. The ideal supported MTU for deploying the Cisco Unified Wireless is 1500, but the solution operates successfully over networks where the MTU is as small as 500 bytes.

The following are some Layer 3 LWAPP packet captures to illustrate LWAPP operation. These three sample decodes of the LWAPP packets use the Ethereal Network Analyzer.

Note Note that the default Ethereal configuration does not decode Cisco LWAPP packets correctly. This can be corrected by using the "SWAP Frame Control" option in protocol preferences.

Figure 2-3 shows the decode of a LWAPP control packet. This is a packet from the WLC, and uses UDP source port 12223, as do all LWAPP control packets from the WLC. The Control Type 12 is a configuration command, where the AP configuration is passed to the LWAPP AP by the WLC. The payload in this LWAPP packet is AES encrypted, using keys derived during the PKI authentication performed between the LWAPP AP and WLC.

Figure 2-3 LWAPP Control Packet

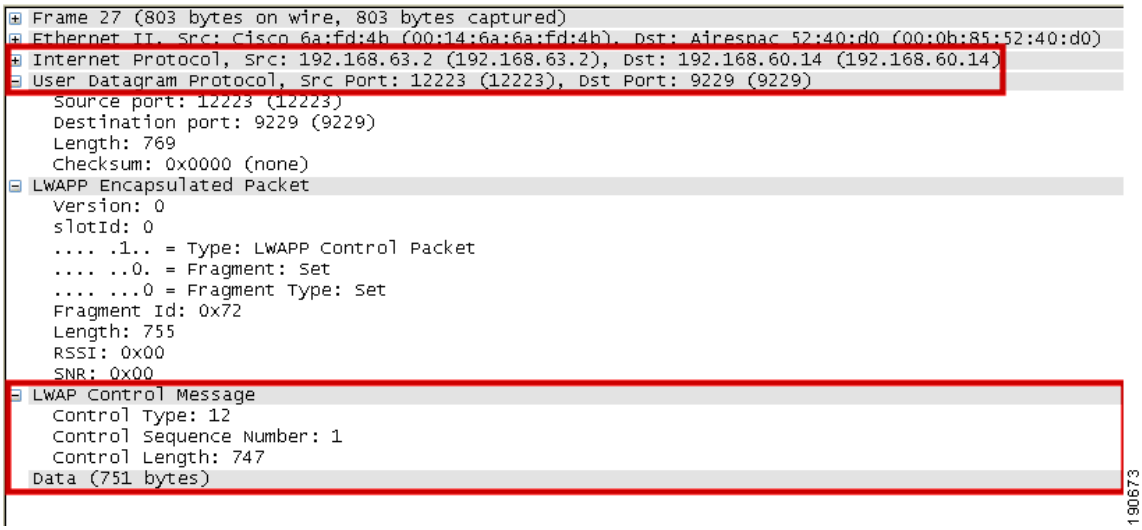

Figure 2-4 shows a decode of an LWAPP packet containing an 802.11 probe request. This packet is from the LWAPP AP to the WLC, and uses UDP port 12222, as do all LWAPP-encapsulated 802.11 frames. In this case, RSSI and SNR values are also included in the LWAPP packet to provide RF information to the WLC.

Figure 2-4 802.11 Probe Request in LWAPP

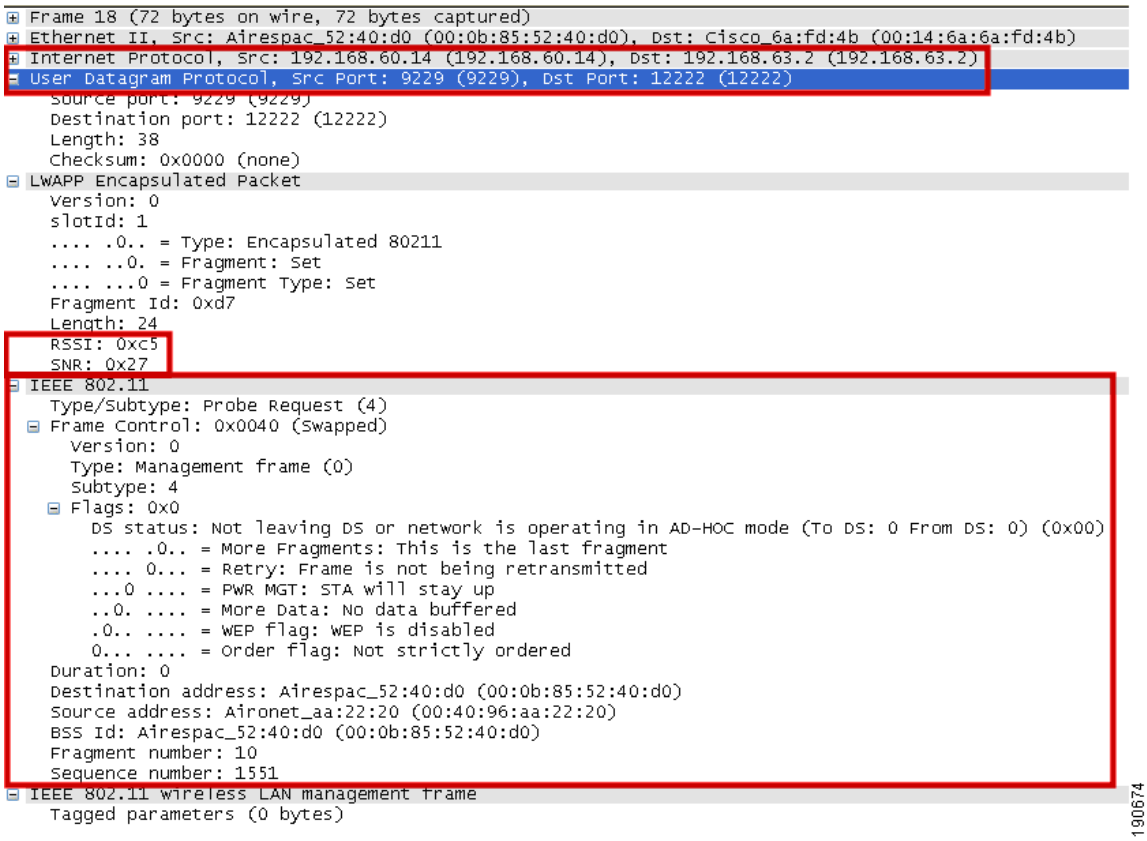

Figure 2-5 shows another LWAPP-encapsulated 802.11 frame, but in this case it is an 802.11 data frame, like that shown in Figure 2-4. It contains the complete 802.11 frame, as well as the RSSI and SNR information for the WLC, and is primarily shown here to demonstrate that the 802.11 data frame is treated the same as other 802.11 frames by LWAPP. Points highlighted in Figure 2-5 are the fragmentation supported by LWAPP, where the LWAPP AP and WLC automatically fragment LWAPP packets to fit the minimum MTU size between the LWAPP AP and the WLC. Note from the Ethereal decode that the frame control decode bytes have been swapped; this is done in the Ethereal protocol decode of LWAPP to take into account that some LWAPP APs swap these bytes.

Figure 2-5 802.11 Data Frame in LWAPP

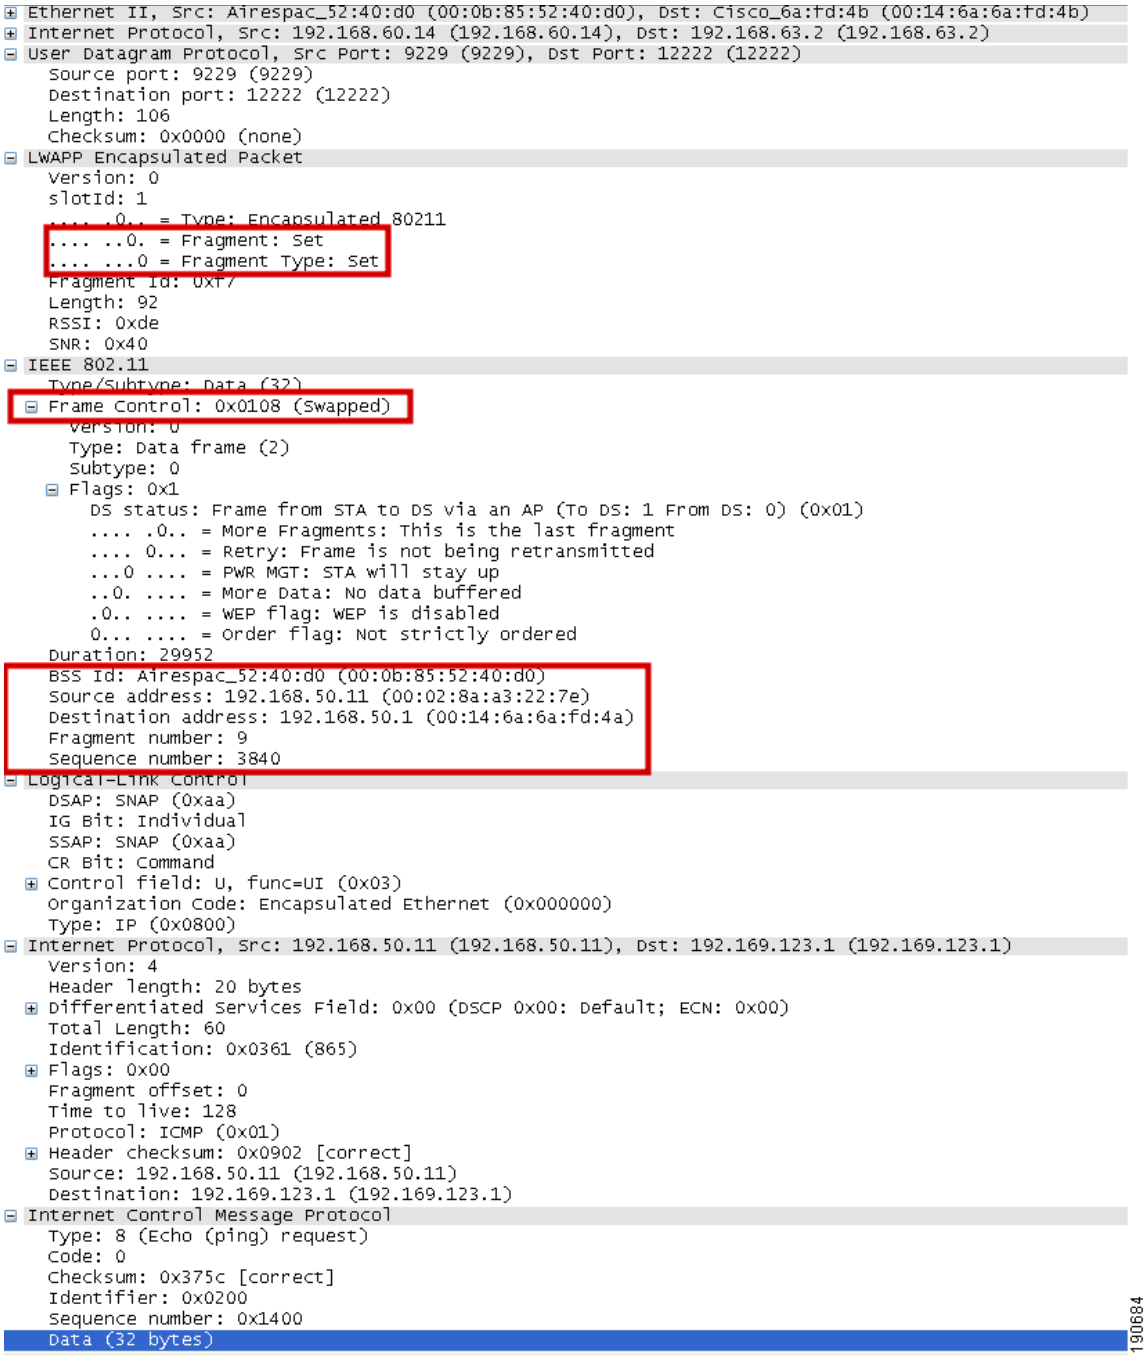

WLC Discovery and Selection

This section discusses the typical Layer 3 LWAPP behavior after a reset of the LWAPP AP, but not the various options that may occur with a new AP deployment.

For a complete description, see the 440X Series Wireless LAN Controllers Deployment Guide at the following URL:
http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html.

The following sequence takes place:

1. The AP broadcasts a Layer 3 LWAPP discovery message on the local IP subnet. Any WLC configured for Layer 3 LWAPP mode that is connected to the local IP subnet receives the Layer 3 LWAPP discovery message. Each of the WLCs receiving the LWAPP discovery message reply with a unicast LWAPP discovery response message to the AP.
2. When a feature called Over-the-Air Provisioning (OTAP) is enabled on a WLC, APs that are joined to the WLC advertise their known WLCs in neighbor messages that are sent over the air. New APs attempting to discover WLCs receive these messages and then unicast LWAPP discovery requests to each WLC. (OTAP is not supported in IOS APs in their initial state; that is, an IOS AP fresh out of the box cannot use OTAP to find a WLC.) WLCs receiving the LWAPP discovery request messages unicast an LWAPP discovery response to the AP.
3. The AP maintains previously learned WLC IP addresses locally in NVRAM. The AP sends a unicast LWAPP discovery request to each of these WLC IP addresses. Any WLC receiving the LWAPP discovery request responds by sending an LWAPP discovery response to the AP. These WLC IP addresses are learned by the AP from previously joined WLCs. The stored WLC IP addresses include all of the WLCs in previously joined WLC mobility groups. (The mobility group concept is discussed in greater detail later in this document.)
4. DHCP servers can be programmed to return WLC IP addresses in vendor specific “Option 43” in the DHCP offer to lightweight Cisco APs. When the AP gets an IP address via DHCP, it looks for WLC IP addresses in the Option 43 field in the DHCP offer. The AP sends a unicast LWAPP discovery message to each WLC listed in the DHCP option 43. WLCs receiving the LWAPP discovery request messages unicast an LWAPP discovery response to the AP.
5. The AP attempts to resolve the DNS name “CISCO-LWAPP-CONTROLLER.localdomain”. When the AP is able to resolve this name to one or more IP addresses, the AP sends a unicast LWAPP discovery message to the resolved IP address(es). Each WLC receiving the LWAPP discovery request message replies with a unicast LWAPP discovery response to the AP.
6. If, after Steps 1 through 5, no LWAPP discovery response is received, the AP resets and restarts the search algorithm.

Typically, the DHCP or DNS discovery mechanism is used to provide seed WLC addresses, and then WLC discovery response provides a full list of WLCs from the mobility group.

An LWAPP AP is normally configured with a list of up to 3 WLCs that are its preferred WLCs. If these WLCs are unavailable or over-subscribed, the AP chooses another WLC from the list of WLCs in the response to its discovery requests and chooses the least-loaded WLC.

Components

The three primary components to the Cisco Unified Wireless Architecture are the APs, the WLC, and the WCS. This section describes the AP and WLC options; the WCS is discussed in detail in another chapter.

WLCs

This document refers to all Cisco Unified Wireless controls as WLCs for convenience, and because of the commonality of features across the various Cisco Unified Wireless WLCs.

The following summarizes various Cisco Unified Wireless WLCs and their features:

- 2006—Standalone WLC that supports up to six APs, with four Fast Ethernet interfaces that can be configured as dot1q trunks to provide connection into the wired network. Ideal for a small-to-medium size office, where an H-REAP would be unsuitable because of the number of users, WAN requirements, or client roaming requirements.
- 4402—Standalone WLC that supports either 12, 25, or 50 APs, with two SFP-based Gigabit Ethernet ports, that can be configured as dot1q trunks to provide connection into the wired network, Gigabit ports can be link aggregated to provide an EtherChannel connection to the wired network. Ideal for medium-size offices or buildings.
- 4404—Standalone WLC that supports 100 APs with four SFP-based Gigabit Ethernet ports that can be configured as dot1q trunks to provide connection into the wired network. Gigabit ports can be link aggregated to provide an EtherChannel connection to the wired network. Ideal for large offices, buildings, and even a small campus.
- WLCM—WLC module for integration into Cisco ISR routers. The WLCM supports up to six APs. The WLCM appears as an interface on the ISR router that can be configured as a dot1q trunk to provide a routed connection to the wired network. Ideal for small-to-medium size offices requiring an integrated solution.
- WS-C3750G—Integrated WLC that supports either 25 or 50 APs, integrated with the 3750 backplane appearing as two Gig Ethernet ports, that can be configured as dot1q trunks to provide connection into the 3750. The Gig ports can be link aggregated to provide an EtherChannel connection to the 3750. Integration with the 3750 provides the WLC with a direct connection into the advanced routing and switching features of the 3750 stackable switch. Ideal for medium-size offices or buildings.
- WiSM—WLC module for integration into a 6500 switch. The WiSM supports up to 300 APs. The WiSM appears as a single link aggregated interface on the 6500 that can be configured as a dot1 trunk to provide connection into the 6500 backplane. Ideal for large buildings or campuses.

Table 2-1 summarizes the Cisco Unified Wireless Controllers.

Table 2-1 Cisco Unified Wireless Controller Summary

Product	Number of APs	Interfaces	Comments
2006	6	4x Fast Ethernet	Cannot be a Mobility Anchor, does not support Layer 2 LWAPP, no H-REAP support
4402	12 or 25	2x Gig Ethernet	
4404	50 or 100	4x Gig Ethernet	

Table 2-1 Cisco Unified Wireless Controller Summary

WLCM	6	ISR backplane	Cannot be a Mobility Anchor, does not support Layer 2 LWAPP, no H-REAP support, and Layer 3 only connection to the network
WS-C3750G	25 or 50	3750 backplane	Full featured 3750 stackable switch with integrated WLC
WiSM	300	6500 backplane	Module directly connecting to the 6500 backplane

APs

Within the Cisco Unified Wireless Architecture, there are two categories of APs: autonomous and lightweight (LWAPP). This section briefly discusses the various models of AP products available within each category, and contrasts features, functionality, and applications.

Cisco Autonomous APs

APs in this category consist of the original Aironet product line. The following select models are available in or are capable of being field upgraded to lightweight (LWAPP) mode of operation. This feature permits an enterprise to standardize on a common AP platform that can be deployed in hybrid topologies.

First generation autonomous APs are as follows:

- AP 1100—This single band AP is orderable as an 802.11g AP or 802.11b AP that is field upgradeable to 802.11g. It possesses an integrated antenna and is considered an entry level AP for enterprise deployments. The part number for the LWAPP AP is AIR-LAP1121G-x-K9 where x= the regional code.
- AP 1200—A single band 802.11b/g AP that is targeted for enterprise deployments. Unlike the 1100 series, the 1200 supports connection to external antennas for more flexibility. It can be field upgraded to support an 802.11a radio as well as upgradeable for lightweight (LWAPP) operation. The part number for the LWAPP AP is AIR-LAP1231G-x-K9 where x= the regional code.
- AP 1230AG—Dual band 802.11a/b/g AP with external connectors for antennas in both bands. It does not possess all of the features (most notably 802.3af PoE) and RF performance of the 1240AG. It also comes in a lightweight (LWAPP) version or can be upgraded later to lightweight mode of operation. The part number for the LWAPP AP is AIR-LAP1232G-x-K9 where x= the regional code.

Second generation autonomous APs are as follows:

- AP 1130AG—The AG version is dual band (a/b/g) AP with integrated antennas. It is designed to be wall-mounted and also uses an integrated dual band antenna. The 1130AG is available in a lightweight (LWAPP) version for implementation in centralized (WLC)-based deployments. The autonomous version can be later upgraded for lightweight operation. The part number for the LWAPP AP is AIR-LAP1131AG-x-K9 where x = the regional code.

- AP 1240AG—A dual band 802.11 a/b/g AP designed for deployments in challenging RF environments such as retail and warehousing. The 1241AG possesses external connections for antennas in both bands. It is the most feature-rich AP in the autonomous category and is also available in a lightweight (LWAPP) version. For greatest flexibility, the autonomous version can be upgraded later to lightweight mode of operation. Other notable features include pre-installed certificates for LWAPP operation mode and the ability to support hybrid REAP. The part number for the LWAPP AP is AIR-LAP1242AG-x-K9 where x = the regional code,
- AP 1300—A single band 802.11b/g AP/bridge designed for outdoor deployments. It comes with an integrated antenna or can be ordered with RP-TNC connectors to support external antenna applications. The LWAPP AP part number is AIR-LAP1310G-x-K9 where x = the regional code.

Cisco Lightweight APs

APs in this category consist of the original Airespace product line, but also include select autonomous AP models above. The following lightweight models can be used only in WLC topologies:

- AP 1010—Dual band, zero touch, 802.11a/b/g AP intended for basic enterprise LWAPP/WLC deployments. The 1010 comes with dual internal sector antennas. The part number is AIR-AP1010-x-K9 where x = the regional code.
- AP 1020—Similar to the 1010, but in addition to its internal sector antennas, it also includes RP-TNC connectors for external 2.4 and 5 GHz antennas. The part is number AIR-AP1020-x-K9 where x = the regional code.
- AP 1030—Also referred to as the REAP AP or Remote Edge AP, the 1030 possesses the same capabilities, features, and performance as the 1020, in addition to being able to be deployed in environments where it is not practical to deploy a WLC, such as in small branch offices. The part number is AIR-AP1030-x-K9 where x = the regional code.
- AP 1500—A dual band AP specifically designed for outdoor, point-to-point, and multipoint MESH deployments. The 802.11a band is used for backhaul while the b/g band is used for wireless client access. The 1500 uses (patent pending) Adaptive Wireless Path Protocol (AWPP) for optimal routing through MESH topologies.

Table 2-2 and Table 2-3 provide a comparison summary of the APs discussed above.

Table 2-2 AP Comparison (1)

Cisco Series	802.11b	802.11g	802.11a	Autonomous	Light weight	# Broadcasted SSIDs	Preinstalled Cert?
1000	YES	YES	YES	NO	YES	16	YES
1100	YES	YES	NO	YES	YES	8	NO
1130AG	YES	YES	YES	YES	YES	8	YES[1]
1200	YES	YES	Optional	YES	YES	8	YES[1]
1230AG	YES	YES	YES	YES	YES	8	YES[1]
1240AG	YES	YES	YES	YES	YES	8	YES[1]
1300	YES	YES	NO	YES	YES	8	NO
1500	YES	YES	YES	NO	YES	16	YES

1. Units shipped prior to Aug 2005 require a Cisco-provided utility to load self-signed certificate, and an 11g radio is required.

Table 2-3 AP Comparison (2)

Cisco Series	Office and similar environments	Challenging Indoor environments	Outdoors
1010	Recommended*	Not Recommended	Not Recommended
1020	Recommended*[1]	Recommended*[1]	Not Recommended
1100	Recommended	Not Recommended	Not Recommended
1130AG	Ideal	Not Recommended	Not Recommended
1200	Recommended***	Recommended	Recommended****
1230AG	Recommended***	Recommended	Recommended****
1240AG	Recommended***	Ideal	Recommended****
1300	Not Recommended	Not Recommended	Ideal
1500	Not Recommended	Not recommended	Ideal*

[1] Or 1030 for Remote offices

* LWAPP Deployments Only

** Autonomous Deployments Only

*** Particularly for deployments above suspended ceilings

**** Can be used outdoors when deployed in weatherproof NEMA rated enclosure

For further detailed information, see the following link:
http://www.cisco.com/en/US/prod/collateral/wireless/ps5679/ps5861/prod_brochure0900aecd8035a015_ps6108_Products_Brochure.html.

Mobility Groups, AP Groups, and RF Groups

Within the Cisco Unified Wireless Architecture, the following are three important concepts in grouping devices:

- Mobility group
- AP groups
- RF groups

This section describes their purpose in the Cisco Unified Wireless Architecture. For more details on operation and configuration these groups, see the following URLs:

- Deploying Cisco 440X Series Wireless LAN Controllers—
http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html
- Cisco Wireless LAN Controller Configuration Guide, Release 4.0—
http://www.cisco.com/en/US/docs/wireless/wcs/4.0/configuration/guide/wcscfg40.html

Mobility Groups

A mobility group is a group of WLCs that acts as one virtual WLC by sharing key client, AP, and RF information. The WLC is able to make decisions based on the data from the entire mobility group domain rather than simply from its own connected APs and clients.

The mobility group forms a mesh of authenticated tunnels between the WLCs in the mobility group, allowing any WLC to directly contact other WLCs in the group, as shown in Figure 2-6.

Figure 2-6 WLC Mobility Group

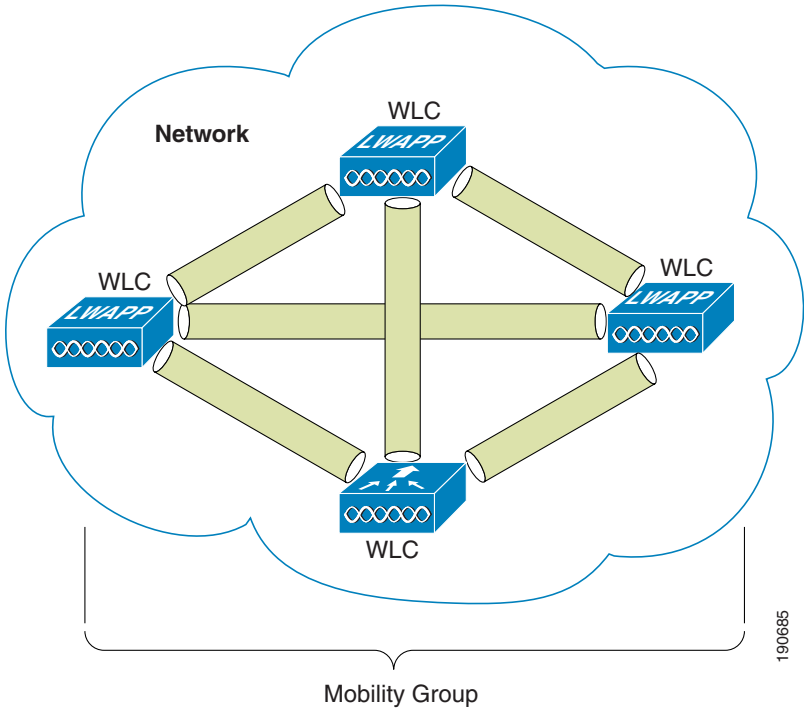

Creating Mobility Groups

Creating mobility groups is simple and well documented, but there are the following important considerations:

- Up to 24 WLAN controllers and 3600 APs are supported per mobility group.
- The WLCs do not have to be the same type to be in the same mobility group; a 4402, 4404, WiSM, WLCM, and 2006 can all be in the same mobility group, but the WLCs should be running the same software revision. Mobility groups do not break because of software differences but they do rely on matching configuration on WLC WLANs.
- A mobility group requires all WLCs in the group to have the same virtual IP address.
- Each WLC has the same mobility group name, and is in the mobility list of each other WLC.
- For a client to seamlessly roam between mobility group members, the client WLANs must match in SSID and WLAN security configuration.

Putting WLCs in Mobility Groups

The primary purpose of a mobility group is the creation of a virtual WLAN domain between multiple WLCs, providing a comprehensive wireless view for client roaming. The creation of a mobility group makes sense only when there is overlapping wireless coverage between APs connected to different WLCs. For example, there is nothing to be gained in having campus and branch WLCs in the same mobility group. Even within the campus, if there is no WLAN coverage between buildings, there is no benefit in having the WLCs of isolated APs within the same mobility group.

Mobility Group Rule Breakers

When using the mobility anchor feature, the anchor WLC can have connections with more than 24 WLCs. Mobility group members of a mobility anchor do not have to have a mobility group connection between each other, but must be in the mobility list of the anchor controller.

For a discussion on mobility anchor configuration, see Chapter 12, “Cisco Unified Wireless Guest Access Services.”

AP Groups

In a default deployment, a WLAN is mapped to a single interface per WLC. Consider a deployment scenario, where you have a 4404-100 WLC supporting the maximum number of APs (100). Now consider a scenario with 25 users associated to each AP. In the default configuration, you have 2500 users on the same VLAN. This is not be a problem because LWAPP is an overlay architecture; there is no spanning tree requirement to all 100 APs. However, there can be broadcast- or multicast-intensive applications running on the wireless LAN end clients, and this leads to a need to break up the number of clients on a single subnet. Also, you may want to distribute the end client load across multiple interfaces in the infrastructure. To create smaller user domains, you should make use of the AP Groups feature and create site-specific VLANs. Figure 2-7 illustrates the AP groups and site-specific VLAN concept.

Note AP groups do not allow multicast roaming across group boundaries; this is discussed in more detail later in this design guide.

Figure 2-7 AP Groups and Site-Specific VLANS

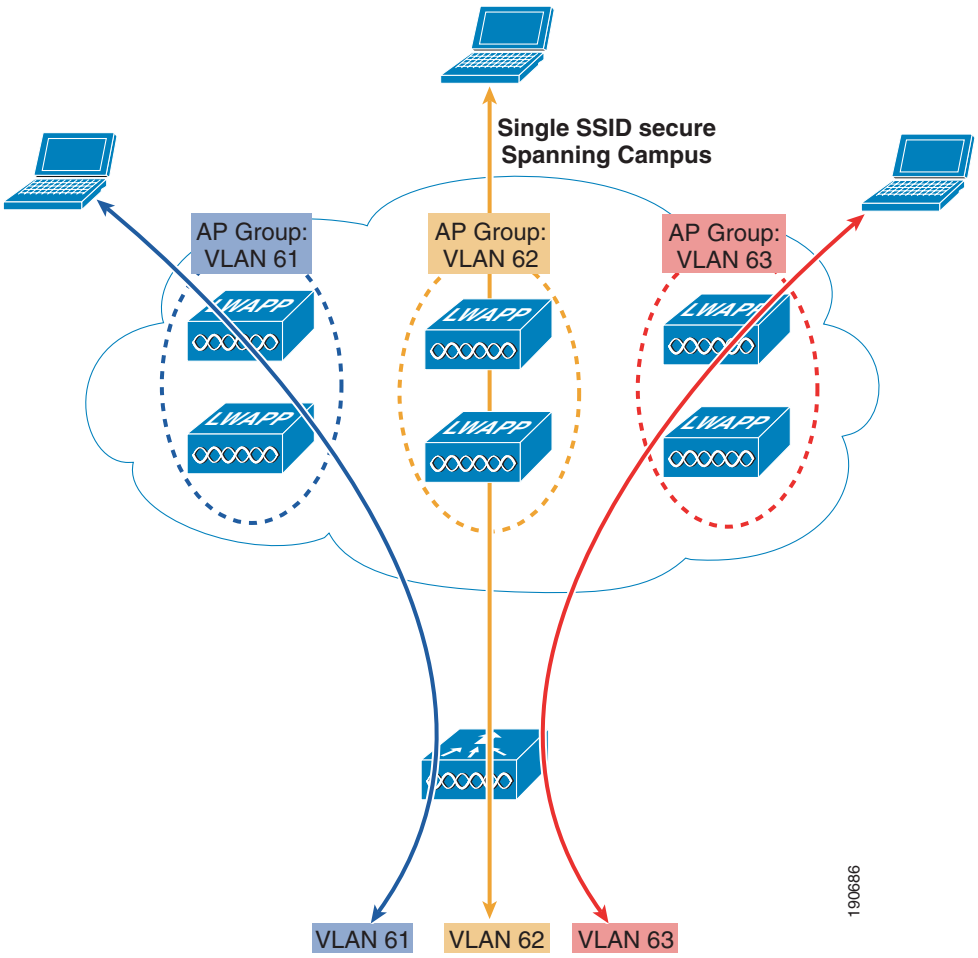

In Figure 2-7, there are three dynamic interfaces configured, mapping to three site-specific VLANs: VLANs 61, 62, and 63. These site-specific VLANs apply to the secure SSID for normal corporate users. A corporate user associating to the secure SSID on an AP in the AP Group corresponding to VLAN 61 gets an IP address on the VLAN 61 IP subnet. A corporate user associating to the secure SSID on an AP in the AP Group corresponding to VLAN 62 gets an IP address on the VLAN 62 IP subnet. A corporate user associating to the secure SSID on an AP in the AP Group corresponding to VLAN 63 gets an IP address on the VLAN 63 IP subnet. Roaming between site-specific VLANs is treated internally by the WLC as a Layer 3 roaming event, so the wireless LAN client maintains its original IP address.

RF Groups

RF groups, also known as RF domains, are another critical deployment concept. An RF group is a cluster of WLCs that coordinate their dynamic radio resource management (RRM) calculations on a per 802.11 PHY type.

An RF group exists for each 802.11 PHY type. Clustering WLCs into RF domains allows the dynamic RRM algorithms to scale beyond a single WLC and span building floors, buildings, and even campuses. RF RRM is discussed in more detail in a later chapter of this document, but can be summarized as follows:

- LWAPP APs periodically send out neighbor messages over the air that include the WLC IP address and a hashed message integrity check (MIC) from the timestamp and BSSID of the AP.
- The hashing algorithm uses a shared secret (the RF Group Name) that is configured on the WLC and pushed out to each AP. APs sharing the same secret are able to validate messages from each other via the MIC. When APs on different WLCs hear validated neighbor messages at a signal strength of -80 dBm or stronger, the WLCs dynamically form an RF group.
- The members of an RF domain elect an RF domain leader to maintain a "master" power and channel scheme for the RF group.
- The RF group leader analyzes real-time radio data collected by the system and calculates the master power and channel plan.
- The RRM algorithms try to optimize around a signal strength of -65 dBm between all APs, and to avoid 802.11 co-channel interference and contention as well as non-802.11 interference.
- The RRM algorithms employ dampening calculations to minimize system-wide dynamic changes. The end result is dynamically calculated, near-optimal power and channel planning that is responsive to an always changing RF environment.
- The RF group leader and members exchange RRM messages at a specified updated interval, which is 600 seconds by default. Between update intervals, the RF group leader sends keepalive messages to each of the RF group members and collects real-time RF data.

Roaming

Roaming in an enterprise 802.11 network can be described as when an 802.11 client changes its AP association from one AP within an ESS to another AP within the same ESS. Depending on the network features and configuration, a lot may occur between the clients, WLCs, and upstream hops in the network, but at the most basic level, it is simply a change of association.

When a wireless client authenticates and associates with an AP, the WLC of the AP places an entry for that client in its client database. This entry includes the client MAC and IP addresses, security context and associations, QoS context, WLAN, and associated AP. The WLC uses this information to forward frames and manage traffic to and from the wireless client.

When the wireless client moves its association from one AP to another, the WLC simply updates the client database with the new associated AP. If necessary, new security context and associations are established as well.

A Layer 2 roam occurs when a client roams from one AP and (re)associates to a new AP, providing the same client subnet. In most cases, the foreign AP can be on the same WLC as the home AP.

This is a very simple roam because the WLC maintains a database with all the information of the client. All upstream network components from the WLC are unaffected by the client moving from home to foreign AP, as illustrated in Figure 2-8.

Figure 2-8 Layer 2 Roam

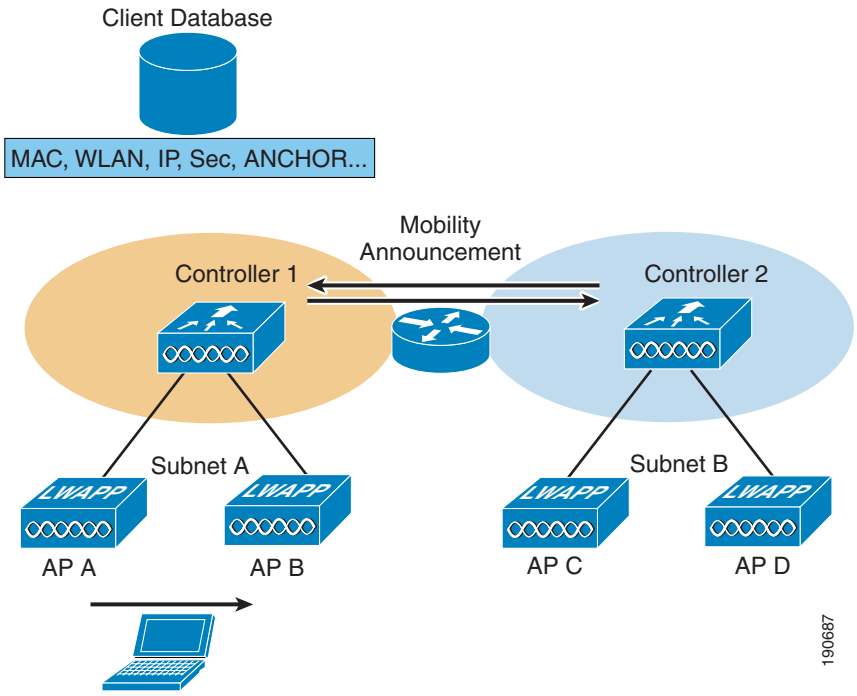

In instances when there are multiple WLCs connected to the same subnet, and therefore a client can roam between WLCs but remain on the same subnet, mobility announcements are passed between the related WLCs to pass client context information between WLCs. This WLC then becomes the anchor WLC for that client.

WLC to WLC, Different Subnet

In instances where the client roams between APs that are connected to different WLCs and the WLC WLAN is connected to a different subnet, a Layer 3 roam is performed, and there is an update between the new WLC (foreign WLC) and the old WLC (anchor WLC) mobility databases.

If this is the case, return traffic to the client still goes through its originating anchor WLC. The anchor WLC uses Ethernet over IP (EoIP) to forward the client traffic to the foreign WLC, to where the client has roamed. Traffic from the roaming client is forwarded out the foreign WLC interface on which it resides; it is not tunneled back. The client MAC address for its default gateway remains the same, with the WLC changing the MAC address to the local interface gateway MAC address when the client traffic is sent to the default gateway.

The example in Figure 2-9 describes a client Layer 3 roam with PMK.

Figure 2-9 Layer 3 Roaming

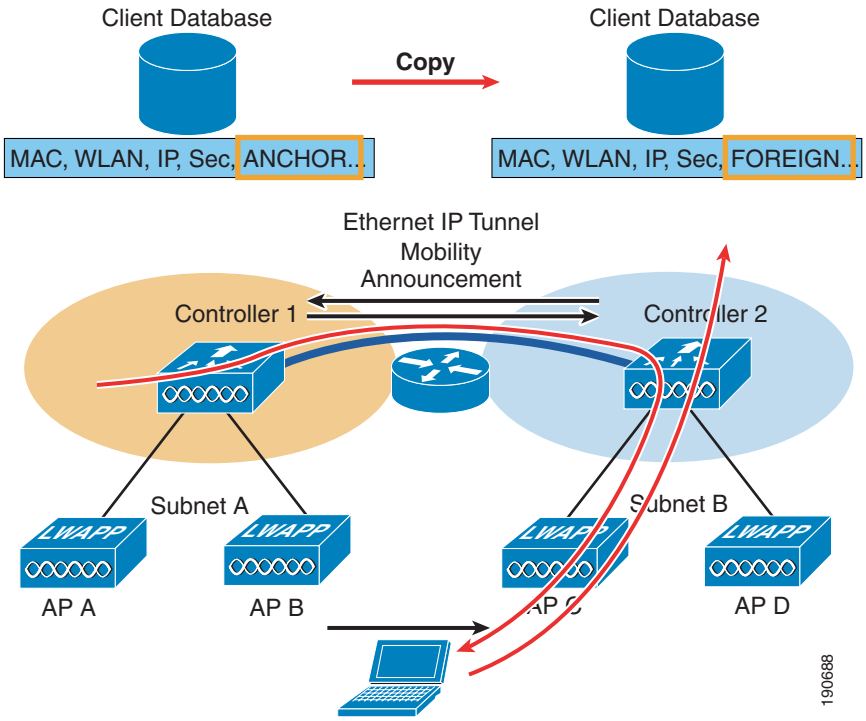

The client begins with a connection to AP B on WLC 1. This creates an ANCHOR entry in the WLC client database. As the client moves away from AP B and makes an association with AP C, WLC 2 sends a mobility announcement to peers in the mobility group looking for the WLC with the client MAC address. WLC 1 responds to the announcement, handshakes, and ACKs. Next the client database entry for the roaming client is copied to WLC 2, and marked as FOREIGN. Included PMK data (master key data from the RADIUS server) is also copied to WLC 2. This provides fast roam times for WPA2/802.11i clients because there is no need to re-authenticate to the RADIUS server.

After a simple key exchange between the client and AP, the client is added to the WLC 2 database and is similar, except that it is marked as FOREIGN.

Points to Remember with Layer 3 Roaming

Layer 3 roaming is a very useful tool, but when deploying with this current software release, remember the following points:

- Traffic is currently asymmetrically routed; that is, roaming client traffic from the anchor WLC are EoIP-tunneled to the foreign WLC, but traffic from the roaming client returns to the network via the foreign WLC. This can be an issue when source address checks or reverse path checks are made within the network or connected systems.
- The EoIP tunnels used to carry roaming traffic between anchor and foreign WLCs are currently DSCP-marked best effort, and not marked with the client traffic DSCP value.

- Multicast group membership is not currently transferred during the client roam; that is, if a client is receiving a multicast stream and roams to a foreign WLC that multicast stream is broken, and must be re-established.
- The basis for Layer 3 roaming is the anchor WLC. The anchor is defined by the subnet of the WLC where a client first associates to the mobility group. This means that Layer 3 roaming assumes a DHCP client where a client gets an appropriate address for the anchor WLC interface, and then roams to a foreign WLC. A client cannot begin its network connection with a static IP address that does not match the subnet of its anchor WLC. In instances where this type of static behavior is required, Mobile IP should be investigated as a solution; for more details concerning Mobile IP and its interaction with the Cisco Unified Wireless architecture, see Chapter 14, "Cisco Unified Wireless and Mobile IP."

Broadcast and Multicast on the WLC

The section discusses the handling of broadcast and multicast traffic by a WLC and its impact on design.

Figure 2-10 shows a schematic of the basic 802.11 broadcast/multicast operation. With a client, such as client 1 in this example, the 802.11 frame is unicast to the AP, and then the AP sends the frame as broadcast out both its wireless and wired interfaces.

Figure 2-10 802.11 Broadcast/Multicast

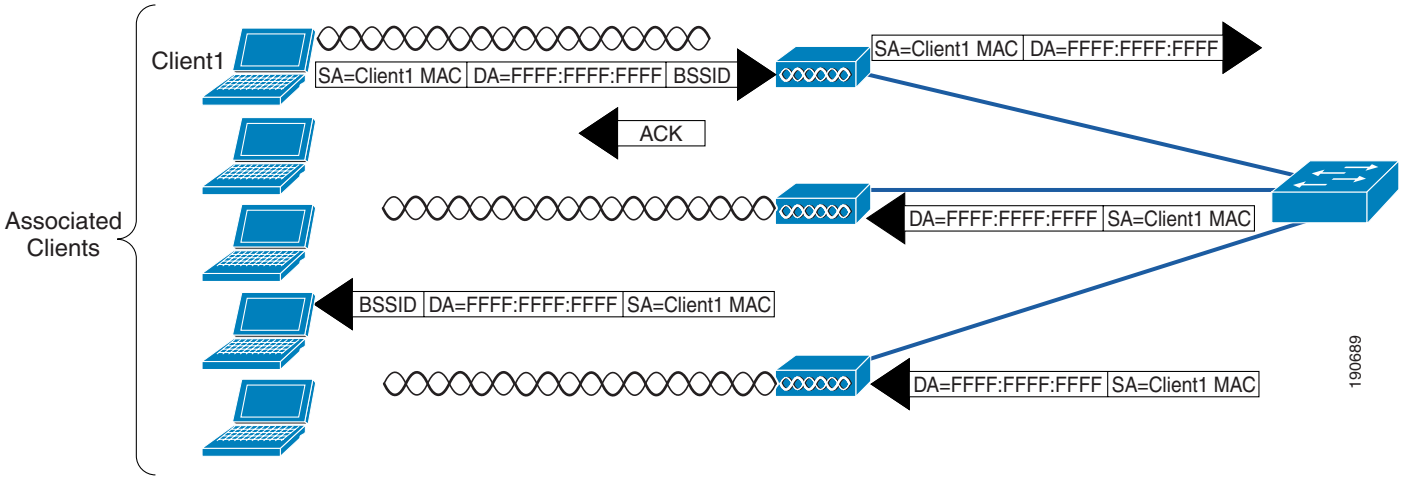

If there are other APs on the same wired VLAN as the AP of Figure 2-10, they forward the wired broadcast packet out their wireless interface.

The WLC split MAC treats broadcast traffic differently, as illustrated in Figure 2-11. In this case, no broadcast traffic is sent back out the WLAN interface, and a limited set of broadcast traffic is sent out the WLAN interface of the WLC.

Note Which protocols are forwarded under which situations is discussed in the following section.

Figure 2-11 Default WLC Broadcast Behavior

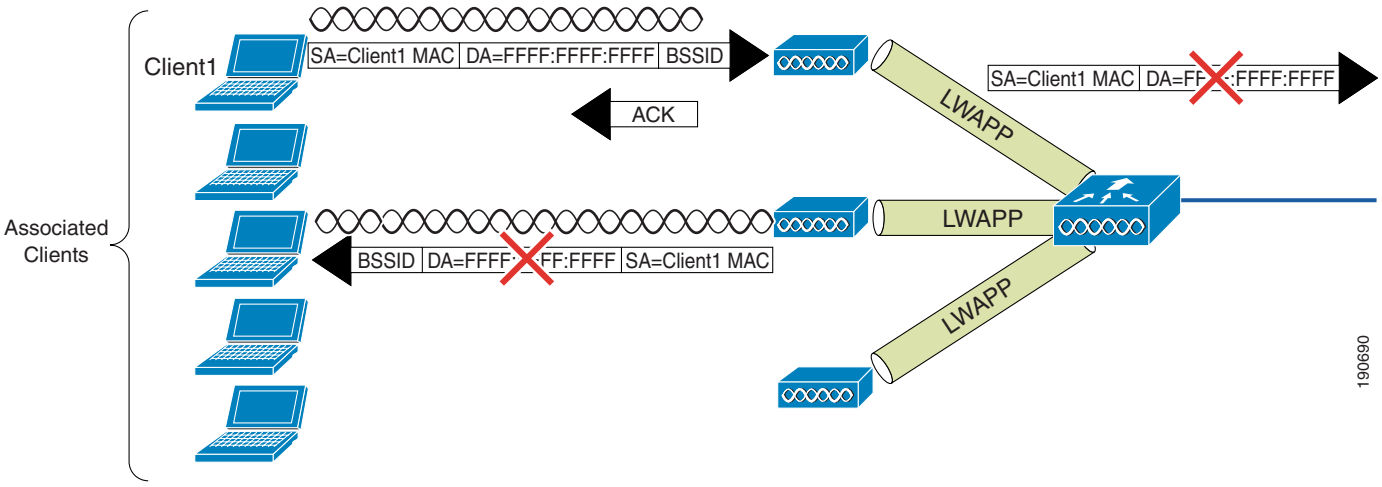

WLC Broadcast and Multicast Details

Broadcast and multicast traffic in WLANs often require special handling in a WLAN network because of the additional load placed on WLANs by broadcasts and multicasts being sent at the lowest available bitrates.

The default behavior of the WLC is not to send any broadcast/multicast traffic out to the WLAN client devices.

The WLC is able to do this without impacting client operation because a typical IP client does not use broadcast/multicast for any other purpose than obtaining network information (DHCP) and resolving a IP address MAC associations (ARP).

DHCP

The WLC acts as a DHCP relay agent for its WLAN clients, unicasting client DHCP requests to the DHCP server configured on the dynamic interface associated with that WLAN, except in roaming as discussed in more detail in this chapter. Because the WLC knows where the DHCP server is, there is no need for it to forward the broadcast DHCP request out its wired or wireless interfaces.

This does a number of things for the WLC and the WLAN:

- It relieves the requirement for the DHCP to broadcast further than the WLC.
- It allows the WLC to be part of the DHCP exchange and to learn the IP address MAC association of its WLAN clients.
- It allows the WLC to send WLAN clients the virtual IP address shared by the WLC mobility group, as the DHCP server answering the DHCP request. This means that a WLC can intercept a DHCP renewal request from a roaming WLAN client, determine if that client has already joined the mobility group, and allow the existing IP address for the client to be renewed even though the client subnet is not native to the WLC.

ARP

Before an IP WLAN client can send an IP packet to any other IP client, it needs to know which MAC addresses to use as the destination MAC address. To do this, the client broadcasts an ARP query, requesting a MAC address to pair with the IP address contained within the ARP request, shown in Figure 2-12.

Figure 2-12 ARP Frame

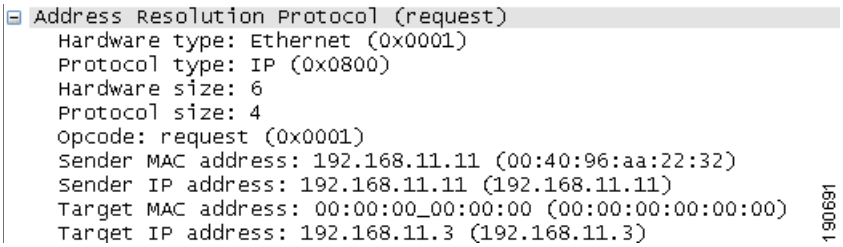

On seeing the ARP request, the WLC either responds directly, acting as an ARP proxy, or forwards the request out the wired interface to have it resolved by another WLC; the WLC does not forward the ARP broadcast back out to the WLAN.

The default behavior of the WLC is to respond to ARP queries directly based on its ARP cache. The **config network arpunicast enable** command can be used to ensure an ARP is sent to the WLAN client, but this ARP request is unicast to the WLAN client, and the primary purpose of this command is to prevent excessive retries by IP clients to a WLAN client that may have roamed from the WLAN network.

Other Broadcast and Multicast Traffic

In its default configuration, no broadcasts and multicasts are forwarded by the WLC. If multicast forwarding is configured as described in Chapter 6, "Cisco Unified Wireless Multicast Design," steps should be taken to minimize the multicast traffic generated at the WLC interface.

The typical steps of limiting multicast addresses groups explicitly supported on the WLAN should be taken, but because enabling multicast allows all multicast traffic including link layer multicasts, multicast is enabled globally on a WLC, and multicast traffic cannot currently be filtered by the WLC, the following steps should be considered:

- Disable CDP on interfaces connecting to WLCs.
- Port filter incoming CDP and HSRP traffic from the WLCs.
- Remember that multicast is enabled on all WLANs on the WLC, including the Guest WLAN, and multicast security including link layer multicast security must be considered.

Design Considerations

In the Cisco Unified Wireless Architecture, the primary considerations are AP connection, and WLC location and connection. This section discusses some of the considerations in these decisions and makes general recommendations where appropriate.

WLC Location

The flexibility of Cisco Unified Wireless LAN solution leads to the following choices about where to locate WLCs:

- Distributed WLC deployment—WLCs are distributed around the campus network, typically on a per building basis, servicing the APs in that building, and connected to the campus network by connecting the WLCs to the distribution routers in that building. In this case, the LWAPP tunnel between the AP and the WLC does not typically leave the building. A schematic of a distributed WLC deployment is shown in Figure 2-13.

Figure 2-13 WLCs Distributed

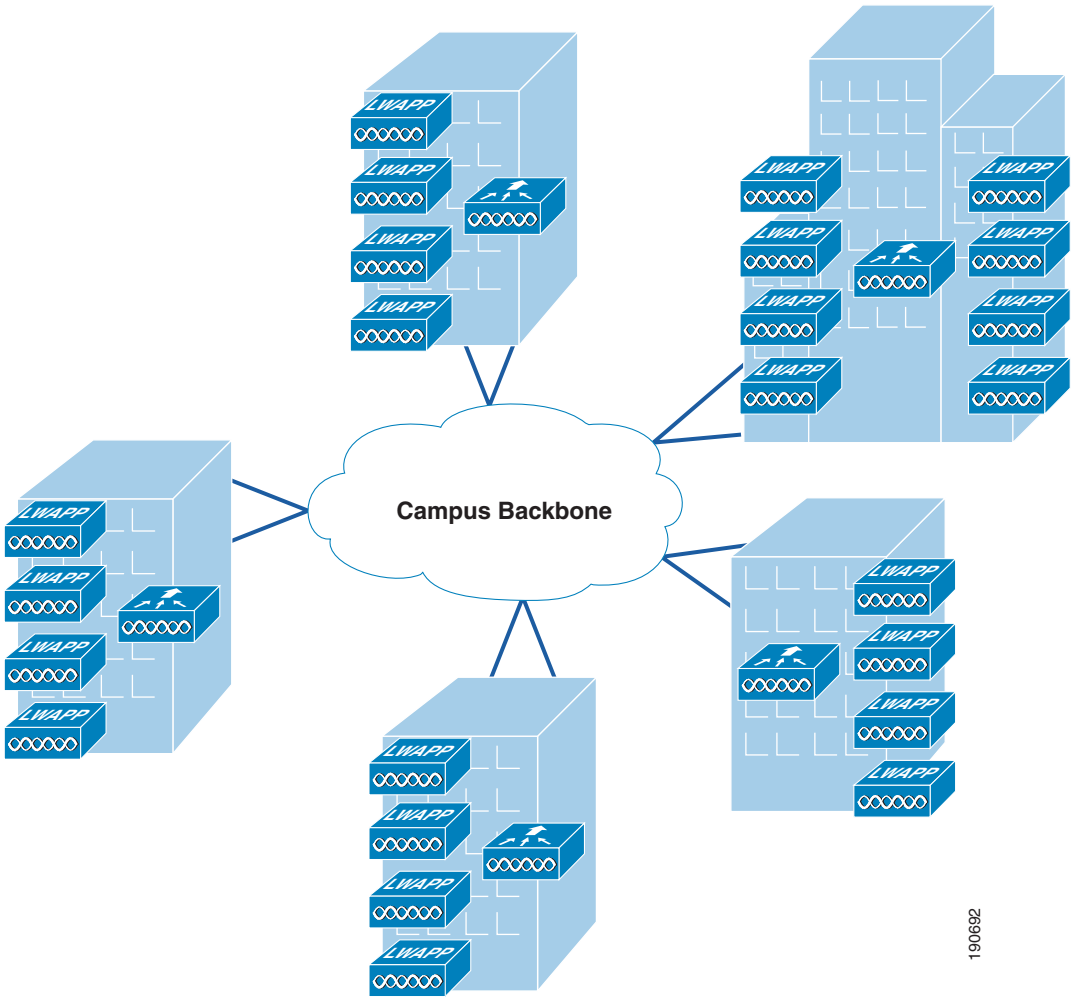

- Centralized WLC deployment—WLCs are placed in a centralized location in the network where most LWAPP tunnels between APs and WLCs must traverse the campus backbone network. A schematic of a centralized WLC deployment is shown in Figure 2-14. Note that the centralized WLC (a pair of WiSMs enabled 6500s in this case) are not shown in a specific building. The centralized WLC cluster would typically be attached to the campus core in the same building as a data center, but not in the data center because the network and security requirements of a data center are generally different to that of WLC cluster.

Figure 2-14 WLCs Centralized

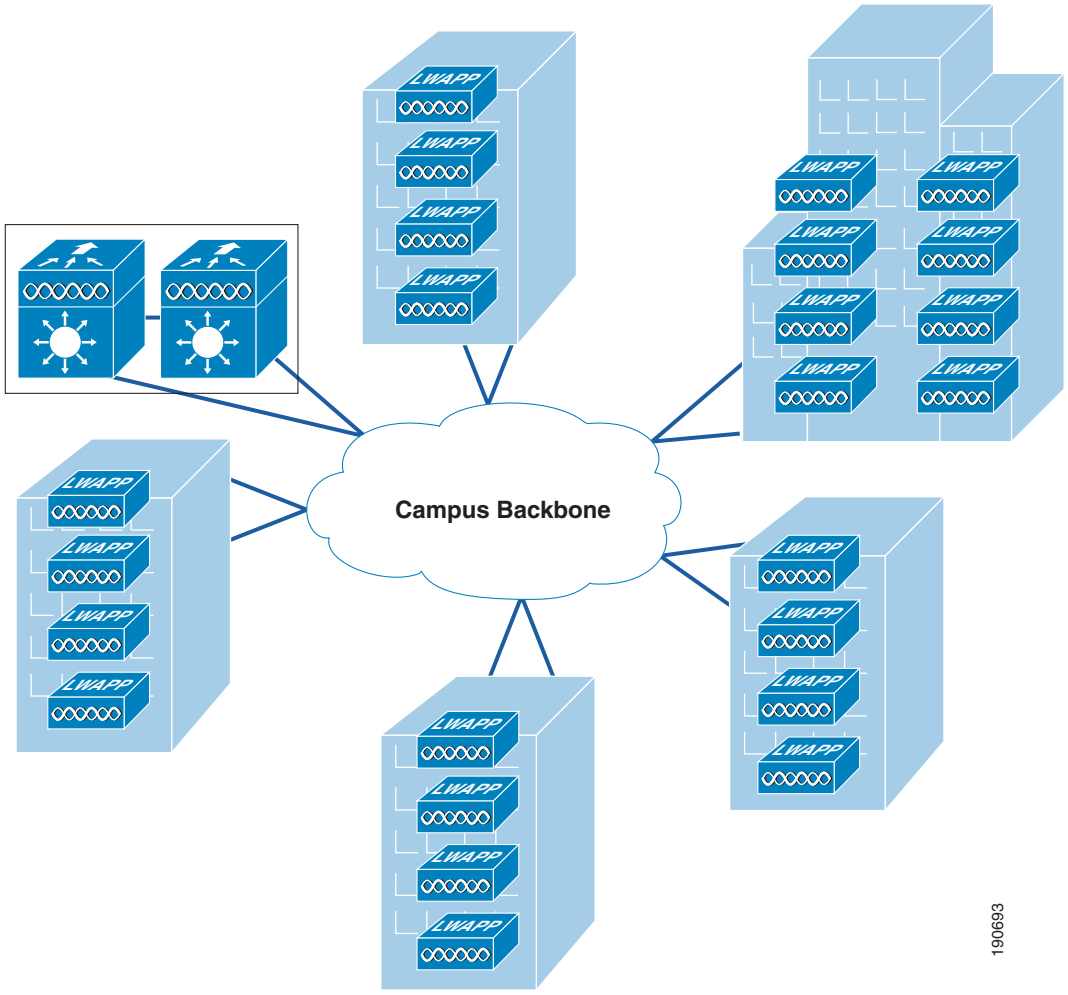

Centralizing WLCs

The general recommendation of this design guide is that WLCs be centralized into a central location in the campus rather than being distributed. The distributed WLC model with mobility groups and Layer 3 roaming is well-proven, and the current gaps in Layer 3 roaming QoS and multicast are expected to be addressed in later software releases. When these are addressed, many of the drivers to centralized are removed.

The best way to address Layer 3 roaming is avoid the issue when possible, the scalability of the WiSM solution, and the broadcast and multicast suppression features of the WLC make the implementation of large mobility subnets practical to implement.

The centralization of the WLC infrastructure makes WLC capacity management simpler and more cost effective, and as WLAN becomes more mission critical, it allows a highly available infrastructure and the capacity to be focused in a small number of locations rather than having to address the same issues in a distributed fashion. The same principle applies with integrating the WLC with other infrastructure; the centralizing of the infrastructure minimizes the number of integration points and integration devices. For example, if the decision is made to implement an inline security component such as a NAC appliance, the centralized WLC would have one integration point, but the distributed solution would have *n* integration points.

The centralization of the WLCs is attractive and is a general recommendation, and the WiSM makes a good choice in this environment. When planning a centralized WLC deployment, consideration should be given to the protection of the network directly connected to the WLC, because the WLC is fundamentally connecting an access network to this network device, and all the security considerations associated with an access layer network device need to be considered. For example, in a WiSM deployment, features such as Denial of Service Protection and Traffic Storm Protection should be considered given the central role of its devices in providing a WLAN service to many users, and the potential for clients with varying levels of security connecting to the switch backplane.

Connecting Distributed WLCs Network

As mentioned earlier in the distributed WLC model, the WLCs are typically at the distribution layer of the campus network. If this is done, Cisco does *not* recommend that the WLC connect to the distribution layer via a Layer 2 connection, as shown in the schematic of Figure 2-15.

Figure 2-15 Layer 2 Connected WLC

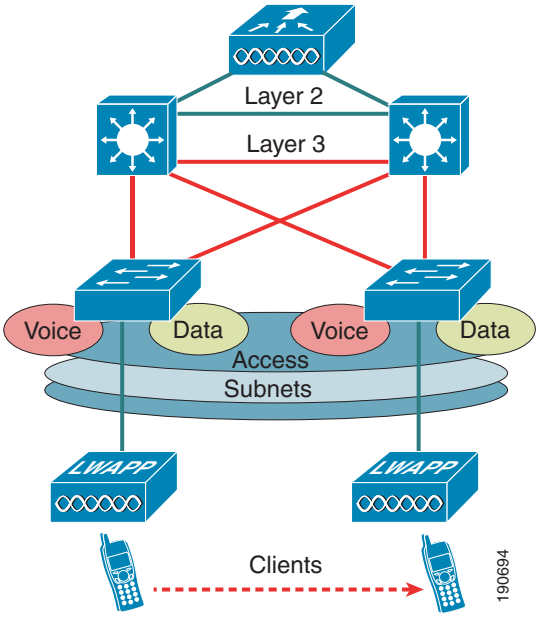

This recommendation is made for a number of reasons, including the following:

- General best practice campus design recommends Layer 3 access and distribute connections to provide fast convergence and simplified operation; inserting a Layer 2 connected WLC breaks this model.
- This requires the introduction of access features at the distribution layer, such as HSRP, and access layer security features. This can be an issue if the distribution does not support all the preferred access switches, or needs to have its software version changed to support access features.
- A Layer 3 connected WLC, as shown in Figure 2-16 (in this case a 3750G), allows the WLAN-related software and configuration to be isolated to a single device, which connects to the network using the same routing configuration as other access layer routing devices; that is, it would typically be configured as a stub router.

Figure 2-16 Layer 3 Connected WLC

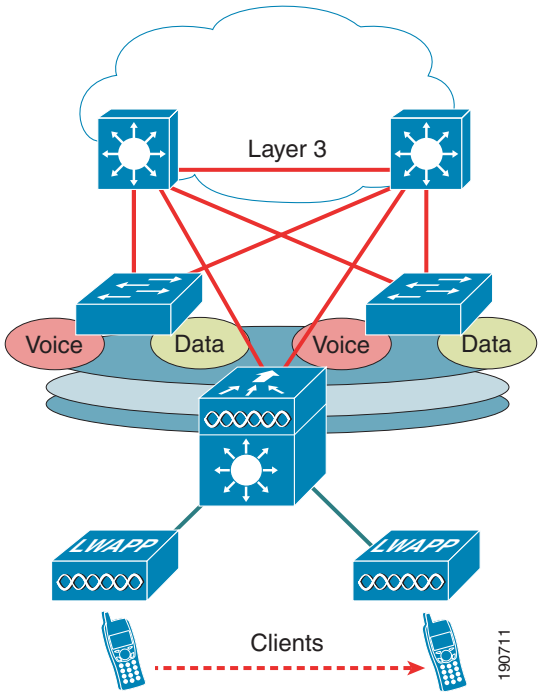

Link Budget and Wired Network Performance

With the use of the Cisco Unified Wireless Architecture where WLAN client traffic is tunneled from the LWAPP AP to the WLC, the question arises concerning the impact upon the backbone wired network, the performance requirements for that network, and the relative benefits of a distributed WLC deployment versus a centralized WLC deployment.

In examining the impact of the LWAPP traffic on traffic volume, there are three main points to consider:

- The volume of LWAPP control traffic—The volume of traffic associated with LWAPP traffic control can vary depending on the actual state of the network; that is, it is higher during a software upgrade or reboot situations. However, traffic studies have found that the average LWAPP WLC traffic load is ~0.35 Kb/sec. In most campuses, this traffic would be considered negligible, and would be neutral when considering a centralized versus distributed WLC deployment.

- The overhead introduced by the tunneling—The Layer 3 LWAPP tunnel adds 44 bytes to a typical IP packet to or from a WLAN client. Given that average packets sizes found on typical enterprises are ~300 bytes, this represents an overhead of approximately 15 percent. In most campuses, this overhead would be considered negligible, and would be neutral when considering a centralized versus distributed WLC deployment.
- Traffic engineering—The tunneling of traffic to a location within the network, and then having it routed to its ultimate destination, rather than having it enter the network at the access layer and then be routed to its ultimate destination, changes traffic flows and volumes within the network. In a distributed WLC model, this impact is minimized because WLC is at the distribution layer and the tunnel is relatively short. In a centralized WLC model, the length of the LWAPP tunnel is longer and the potential to taking traffic off its most efficient path increases. The longer path and the potentially inefficient traffic flows can be mitigated, ensuring that the centralized WLCs are close to the part of the campus the network that has the most client traffic. For example, having the centralized WLC adjacent to the data center would generally be an efficient location because the majority of the client traffic would typically be to and from servers located in the data center. Given that most enterprise client traffic is to and from servers in the data center, and that the enterprise backbone network is of low latency, the overhead associated with inefficient traffic flow would be considered negligible, and would be neutral when considering a centralized versus distributed WLC deployment.

For most enterprises, the introduction of a WLAN does not introduce new applications, at least not immediately. The addition of a Cisco Unified Wireless LAN network is unlikely to have a significant impact on campus backbone traffic volumes.

AP Connection

APs should be on a separate network to the end users. This is in line with the general best practice that infrastructure management interfaces be on a separate subnet from end users. In addition, Cisco recommends that Catalyst Integrated Security Features (CISF) be enabled on the LWAPP AP switch ports (REAP and H-REAP APs, which are discussed in a later chapter) to provide additional protection to the WLAN infrastructure.

DHCP is the generally recommended mechanism for address assignment, because it provides a simple mechanism for providing up-to-date WLC address information and ease of deployment.

A static IP address can be assigned APs and requires more planning and individual configuration. APs with console ports allow the setting of IP address information through the console.

To effectively provide WLAN QoS features in the Cisco Unified Wireless Architecture, QoS should be enabled on the network between the LWAPP APs and the WLCs.

Operation and Maintenance

This section focuses on general deployment considerations and recommendations for easy operation and maintenance of a Cisco Unified Wireless deployment.

WLC Discovery

The multiple WLC discovery mechanisms for APs make the initial deployment of LWAPP APs very simple, with a range of options from staging LWAPP APs with a WLC in a controlled environment to deploying them straight out of the box, and using one of the discovery mechanisms to find a WLC.

Although this flexibility in finding a WLC is very useful, an enterprise deployment generally wants to be able to predict which WLC is used when an AP is first connected to the network, which WLC will be the primary WLC used in the normal operation of an AP, and which WLC will be the secondary and alternate WLC by an AP.

AP Distribution

The WLC discovery process was discussed earlier in this chapter. In a standard initial deployment, the APs automatically distribute themselves across the available WLCs based on the load on each WLC. Although this process makes for an easy deployment, there are a number of operational reasons not to use the auto distribution of APs across WLCs.

APs in the same location should use the same WLC. This makes it easier for general operations and maintenance, allowing staff to know which operations impact which locations, and to be able to quickly associate WLAN issues with specific WLCs, roaming within a WLC, or roaming between WLCs.

The tools that are used to manage AP distribution across WLCs are as follows:

- Primary, secondary, and tertiary WLCs—Each AP can be configured with primary, secondary, and tertiary WLC names that determine the first three WLCs in the mobility group with which the AP will prefer to partner, regardless of the load differences between WLCs in the mobility group.
- Master WLC—When an AP initially partners with a WLC in the mobility group, it has not been configured with a preferred primary, secondary, and tertiary WLC, so it can partner with any WLC based upon the perceived WLC load; or if a WLC is configured as a Master WLC, all APs without primary, secondary, and tertiary WLCs configured will partner with the Master WLC. This allows operations staff to know where to find new APs, and to control when the APs go into production and which WLCs will be the primary, secondary, and tertiary WLCs.

Firmware Changes

One key consideration in the Cisco Unified Wireless operation is how to upgrade WLC firmware with minimal disruption to the WLAN network, because the simple upgrade and reboot of a WLC can result in a general loss of WLAN coverage in some locations while all the APs in that area download new software.

A better option is to move the APs to their secondary WLC, upgrade their primary WLC, and then move the APs to the now upgraded WLC in a controlled manner.

The process can vary slightly, depending on the failover infrastructure, in 1+1 scenario:

- APs are moved off the primary WLC to the secondary
- The primary WLC is upgraded
- All APs are then moved to the primary WLC
- The secondary WLC is upgraded
- Secondary APs are moved back to the secondary AP.

In an N+1 scenario:

- Each WLC moves its APs to the +1 WLCs while the WLC is upgraded.
- APs are moved back to their primary WLC after it is upgraded.
- After all N WLCs are upgraded, the +1 WLC is upgraded.

Note AP Failback should be disabled to ensure that the APs return to their primary WLC in a controlled manner.