CHAPTER 2

Cisco Unified Wireless Technology and Architecture

This chapter discusses the key design and operational considerations associated with 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 underlying protocol used in Cisco's centralized WLAN architecture. It provides for the configuration and management of WLAN(s), in addition to tunneling WLAN client traffic to and from a centralized WLAN controller (WLC). Figure 2-1 shows a high level diagram of a basic centralized WLAN architecture, where LWAPP APs connect to a WLC via LWAPP.

Note Because the foundational WLAN features are the same, the term WLC is used generically to represent all Cisco WLAN Controllers, regardless of whether the controller is a standalone appliance, an ISR with a WLC module; or a Catalyst switch with a service module or integrated WLC.

Figure 2-1 LWAPP APs Connected to a WLC

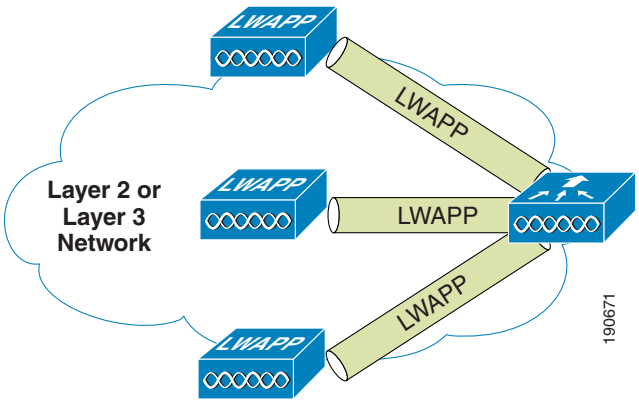

The LWAPP protocol comprises of a number of functional components; however, only those that influence the design and operation of a centralized WLAN network are discussed in this document.

The key features of LWAPP are:

- Split MAC tunnel
- L2 or L3 based tunnels
- WLC discovery process.

Split MAC

A key component of the LWAPP protocol is the concept of split MAC, where part of the 802.11 protocol operation is managed by the LWAPP AP, while the remaining parts are managed by the WLC. A diagram of the split MAC concept is shown in Figure 2-2c.

A generic 802.11 AP, at the simplest level, is nothing more than an 802.11 MAC-layer radio that bridges WLAN clients to a wired network based on association to a Basic Service Set Identifier (BSSID). See Figure 2-2a. The 802.11 standard extends the single AP concept (above) 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 a common network via more than one AP. See Figure 2-2b.

The LWAPP split MAC concept takes all of the functions normally performed by individual APs and distributes them between two functional components: an LWAPP AP and a WLC. The two are linked across a network by the LWAPP protocol and together provide equivalent radio/bridging services in a manner that is simpler to deploy and manage than individual APs.

Note Although 'split MAC' facilitates Layer 2 connectivity between the WLAN clients and the wired interface of the WLC; this does not mean that the LWAPP tunnel will pass all traffic. The WLC forwards only IP Ethertype frames, and its default behavior is to not forward broadcast and multicast traffic. This is important to keep in mind when considering multicast and broadcast requirements in a WLAN deployment.

Figure 2-2 Split MAC Concept

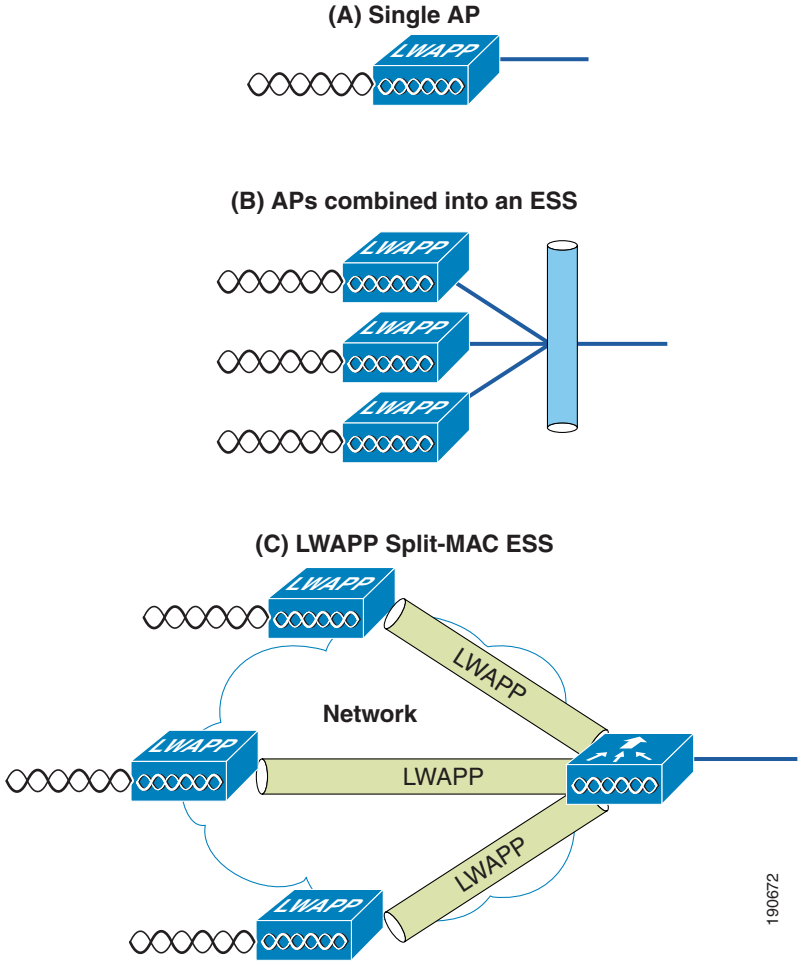

The simple timing-dependent operations are generally managed locally on the LWAPP AP, while more complex, less time-dependent operations are managed on the WLC.

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 in the case of REAP and H-REAP configured APs, which are discussed later in this guide

An LWAPP tunnel supports two categories of traffic:

- LWAPP control messages—Used to convey control, configuration, and management information between the WLC and APs.
- Wireless client data encapsulation—Transports Layer 2 wireless client traffic in IP Ethertype encapsulated packets from the AP to the WLC.

When encapsulated client traffic reaches the WLC, it is mapped to a corresponding VLAN interface/port at the WLC. This interface mapping is defined as part of a WLAN's configuration settings on the WLC. The interface mapping is usually static, but a WLAN client can be dynamically mapped to a specific VLAN based on parameters sent by an upstream AAA server upon successful EAP authentication. In addition to the VLAN assignment, other WLAN configuration parameters include: SSID, operational state; authentication and security method; and QoS.

Layer 2 and Layer 3 Tunnels

LWAPP allows tunneling within Ethernet frames (Layer 2) or within UDP packets (Layer 3). This is configurable on the WLC. Only one method can be supported at a time and not all WLCs support the Layer 2 method.

Layer 2 Tunnel

When deploying Layer 2 LWAPP, the WLC and the LWAPP APs require IP addresses even though the LWAPP tunnel uses Ethertype 0xBBBB to encapsulate traffic between them. All communication between the LWAPP AP and the WLC is encapsulated using Ethertype 0xBBBB.

Although Layer 2 LWAPP is one of the simplest ways to establish AP connectivity and configuration, it is generally not recommended for enterprise deployments, and therefore will not be discussed further in this document.

The primary reasons why the Layer 2 method is not a current Cisco best practice recommendation:

- Layer 2 connectivity between the LWAPP APs and the WLC potentially limits the location of where the APs or WLC can be positioned within the overall network. Extending Layer 2 transport across an enterprise network to get around this limitation is not a current Cisco best practice recommendation.
- Layer 2 LWAPP is not supported on all LWAPP APs and WLC platforms.
- Even though client traffic DSCP values are preserved within the tunnel, Layer 2 LWAPP does not provide corresponding CoS marking for the Ethertype frames, and therefore is not able to provide transparent, end-to-end QoS for the tunneled traffic.

Layer 3 Tunnel

Layer 3 LWAPP is the recommended tunnel type. This method uses IP UDP packets to facilitate communication between the LWAPP AP, and the WLC. L3 LWAPP is able to perform fragmentation and reassembly of tunnel packets; thereby allowing client traffic to make use of a full 1500 byte MTU and not have to adjust for any tunnel overhead.

Note In order to optimize the fragmentation and reassembly process, the number of fragments that the WLC or AP expect to receive is limited. The ideal supported MTU size for deploying the Cisco Unified Wireless network 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. The sample decodes were captured using a Wireshark Network Analyzer.

Note The Wireshark's default configuration does not decode Cisco LWAPP packets correctly. This can be corrected by using the "SWAP Frame Control" option under protocol preferences.

Figure 2-3 shows a decode of an LWAPP control packet. This packet originates from the WLC using UDP source port 12223 (as do all LWAPP control packets from the WLC). Control Type 12 represents a configuration command used to pass AP configuration information to the LWAPP AP by the WLC. Control packet payloads are AES encrypted, using keys derived from the PKI authentication process that is performed when an LWAPP AP first establishes a connection with the WLC.

Figure 2-3 LWAPP Control Packet

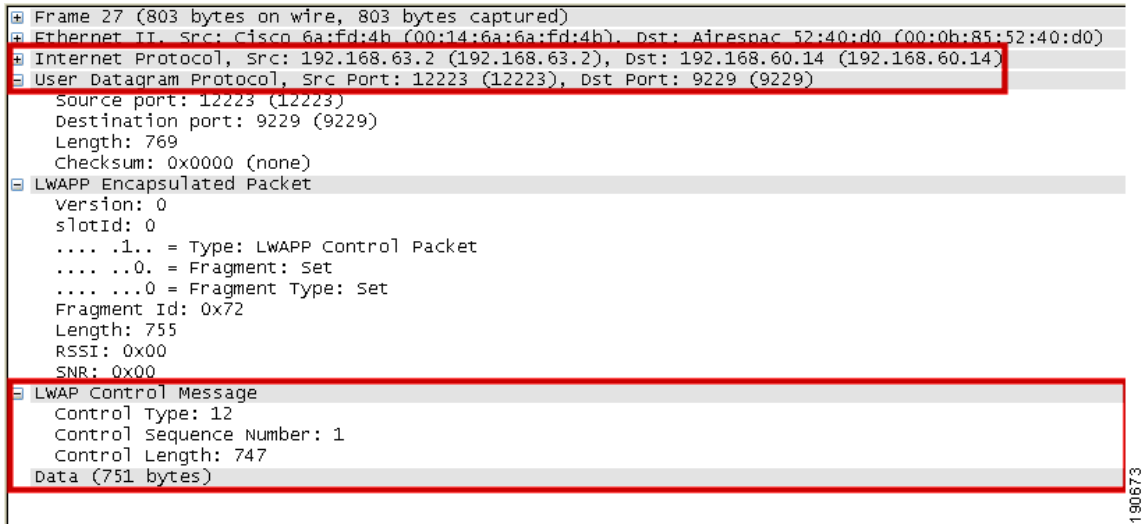

802.11 Probe Request in LWAPP shows a decode of an LWAPP packet containing an 802.11 probe request. This packet originates from the LWAPP AP to the WLC using UDP port 12222, as do all LWAPP-encapsulated 802.11 frames. In this example, RSSI and SNR values are also included in the LWAPP packet to provide RF information to the WLC.

Figure 2-4 802.11 Probe Request in LWAPP

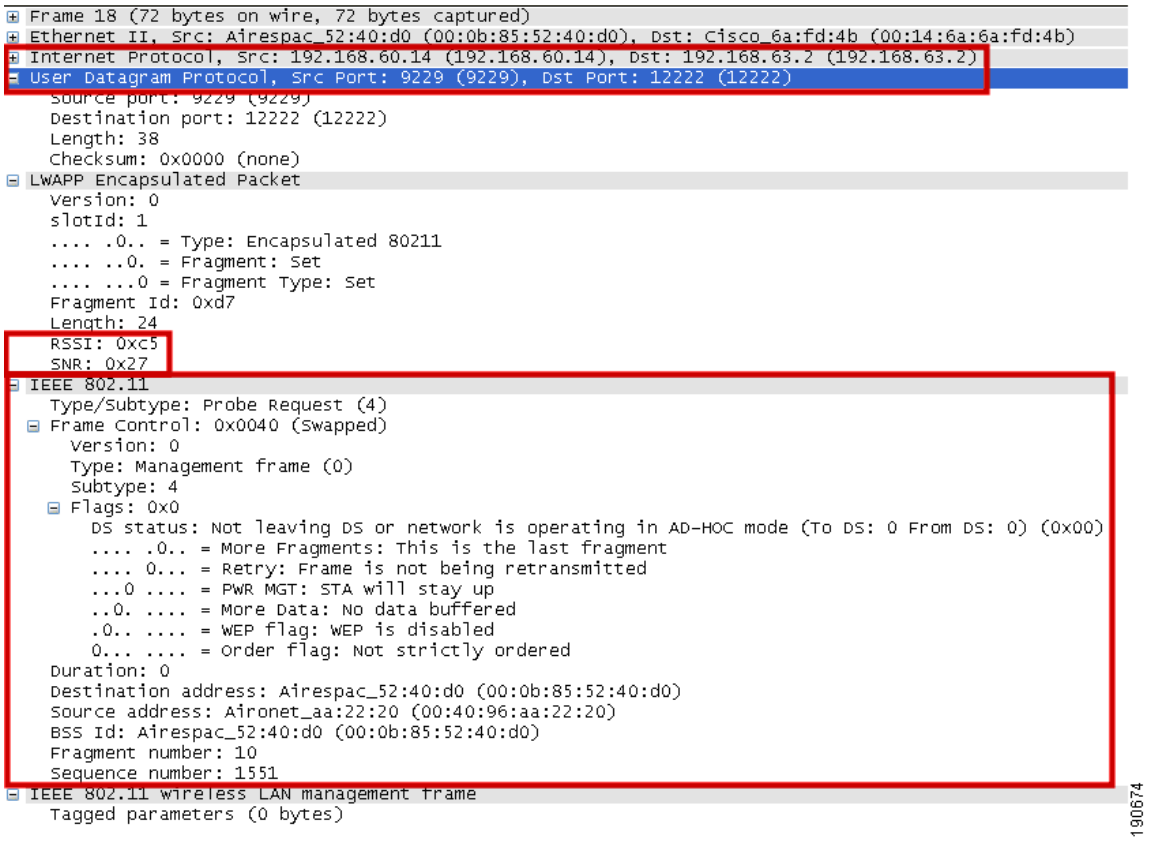

Figure 2-5 shows another LWAPP-encapsulated 802.11 frame, but in this case it is an 802.11 data frame, like that shown in Figure 2-4. It contains a complete 802.11 frame, as well as RSSI and SNR information for the WLC. This capture is being shown to illustrate that an 802.11 data frame is treated the same by LWAPP as the other 802.11 frames. Figure 2-5 highlights that fragmentation is supported, in order for LWAPP packets to accommodate the minimum MTU size between the LWAPP AP and the WLC. Note in the Wireshark decode that the frame control decode bytes have been swapped; this is accomplished during Wireshark's protocol analysis of the LWAPP packet to take into account that some LWAPP APs swap these bytes.

Figure 2-5 802.11 Data Frame in LWAPP

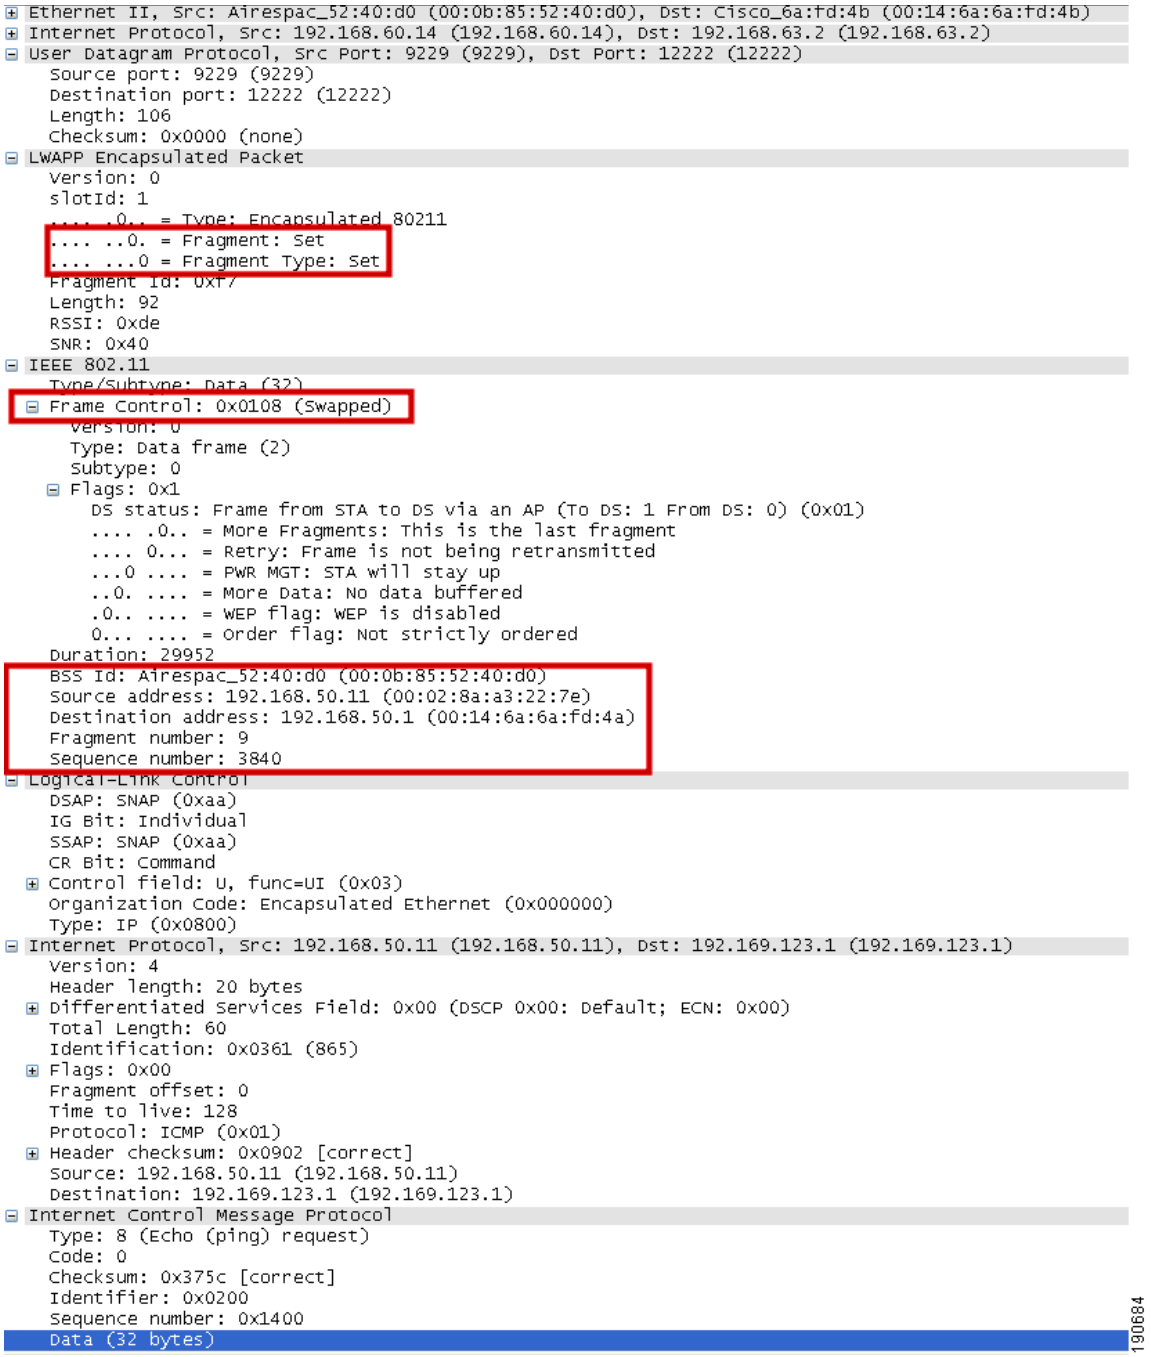

WLC Discovery and Selection

The following section highlights the typical behavior of a Layer 3 LWAPP AP upon being reset. For a comprehensive description of the discovery/join process, 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.

Upon reset, the following sequence takes place:

Step 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 same IP subnet will see the discovery message. Each of the WLCs receiving the LWAPP discovery message will in turn reply with a unicast LWAPP discovery response message to the AP.

Step 2 If a feature called 'Over-the-Air Provisioning' (OTAP) is enabled on a WLC, APs that are already joined to that WLC will advertise their known WLCs in neighbor messages sent to other APs 'over the air'. Any new AP attempting to 'discover' WLCs for the first time will receive these messages and in turn unicast an LWAPP discovery request to each WLC advertised in the OTAP message. (OTAP is not supported by IOS APs in their initial state. In other words, a new IOS-based AP cannot use OTAP to discover a WLC.) WLCs that receive LWAPP discovery request messages unicast an LWAPP discovery response to the AP.

Step 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 an LWAPP discovery request responds by sending an LWAPP discovery response to the AP. As stated earlier, WLC IP addresses can be learned via OTAP messages sent from existing APs already joined to WLCs. The information stored in NVRAM also includes address information for any previously joined WLC that was a member of another mobility group. (The mobility group concept is discussed in greater detail later in this document.)

Step 4 If OTAP is not used, DHCP servers can be programmed to return WLC IP addresses using vendor specific DHCP options. In this case "Option 43" is used in a "DHCP offer" to "advertise" WLC addresses to LWAPP APs. When an AP receives its IP address via DHCP, it checks for WLC IP address information in the Option 43 field of 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.

Step 5 In lieu of Option 43 information, the AP attempts to resolve the following DNS name: "CISCO-LWAPP-CONTROLLER.localdomain". If the AP is able to resolve this, it sends a unicast LWAPP discovery message to each IP address returned in the DNS reply. As described above, each WLC that receives an LWAPP discovery request message replies with a unicast LWAPP discovery response to the AP.

Step 6 If after Steps 1 through 5 no LWAPP discovery response is received, the AP resets and restarts the search algorithm.

Typically, either the DHCP or DNS discovery mechanism is used to provide one or more seed WLC addresses, and then a subsequent WLC discovery response provides a full list of WLC mobility group members.

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

Components

There are three primary components that make up Cisco's Unified Wireless Architecture: the Lightweight APs, the WLC, and the WCS. This section describes the AP and WLC product options.

The Cisco WCS is an optional network component that works in conjunction with Cisco Aironet Lightweight APs, Cisco wireless LAN controllers and the Cisco Wireless Location Appliance. With Cisco WCS, network administrators have a single solution for RF prediction, policy provisioning, network optimization, troubleshooting, user tracking, security monitoring, and wireless LAN systems management. Robust graphical interfaces make wireless LAN deployment and operations simple and cost-effective. Detailed trending and analysis reports make Cisco WCS vital to ongoing network operations. More information on Cisco WCS can be found at the following URls:

http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd802570d0.html

http://www.cisco.com/en/US/products/ps6305/products_installation_and_configuration_guides_list.html

WLCs

For convenience, this document refers to all Cisco Unified Wireless controllers as WLCs due to the general uniformity and commonality of features across all of Cisco's WLC platforms.

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

- 2106—Is a standalone WLC that supports up to six APs, with eight Fast Ethernet interfaces. Two of the Fast Ethernet interfaces can be used to power (802.3af) directly connected APs. The interface can be configured as dot1q trunks to provide connection into the wired network. The 2106 is ideal for a small-to-medium size offices, where an H-REAP would otherwise be unsuitable because of the number of users, WAN requirements, and/or client roaming requirements.
- 4402—Is a standalone WLC that supports either 12, 25, or 50 APs. It comes with two SFP-based Gigabit Ethernet ports that can be configured as dot1q trunks to provide connection into the wired network, or the Gigabit ports can be link-aggregated to provide an EtherChannel connection to the switched network. This is ideal for medium-size offices or buildings.
- 4404—Is a standalone WLC that supports 100 APs. It comes with four SFP-based Gigabit Ethernet ports that can be configured as dot1q trunks to provide connection into the wired network. The Gigabit ports can be link aggregated to provide an EtherChannel connection to the switched network. This is ideal for large offices, buildings, and even a small campus.
- WLCM—The WLC module is specifically designed for Cisco's Integrated Service Router (ISR) series. It's currently available in a 6, 8 or 12 AP version. The WLCM appears as an interface on the ISR router that can be configured as a dot1q trunk to provide a routed connectivity to the wired network. This is ideal for small-to-medium size offices requiring an integrated solution.
- WS-C3750G—Is a WLC that supports either 25 or 50 APs that comes integrated with the Catalyst 3750 switch. The WLC's backplane connections appear as two Gig Ethernet ports, that can be configured separately as dot1q trunks to provide connection into the 3750. Or, the Gig ports can be link aggregated to provide a single EtherChannel connection to the 3750. Because the WLC is integrated directly it has access to all of the advanced routing and switching features available in the 3750 stackable switch. It is ideal for medium-size offices or buildings. The '50 AP' version can scale up to 200 APs when four 3750s are stacked together as a virtual switch.

- WiSM—Is a WLC module that is designed specifically for Cisco's Catalyst 6500 switch series. It supports up to 300 APs per module. Depending on the 6500 platform, multiple WISMs can be installed to offer significant scaling capabilities. The WiSM appears as a single aggregated link interface on the 6500 that can be configured as a dot1 trunk to provide connection into the 6500 backplane. This is 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
2106	6	8x Fast Ethernet	Cannot be a Mobility Anchor, does not support Layer 2 LWAPP, 2 of the Fast Ethernet interfaces support 802.3af for PoE.
4402	12 or 25	2x Gig Ethernet	
4404	50 or 100	4x Gig Ethernet	
WLCM	6, 8 or 12	ISR backplane	Cannot be a Mobility Anchor, does not support Layer 2 LWAPP. Layer 3 sub-interface termination of static and dynamic WLC interfaces only, no support for dot1q trunking.
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: standalone and LWAPP. This section briefly discusses the various models of AP products available within each category, and contrasts features, functionality, and applications. Cisco's 1500 series MESH APs are mentioned briefly below; however, this design guide does not address wireless MESH applications or deployment guidelines. Refer to the following guides for further information about the Cisco MESH solution:

- Cisco Mesh Networking Solution Deployment Guide: http://www.cisco.com/en/US/docs/wireless/access_point/mesh/4.0/deployment/guide/overview.html
- Cisco Aironet 1500 Series Wireless Mesh AP Version 5.0 Design Guide: http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP.html

Cisco Standalone 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 LWAPP mode of operation. This feature permits an enterprise to standardize on a common AP platform that can be deployed in mixed wireless topologies.

First generation standalone APs are as follows:

- AP 1100—This single band 802.11b/g AP. It has 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 connections 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 have all of the features (most notably 802.3af PoE standard) 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 standalone 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 makes use of an integrated dual-band antenna. The 1130AG is available in a lightweight (LWAPP) version for implementation in centralized (WLC)-based deployments. The standalone version can be 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 standalone category and is also available in a lightweight (LWAPP) version. For greatest flexibility, the standalone 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.

A new third generation AP, the Cisco 1252, is a business class AP that supports draft 2 of the emerging 802.11n standard. 802.11n offers combined data rates up to 600Mbps using Multiple-Input Multiple-Output (MIMO) technology. The Cisco 1252 is available in a dual-band a/b/g or a single-band b/g radio configuration and can be deployed as a stand alone AP (standalone) or as part of a unified (controller) wireless deployment. In order to offer maximum deployment flexibility, the Cisco 1252 is equipped with RP-TNC connectors for use with a variety of external 2.4 and 5Ghz antennas. In order to support the greater throughput rates offered by 802.11n, the Cisco 1252 incorporates a gigabit 10/100/1000 interface. The Cisco 1252 is designed to be deployed in challenging RF environments where high bandwidths are needed. Part numbers for the standalone version include: AIR-AP1252AG-x-K9 (Dual Band) and AIR-AP1252G-x-K9 (single band). Part numbers for the Cisco Unified Wireless versions include: AIR-LAP1252AG-x-K9 (dual band) and AIR-LAP1252G-x-K9 (single band).

Cisco LWAPP APs

APs in this category consist of the original Airespace product line, but also include the standalone AP models noted above. The following 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 internal dual 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 number is 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/g	802.11a	802.11n	Standalone	LWAPP	# Broadcasted SSIDs	Preinstalled Cert?
1000	YES	YES	NO	NO	YES	16	YES
1100	YES	NO	NO	YES	YES	8[1]	NO
1130AG	YES	YES	NO	YES	YES	8[1]	YES[2]
1200	YES	Optional	NO	YES	YES	8[1]	YES[2]
1230AG	YES	YES	NO	YES	YES	8[1]	YES[2]
1240AG	YES	YES	NO	YES	YES	8[1]	YES[2]
1252AG	YES	YES	YES	YES	YES	8[1]	YES
1252G	YES	NO	YES	YES	YES	8[1]	YES
1300	YES	NO	NO	YES	YES	8[1]	NO
1500	YES	YES	NO	NO	YES	16	YES

1. 16 BSSIDs to be supported in future Releases.
2. Units shipped prior to August 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[1]	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[2]	Recommended	Recommended[2]
1230AG	Recommended	Recommended	Recommended[2]
1240AG	Recommended[2]	Ideal	Recommended[2]
1300	Not Recommended	Not Recommended	Ideal[3]
1500	Not Recommended	Not recommended	Ideal[1]

1. Or 1030 for Remote offices. LWAPP Deployments Only.

2. Can be used outdoors when deployed in weatherproof NEMA rated enclosure. Particularly for deployments above suspended ceilings.
3. Standalone Deployments Only.

For further detailed information, see the following link:
http://www.cisco.com/en/US/products/ps6108/prod_brochure0900aecd8035a015.html

Mobility Groups, AP Groups, and RF Groups

Within the Cisco Unified Wireless Architecture, there are three important 'group' concepts:

- Mobility groups
- AP groups
- RF groups

This section describes the purpose and application of these groups within the Cisco Unified Wireless Architecture. For more details on operation and configuration, 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 4.1—
http://www.cisco.com/en/US/docs/wireless/controller/4.1/configuration/guide/ccfig41.html

Mobility Groups

A mobility group is a group of WLCs that together, act as a single virtual WLC by sharing essential end client, AP, and RF information. A given WLC within a mobility domain, is able to make decisions based on data received from other members of the entire mobility group, rather than relying solely on the information learned from its own directly connected APs and clients.

A mobility group forms a mesh of authenticated tunnels between member WLCs, thereby allowing any WLC to directly contact another WLCs within the group, as shown in Figure 2-6.

Figure 2-6 WLC Mobility Group

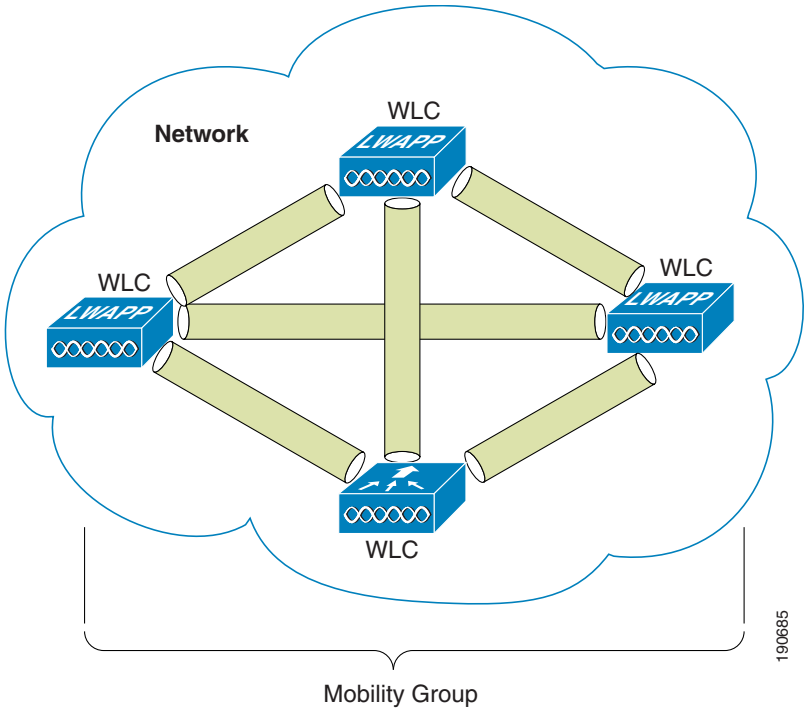

Mobility Group Definition

Creating a mobility group is simple and well documented; however, there are some important considerations to keep in mind:

- Up to 24 WLAN controllers and 3600 APs are supported per mobility group. An enterprise may consist of more WLAN controllers and APs, but they must be configured as members of another mobility group.
- The WLCs do not have to be of the same model/type to be a member of a mobility group. For example, a group may comprise of any combination of the following: 4402, 4404, WiSM, WLCM, 3750G, and 2106; however, they should all be running the same software version. With that said, a mobility group will not be broken simply because of software differences, but a common software version is strongly recommend in order to ensure feature and functional parity across a unified wireless deployment.
- A mobility group requires all WLCs in the group to use the same virtual IP address.
- Each WLC must use the same ‘mobility domain name’ and be defined as a peer in each others ‘Static Mobility Members’ list.
- In order for a wireless client to seamlessly roam between mobility group members (WLCs), a given WLAN's SSID and security configuration must be configured identically across all WLCs comprising the mobility group.

Mobility Group Application

Mobility groups are used to help facilitate seamless client roaming between APs that are joined to different WLCs. The primary purpose of a mobility group is to create a virtual WLAN domain (across multiple WLCs) in order to provide a comprehensive view of a wireless coverage area. The use of mobility groups are beneficial only when a deployment comprises of 'overlapping' coverage established by two or more APs that are connected to different WLCs. A mobility group is of no benefit when two APs, associated with different controllers, are in different physical locations with no overlapping (contiguous) coverage between them (for example, Campus and Branch or between two or more buildings within a campus).

Mobility Group—Exceptions

The Cisco Unified Wireless solution offers network administrators the ability to define static mobility tunnel (Auto Anchor) relationships between an ‘anchor’ WLC and other WLCs in the network. This option, among other things, is used when deploying wireless guest access services.

If the auto anchor feature is used, no more than 24 (foreign) WLCs can be mapped to a designated anchor WLC. Foreign WLCs do not, by virtue of being connected to the auto anchor, establish mobility relationships between each other. The anchor WLC must have a ‘static mobility group member’ entry defined for each foreign WLC where a static mobility tunnel is needed. Likewise for each foreign WLC where a static mobility tunnel is being configured, the anchor WLC must be defined as a ‘static mobility group member’ in the foreign WLC.

A WLC can only be member of one mobility group for the purpose of supporting dynamic inter-controller client roaming. A WLC that is configured as an ‘auto anchor’, does not have to be in the same mobility group as the foreign WLCs. It is possible for a WLC to be a member of one mobility group whilst at the same time, act as an auto anchor for a WLAN originating from foreign WLCs that are members of other mobility groups.

For a discussion on mobility anchor configuration, see Chapter 10, “Cisco Unified Wireless Guest Access Services.”

AP Groups

In typical deployment scenarios, each WLAN is mapped to a single dynamic interface per WLC. However, consider a deployment scenario where there is a 4404-100 WLC supporting the maximum number of APs (100). Now consider a scenario where 25 users are associated to each AP. That would result in 2500 users sharing a single VLAN. Some customer designs may require substantially smaller subnet sizes. One way to deal with this is to break up the WLAN into multiple segments. The WLC’s AP grouping feature allows a single WLAN to be supported across multiple dynamic interfaces (VLANs) on the controller. This is done by taking a group of APs and mapping them to a specific dynamic interface. APs can be grouped logically by employee workgroup or physically by location. Figure 2-7 illustrates the use of AP groups based on site-specific VLANs.

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

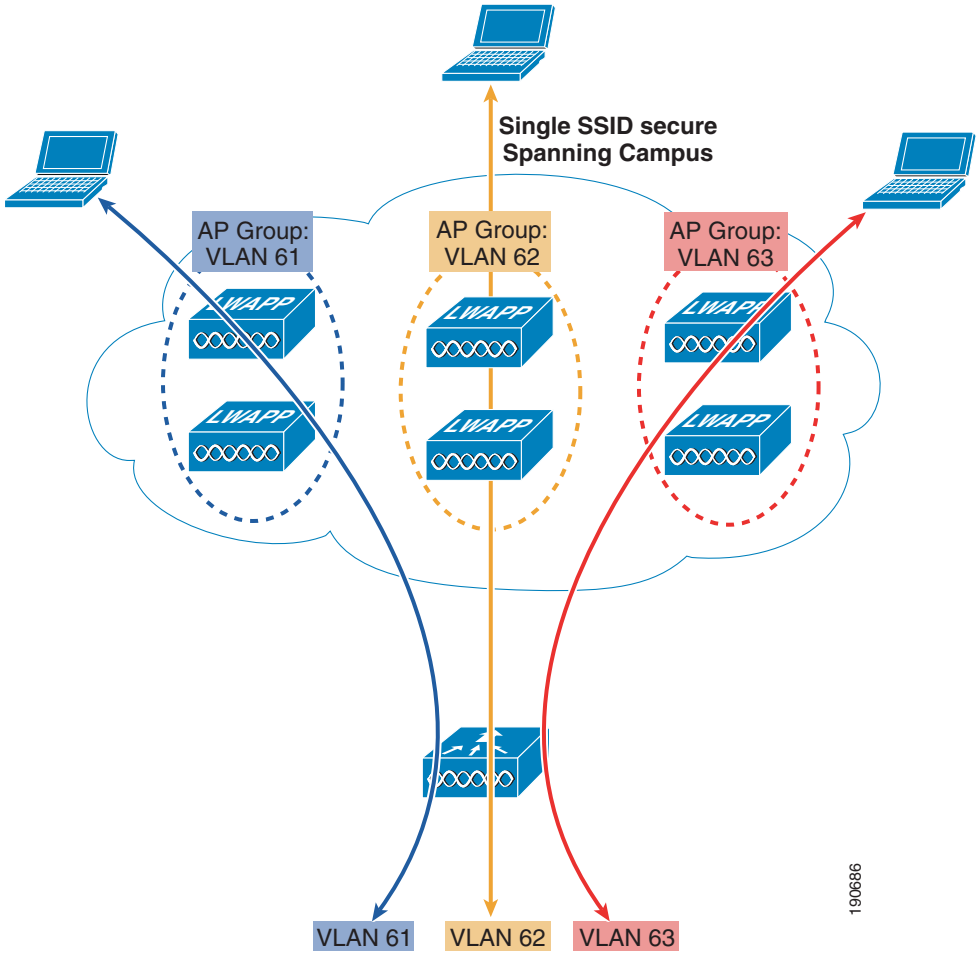

In Figure 2-7, there are three dynamic interfaces configured, each mapping to a site-specific VLAN: VLAN 61, 62, and 63. Each site specific VLAN and associated APs are mapped to the same WLAN SSID using the AP grouping feature. A corporate user associating to the WLAN on an AP in the AP Group corresponding to VLAN 61 gets an IP address on the VLAN 61 IP subnet. Likewise, a corporate user associating to the WLAN on an AP in the AP Group corresponding to VLAN 62 gets an IP address on the VLAN 62 IP subnet and so on. Roaming between the site-specific VLANs is handled internally by the WLC as a Layer 3 roaming event and as such, the wireless LAN client maintains its original IP address.

RF Groups

RF groups, also known as RF domains, represent another important deployment consideration. An RF group is a cluster of WLCs that collectively coordinate and calculate their dynamic radio resource management (RRM) settings based on 802.11 PHY type (for example, 802.11b/g and 802.11a).

An RF group exists for each 802.11 PHY type. Grouping WLCs into RF domains allows the solution's dynamic RRM algorithms to scale beyond a single WLC, thereby allowing RRM for a given RF domain to extend between floors, buildings, and even across campuses. RF Groups and 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 includes the WLC's IP address and a hashed message integrity check (MIC) derived from a timestamp and the BSSID of the AP.
- The hashing algorithm uses a shared secret (the RF Group Name) that is configured on the WLC and is pushed out to each AP. APs sharing the same secret are able to validate messages from each other using the MIC. When APs belonging to other WLCs hear validated neighbor messages at a signal strength of -80 dBm or stronger, their WLCs dynamically become members of the RF group.
- Members of an RF group 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 a master power and channel plan.
- The RRM algorithms:
  - Try to achieve a uniform (optimal) signal strength of -65 dBm across all APs
  - Attempt to avoid 802.11 co-channel interference and contention
  - Attempt to avoid 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 ever changing RF environment.
- The RF group leader and members exchange RRM messages at a specified update interval, which is 600 seconds by default. Between update intervals, the RF group leader sends keep alive messages to each of the RF group members and collects real-time RF data. Note that the maximum number of controllers per RF group is 20.

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 in the same ESS. Depending on network features and configuration, several events can occur between the client, WLCs, and upstream hops in the network, but at the most basic level, roaming is simply a change in AP association.

When a wireless client authenticates and associates with an AP, the corresponding WLC (to which the AP is connected) creates 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 information about the new AP. If necessary, new security context and associations are established as well.

A Layer 2 roam occurs when a client leaves one AP and re-associates with a new AP, in the same client subnet. In most cases, the 'roamed to' AP is connected to the same WLC as the original AP.

The description above represents the simplest roaming scenario because a single WLC database maintains all information about the client. Network elements upstream from the WLC are unaffected by the client moving from one AP to another as illustrated in Figure 2-8.

Figure 2-8 Layer 2 Roam

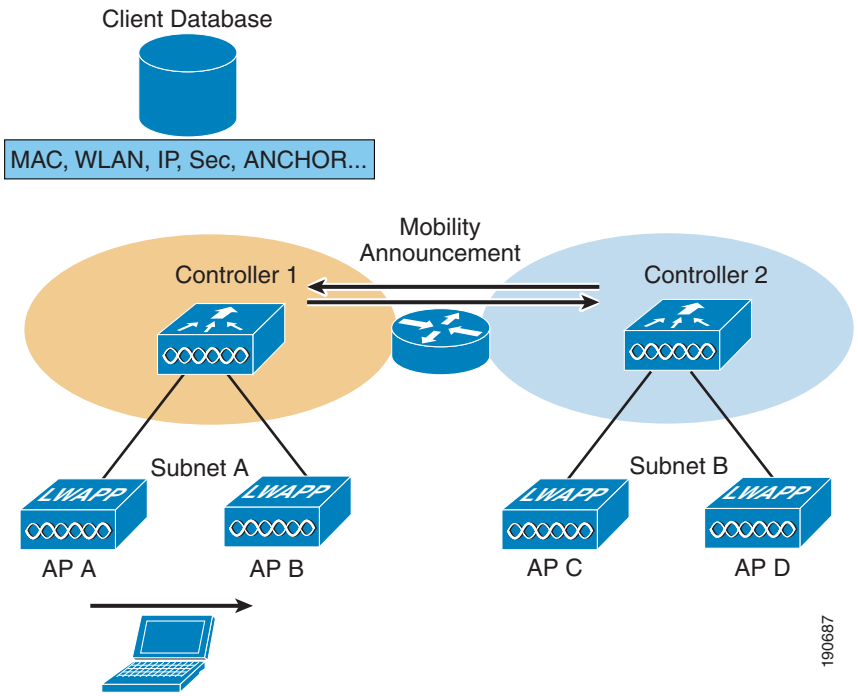

When there are multiple WLCs connecting a WLAN to the same subnet and a client roams between APs connected to different WLCs, a mobility announcement is exchanged between the WLCs. The mobility announcement passes client-context information between WLCs.

WLC to WLC Roaming Across Client Subnets

In cases where a client roams between APs that are connected to different WLCs and the client subnet/VLAN is not the same between the WLCs, then a Layer 3 roam is performed. A mobility announcement is exchanged between the 'roamed to' (foreign) WLC's mobility database and the home (anchor) WLC's mobility database.

A Layer 3 roam is more complex because the wireless client is moving from one VLAN/subnet to another. Unless the WLAN system takes action to make the client subnet change transparent, the Layer 3 roam event has an adverse impact on client communication with upstream services. Existing client sessions will either hang or eventually timeout and disconnect. The Cisco Unified Wireless solution uses mobility tunnels to facilitate Layer 3 roaming that is transparent to the upstream network There are two types of mobility tunnels:

- Asymmetrical (default behavior – WLC Releases 4.0 and earlier)
- Symmetrical (new option beginning with WLC Releases 4.1 and later)

Note In WLC Release 4.1, asymmetrical tunneling is still the default behavior. Administrators must explicitly configure symmetrical tunnel behavior.

Layer 3 Roam—Asymmetrical Mobility Tunnel

In a Layer 3 roaming scenario, traffic returning to the wireless client goes through the anchor WLC. The anchor WLC establishes an Ethernet-over-IP (EoIP) tunnel to forward client traffic to the foreign WLC where it is then delivered to the client. All traffic originated by the client is forwarded out the corresponding VLAN interface to which the WLAN is mapped to at foreign WLC. The client's original IP address and default gateway IP (MAC) address remain the same. All traffic, other than that which is destined for the local subnet, is forwarded to the default router where the foreign WLC substitutes the client's default gateway MAC address with the MAC address of the default gateway associated with dynamic interface/VLAN at the foreign controller.

Figure 2-9 illustrates a client Layer 3 roam using an asymmetrical mobility tunnel.

Figure 2-9 Layer 3 Roaming

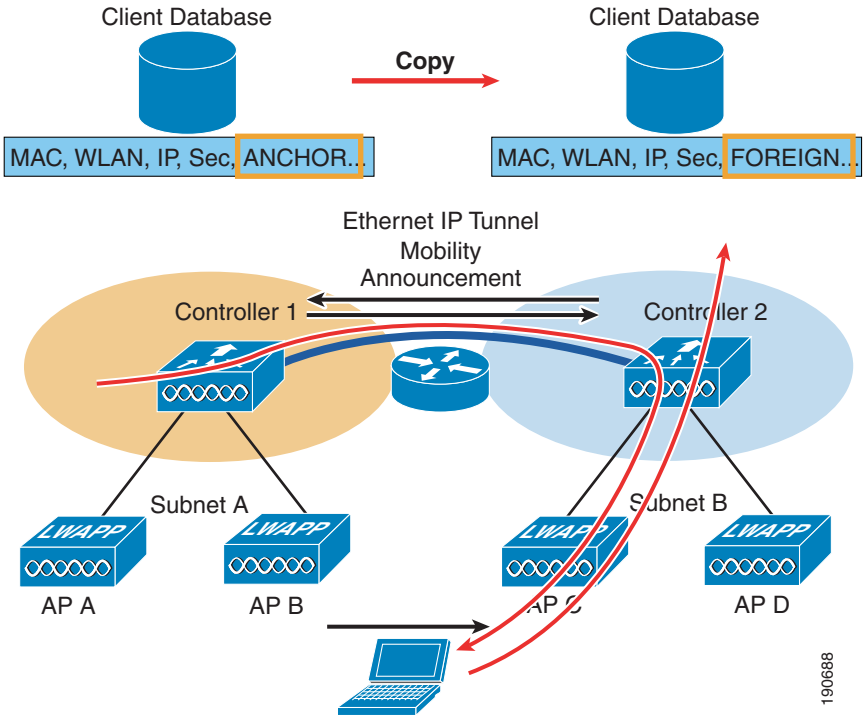

Using Figure 2-9, the following occurs when a client roams across a Layer 3 boundary:

1. The client begins with a connection to AP B on WLC 1.
2. This creates an ANCHOR entry in WLC 1's client database.
3. As the client moves away from AP B and begins association with AP C, WLC 2 sends a mobility announcement to its peers in the mobility group looking for the WLC with information for the client MAC address.
4. WLC 1 responds to the announcement, handshakes, and ACKs.
5. The client database entry for the roaming client is copied to WLC 2, and marked as FOREIGN. PMK data (master key data from the RADIUS server) is also copied to WLC 2. This facilitates fast roaming for WPA2/802.11i clients because there is no need to undergo full re-authentication with the RADIUS server.

6. A simple key exchange is made between the client and AP, the client is added to WLC 2's database, which is similar to the anchor controller's entry, except that the client entry is marked as FOREIGN.
7. Data being sent to the WLAN client is now EoIP tunneled from the anchor WLC to the foreign WLC.
8. Data sent by the WLAN client is sent out a local interface VLAN at the foreign controller.

The 'asymmetrical' Layer 3 roaming procedure described above solves the challenge of roaming transparently across Layer 3 boundaries; however, the asymmetric flows can cause other issues in the upstream network. This is especially true if wireless client traffic is expected to flow bi-directionally through adjacent appliances or modules such as firewalls, NAC and or IPS/IDS appliances. Or, for example, if uRPF checks are enabled on next hop routed interfaces, traffic is dropped after the client roams to a different subnet. This is the reason why a symmetrical mobility tunnel capability was introduced to the Cisco Unified Wireless solution.

Layer 3 Roam—Symmetrical Mobility Tunnel

Beginning with WLC Release 4.1 and later, the WLCs can be configured to support dynamic, bi-directional tunneling between the foreign AP/WLC and the anchor WLC as shown in Figure 2-10.

Figure 2-10 Layer 3 Roam—Symmetrical Mobility Tunnel

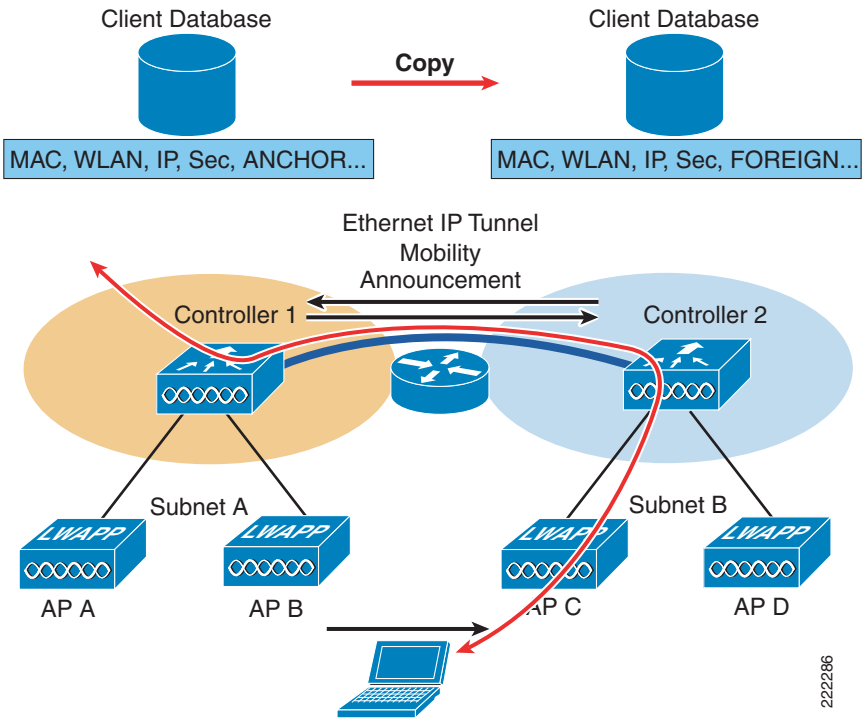

The WLC's Layer 3 mobility handoff procedure remains unchanged. However, WLC Release 4.1 makes use of existing capabilities associated with the solution's auto anchor tunneling mechanism to create a dynamic symmetrical tunnel when a client performs a Layer 3 roam.

Symmetrical tunneling is not enabled by default. It must be explicitly configured either through the controller's web configuration interface, WCS template or the controller's CLI. Symmetrical mobility tunnel operation must be enabled for each controller that is a member of a given mobility group, otherwise unpredictable behavior can occur.

Figure 2-11 and Figure 2-12 show Wireshark protocol traces of a bidirectional mobility tunnel.

Figure 2-11 Bi-directional Mobility Tunnel(1)

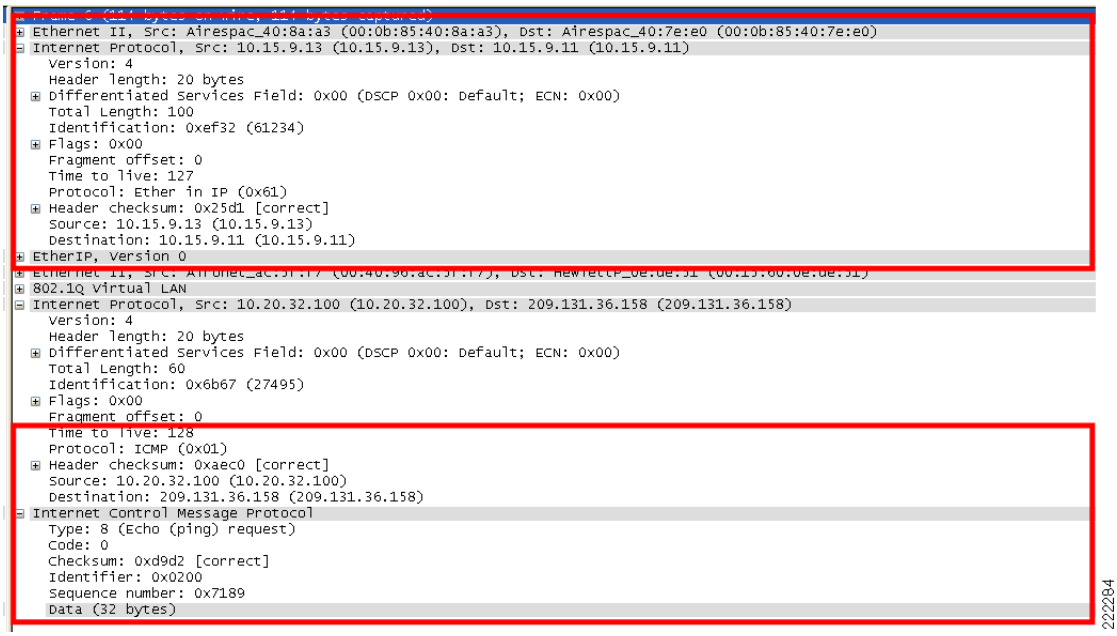

Figure 2-12 Bi-directional Mobility Tunnel(2)

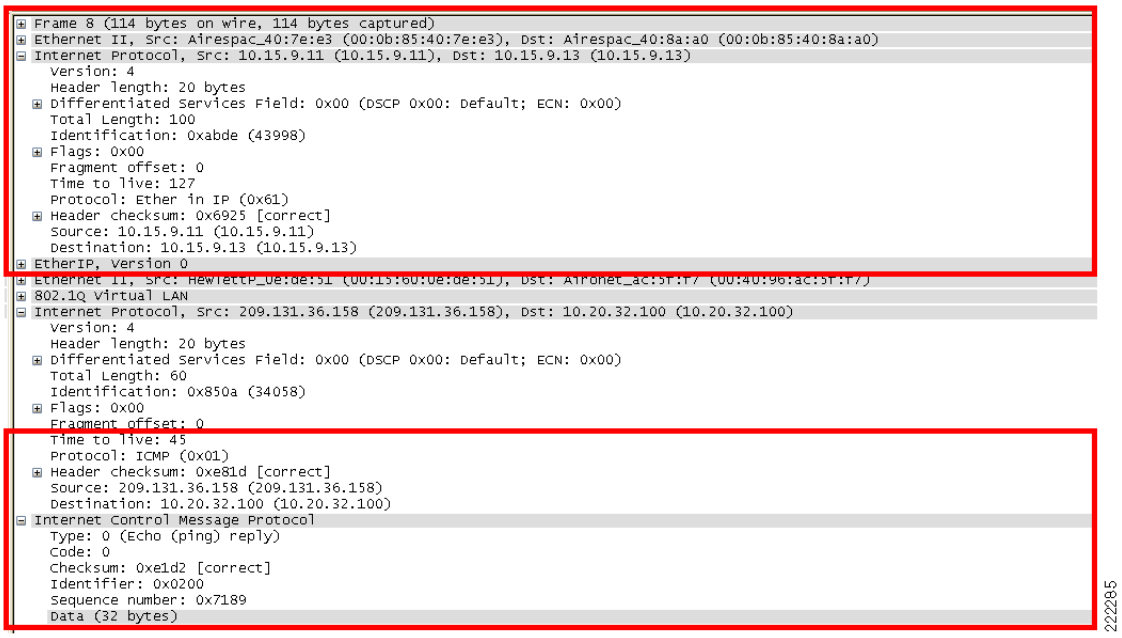

In the protocol traces above, a symmetrical mobility tunnel (EtherIP) is established between two WLCs, 10.15.9.11 (anchor) and 10.15.9.13 (foreign). In Figure 2-11, client 10.20.32.100, which has roamed to an AP on controller 10.15.9.13, is sending an ICMP ping request to Internet site 208.131.36.158 (yahoo.com). Note that the foreign controller tunnels the client's packet to the anchor controller. If the controllers were configured for asymmetrical mobility tunneling, this packet would not appear in the

trace because the foreign controller would have forwarded it locally out the VLAN interface associated with the WLAN. In Figure 2-12, the ping reply is received by the anchor controller and forwarded to the foreign controller via same the mobility tunnel, which is the same as the asymmetrical tunnel.

Important Notes About Layer 3 Roaming

Layer 3 roaming is a highly useful capability, but when deploying with the 4.1 software release, remember the following points:

- 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 foundation for facilitating Layer 3 roaming within the Unified Wireless solution is based on the concept of mobility anchors and EoIP tunnels. An 'anchor WLC' is that WLC through which a client first associates to a WLAN. The client is then assigned an address, via DHCP, that corresponds to the interface/subnet assigned to the WLAN at the anchor controller. Currently, the Unified Wireless solution does not permit clients to connect to a WLAN with a static IP address that is outside the subnet defined for the WLAN. In deployment scenarios where static client addressing is necessary, Mobile IP should be investigated as a potential solution. For more details concerning Mobile IP and its compatibility with the Cisco Unified Wireless architecture, see Chapter 12, "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 2-13 depicts basic 802.11 broadcast/multicast behavior. When client 1 in this example sends an 802.11 broadcast frame, it is unicasted to the AP. The AP then sends the frame as a broadcast out both its wireless and wired interfaces.

Figure 2-13 802.11 Broadcast/Multicast

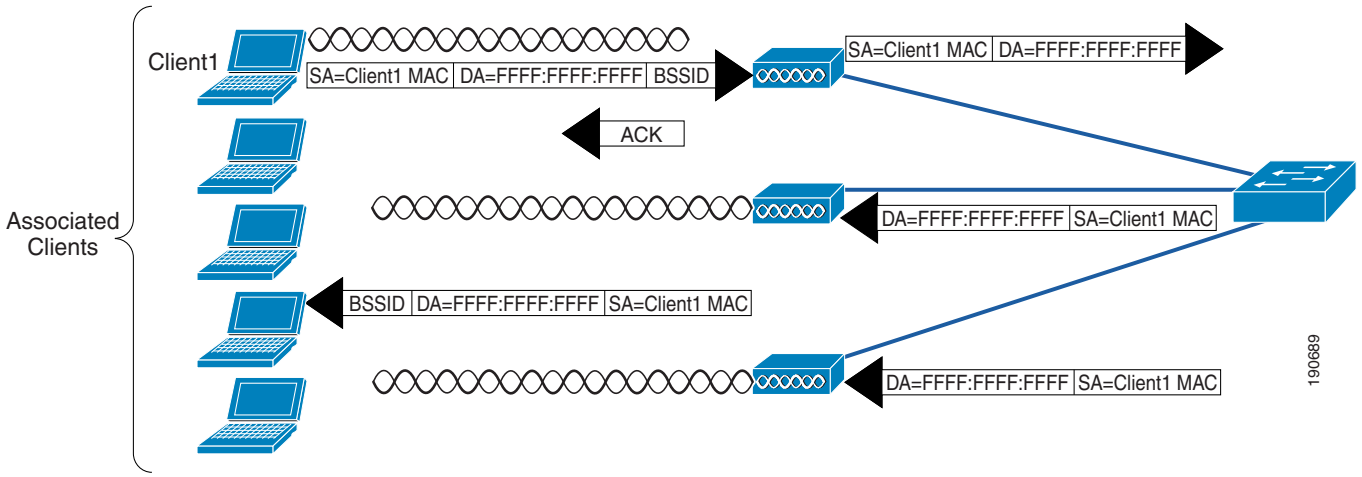

If there are other APs on the same wired VLAN as the AP as depicted in Figure 2-13, they forward the wired broadcast packet out their wireless interface.

The WLC's LWAPP split MAC method treats broadcast traffic differently, as shown in Figure 2-14. In this case, when a broadcast packet is sent by a client, the AP/Controller does not forward it back out the WLAN, and a only subset of all possible broadcast messages are forwarded out a given WLAN's wired interface at the WLC.

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

Figure 2-14 Default WLC Broadcast Behavior

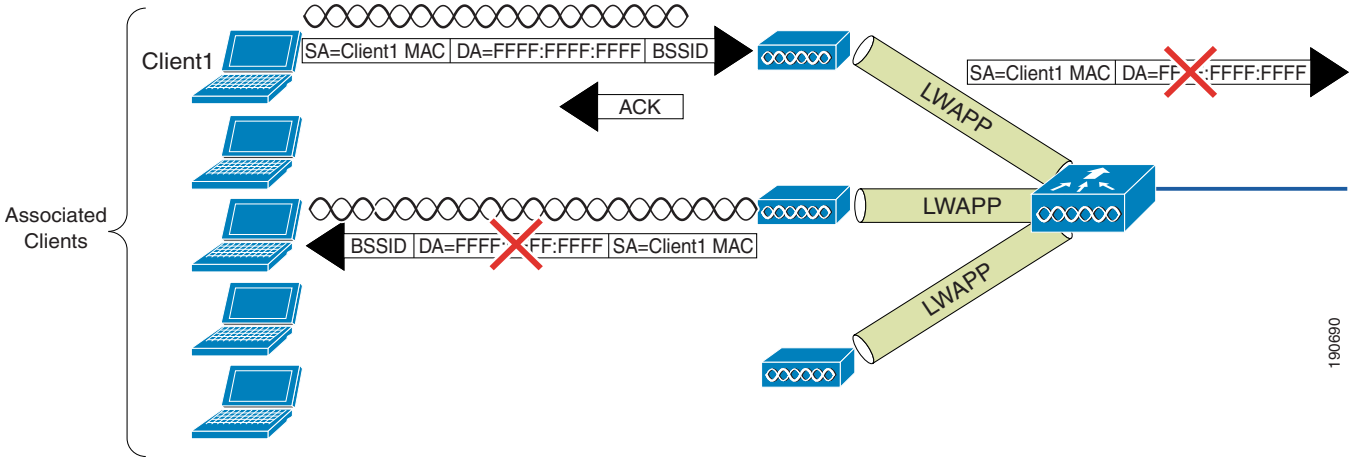

WLC Broadcast and Multicast Details

Broadcast and multicast traffic often require special treatment within a WLAN network because of the additional load placed on then WLAN as a result of this traffic having to be sent at the lowest common bitrate. This is done to ensure that all associated wireless devices are able to receive the broadcast/multicast information.

The default behavior of the WLC is to block broadcast and multicast traffic from being sent out the WLAN to other wireless client devices. The WLC can do this without impacting client operation because most IP clients do not send broadcast/multicast type traffic for any reason other than to obtain network information (DHCP) and resolve IP addresses to MAC addresses (ARP).

DHCP

The WLC acts as a DHCP relay agent for associated WLAN clients. It unicasts client DHCP requests to a locally configured or upstream DHCP server except during L3 client roaming, which will be discussed in more detail below. DHCP server definitions are configured for each dynamic interface, which in turn is associated with one or more WLANs. DHCP relay requests are forwarded via the dynamic interfaces using the source IP address of a given dynamic interface. Because the WLC knows which DHCP server to use for a given interface/WLAN, there is no need to broadcast client DHCP requests out its wired and wireless interfaces.

The method above accomplishes the following:

- It eliminates the need for DHCP requests to be broadcasted beyond the WLC.

- The WLC becomes part of the DHCP process, thereby allowing it to learn the MAC / IP address relationships of connected WLAN clients, which in turn allows the WLC to enforce DHCP policies and mitigate against IP spoofing or denial-of-service (DoS) attacks.
- It allows the WLC to relay DHCP reply messages using a virtual DHCP server IP address rather than the actual IP address of a DHCP server The aforementioned behavior is configured via the WLC's CLI, and is enabled by default. The virtual address is shared by all WLCs that comprise a mobility group. The benefit of DHCP proxy is realized during L3 client roaming or when a client roams across an AP group boundary. In these cases, the WLC will receive a client DHCP renewal request upon which it will verify the client is roaming within the mobility group and allow the client to renew (keep) it's existing IP address/subnet assignment even though the client roamed to a new subnet on the foreign WLC. See Roaming, page 2-17.

Note The virtual IP/Proxy DHCP behavior described above is required if the asymmetrical mobility tunnel method is configured (default), see Roaming section above. Otherwise, if the symmetrical tunnel method is used, WLC based DHCP proxy is not necessary because client traffic and DHCP requests are always tunneled back to the anchor controller.

ARP

Before a WLAN client can send IP packets to any other IP address, it needs to know the MAC address of the target client to forward the frame to. To accomplish this, a client will broadcast an ARP query, requesting the MAC address for the IP host that it wishes to communicate with, see Figure 2-15.

Figure 2-15 ARP Frame

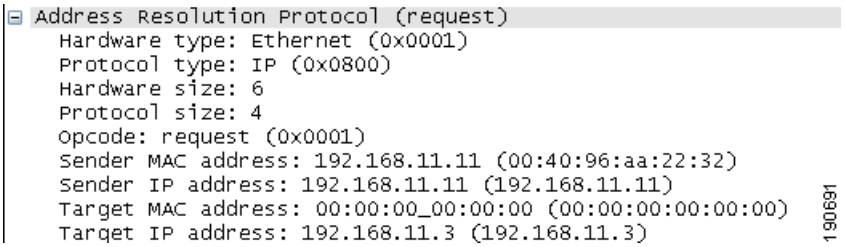

Upon seeing a wireless client ARP request, the WLC will either respond directly, acting as an ARP proxy in behalf of the other wireless clients, or it will forward the request out it's wired interface to have it resolved by another WLC. The WLC will not forward the ARP broadcast back out to the WLAN.

The default behavior of the WLC is to respond to ARP queries directly based on its local ARP cache. The WLC CLI command**: network arpunicast enable** can be used to override this behavior. In this case the WLC will unicast an ARP request directly to the target host rather than responding in behalf of the target. The target will unicast its ARP reply back to the requesting host. The purpose of this command is to avoid excessive retries by IP clients looking for a WLAN client that may have roamed from the WLAN network.

Other Broadcast and Multicast Traffic

As mentioned earlier the WLC (by default) will not forward broadcasts or multicasts toward the wireless users. If multicast forwarding is explicitly enabled as described in Chapter 6, "Cisco Unified Wireless Multicast Design," steps should be taken to minimize the multicast traffic generated on those interfaces that the WLC connects to.

All normal precautions should be taken to limit the multicast address groups explicitly supported by a WLAN. When multicast is enabled, it is global in nature, meaning it is enabled for every WLAN configured regardless if multicast is needed by that WLAN or not. The unified wireless solution is not able to distinguish between data link layer versus network layer multicast traffic neither is the WLC capable of filtering specific multicast traffic. Therefore, the following additional steps should be considered:

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

Design Considerations

Within a Cisco Unified Wireless deployment, the primary design considerations are: AP connectivity, and WLC location and connectivity. This section will briefly discuss these topics and make 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—In the distributed model, WLCs are located throughout the campus network, typically on a per building basis, managing the APs that are resident in a given building. The WLC(s) is connected to the campus network using the distribution routers within the building. In this scenario the LWAPP tunnels, between APs and the WLC typically stay within the building. Figure 2-13 shows a distributed WLC deployment.
- Each of the distributed WLCs could be configured with as a separate RF group and mobility group, so long as the WLAN coverage is not overlapping between buildings.

Figure 2-13 WLCs Distributed

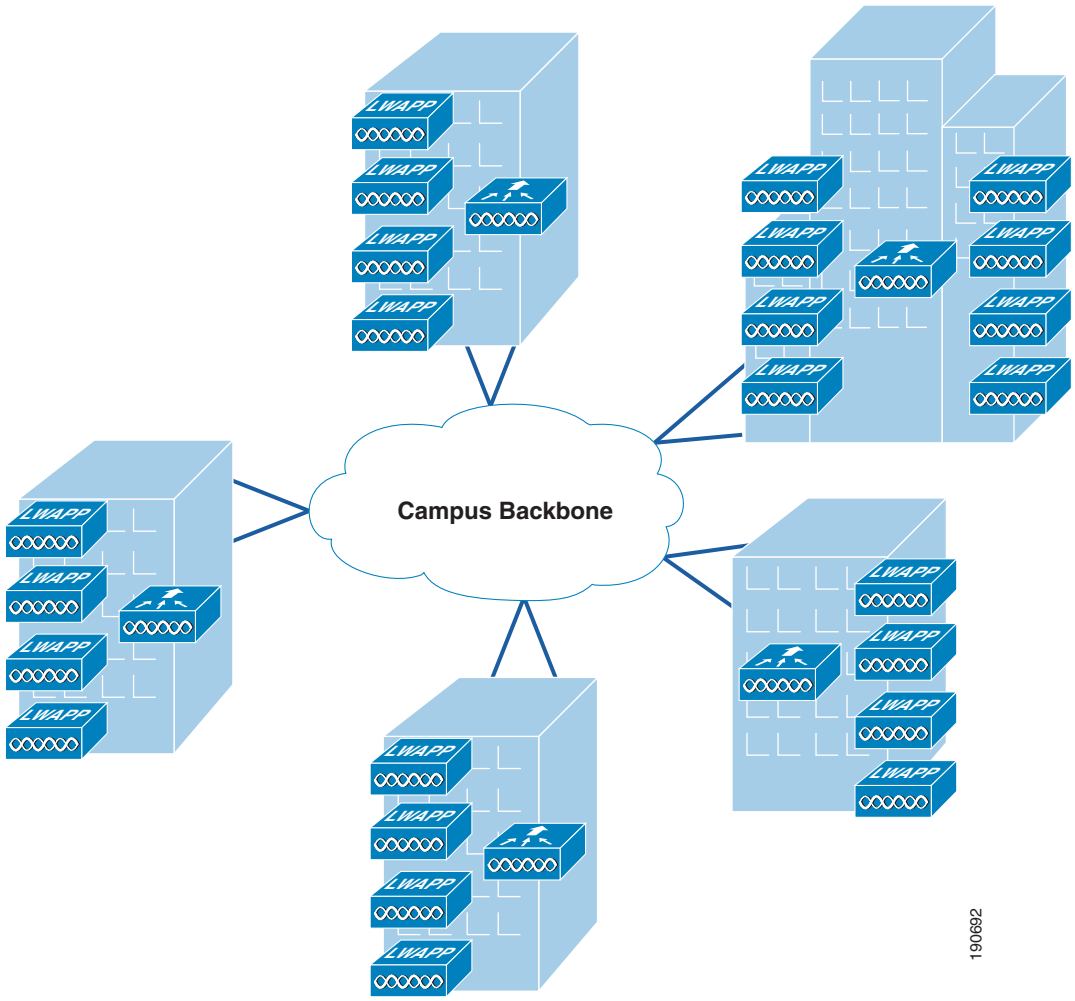

- Centralized WLC deployment—In this model, WLCs are placed at a centralized location in the enterprise network. This deployment model requires the AP/WLC LWAPP tunnels to traverse the campus backbone network. An example of a centralized WLC deployment is shown in Figure 2-14. Note in the example below that the centralized WLCs (a pair of WiSMs in Catalyst 6500's) are not shown in a specific building. A centralized WLC cluster is connected via a dedicated switch block to the campus core, which is typically located in the same building where the data center resides. The WLCs should not be connected directly to the data center's switching block because the network and security requirements of a data center are generally different then that of a WLC cluster.

Figure 2-14 WLCs Centralized

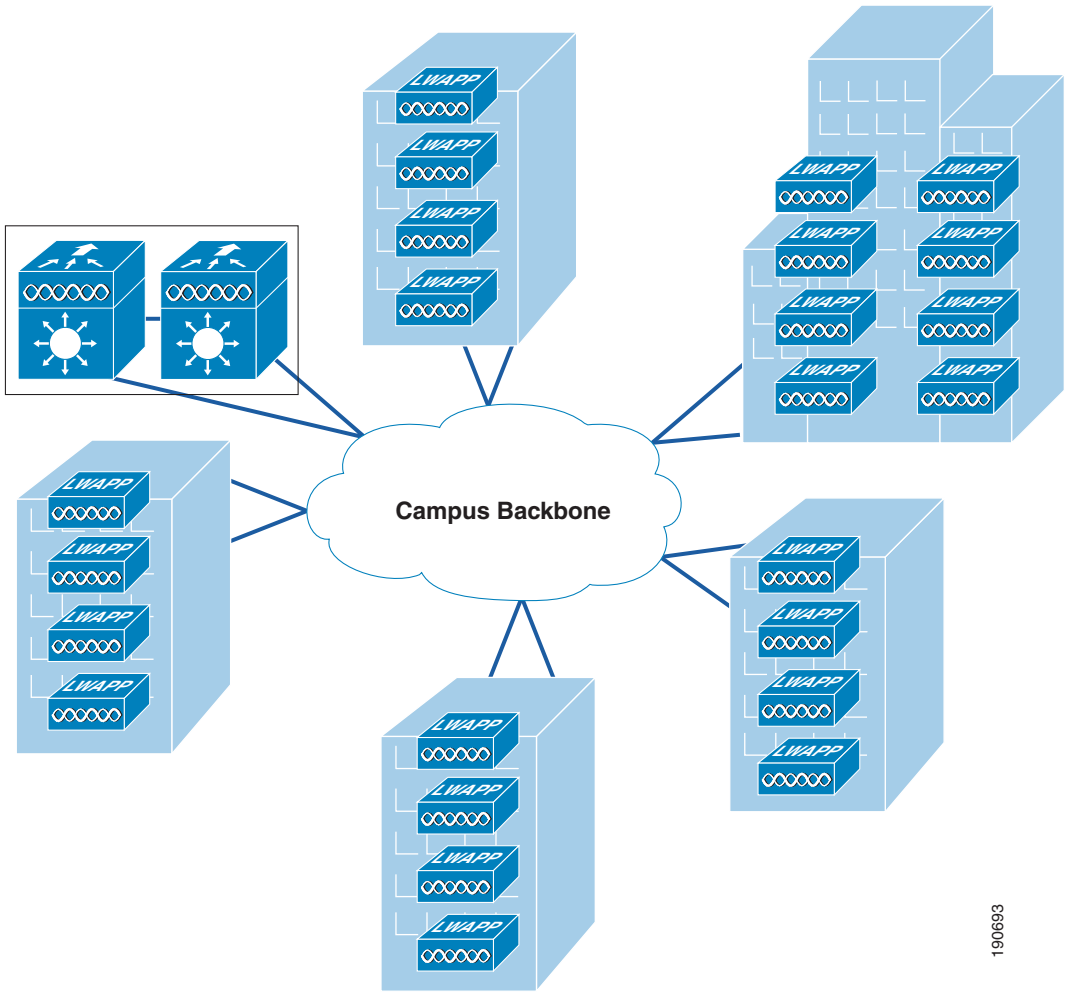

Centralizing WLCs

The general recommendation of this design guide is to deploy the WLCs at a central location within the overall campus environment. The distributed deployment model (which would require mobility groups and Layer 3 roaming) is well proven, but it is not recommended because of current shortcomings with multicast support associated with Layer 3 roaming. When these are addressed, most of the barriers preventing consideration of a distributed deployment model will be removed. Prior to Release 4.1, there were other functionality shortcomings (tunnel QoS and asymmetrical tunneling) that made distributed deployments impractical, but these have since been resolved.

The best way to address Layer 3 roaming is to avoid deployment scenarios that would otherwise necessitate it. Currently, large mobility subnets are more feasible to implement due to the scaling capabilities of the WISM module coupled with the broadcast/multicast suppression features offered by the WLC.

By centralized the WLC infrastructure, capacity management becomes simpler and more cost effective. Also, as WLANs become more mission critical, centralized deployments make it easier to create a high availability WLC topology. Centralization reduces the number of locations where capacity management and high availability issues must be dealt with.

The same principle applies when integrating the WLC with other infrastructure components. Centralized WLCs minimize the number of integration points and integration devices. For example, if a decision is made to implement an inline security device such as a NAC appliance, the centralized WLC will have one integration point, whereas a distributed solution will have '*n'* integration points, where *n* equals the number of locations where WLCs are deployed.

In summary, a centralized WLC deployment is the preferred and recommended method. When planning any centralized WLC deployment, consideration must be given to the protection of the wired network infrastructure that directly connects to the WLC. The reason is because the WLC essentially attaches an 'access' network at a location within the overall enterprise topology that would not otherwise be exposed to 'access network' and its associated vulnerabilities. Therefore, all security considerations normally associated with an access layer network device must be considered. For example, in a WiSM based deployment, features such as denial-of-service protection and traffic storm protection should be considered because of the large scale role the WISM plays in providing diverse WLAN services to large numbers of end users while at the same time being directly connected to the backplane of a core multi-layer, multi-function Catalyst 6500 switching platform.

Distributed WLC Network Connectivity

As mentioned above, distributed WLCs are typically connected to the distribution layer router within the campus network. If this is the case, Cisco does *not* recommend the WLC connect to the distribution layer via a Layer 2 link, as shown in Figure 2-16.

Figure 2-16 Layer 2 Connected WLC

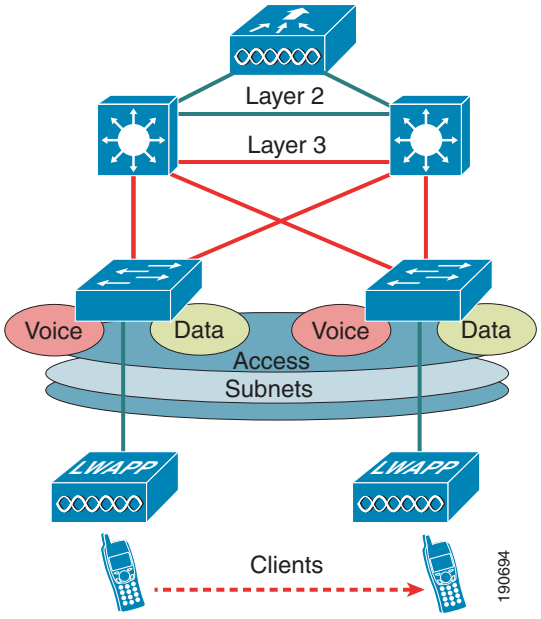

This recommendation is made for a number of reasons, including the following:

- General best practice campus design recommends Layer 3 access and distribution connectivity to provide fast convergence and simplified operation; inserting a Layer 2 connected WLC breaks this model.
- Layer 2 WLC connectivity requires the introduction of access layer features at the distribution layer, such as HSRP, and access layer security features. This may be an issue if the distribution layer 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 2-17 (in this case a 3750G), allows the WLAN-related software and configuration to be isolated to a single device and connects to the distribution layer using the same routing configuration as other the access layer routing devices.

Figure 2-17 Layer 3 Connected WLC

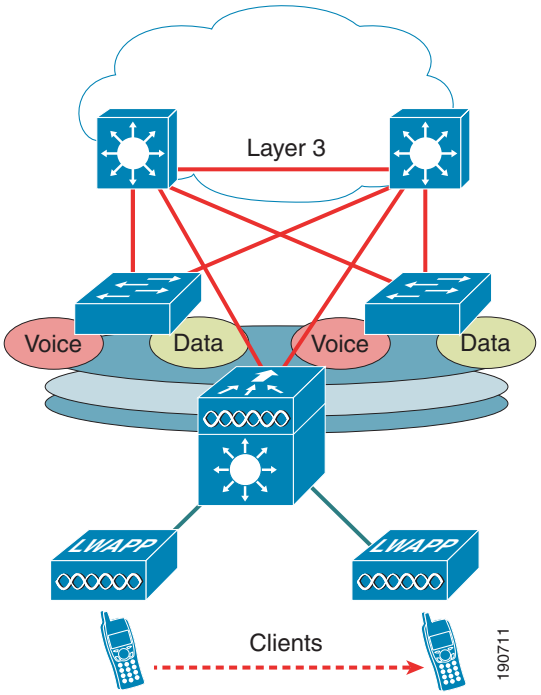

Traffic Load and Wired Network Performance

When deploying a Unified Wireless solution, questions often arise concerning:

- LWAPP traffic impact/load across the wired backbone.
- Minimum performance requirements to support a Unified Wireless deployment.
- Relative benefits of a distributed versus centralized WLC deployment in the context of traffic load on the network.

In examining the impact of the LWAPP traffic in relation to overall network traffic volume, there are three main points to consider:

- The volume of LWAPP control traffic—The volume of traffic associated with LWAPP control can vary depending on the actual state of the network. That is to say, it is usually higher during a software upgrade or WLC reboot situations. With that said, traffic studies have found that the

average load LWAPP control traffic places on the network is approximately 0.35 Kb/sec. In most campuses, this would be considered negligible, and would be of no consequence when considering a centralized deployment model over a distributed one.

- The overhead introduced by tunneling—A Layer 3 LWAPP tunnel adds 44 bytes to a typical IP packet to and from a WLAN client. Given that average packets sizes found on typical enterprises are approximately 300 bytes, this represents an overhead of approximately 15 percent. In most campuses, this overhead would be considered negligible, and again would be of no consequence when considering a centralized deployment model over a distributed one.
- Traffic engineering—Any WLAN traffic that is tunneled to a centralized WLC is then routed from the location of the WLC to its end destination in the network. Depending on the distance of the tunnel and location of the WLC, WLAN client traffic may not otherwise follow an optimal path to a given destination. In the case of a traditional access topology or distributed WLC deployment, client traffic enters the network at the edge and is optimally routed from that point based on destination address.

With that said, the longer tunnels and potentially inefficient traffic flows associated with a centralized deployment model can be partially mitigated by positioning the WLCs in that part of the network where most of the client traffic is destined (for example, a data center). Given the fact that most enterprise client traffic goes to servers in the data center and the enterprise backbone network is of low latency, any overhead associated with inefficient traffic flow would be negligible, and would be of no consequence when considering a centralized deployment model over a distributed one.

For most enterprises, the introduction of a WLAN does not result in the introduction of new applications, at least not immediately. Therefore, the addition of a Cisco Unified Wireless network alone is not likely to have a significant impact on campus backbone traffic volumes.

AP Connectivity

APs should be on different subnets from the end users. This is consistent with general best practice guidelines that specify that infrastructure management interfaces should be on a separate subnet from end users. Additionally, Cisco recommends that Catalyst Integrated Security Features (CISF) be enabled on the LWAPP AP switch ports to provide additional protection to the WLAN infrastructure. (H-REAP AP connectivity is discussed in Chapter 7, “FlexConnect.”)

DHCP is generally the recommended method for AP address assignment, because it provides a simple mechanism for providing up-to-date WLC address information for ease of deployment. A static IP address can be assigned to APs, but requires more planning and individual configuration. Only APs with console ports permit static IP address configuration.

In order to effectively offer WLAN QoS within the Cisco Unified Wireless network, QoS should also be enabled throughout the ‘wired’ network that provides connectivity between 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 different WLC discovery mechanisms for APs (discussed earlier) make initial deployment of LWAPP APs very simple. Options include:

- Staging (priming) LWAPP APs in advance using a WLC in a controlled environment
- Deploying them straight out of the box by using one of the auto discovery mechanisms (DHCP, DNS or OTP)

Although auto discovery is very useful, a network administrator will generally want to be able to control which WLC an AP will join once it is connected to the network for the first time. Subsequently then, an administrator will want to define which WLC will be the 'primary' for a given AP during normal operation in addition to configuring secondary and tertiary WLCs for backup purposes.

AP Distribution

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

APs in the same physical location should be joined to the same WLC. This makes it easier for general management, operations and maintenance, allowing staff to control the impact that various operational tasks will have on a given location, and to be able to quickly associate WLAN issues with specific WLCs, whether it be roaming within a WLC, or roaming between WLCs.

The tools that are used to control AP distribution across multiple WLCs are:

- Primary, secondary, and tertiary WLC Names—Each AP can be configured with a primary, secondary, and tertiary WLC name, which in turn determine the first three WLCs in the mobility group that the AP will prefer to join, regardless of the load variations across WLCs in the mobility group.
- Master WLC—When an AP joins a WLC for the first time in the mobility group, it is not yet configured with a preferred primary, secondary, and tertiary WLC; therefore, it will be eligible to partner with any WLC (within the mobility group) depending upon the perceived WLC load. If a WLC is configured as a master WLC, all APs without primary, secondary, and tertiary WLC definitions will join with the master WLC. This allows operations staff to easily find newly joined APs and control when they go into production by defining the primary, secondary, and tertiary WLCs name parameters.

Firmware Changes

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

A better option is to migrate the APs to their secondary WLC, upgrade their primary WLC, and then migrate the APs back to the primary (upgraded) WLC in a controlled manner.

The process will vary slightly, if a deployment has been designed for high availability, 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 WLC.

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 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.